BOOK

# The Right Policy at the Wrong Time

*Carbon Pricing and the Order of Things*

**Stephen J. Stretton**

## Contents

1 Preface . . . 2
2 A Note on Terms . . . 3
2.1 1. Why Feebates Beat Carbon Taxes . . . 4
2.2 2. What a Feebate Is . . . 6
2.3 3. The Energy Price Crisis and the Case for Feebates . . . 8
2.4 4. Scarcity Rent: Prices, Taxes, and Rationing . . . 10
2.5 5. Pricing Across Borders . . . 11
2.6 6. The Price Set by a Technology Competition . . . 13
2.7 7. Carbon Markets Under Net Zero . . . 16
2.8 8. Positive Framing . . . 18
2.9 9. Innovation in Waves . . . 20
2.10 10. Three Phases of Collective Action . . . 21
2.11 11. The Sequence . . . 22
2.12 12. Dynamic Incentives . . . 24
2.13 13. The Order of Things . . . 26
3 References . . . 27

*The views expressed are the author's own and do not represent those of any institution with which he is or has been affiliated.*

## 1 Preface

On the fourteenth of March 2025, a few hours after being sworn in, the new prime minister of Canada signed a directive abolishing the country's consumer carbon tax. It had been in place for six years. It returned its revenue to households through a quarterly rebate — the design economists most often recommend — and it was, by every study that examined it, doing what it was meant to do. It was also the single most effective line of attack the opposition had found in a decade, and the man who abolished it was a former central bank governor who had spent years arguing for carbon pricing. He kept the industrial carbon price, which charges large emitters only for emissions above a facility benchmark. Nobody campaigned against that one. Most voters did not know it existed.

This book is about why that happened, and what it means for the design of climate policy.

The conventional reading is that economists are right and politicians are weak. Carbon pricing is the cheapest way to cut emissions; the public will not accept it; therefore the public must be persuaded, or the policy must be smuggled in under another name. The reading offered here is different. Carbon pricing as usually designed has a structural political flaw — its costs are concentrated on people who can organize, its benefits diffuse and deferred — and that flaw is not a failure of communication. It is a property of the instrument, and it can be designed out.

Two claims follow.

The first is that the *marginal signal* a carbon price sends can be separated from the *net burden* it imposes. A feebate — a charge on emitters above a benchmark and a rebate to those below it, designed so the two roughly cancel — delivers the same incentive at the margin as a tax while raising little or no net revenue. It therefore creates winners and losers inside the same industry rather than a united industry opposed to the policy. Canada's industrial system is a feebate. So is the European Union's emissions trading scheme, once free allocation is taken into account. So is France's vehicle bonus-malus, which has survived four presidents. Part I sets out why these survive when taxes do not, and tests the argument against the case where every government in Europe reached for the wrong instrument at once: the gas crisis of 2022.

The second claim is about what a carbon price *is*. The conventional answer is that it is a number computed from a damage function and a discount rate. Part II argues it should be a price *discovered by competition* — specifically, the competitively bid cost of removing carbon from the air — and that the consequence is a much higher price than today's offset markets produce, together with a different accounting: one ton emitted would be offset not by one ton removed but by however much mitigation the full price can buy, which is usually several tons of forest protection. That change would redirect a great deal of money toward conservation while funding the removal technology that will eventually be needed.

Part III puts the two together as a *sequence*. Innovation proceeds in phases, and so does collective action, and the right instrument differs at each stage. Subsidy is close to optimal early, when the alternatives are small and expensive and no one is organized against them. Feebates suit the middle, when volumes have grown too large to subsidize but the incumbents remain strong enough to block a tax. A conventional carbon tax with general revenue recycling is right last, when the incumbents have shrunk and the tax can do double duty as a revenue instrument. The characteristic error of climate policy — the one the title names — is deploying a late-phase instrument in an early-phase world.

A note on what this book does not cover. Its companion, *Time and Chance: The Elements of Scale in the Green Transition*, argues that the binding constraint in most of the world is the cost of capital, and sets out the financial machinery for lowering it. The two arguments meet at one point, which each volume names so that the other can stand alone. A cheap cost of capital requires a predictable revenue stream; such streams are created by policy; and the policies that create them best are quantity instruments — auctions, power purchase agreements, contracts for difference — rather than prices. On that argument, the pricing question this book addresses is not an alternative to the financing agenda but its foundation. Readers who want the financing followed through will find it there. Readers who want to know why the instrument should be a feebate, and why the price should be set by a competition, will find it here.

## 2 A Note on Terms

**Carbon tax.** A charge per ton of carbon dioxide emitted, usually levied on fuels in proportion to their carbon content. Its revenue goes to the general budget unless a use is specified.

**Emissions trading scheme (ETS).** A cap on total emissions with tradeable permits. The quantity is fixed; the price emerges from trading and can be volatile.

**Feebate.** A charge on emitters above a benchmark and a rebate to those below it, calibrated so that the flows roughly cancel. Same marginal signal as a tax; little or no net revenue.

**Output-based pricing.** A feebate applied to industry, where the rebate is paid per unit of output — per ton of steel, per megawatt-hour — so that a producer is charged only on emissions above the benchmark for its sector.

**Benchmark.** The level of emissions per unit of output, or per household, against which a feebate charges and rebates. Where the benchmark sits determines who wins and who loses.

**Revenue recycling.** What is done with a carbon tax's proceeds: cutting other taxes, paying a dividend, funding investment.

**Offset.** A credit representing a ton of emissions avoided or removed elsewhere, purchased to compensate for a ton emitted. The market for offsets is the subject of Part II.

**Direct air capture (DAC).** Machines that remove carbon dioxide from the ambient air. Currently expensive; the price of doing it is this book's proposed benchmark for the social cost of carbon.

**Social cost of carbon.** The monetary value of the damage caused by one additional ton of carbon dioxide. Conventionally estimated from models; Part II proposes discovering it by competition instead.

**Carbon border adjustment.** A charge on imports proportional to their embedded carbon, intended to prevent a domestic carbon price from simply relocating emissions abroad.

**Feed-in tariff.** A guaranteed price per kilowatt-hour for renewable electricity, fixed for a term of years. The instrument that started the solar boom.

**Reverse auction.** A procurement in which the government fixes the quantity it wants and bidders compete on price.

**Logic of collective action.** Mancur Olson's observation that a policy is politically feasible when its benefits are concentrated and its costs diffuse, and infeasible when the reverse.

## 2.1 1. Why Feebates Beat Carbon Taxes

Economists have spent thirty years explaining that the cheapest way to cut emissions is to price them. Governments have spent thirty years discovering how hard that is to do. The gap between the two is not a failure of economics. It is a failure to take the politics of pricing as seriously as the arithmetic.

*The Record*

Start with what has happened to carbon taxes that people could see.

Australia introduced a carbon price in 2012 and repealed it in 2014, after an election fought substantially on the promise to do so. France froze its carbon tax in December 2018, cancelling a scheduled rise from about 45 euros a ton toward 86, after the fuel-price protests of the *gilets jaunes* made the increase impossible to sustain. The state of Washington put a carbon tax to its voters twice, in 2016 and 2018, and lost both times. And Canada, having built the textbook system — a rising consumer fuel charge with the revenue returned as a per-household rebate that left most families better off — abolished it in 2025 in the first act of a new government, because it had become the most effective political weapon in the country.

Now look at what survived. Canada kept its industrial carbon price, which charges large emitters only for emissions above a benchmark for their sector and is, in the language of this book, a feebate. The European Union's emissions trading scheme has run since 2005 and now prices carbon at levels France's tax never reached — but for most of its life, and still for exposed industry, the bulk of allowances were allocated free, so that the effective design is a charge on emissions above an allocation with the allocation functioning as a rebate. France's bonus-malus on vehicle purchases, introduced in 2008, charges buyers of high-emission cars and pays buyers of low-emission ones, and has been extended and tightened by governments of both left and right. Nobody has marched against it.

The pattern is not an accident, and it is not about the size of the price. The EU price has been higher than any of the repealed taxes. It is about the *shape* of the instrument, and the shape is what this chapter explains.

*Four Claims*

**First, a carbon tax is dominated by a feebate on political-economy grounds.** A feebate charges emitters above a benchmark and pays those below it, with the two flows designed to roughly cancel. It delivers the same marginal signal as a tax — a ton avoided is worth the same either way — while raising little or no net revenue. That sounds like a weakness. It is the source of its strength. A tax creates a visible class of losers who can organize; a feebate creates winners and losers inside the same industry, and the winners have every reason to defend the scheme. In Canada, the steel mill that beat its benchmark was a beneficiary of the industrial system and had no interest in its abolition. Every household was a payer of the consumer tax, and the rebate — paid quarterly, into a bank account, unlabelled — did not feel like a return of anything.

