

Karp, Larry S.; Träger, Christian

Working Paper
Smart Cap

CESifo Working Paper, No. 8917

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Karp, Larry S.; Träger, Christian (2021) : Smart Cap, CESifo Working Paper, No. 8917, Center for Economic Studies and Ifo Institute (CESifo), Munich

This Version is available at:

<https://hdl.handle.net/10419/235287>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

CESifo WORKING PAPERS

8917
2021

February 2021

Smart Cap

Larry Karp, Christian Traeger

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editor: Clemens Fuest

<https://www.cesifo.org/en/wp>

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: <https://www.cesifo.org/en/wp>

Smart Cap

Abstract

We introduce a “smart” cap and trade system that eliminates the welfare costs of asymmetric information (“uncertainty”). This cap responds endogenously to technology or macroeconomic shocks, relying on the market price of certificates to aggregate information. It allows policy makers to modify existing institutions to achieve more efficient emission reductions. The paper also shows that the efficient carbon price is more sensitive to technological innovations than usually assumed. The lasting impact and slow diffusion of these innovations typically make the optimal carbon price a much steeper function of emissions than suggested by the social cost of carbon.

JEL Codes: Q000, Q500, H200, D800.

Keywords: pollution, climate change, taxes, quantities, regulation, smart cap, uncertainty, technology diffusion, dynamic programming, integrated assessment, DICE.

Larry Karp
Department of Agricultural & Resource
Economics, UC Berkeley
Berkeley / CA / USA
karp@berkeley.edu

Christian Traeger
Department of Economics
University of Oslo / Norway
christian.traeger@econ.uio.no

Version of February 2021

We are grateful for feedback from Grischa Perino, Christian Gollier, Severin Borenstein, Paolo Piacquadio, Hiroaki Sakamoto, Robert Stavins, Jean Tirole, Marty Weitzman, Karen Pittel, Till Requate, Wilfried Rickels, Alex Schmidt, Lint Barrage, Inge van Den Bijgaart, Rolf Golombek, Bard Harstad, Michael Hoel, Eric Naevdal, Armon Rezai, our RA Kristen Vamsaeter, and participants at the EAERE 2017, CESifo 2017, EEA 2018, and the CEPR 2020 Climate Change Workshop. We thank the INNOPROD project for providing us with access to the PATSTAT data. Christian Traeger gratefully acknowledges support by the National Science Foundations as part of the SCRiM network on Sustainable Climate Risk Management.

1 Introduction

Forty countries have implemented either a tax or a cap and trade system to regulate greenhouse gas emissions and mitigate climate change. The European Emissions Trading System (ETS) has an annual market value of about 60 billion USD. Following the Paris Climate Agreement, another 88 countries are currently considering introducing either a tax or a cap and trade system. These instruments are also used to regulate other pollutants, and there are many subnational initiatives. The high cost of reducing greenhouse gas (GHG) emissions, and the potentially enormous costs of failing to deal with the climate problem, make it important to use efficient policies. The problem is too big, and the remedy too expensive, to waste effort. Asymmetry of information between firms and the regulator is central to the policy design problem. We introduce a “smart cap” that efficiently aggregates information and responds to technological innovations or other macroeconomic shocks, thereby reducing the cost of climate change mitigation.

Cap and trade is currently the most widely used market-based policy to control GHG emissions. California’s ETS, with a market value of 6 billion USD, is the largest in the U.S., and China is currently introducing the largest overall. Due largely to technological and macroeconomic shocks, these policies have resulted in low carbon prices and modest emission reductions in the European ETS, the Regional Greenhouse Gas Initiative (RGGI), and California’s carbon market. As a result, these policies miss low-cost emissions reduction opportunities and undermine the public’s confidence in market-based regulation. A smart cap endogenously contracts the emission cap.