**Second, feebates are better from the international perspective.** A country that taxes carbon unilaterally hands a cost advantage to its competitors and invites carbon leakage. A country that runs a feebate does not: because the average firm pays close to nothing, the sector-wide cost burden is small, while the incentive to be cleaner than the average is undiminished. This makes the instrument far easier to adopt without waiting for everyone else — which matters when the alternative is waiting for a global agreement that has not arrived in thirty years. Chapter 5 takes this further.

**Third, procurement auctions coordinate better across time than a given price.** A carbon price tells an investor what emitting costs today. It does not tell them what the price will be in 2040, which is the number that matters when the asset lasts twenty-five years. A reverse auction for a fixed quantity of renewable capacity does the opposite: it fixes the quantity and lets the market discover the price. Where the objective is to build a certain amount of something by a certain date, quantity instruments coordinate the investment path more reliably than price instruments do.

**Fourth, and consequently, renewables financed through procurement auctions carry a lower cost of capital.** An auctioned power purchase agreement converts a merchant project exposed to volatile wholesale prices into a contracted project with a predictable revenue stream. Lenders price that difference. This is the thread the companion volume follows; here it matters because it means the instrument that is best on political grounds is also, for capital-intensive technologies, best on financial ones.

*The Logic of Collective Action*

Behind all four claims sits a single principle of political economy, which is worth stating explicitly because the rest of the book leans on it. Following Mancur Olson, call it the logic of collective action.

A necessary condition for a policy to be politically feasible is that its benefits be concentrated and its costs diffuse, rather than the reverse. Where costs fall on a group that can organize, incentivize and coerce its own members — a trade association, a utility, a regional industry, a farmers' union — that group will lobby to block the policy. Where benefits fall on such a group, it will lobby for it. Costs or benefits spread thinly across millions of households are individually too small to be worth acting on, and the cost of organizing those millions is prohibitive.

This is uncomfortable for climate policy, because a conventional carbon tax has exactly the wrong shape. Its costs are concentrated on identifiable emitters in power, industry, transport and buildings — and, through fuel prices, on every driver and every household with a gas boiler, who may not organize but who vote. Its benefits are diffuse: a slightly better climate for everybody, decades hence, and some revenue that flows into the general budget where nobody can point to it. Green subsidies, by contrast, are naturally compliant with the logic — concentrated beneficiaries, diffuse funding — which is one reason the world has rather more of them than it has carbon taxes, and why the largest climate legislation in American history, the Inflation Reduction Act of 2022, contained no carbon price at all.

The usual reply is that the revenue can be recycled. But the standard recycling proposals are macroeconomic in flavor — cuts in labor taxes, reductions in corporate rates, a uniform per-capita dividend — and are insufficiently targeted to solve the problem. Canada tried the dividend and it did not work, for a reason worth understanding. A per-capita rebate leaves the heavy emitters comprehensively out of pocket, and it leaves every household seeing a charge at the pump each week and a deposit in the bank each quarter, with no visible connection between the two. Losers are reliably more motivated to resist than winners are to support, and a benefit that does not feel like one does not recruit defenders.

*A Rebate Design That Passes the Test*

Here is a design that does pass. Rank sectors by emissions intensity per unit of gross value added. For those above the average — power first, then heavy industry — do not impose a tax at all. Impose a feebate, with the neutral point set as close as possible to recent historical usage. Do the same for domestic and commercial buildings: not a charge on energy, but an incentive calibrated against what each user has recently consumed.

A system of incentives benchmarked to recent history, running at a slightly negative net revenue, has three properties worth having.

It gives everybody a large marginal incentive to cut emissions, because the fee and the rebate are both live at the margin even when they cancel on average.

It creates no concentrated losers, because the benchmark is each participant's own recent behavior rather than an abstract standard that happens to favor whoever is already clean.

And — this is the subtle one — it is slightly unfair, which is a virtue. Grandfathering to historical usage hands different participants different endowments of rebate for no defensible reason. That arbitrariness creates its own pressure, over time, to move toward a fairer common benchmark. The sequence matters: the climate incentive goes in immediately, and the substantial revenue, along with the opportunity for genuine tax reform, arrives later. This is close to the reverse of what is normally expected of a carbon tax, and that is the point.

## 2.2 2. What a Feebate Is

The previous chapter used the word as if its meaning were obvious. This one takes it apart, because the design choices inside a feebate decide whether it works, and because the instrument already exists in several forms that most people have never noticed.

*The Mechanism*

A feebate has three parts: a base on which the charge is levied, a benchmark that divides payers from recipients, and a rebate rule that returns the money.

The *base* is what is being priced — carbon dioxide emitted, fuel consumed, or some proxy such as a vehicle's rated emissions per kilometre. The *benchmark* is the level above which a participant pays and below which they receive. And the *rebate rule* determines how the money collected from those above the line is returned to those below it: per unit of output, per capita, per vehicle, or in proportion to historical usage.

Two things follow from this construction that a carbon tax does not share. The total flow is small relative to the total emissions priced, because most participants sit near the benchmark and pay or receive little. And the net effect on any sector's cost base is close to zero, so the instrument does not raise the sector's prices or push its production abroad — while every participant still faces the full price at the margin, because moving one ton across the benchmark changes what they pay by the full amount.

That second property is the one economists sometimes miss. A feebate is not a weaker carbon price. At the margin it is exactly the same carbon price. What it removes is the *average* burden, and it is the average burden, not the marginal signal, that generates the political opposition.

*Vehicles: The Instrument That Already Exists*

France introduced its bonus-malus in January 2008. Buyers of cars emitting more than a threshold of carbon dioxide per kilometre paid a malus, escalating with emissions; buyers of cars below a lower threshold received a bonus. The thresholds have been tightened almost every year since, and the malus at the top of the scale has risen from a few hundred euros to tens of thousands. The effect on the composition of sales was immediate and large — the share of low-emission vehicles rose sharply in the first year — and the scheme has survived four presidents and every change of government.

Sweden adopted a bonus-malus in 2018 on the same principle. Norway's vehicle taxation, though not labelled a feebate, has the same structure and produced the highest electric vehicle share in the world. And in the United States, the corporate average fuel economy standard functions as a feebate among manufacturers, with credits tradeable between them

— a design that created the market for compliance credits on which Tesla's early profitability depended.

None of these has generated a *gilets jaunes*. The reason is the one Chapter 1 gave: the buyer of a large car pays, the buyer of a small one is paid, and neither can point to the other as the beneficiary of a transfer from themselves to the state.

*Industry: Output-Based Pricing*

Applied to industry the rebate rule is per unit of output, and the instrument is usually called output-based pricing. A steel mill pays for its emissions above a benchmark for steel, per ton of steel it produces. A cement plant does the same against a benchmark for cement. A producer that is cleaner than its sector's benchmark receives credits; one that is dirtier pays.

Canada's industrial system, which survived the 2025 repeal untouched, works this way, and so — in substance — does the European scheme for exposed sectors. In the EU, the free allocation of allowances to industries at risk of carbon leakage is set by a product benchmark, so that a plant emitting at the benchmark receives allowances covering its emissions and pays nothing net, while one above the benchmark must buy the difference at the market price. Since 2013 the benchmark has been set at the average of the best-performing ten percent of installations, which means most plants are payers at the margin and the cleanest are net sellers. Europe has run a feebate on its heavy industry for over a decade under another name, and it is the part of the scheme that has never been seriously challenged.

The design choice that matters most is the benchmark. Set it at the sector average and the sector as a whole pays nothing net, which is politically easiest and gives the largest incentive to the dirtiest half. Set it at best practice and the sector pays net, which raises revenue and raises resistance. Chapter 1's proposal — benchmark to recent historical usage, then tighten — is the version that gets the incentive in first and argues about the revenue later.

*Households: Benchmarks and Fairness*

For households the base is energy consumption and the rebate rule is the contested part. A pure per-capita rebate is fair but blunt: it rewards the small household in a well-insulated flat and penalizes the large family in a leaky house they did not build. A rebate keyed entirely to historical usage is the reverse: it locks in the profligate.

The workable design is a blend — part per capita, part historical — that moves toward the per-capita share over time. The blend does two things. It avoids punishing people for a housing stock they did not choose, in the first years when they cannot change it. And it sets a known trajectory toward a fairer benchmark, so that the incentive to insulate, to replace the boiler, and to move is present from the start even though the bill does not bite at once.

The instrument can also be pointed at the durable decision rather than the flow. A feebate on heating equipment — a charge on the new gas boiler, a bonus on the heat pump — settles a household's next twenty years of gas demand in a single purchase, and it is the design most likely to work in a crisis, for the reasons the next chapter sets out.

*What a Feebate Cannot Do*

It should be said what the instrument does not do, because its advocates sometimes forget.

A feebate raises little revenue by design, so it cannot fund the things a carbon tax's revenue might have funded. That is a real cost where a government needs the money. It is the reason Chapter 11 places the feebate in the *middle* of the sequence, and the revenue-raising tax at the end.

A feebate's benchmark is a political object. Where it sits, who sets it, and how fast it tightens are decisions that the affected industry will lobby over, and the history of EU free allocation

— repeatedly loosened under pressure — shows that a benchmark can be captured. The defense against capture is a published tightening schedule set in statute, which is the same commitment-device logic that the companion volume applies to contracts for difference.

And a feebate does nothing for the sectors where the base cannot be measured. It works where emissions per unit of output can be counted — power, steel, cement, vehicles, buildings. It does not work for agriculture, land use, or the diffuse emissions that no meter records, and those are the sectors where Part II's approach to carbon markets has to take over.

### 2.3 3. The Energy Price Crisis and the Case for Feebates

The argument of the previous chapters can look abstract until a real crisis arrives and governments have to choose an instrument in a hurry. Europe's gas crisis of 2022 was such a moment, and it is worth walking through, because almost every government reached first for the one tool the logic of collective action would have warned them against — and then, under pressure, discovered the alternative.

*The Shape of the Problem*

The immediate trigger was the war in Ukraine and Europe's dependence on Russian gas. That dependence was not uniform. It differed sharply by member state, and it differed by end use: gas went into power generation, into industrial process heat, and into household heating, and each had a different substitution possibility on a different timescale. Any serious policy response had to disaggregate.

Europe also faced what might be called the Gazprom double bind. Sign long-term contracts, and you lock yourself into Russian gas for a decade. Buy on the spot market instead, and you expose yourself to the dominant supplier's market power — the supplier can throttle flow, drive up the hub price, and collect more revenue on less gas. Neither door was safe. And the composition of Europe's purchases had shifted: more than 80 percent of Gazprom's delivered volumes had a direct link to trading hubs, with only around 13 percent still linked to oil prices, so the spot exposure was larger than the long-term-contract framing suggested.

The price moves were extraordinary. Natural gas reached record highs. The World Bank's energy price index rose 34 percent between January and March 2022, on top of a 50 percent increase over the two preceding years. Crude oil rose roughly 350 percent in nominal terms from its pandemic low in April 2020 to April 2022. Modelling of the baseline upward revision to oil prices, if driven by supply shocks, suggested a reduction in global output of about 0.3 percent after two years.

*Why Subsidy Makes It Worse*

The instinctive response — hold the consumer price down — is the wrong one, and not only for the usual fiscal reasons. In a shortage, price is doing necessary work. Suppressing it makes the scarcity worse.

There are four distinct reasons to resist the subsidy reflex.

It slows the transition away from the very dependence that caused the crisis. Cheap gas is an argument for more gas boilers.

It pushes governments toward high-carbon substitutes. Several countries, France included, moved to reopen coal plants to cover the coming winter's power shortfall — trading an energy problem for a climate one.

It consumes fiscal space that is not there. The conversation about constrained public finance usually concerns developing countries, but after the pandemic several European states, France and Italy among them, were carrying very high debt. Italian debt came under

visible market tension around Mario Draghi's resignation. Broad energy subsidies are an expensive way to spend borrowed money.

And it fails to change long-run household behavior, which is precisely what a durable answer to gas dependence requires. A subsidy tells households the problem is temporary. Insulation and heat pumps are decisions taken on the assumption that it is not.

France illustrates both the temptation and the partial escape: the government moved to grant allowances targeted at low-income households as part of a broader cost-of-living package, rather than holding down the price for everyone.

There is also a definitional trap worth flagging. A tax on energy shows up in the inflation index. A revenue-neutral incentive largely does not, since the fee and the rebate offset. Two policies with the same marginal signal therefore look very different in the headline number that finance ministers are judged on. That is an argument for the feebate, not against it.

*Direct Assistance, Properly Targeted*

If prices must be allowed to do their work, the distributional problem has to be handled directly rather than through the price. The tools are targeted transfers to vulnerable households and output-based rebating for exposed industry. Both preserve the marginal incentive to economize — the household still faces the true price of the next unit of gas, the factory still faces the true price of the next unit of process heat — while protecting the level of income or competitiveness. This is the structural insight of the feebate: separate the incentive from the burden.

Alongside this sat quantity and efficiency measures that were genuinely useful. The EU's "Save gas for a safe winter" plan set demand reduction targets. The IEA's ten-point plan emphasized heat pumps, thermostat settings and efficiency retrofits. National measures followed in France, Germany and Italy. Lessons from the shocks of the 1970s point the same way: efficiency standards and renewable mandates proved highly effective, while price controls and the promotion of coal for power produced distortions that took decades to unwind.

*The Feebate Answer*

The instrument that fits this problem is the one from Chapter 1: an incentive that does not raise net revenue.

Consider the power sector first. A charge on gas used for generation gives every generator a reason to move away from gas — or, if they cannot, to run their plant more efficiently. Refund the proceeds per kilowatt-hour generated. Consumers' bills are relieved directly, the sector as a whole is roughly whole, and the relative price of gas-fired against non-gas-fired generation has moved decisively.

In heavy industry, the same structure works with a different rebate base: charge on gas use, refund per physical unit of output. A steel plant that makes a ton of steel with less gas than its competitors comes out ahead. A plant that does not, pays. Neither faces a sector-wide levy that would simply relocate production outside Europe.

For households, a benchmark feebate can be built from a combination of historical usage and a per-capita allocation — the first to avoid punishing people for a housing stock they did not choose, the second to stop the benchmark from permanently rewarding the profligate. And feebates can be pointed at the durable decision: disincentivize new gas boilers, reward heat pumps, and the household's next twenty years of gas demand is settled by one purchase.

A feebate can be levied directly on the CO2 content of a product, with the proceeds rebated per unit produced. In the European context the fee base could be gas consumption alone,

all fossil fuels weighted by carbon content, or all fossil fuels weighted by full greenhouse-gas content. The last is the right choice. If the objective is both to meet climate goals and to curtail dependence on Russian gas, it is past time for Europe to run a greenhouse gas policy rather than a carbon dioxide policy.

*What Europe Actually Did*

The record is instructive, because Europe ended up somewhere close to the feebate without calling it one.

In October 2022 the Council adopted an emergency regulation that capped the revenues of "inframarginal" electricity generators — those, such as renewables and nuclear, whose costs had not risen but whose prices had — at 180 euros per megawatt-hour, with the proceeds above the cap returned to consumers. It imposed a "solidarity contribution" of at least 33 percent on the excess profits of fossil fuel companies. And it set a binding target to cut peak electricity demand by 5 percent. That package charged those who were collecting the scarcity rent and returned the money to those paying it, while leaving the marginal price intact. It was a feebate on the power sector, arrived at under duress.

The demand response was larger than most forecasts. EU gas consumption fell by around 18 percent between August 2022 and March 2023 against the five-year average, and the reduction persisted into the following winter. Prices did their work. What the emergency package added was a way of letting them do it without the distribution of the rent becoming politically intolerable.

*Burden Sharing Across the Union*

Any Europe-wide action immediately raises the question of who bears what. Here too the feebate has an intergovernmental form. Member states could be allocated a quantity of gas at market prices, with imports of Russian gas taxed across the Union and states consuming beyond their quota paying more — a structure closely analogous to the double quota systems long used in agricultural production control.

A common financial facility for heat pumps would work on the same principle at the level of investment rather than consumption: countries that act more aggressively draw more from the pot. And there is no reason to stop at the Union's borders. A burden-sharing structure designed to include non-EU participants becomes a prototype for exactly the kind of international arrangement that Chapter 5 proposes. A crisis instrument, if designed properly, is also a climate instrument.

## 2.4 4. Scarcity Rent: Prices, Taxes, and Rationing

The gas crisis raised a question that pricing theory answers unusually cleanly: when a commodity's supply cannot adjust and its demand will not, who should capture the resulting windfall?

*Why Prices Went Where They Went*

How does one explain prices eight times higher in one year than the last? Supply was constricted and could not adjust quickly. LNG terminals were being built, but terminals take years. With supply effectively fixed in the short run, the only remaining equilibrating mechanism is demand destruction — and if demand is highly inelastic, as heating and industrial process demand are over a single winter, prices must rise to extraordinary levels before enough demand goes away.

This is not a market failure. It is a market doing the only thing available to it. But it has a striking distributional consequence: the enormous gap between the cost of producing the gas and the price paid for it — the scarcity rent — accrues to whoever holds the gas. In

Europe's case, substantially to Russia, and to every generator whose price was set by gas but whose fuel was not.

*The Ramsey Argument*

An inelastic commodity is a tempting object of taxation, and the observation is old. Frank Ramsey's rule holds that commodities with the lowest elasticity of supply and demand should be taxed the most, since the deadweight loss of a tax depends on how much the tax changes behavior, and by assumption here it changes it very little.

The Ramsey logic is doubly compelling in a scarcity episode. Prices have to climb to astronomical levels before demand falls enough to clear the market. That money is going somewhere regardless. If a tax were used to limit demand instead, the same demand reduction would occur, but the funds would stay inside the importing bloc and could be refunded to domestic industry and households — which is to say, the competitiveness problem and the cost-of-living problem could both be addressed with the rent that is currently being exported.

Put plainly: the choice is not between high prices and low prices. Supply is fixed; the price that clears the market is what it is. The choice is about who collects the difference.

*Why This Is Harder Than It Sounds*

Two obstacles stand in the way, one economic and one political.

The economic one is market geography. Coordinated demand-side taxation of oil is genuinely difficult, because oil trades globally and a bloc that taxes it simply reallocates supply toward those who do not. Gas is different. It is a regional market, linked only partially and expensively through LNG, which means a large importing bloc has real monopsony power. And the European Union is the rare importing bloc that pools sovereignty and can coordinate policy across its members. If the Ramsey argument is ever going to be applied in practice, gas in Europe is the case — and the 2022 revenue cap, which captured the rent of generators whose price was set by gas, was a first application of exactly this logic.

The political obstacle is the one this book keeps returning to. Reducing demand through flexible taxes or a cap-and-trade system on individual fuels is analytically attractive and politically very hard, for exactly the reasons set out in Chapter 1: the costs are concentrated and visible, the benefits diffuse and delayed.

Which brings the argument back to the same place. It makes obvious sense for Europe to capture more of the scarcity rent associated with restricted supply rather than pay extortionate prices to the dominant supplier. Doing it through taxation alone is, at present, close to politically impossible. Doing it through a feebate — same marginal signal, no net revenue, no concentrated losers — is not, and in the event Europe did something very like it.

*Author's notes for this chapter: the oil price cap, or consumer cartel, idea is easiest to implement through prices, taxes or a permit cap, and easiest of all with gas. A fee on gas imported by pipeline could also be justified on greenhouse-gas emissions grounds.*

## 2.5 5. Pricing Across Borders

Every argument so far has been about a single jurisdiction. This chapter asks what the feebate looks like between countries — and whether the same logic that makes it survive domestically can make it work where carbon pricing has failed most completely, which is internationally.

*The International Problem*

Thirty years of negotiation have not produced a global carbon price, and the reason is the logic of collective action applied to sovereign states. A global tax would concentrate costs

on the largest emitters, who are also the largest economies, and diffuse the benefits across everyone and the future. Suppose the required incentive is 100 dollars a ton. A putative international tax at that level would raise nearly 5 trillion dollars and require transfers, before any use of the revenue, of roughly a trillion dollars a year from China alone. No amount of mechanism design makes that signable.

Direct global subsidies fail the other way: they need a pot of money that does not exist, and they need a baseline against which to pay, which invites every country to inflate its baseline and delay its action until it is paid to act. The additionality problem that Part II examines in offset markets is the same problem at the level of states.

*The Feebate Between Countries*

The country-level version of the feebate is a fee-and-rebate on national carbon intensity, with the overall pot netting to zero. Countries that improve their intensity faster than the benchmark receive; countries that improve more slowly pay. Better performers help worse ones, and the transfers required are a fraction of what a tax would move, because most countries sit near the benchmark.

The double-quota structure sketched in Chapter 3 is one form of this. Each participant is allocated a quantity at a reference price; use beyond the quota is charged; the charges fund the rebates to those below. It is familiar from agricultural policy, where it has been used to manage production for decades, and its virtue is that it does not require anyone to agree the level of a tax — only the allocation of quotas, which is a negotiation countries know how to have.

The difficulty is the negative transfer. A country that must *pay* into the pot is being taxed by its peers, and sovereign states do not accept that. This is the point at which the companion volume's Climate Incentives Trust becomes relevant: if the rebates are held in accounts that grow with performance and shrink with its reverse, the negative transfer becomes a deduction from money a country has not yet received rather than a payment from money it has. Deducting from an account is politically achievable in a way that taxing a sovereign is not. The international feebate, in other words, needs a trust to hold the pot.

*Border Adjustment as a Feebate*

The European Union's carbon border adjustment mechanism, in transitional operation since 2023 and charging from 2026, is usually described as a tariff. It is more usefully described as the international extension of the EU's domestic feebate.

The logic is this. EU industry pays the ETS price on emissions above its free allocation. Imports, until the mechanism existed, paid nothing, which gave foreign producers the advantage that Chapter 1 warned a unilateral price creates. The border adjustment charges imports the same price on the same basis — embedded emissions above the EU benchmark — so that a foreign steel mill faces the same marginal signal as a European one. As free allocation is phased out, the border charge phases in.

Two features make this a feebate rather than a tariff. The charge is on emissions above a benchmark, not on the product; a clean foreign producer pays nothing. And the benchmark is the same one European producers face, so the mechanism does not protect European industry as such — it protects clean industry, wherever it is. A foreign producer that beats the benchmark is better placed than a European one that does not.

The limits are real. The mechanism covers only the most exposed sectors — iron and steel, cement, aluminium, fertilisers, electricity, hydrogen — and can therefore only play a part. It requires measurement of embedded emissions that many exporters cannot yet provide. And it has been received in much of the developing world as protectionism dressed up,

which is a political cost the EU has chosen to bear. But the underlying structure is the right one, and it is the structure of this book: a marginal signal that follows the product across borders, with no net burden on those who meet the benchmark.

*Clubs*
The alternative route to international pricing is the climate club: a group of countries that agree a common minimum price among themselves and apply a border adjustment to non-members. The idea is that the club grows because the border charge makes membership cheaper than exclusion.

The club logic and the feebate logic are the same. Members face a common benchmark; non-members are charged at the border on emissions above it; a non-member that meets the benchmark loses nothing by staying out and gains nothing by joining, so the club's growth depends on the benchmark being one that most countries will eventually want to meet anyway. Set the benchmark at a level that rewards what clean producers are already doing, and the club is a coalition of the willing. Set it at a level that punishes, and it is a trade war.

This is why the domestic argument of Chapter 1 and the international argument of this chapter are the same argument. The instrument that creates winners inside an industry rather than a united industry against it, applied across borders, creates member states inside a club rather than a united world against it. The feebate is the form of carbon pricing that can be adopted one jurisdiction at a time and still add up. # Part II — What Carbon Should Cost

## 2.6 6. The Price Set by a Technology Competition

Carbon markets rest on an accounting convention so intuitive that it is rarely examined: one ton emitted here can be cancelled by one ton removed there. This chapter argues that the convention is wrong, that it is holding the price of carbon far below where it should be, and that replacing it would unlock a great deal more mitigation from the same money.

*The State of the Market*
Before the proposal, the record, because the argument depends on it.

In January 2023 an investigation by *The Guardian*, *Die Zeit* and SourceMaterial examined the rainforest offsets certified by Verra, the largest standard in the voluntary market, and concluded that more than 90 percent of them represented no real emissions reduction — the forests would not have been cut down anyway, or the baselines against which "avoided deforestation" was measured had been inflated. The findings were contested in their details and confirmed in their thrust by later academic work. The voluntary carbon market, which had reached about 2 billion dollars in 2021, fell below 1 billion by 2023, and the price of a nature-based credit settled in the range of 5 to 10 dollars a ton. In some categories it fell below 2.

Meanwhile the cost of actually removing a ton of carbon dioxide from the air — by direct air capture — ran at 600 to 1,000 dollars a ton at the first commercial plants. The United States now pays a tax credit of 180 dollars per ton for direct air capture. A coalition of companies led by Stripe, Alphabet and Shopify committed over a billion dollars through the Frontier advance market commitment to buy permanent removals at whatever price the technology could deliver.

So there are two prices for a ton of carbon in the world. One is what it costs to claim you have offset a ton, and it is a few dollars. The other is what it costs to actually remove one, and it is a few hundred. The argument of this chapter is that the first price is fictitious and the second is the one that should govern.

*The Proposal in Brief*

The argument is for a social cost of carbon defined not by a committee or a discount-rate assumption but by a technology competition — by the observed, competitively bid cost of removing carbon from the air inorganically. A ton of emissions would then be certified as offset only if the buyer pays that full global social cost.

Crucially, the money does not have to be spent on air capture. A second market allocates the funds to whatever delivers the most mitigation per dollar. So one ton of emissions, offset at the air-capture price, might purchase five tons of avoided deforestation. Net-zero commitments in the developed world, by companies and by countries, would pay for the conservation in the developing world needed to preserve carbon stocks that are genuinely irreplaceable — while simultaneously funding the technology that will eventually be needed for drawdown at scale.

Achieving this requires three institutional pieces:

1. Global auctions for inorganic carbon removal, which set the price.
2. A process for determining the cost-effective use of the funds: a global assessment of organic carbon protection and removal cost-effectiveness.
3. A consistent standards body that certifies offsets as gold standard only if the buyer has paid the price established in the auction.

*The Context*

Compliance carbon markets and the voluntary offset market are a promising source of funds for mitigation, and particularly for preserving and enhancing forests. The basic transaction is sound: corporations or countries pay others to make reductions they cannot economically make themselves. That is what a market is for.

The problems are real too, and they are widely recognized. The price is too low to give a good incentive, and the effectiveness of the resulting mitigation is too uncertain to justify the claims made for it. These are usually treated as two separate scandals — a price scandal and a quality scandal. They are in fact the same scandal, and the accounting convention is what connects them. A market that lets a buyer claim a ton for 5 dollars will find sellers willing to supply tons that are not really there, because at 5 dollars nothing real can be supplied.

*The Current Logic and Why It Fails*

At present, carbon markets work one ton for one ton. One ton of emissions is deemed offset by one ton of drawdown, typically from forests. The buyer's emissions and the seller's sequestration are treated as fungible.

This chapter argues against that idea, and for a global mitigation price set by the cost of inorganic removal. Offsets purchased at that price could then pay for multiple tons of, say, forest protection. Concretely, that price might be somewhere around 200 dollars a ton — well below the current cost of direct air capture, roughly where a mature auction might settle, and forty times the price of a rainforest credit today.

*The Argument from Economic Theory*

The theoretical case starts from an observation that is uncomfortable but hard to dispute. Economic theory tells you to pay the marginal cost at the social optimum. It does not tell you to pay the marginal cost in whatever market happens to exist. And the market for emissions and emissions removal is nowhere near the social optimum, and shows no sign of approaching it.

Look at where we actually are. There is far too little removal, which means the price is too low. There is far too little emissions reduction, which also means the price is too low. Both

margins point the same direction. The prices observed in today's offset markets are prices from deep inside a badly under-abated equilibrium, and using them to value a ton of carbon is like using famine-era rationing prices to value food.

There is a second, more formal way to see it. Take the global social optimum to be the intersection of a global marginal abatement cost curve and a global marginal damage cost curve. Since we are demonstrably not at that intersection — emissions are higher than optimal and the price paid is below the optimal level — a gold-standard offset market should pay the social cost *at* the optimum, not the shadow price of our current position. The cost of removing carbon from the atmosphere is a reasonable upper bound on that social cost. And given the economics of catastrophic risk — the property losses from flooding coastal cities alone would run to trillions — it may well be the lower bound too. When the plausible upper and lower bounds converge, that is a good sign the number is usable.

*The Argument from Practice*

The practical case is about trajectories. A 1.5C-compliant path requires that developed countries reach *financial* net zero well before the date at which everyone has committed to territorial net zero. Financial net zero means every country and every consumer group commits to offsetting all the emissions from each of the sectors that emit — and then that money is directed to protecting the carbon stocks we already have, creating more, and funding the technology markets that will eventually take carbon out of the atmosphere altogether.

The theory and the practice happen to agree, which is reassuring. It is not right that emitters carry on emitting while paying a paltry carbon price. The credits they buy are, on the evidence above, frequently of no value at all, and even the good ones may be impermanent — a protected forest can burn, or be logged after the contract expires. Paying for multiple tons is a far better response to that uncertainty than assuming every purchased ton was permanently sequestered. And genuine conservation needs more money than it is getting, not more credits.

It is worth being clear that this is an argument *for* carbon markets, not against them. It would be much better for mitigation if there were a proper carbon market in the world. At some abatement cost, offsetting genuinely beats eliminating — aviation is the standard example, and probably the right one, since the alternative fuels are expensive and the emissions are hard to abate. The objection is not to trade. It is to the price.

*What Is Needed to Make It Work*

Three institutions, each recognizable rather than a wish.

**A technology market for carbon capture, to set the global social cost of carbon.** This is a worldwide competition to determine the cheapest way to draw carbon out of the atmosphere: a global reverse auction, with each contract covering a plant sized at roughly 20 million tons of lifetime removal — a million tons a year for twenty years. The winning bid sets the price. This is the instrument that drove solar costs down through a decade of renewable auctions, applied to a different technology; and it has the same property, that the price is discovered rather than asserted. The Frontier commitment and the US tax credit are early, partial versions of this — a demand signal without yet an auction. The proposal is to make the auction the mechanism and its clearing price the benchmark.

**A mechanism for the most cost-effective use of funds.** This is a scientific and economic assessment exercise, and it should be designed by scientists rather than by market participants. If we know that deforesting the next 2 percent of the Amazon would trigger the loss of 75 percent of the remainder — through the forest-feeds-rain-feeds-forest dynamic — then

that particular carbon is immensely more valuable than carbon at the margin elsewhere. Having identified it, one designs a cost-effective preservation mechanism for it, potentially through competition, and pays for that mechanism in different parts of the forest. The fund allocation is deliberately non-uniform: the whole point is that not all tons are equal, and the Verra scandal was in large part the consequence of pretending they were.

**A gold standard for offsets: pay the full price, and spend it well.** Only this logic would be certified as genuine net zero. Countries and companies would define a date for financial net zero, ahead of global net zero. Reaching it means paying the full auction-set price per ton of your emissions — but you receive several tons of mitigation for that money, because most of the time protecting rainforest is more cost-effective than air capture. The buyer's claim is honest, the price is defensible, and the money goes where it does the most good. The Integrity Council's core carbon principles, adopted in 2023, are a step toward the standards half of this; what they lack is the price.

## 2.7 7. Carbon Markets Under Net Zero

The previous chapter proposed a new pricing logic. This one explains why the existing logic is not merely underpriced but structurally running out of sellers — and what has to replace it.

*The Sellers Are Disappearing*

With the Paris Agreement and the definition of net-zero strategies for Europe, China, the United States and India, long-term strategies have become an established norm. Some 140 countries, accounting for around nine-tenths of global emissions and output, have now pledged net zero by mid-century or thereabouts. Meeting the Paris temperature targets requires *global* net zero by midcentury for a 2C target, and earlier for 1.5C.[1]

The distinction between national and global net zero is doing a great deal of quiet work here. A standalone *national* net zero target admits the possibility that emissions in one country are compensated by counterfactual emissions reductions — offsets — in another. A *global* net zero target does not. It requires that gross emissions be compensated by actual *sinks*: the permanent drawing-down of carbon from the atmosphere. And because the 1.5C budget in particular may well be overshot, there is the further prospect of needing to fund removal simply to come back down again.

This creates a problem for carbon markets that has not been widely absorbed. Where will the *sellers* of credits come from — which is to say, which countries will overachieve their targets? The offset model implicitly assumes some underachievement of stringent targets in developed countries, matched by overachievement of less stringent targets in developing ones. But under a global net-zero norm, developing countries will themselves need stringent targets, which are by construction very hard to overachieve. The supply side of the offset market is being legislated out of existence.

The Paris Agreement's Article 6, whose rules were finally settled at the Baku conference in 2024, recognizes this in its requirement for "corresponding adjustments": a country that sells a credit must add the ton back to its own inventory, so that the ton is counted once. That is correct accounting, and it is fatal to the old model. A country with a net-zero target that sells a ton it has reduced must find another ton to reduce at home. The credit is no longer free to the seller, and the seller's price must reflect that.

[1]These targets matter: latest evidence suggests fundamental instability of Greenland ice sheets well below the 2C level, and perhaps even below current levels. Indeed the rapid warming of the arctic is already giving considerable global effects.

*What Carbon Markets Were Also For*

That matters more than it first appears, because carbon markets were never only an efficiency device. They were conceived as a burden-sharing mechanism: a way to transfer resources from the richer developed world, which carries the historical responsibility, to the developing world, and thereby to give developing countries a reason to participate in global climate action at all.

Remove offset-based markets and the efficiency loss is manageable. The loss of the transfer channel is not. If offsets cannot do this work, the participation incentive has to come from somewhere, which means side payments in some form.

So the general question about how carbon markets should be structured for long-term strategies resolves into five requirements:

a. a carbon *sink* logic rather than an offset logic;
b. side payments to poor countries in return for stringency;
c. an additional source of funding beyond offsets;
d. a market-based logic for efficient selection of projects and policies; and
e. fair burden sharing.

The technology-competition price of Chapter 6 satisfies (a) and (d). The rest of this book and its companion are largely an attempt to satisfy (b), (c) and (e).

*Funding*

Side payments and drawdown both imply significant funding. Where might it come from? Three sources are worth examining.

The first is gold-standard corporate net-zero commitments taken sometime before 2050 — in the 2030s, say. This is the demand side of Chapter 6′s proposal, and its attraction is that the money is voluntary and already notionally committed. Its weakness is the one the Verra episode exposed: voluntary demand collapses when the product is shown to be worthless, and it will return only when the product is credible.

The second is sovereign co-guaranteed long-term borrowing — hundred-year maturities — repaid in proportion to historical emissions. This is the mechanism sketched in the companion volume's preface: the present borrows from the future to prevent damage to the future, with the repayment key being cumulative responsibility.

The third is the concessionality component of scaled-up development bank lending. This could plausibly fund side payments, though probably not drawdown, and it is the subject of the companion volume.

*Assessing These by Political Economy*

The right way to compare these options is not by their elegance but by whether they can survive contact with politics. Three components of political economy matter.

*Between nations.* States are sovereign and face something close to a prisoner's dilemma. Any mechanism requiring simultaneous sacrifice by all parties will fail; mechanisms that reward early movers, or that make defection visible and costly, have a chance. Chapter 12 develops this.

*Within nations.* Concentrated interests — fossil fuel interests above all — have more influence than diffuse ones, and every country has its own entrenched set. This is the logic of collective action from Chapter 1, applied to international mechanism design, and it explains why instruments that need no net revenue travel better than instruments that need a lot.

*Across time.* Sacrifice later is considerably easier to agree to than sacrifice now. This is usually deplored as short-termism. It is more useful to treat it as a design parameter: a

mechanism that front-loads the incentive and back-loads the payment is exploiting a real feature of political preference rather than fighting it.

*The Legal Character of an Emission*
There is a legal dimension that carbon market design tends to ignore, and it strengthens the pricing argument of Chapter 6.

Emitting CO2 is not simply the consumption of a shared budget. It inflicts damage on identifiable agents across the world, in the future. Where there is potential legal liability for future damage, the appropriate response is to set aside money now to pay for it later — which is to say, it is a *capital* requirement, not an expense. This is the concept underlying both corporate emissions liability and national cumulative responsibility. An emitter who has not provisioned against the damage they are causing is running an unreserved liability, and no other industry is permitted to do that.

*Three Allocations*
Finally, there are three distinct allocation questions, which are often run together and should not be.

**The allocation of money to its most efficient use.** This requires market design. Renewables are procured through reverse auctions; the same instrument can procure removal, as Chapter 6 proposes. What matters here is *intertemporal* rather than static optimality — which may mean deliberately paying more for a particular technology now, because doing so moves it down its learning curve and lowers the cost of everything that follows. Static cost-effectiveness would never have funded solar in 2005, and it will not fund direct air capture in 2026.

**The burden-sharing arrangement.** One approach recognizes the need for low-cost financing at scale in middle-income countries. What form the subsidy to middle-income countries should take is unclear, and the numbers involved imply that the available funds are insufficient — which suggests the arrangement will need both carrots and sticks rather than carrots alone.

**The intergenerational allocation.** Here there is the possibility of a genuine Pareto improvement. If borrowing from the future funds the prevention of damage to the future, and the damage prevented exceeds the debt service, then every generation is better off. This is not redistribution. It is an investment that happens to cross a generational boundary, and the only reason it does not happen automatically is that the beneficiaries cannot yet sign the contract.

## 2.8 8. Positive Framing

The earliest draft of this project listed five ideas, and the fourth was called positive framing. It was described in one sentence: agree on something that everyone can agree on; don't lead with what is tricky or unpopular. This chapter is what that sentence means, because it turns out to be the principle that connects everything else in this part.

*The Principle*
Every climate policy can be described in more than one way. A carbon tax is a charge on pollution; it is also a rise in the price of petrol. A renewable auction is a subsidy to wind farms; it is also the cheapest way to buy electricity. A feebate on vehicles is a penalty on large cars; it is also a discount on small ones. The description is not neutral. It determines who thinks the policy is aimed at them.

Positive framing is the discipline of choosing, among the true descriptions of a policy, the one that identifies beneficiaries rather than payers — and then, more importantly, *designing*

*the policy so that the positive description is the accurate one*. It is not spin. A policy that is genuinely a discount on small cars is different from a policy that is genuinely a penalty on large ones, even if their marginal incentives are identical; the difference is where the benchmark sits and where the money flows, and those are design choices.

*Why It Is the Same Argument*

Read that way, positive framing is the logic of collective action turned into a design rule. Olson says a policy passes when its benefits are concentrated and visible and its costs diffuse and deferred. Positive framing says: build the policy so that it has visible beneficiaries, and describe it by naming them.

The feebate is the instrument that does this for carbon pricing. It has beneficiaries — the clean half of every sector — who can see what they gained. A tax with a dividend has beneficiaries too, on paper; the Canadian rebate left most households better off. But the rebate was not *framed* as a benefit, because it was not designed as one: it arrived unlabelled, disconnected from the charge, and it could not be pointed to. The positive frame was available and the design did not support it.

*Evidence from the Largest Climate Law*

The Inflation Reduction Act of 2022 is the largest climate measure any government has passed, and it contains no carbon price. It consists almost entirely of tax credits — for clean electricity, for manufacturing, for hydrogen, for carbon capture, for electric vehicles, for home efficiency. It was framed, in its title and its politics, as a measure against inflation and for domestic manufacturing. Its climate content was real and its climate framing was almost absent.

It passed the Senate by a single vote, which is to say it could not have passed as a carbon tax. And it has proved durable in a way a tax would not have, because its beneficiaries — the factories, the developers, the counties where the plants were built, many of them in districts represented by the party that opposed it — are concentrated and visible and organized. The credits have been trimmed since; they have not been repealed, and the trimming has been resisted by the beneficiaries. That is what positive framing looks like when it is built into the design.

The cost is the one Chapter 2 named. The act raises no revenue and spends a great deal, and its price signal is indirect. It is a phase-one instrument in the language of Part III, and its success is evidence for the sequence argument rather than against pricing. But it is also evidence that the frame decides what survives.

*Framing a Feebate*

Applied to the instruments of this book, positive framing yields three practical rules.

Name the rebate before the fee. A feebate on heating is a heat pump bonus, funded by a boiler charge; the order of the description matters because the first term is what people remember.

Make the rebate visible at the point where the fee is paid. Canada's rebate failed partly because it arrived three months after the charge and through a different channel. A feebate whose rebate appears on the same bill as the fee, or at the same till, is understood as a single instrument rather than as a tax and a separate handout.

Set the benchmark so that a majority receives. A feebate whose benchmark sits at the median has half the participants paying and half receiving, and the payers are the more motivated. A benchmark set slightly above the median puts a majority on the receiving side, at the cost of a small net fiscal outlay — which is Chapter 1's "slightly negative net revenue"

and its "slightly unfair" dynamic. The book's design choice there was not an accident of arithmetic. It was positive framing applied to the benchmark.

*The Limit*
There is one thing positive framing cannot do, and it should be stated. It cannot make a sacrifice popular. Where a policy genuinely requires most people to be worse off — a carbon tax whose revenue is needed for the budget, in the last phase of the sequence — no framing rescues it, and the honest course is to say so and to time the policy for the moment when the sacrifice is smallest. That moment, Part III argues, is later than most economists have wanted it to be. # Part III — Transitioning Pricing

## 2.9 9. Innovation in Waves

The remaining chapters of this book are about sequence — about the claim that the right instrument depends on the moment. That claim rests on a view of how innovation proceeds, and this chapter sets it out.

Governments seek to develop rapidly and in a low-carbon fashion. Can the two objectives be reconciled? Since climate action requires the turnover of older high-carbon ways of doing things and their replacement by newer low-carbon means, could that capital replacement — and the innovation and reconfiguration it demands — be a spur to growth rather than a drag on it?

*The History of Innovative Growth*
Historical examples abound of innovation driven by the state, by the market, and by both together. Looking back, such innovations have often been driven by energy and communications: the steam engine, the coming of electrification, the internet. Major state projects have been drivers too — the Second World War, the Apollo program, the search for a covid vaccine.

The mechanism in each case is recognizable. A new challenge becomes a spur to entrepreneurial initiative, and the inventions created threaten the interests of incumbents, who must react or be displaced. These new configurations also lead to more innovative firms replacing less creative peers, which means the process may generate not only better outcomes but cultural change over time. The countervailing tendency should be noted: market concentration, and the political influence of concentrated interests, can work against this — the logic of collective action from Chapter 1, appearing here as a brake on innovation rather than on policy.

Climate shows both dynamics in action, and policy action has transformed the story. Subsidies for renewable installation in Europe, North America and elsewhere, together with R&D and production support in China — preferential strategic industry status, free land, low-cost credit and other production subsidies — have combined with an intrinsic learning curve effect to drive costs down. The price of a solar module fell by roughly 90 percent between 2010 and 2020. Germany's feed-in tariff, funded by a surcharge on electricity bills that cumulatively exceeded 200 billion euros before it was abolished in 2022, is the single largest reason that solar is cheap in India today. It was, by any static measure, an extraordinarily expensive way to buy German electricity. It was, by the measure that matters, the cheapest investment in the transition anyone has made.

Reducing the cost of renewables has in a real sense transformed the outlook for climate policy, because it finally provides a cost-effective alternative energy supply at scale. When the state provides incentives, universities do primary research, and firms do the innovating, the combination can produce results that none of the three would generate alone.

*Externalities, Market Structure, and Phases*

According to Schumpeter, economic growth is driven by successive waves of innovation. Critical to the Schumpeterian view is the sense in which an externality loop can be *closed* — where and whether productive investment can be given a return. Innovation requires R&D, and R&D requires that someone capture enough of the gain to justify it.

Market structure therefore matters. If a monopolist controls an industry, Schumpeter argued, it may invest more in R&D, because it captures more of the rent from innovation. A competitive market might invest less. Arrow argued the opposite: that more competitive markets would be more innovative, since monopolists grow lazy and lack the drive to innovate. Recent empirical work suggests both effects are in play, with oligopoly the ideal balance between pure monopoly and perfect competition from the standpoint of innovation.

The *phases* of investment matter equally. Introducing new ways of doing things follows an S-shaped curve, like an epidemic: an early R&D and innovation phase, then a diffusion phase, then maturity. Climate change requires a dynamic process of creative reconfiguration across the whole economy, and each technology needed will sit somewhere on such a curve. Solar is well into diffusion. Batteries are entering it. Green steel and direct air capture are in the early phase. Aviation fuel is barely past R&D.

This is the structure the rest of Part III relies on, and its practical implication is the book's thesis: the right intervention depends on the phase, and applying a maturity-phase instrument to an early-phase technology accomplishes nothing.

*Drivers of Innovation-Led Growth*

Several distinct drivers operate: innovation policy, research clusters, culture, industrial structure, technological dynamics, and the potential for new markets or the destruction of old capital.

Innovation requires secure property rights, in the sense of sufficient certainty about the business environment. This links directly to the companion volume — the need for low-risk innovation possibilities and safe business environments is the same need that commitment devices and guarantee structures address. An investor who cannot predict the policy environment will not fund the plant that embodies the innovation.

What of other growth factors? Climate action also requires vast quantities of capital, and such capital is available — the companion volume's argument is that the constraint is risk-absorbing capital rather than capital as such. Growth theory emphasizes both innovation and accumulation as significant drivers. Green development involves both, which is the strongest theoretical reason to expect it to be growth-enhancing rather than growth-reducing.

### 2.10 10. Three Phases of Collective Action

The last thirty years of climate policy can be described in a single sentence: an extraordinary innovation success combined with a massive collective action failure. Solar modules fell in price by ninety percent. Global emissions rose. Any theory of what to do next has to explain how both things happened at once.

*Borrowing from Schumpeter*

Schumpeter's insight about innovation was that market structure and innovation interact: monopolists, contrary to the naive competitive story, may innovate more, because they can capture the returns. The suggestion here is to approach the green transition from the same angle but applied to policy — to ask whether competitive dynamics could inspire policy

and collective action, and whether better and more dynamic policy could in turn generate innovation gains.

What this points toward is a *phase* model of policy, parallel to the phase model of technology in the previous chapter. Technologies pass through invention, demonstration, early deployment, mass deployment and maturity, and the right intervention differs at every stage. Policies plausibly do the same: policy innovation leads to coordinated policy diffusion, which leads to the creation of new norms.

*Elster's Three Phases*

Jon Elster provides a compelling account of how collective action actually gets going, and it maps onto the climate problem unusually well.

The first phase is driven by what he calls everyday Kantianism. Early movers act unilaterally, according to a rule they would wish to see universalized — they are, in the popular formulation, the change they want to see in the world. Such actors are not making a cost-benefit calculation, and criticizing them for irrationality misses what they are doing: they are creating the norm that later actors will respond to. Germany's feed-in tariff was everyday Kantianism at national scale. It made no sense as a German energy policy. It made sense as a rule Germany wished the world would follow, and it cost Germany a great deal to follow it first.

The second phase belongs to the utilitarians. These more pragmatic actors look for the diffusion of the new norm through local win-win arrangements. At the level of a state, the argument is about co-benefits and competitive advantage; inside a business, innovation is argued for as a business proposition. The essential move in this phase is to close the loop — to find a network of beneficiaries who can actually capture the gains, because a win-win that nobody can appropriate does not diffuse. The spread of renewable auctions from a handful of countries in 2010 to well over a hundred by 2020 is diffusion in this sense: no country adopted an auction because it wished the world would; each adopted it because it was the cheapest way to buy power, and the neighbours could see that.

The third phase is the norm of fairness. By this point the new standard has been established as a social norm, and the question has changed from whether to comply to whether one is doing one's share. Enforcement becomes largely social rather than administrative, which is what makes it cheap. Net-zero pledges are the climate norm in this phase — adopted by some 140 countries not because each calculated the benefit but because not having one became, after about 2019, a thing that had to be explained.

The practical value of this scheme is that it tells you the sequence of instruments is not arbitrary. Attempting a phase-three instrument in a phase-one world is the characteristic error of climate policy. A carbon tax with revenue recycling is a phase-three instrument: it assumes the norm of doing one's share is already established, so that the tax is experienced as fair. Australia introduced one in 2012, France tried to raise one in 2018, and Canada abolished one in 2025, in each case in a polity where the norm had not been established and the tax was experienced as an imposition. The instrument was not wrong. The timing was.

## 2.11 11. The Sequence

Putting the two previous chapters together yields a specific and somewhat heretical conclusion about carbon pricing: that it is the right policy at the wrong time, and that the sequence matters more than the instrument.

*The Missing Middle*

Looking at climate strategies as they exist, we have different norms operating at different levels but no coherent process for getting from one stage to the next. The gap is most

conspicuous in the middle. We lack policies for the *mass deployment* of technologies during the phase before carbon pricing can become dominant. Subsidy handles the early phase; carbon taxation handles the mature phase; the crossing between them is largely unmanaged, and it is where most of the world's emissions currently sit.

A sequence theory of carbon pricing can be constructed to fill it. The Schumpeterian and game-theoretic grounds are the same: existing high-carbon interests will block carbon taxes unless incentives are tailored to prevent resistance and, better, to transform those interests into stakeholders in the new system.

*Three Phases, Three Instruments*

In the early phase, green alternatives are small but expensive. Subsidies are close to optimal from a collective action perspective — concentrated beneficiaries, diffuse funding, no organized opposition. This is the pattern Chapter 1 identified as politically feasible, and it explains why the world has so much subsidy and so little taxation. It also explains the Inflation Reduction Act, which is a phase-one instrument for the technologies — hydrogen, capture, green manufacturing — that are in phase one, deployed alongside phase-one instruments for solar and wind that were, by 2022, arguably no longer in it.

In the intermediate phase, low-carbon assets have a foothold and are growing fast. Subsidies now become too expensive, precisely because they are succeeding: the volume being subsidized has grown. Germany's surcharge reached its political limit for exactly this reason — it was funding a success, and the success had become the largest item on the electricity bill. Here a revenue-neutral tax with targeted sectoral rebating — a feebate — is the better instrument. It maintains the marginal incentive while the fiscal cost falls away, and it does so without creating the concentrated losers that would block a tax. Reverse auctions belong here too: they replaced feed-in tariffs across the world in the 2010s for the same reason, converting an open-ended subsidy into a competitive procurement.

Later, and also earlier for sectors that are not carbon-intensive to begin with, carbon taxation with general revenue recycling is the ideal. By this stage the high-carbon incumbents are diminished, the alternatives are cheap, and the tax can do double duty as a revenue instrument and a climate one. The European emissions trading price, which languished below 10 euros for most of a decade and rose to between 80 and 100 after 2021, is the sequence in a single series: the price could rise once the alternatives were cheap enough that a high price no longer meant a high electricity bill.

*The Sequence in Practice*

Solar is the proof of concept, and it is worth stating the whole arc.

*Phase one, 2000–2010.* Germany, Spain and Italy pay feed-in tariffs far above wholesale prices. The cost is large and the volume is small. Nobody organizes against a subsidy to a few thousand rooftops. Chinese manufacturers, supported by production subsidies, build capacity to serve the European demand. Module prices fall by more than half.

*Phase two, 2010–2020.* Volumes explode; Germany's surcharge becomes politically toxic; Spain retroactively cuts its tariff, with lasting damage to its credibility. Governments switch to auctions — South Africa in 2011, India, Chile, Brazil, Mexico, then dozens more — which maintain the incentive while letting competition drive the price. Module prices fall by another 80 percent. The instrument has changed from a subsidy to a quantity contract, and the political opposition has largely dissolved because the electricity is now cheap.

*Phase three, 2020 onward.* Solar is the cheapest source of new power in most of the world. The European carbon price rises to 80 euros and above without a fuel-price revolt, because

the alternative to paying it is now cheaper than paying it was. Coal is being priced out. The tax works, because the sequence made it survivable.

The argument of this book is that this arc was not luck. It is the shape that any successful transition takes, and it can be planned for rather than stumbled into. The failure mode is trying to start at the end: a carbon tax on transport fuel in 2018, when the electric alternative was still expensive and the tax was experienced not as a nudge toward a cheaper option but as a pure cost. That is Australia, and France, and Washington, and Canada. The right policy at the wrong time.

*Conclusions*

Six propositions follow from this part of the argument.

1. Green development is concrete. It requires frameworks and platforms, not just prices.
2. Research and development remain key, because the sequence starts with something to deploy.
3. Secure property rights and derisked investment are prerequisites, not refinements — the subject of the companion volume.
4. Supporting diffusion is critical, and not only for technologies. Policies diffuse too, and can be designed to diffuse faster.
5. Policy forms and financial forms plausibly share the same dynamic diffusion processes, which means the same phase logic applies to instrument design as to technology support.
6. Policy and norm diffusion could both lead to and be led by Schumpeterian green growth. The causation runs in both directions, which is what makes the virtuous circle possible — and what makes getting the sequence right so valuable.

### 2.12 12. Dynamic Incentives

The fifth idea in the earliest draft of this project was described as "dynamic incentives and collective action voting schemes." The first half of that phrase has run through every chapter of this book. The second half was never developed in the source material, and this chapter says what can be said about it honestly: the dynamic incentive is a real and working idea, and the voting scheme is a sketch.

*Act Now or Miss Out*

A static incentive pays the same amount for the same action whenever it is taken. A dynamic incentive pays more for acting sooner. The difference sounds small and is not, because it changes the strategic calculation of waiting.

Under a static incentive, waiting is free: a country or a firm that delays loses nothing, and may gain, because the technology gets cheaper and the incentive is still there. Under a dynamic incentive, waiting has a price: the reward available today is larger than the reward available next year, so delay is a decision to forgo something. The additionality problem that Chapter 7 identified in offset markets — that paying for action invites countries to delay until they are paid — is a static incentive problem, and a dynamic incentive is its solution.

Renewable auctions are the working example, and they were not designed as dynamic incentives; they became one. Because auction clearing prices fall over time as the technology matures, a developer who wins a contract in this year's round locks in a higher tariff than one who waits for next year's. The early mover is rewarded for moving early, which drives the competition that pushes the next round's price down. The auction rewards speed as a by-product of rewarding cheapness.

The companion volume builds two mechanisms on this principle explicitly. Its policy-lending fund gives countries that act early on policy a larger share of the fund's equity than late movers. Its Climate Incentives Trust holds accounts that grow with a country's action and shrink with its reversal, so that delay is visibly a loss and reversal visibly a cost. Both convert the question "will you act?" into the question "will you act *now*?", and the second question has a different answer.

*Time Consistency*

There is a structural reason dynamic incentives work where static ones do not, and it is about time consistency.

A government promising a reward for future action faces a credibility problem: once the action is taken, the reward is a pure cost, and the government has every reason to reduce it. Spain's retroactive tariff cut and Vietnam's, described in the companion volume, are this problem in action. A dynamic incentive is more credible because it is front-loaded: the reward for early action is paid early, before the government has had time to regret it, and the reward for late action is smaller by design, so a later government reducing it is doing what the schedule already said.

This is why Chapter 7's "across time" observation — that sacrifice later is easier to agree to than sacrifice now — is a design parameter rather than an obstacle. A dynamic incentive front-loads the reward and back-loads the obligation, which is the shape of a promise that political preferences will let a government keep.

*The Voting Scheme*

The original notes said no more than the phrase. Here is what it might have meant, offered as a sketch and not as a proposal.

International incentives require someone to decide the allocation — who receives what for which action. Under a treaty, that decision is made by negotiation, which is slow and rewards obstruction. Under a fund, it is made by the fund's managers, which is faster and rewards whoever controls the managers. A voting scheme is a third option: the participants themselves vote on allocation, with votes weighted by something other than emissions or population.

The weighting that makes this a collective-action device rather than a plebiscite is *contribution*. A participant's vote on how the pool is allocated is proportional to what it has put in — in money, or in verified action. Early movers therefore control the allocation rules, which is a dynamic incentive to move early; and the rules they set, being made by those who have already acted, are likely to reward action rather than promises. This is the governance structure of a mutual society or a cooperative, applied to a climate fund, and it is the natural complement to the trust structure the companion volume proposes.

Whether it would work depends on questions the source material did not answer: how contribution is measured, how new members enter, and how the scheme avoids becoming a club of the rich. Those questions are left open. What can be said is that the scheme, if it can be made to work, is the mechanism by which the dynamic incentive becomes self-governing — and that a self-governing incentive is one that no future government can quietly repeal.

## 2.13 13. The Order of Things

This book opened with a prime minister abolishing a carbon tax on his first day in office, and keeping another one that nobody noticed. It is time to say what that meant.

*What the Argument Has Been*

The argument has been that carbon pricing is right and is being done wrong — wrong in shape, wrong in price, and wrong in order.

*Wrong in shape.* A tax concentrates its costs on those who can organize and diffuses its benefits across those who cannot. It fails for that reason, not because the public is irrational, and it fails in rich democracies with sophisticated electorates as readily as anywhere. The instrument that survives — in Canada's industry, in Europe's exposed sectors, in France's car showrooms — is the one that creates winners and losers inside the same group, so that the winners defend it. That instrument delivers the same marginal signal as a tax. What it removes is the average burden, and the average burden is the politics.

*Wrong in price.* The price of a ton of carbon in today's offset markets is a few dollars, and the evidence of 2023 is that at a few dollars nothing real is supplied. The price of actually removing a ton is a few hundred. A market that let the second price govern — discovered by auction, not asserted by committee — and that let one ton bought at that price purchase several tons of the cheapest real mitigation, would move a great deal of money toward standing forests while funding the removal technology the second half of the century will need.

*Wrong in order.* The transition proceeds in phases, and so does collective action. Subsidy fits the early phase, when the alternative is small and nobody organizes against helping it. The feebate fits the middle, when the alternative has grown too large to subsidize but the incumbent is still strong enough to block a tax. The tax fits the end, when the alternative is cheap and the tax is experienced as a nudge rather than a cost. Solar walked that arc in twenty years. Transport fuel was taxed as if it had, and had not.

*Why Timing Is the Argument*

The companion volume argued that the cost of capital decides which of two worlds a country lives in. This book has argued that the timing of the pricing instrument decides whether a country can stay in the better one.

The two claims are the same claim. A cheap cost of capital needs a contracted revenue stream; a contracted revenue stream is created by a quantity instrument — an auction, a tariff, a contract for difference; and a quantity instrument is a phase-two policy, the one that fits the moment after subsidy and before tax. The instrument that is right on political grounds turns out to be the one that is right on financial grounds, and it is right at the same time. That coincidence is not a coincidence. It is what a well-sequenced transition looks like from two sides.

The phrase in the title is not a complaint about carbon pricing. It is a statement of when it belongs. The right policy at the wrong time is still the right policy — and the whole of this book has been an attempt to say when the right time is, and what to do until then.

## 3 References

Works cited or drawn on in the text. The wider bibliography of the project is in the working papers volume.

Council of the European Union (2022). *Regulation (EU) 2022/1854 on an emergency intervention to address high energy prices*.

Elster, J. (1989). *The Cement of Society: A Study of Social Order*. Cambridge University Press.

Government of Canada (2025). *Removing the consumer carbon price, effective April 1, 2025*. Department of Finance backgrounder.

*The Guardian*, *Die Zeit* and SourceMaterial (2023). Investigation into Verra-certified rainforest offsets, January 2023.

I4CE (2025). "In the absence of a carbon tax in Canada, measures to fill the gap are essential."

IEA (2022). *A 10-Point Plan to Reduce the European Union's Reliance on Russian Natural Gas*.

IEA (2022). *World Energy Outlook 2022*.

Integrity Council for the Voluntary Carbon Market (2023). *Core Carbon Principles*.

Olson, M. (1965). *The Logic of Collective Action*. Harvard University Press.

Pahle, M., Tietjen, O., Osorio, S., Egli, F., Steffen, B., Schmidt, T. S. and Edenhofer, O. (2022). "Safeguarding the energy transition against political backlash to carbon markets." *Nature Energy*, 7(3), 290–296.

Ramsey, F. P. (1927). "A Contribution to the Theory of Taxation." *Economic Journal*, 37, 47–61.

Schumpeter, J. A. (1942). *Capitalism, Socialism and Democracy*. Harper.

UNFCCC (2024). Decisions on Article 6.2 and 6.4 of the Paris Agreement, COP29, Baku.

Van der Ploeg, F. (2021). "Climate Policies: Challenges, Obstacles and Tools." *National Institute Economic Review*, 258.

World Bank (2022). *Global Economic Prospects*, June 2022.

World Bank (annual). *State and Trends of Carbon Pricing*.