Our main innovation shows how a policy that endogenously adjusts the emission cap in response to the certificate price overcomes current policies’ limitations. Under this smart cap, the regulator auctions or gives away certificates at the beginning of each compliance period, and simultaneously announces a “redemption function” that depends on the equilibrium certificate price. The redemption function determines the number of allowable units of emissions per certificate. The equilibrium certificate price, and thus the equilibrium aggregate emissions level, responds to technological innovations or macroeconomic shocks. In contrast, under a standard cap, certificates are issued in fixed units of emissions and aggregate emissions do not respond to prices and shocks.

We incorporate a tractable model of innovation and diffusion into a dynamic analytic model of emissions and climate response. The social cost of carbon (SCC) equals

the present discounted costs from releasing a ton of carbon dioxide (CO_2) today. The SCC is the stock-pollution analogue of marginal damages in a static model. It is generally accepted that damages are only moderately convex in the pollution stock, so the slope of the SCC is small.

The resulting model gives rise to a conceptual insight that is relevant well beyond the smart cap. It is widely believed that, because the slope of the SCC is small, the optimal carbon price response to emission fluctuations is also small. We show instead that the optimal price response is generally much steeper than the SCC. In extreme cases, the sign of the optimal price response can even reverse the sign of the slope of the SCC. In knife-edge situations, the standard cap and trade system can be first best. A technological innovation that reduces current abatement costs also reduces future costs, thereby reducing future emissions, future damages and today's SCC. The resulting positive correlation between marginal abatement costs and the SCC causes the optimal equilibrium price to respond much more strongly to emission shocks than the SCC suggests. The slope of the price response function is even steeper if an innovation in abatement technology is adopted over the course of several compliance periods. Then, a change in abatement cost during the current compliance period implies an even higher (persistent) long-term impact on future emissions and the SCC. Thus, efficient mitigation policies are sensitive to the speed of technology diffusion. A regression of emissions on green patents suggests moderately slow diffusion that turns out highly policy-relevant for efficient mitigation policies.

The smart cap's trading system uses the market to resolve the information aggregation problem discussed in Kwerel (1977), Dasgupta, Hammond & Maskin (1980), and Boleslavsky & Kelly (2014). The tax analogue of our smart cap is a nonlinear emission-dependent tax that we refer to as a "smart tax". In principle, this smart tax can also solve the problem caused by asymmetric information between firms and the regulator. However, the smart tax requires an agency to keep track of and publicize cumulative emissions levels in real time, enabling firms to base their emissions decision on the equilibrium unit tax. In contrast, with the smart cap, the market aggregates and reveals information via the certificate price.

The smart cap is a smooth first-best improvement over hybrid trading systems that add a price floor and ceiling to a standard cap and trade system (Roberts & Spence 1976, Weitzman 1978, Pizer 2002, Hepburn 2006, Fell & Morgenstern 2010, Grull & Taschini 2011, Fell, Burtraw, Morgenstern & Palmer 2012). In the hybrid system,

partly implemented in California, the policy maker commits to buying and selling certificates to keep the abatement cost within a pre-defined price window, making it effectively a tax when the price reaches these boundaries. The smart cap smoothly responds to price changes, eliminating the need for a regulator to buy or sell permits to maintain the price floor or ceiling.

Many papers discuss emissions regulation with asymmetric information for flow pollutants, i.e., pollutants that do not cause damages beyond the period in which they are emitted.¹ Requate & Unold (2001) explain how the issuance of options on emission certificates implements a step function approximation to the marginal damage curve, resulting in the static smart tax. Newell, Pizer & Zhang (2005) show how a committed agency can manage allowances to use a standard cap for direct price control. Taking this idea a step further, Kollenberg & Taschini (2016) show that an appropriate management of banking reserves can transform a standard cap with banking and borrowing into a hybrid mechanism that continuously interpolates between a standard cap and a standard tax. Pizer & Prest (2020) note that with banking and borrowing, adjustment of the intertemporal exchange rates enables the regulator to achieve the first best, provided that all uncertainty is resolved in the last period.²

It is widely understood that policies should be conditioned on available information (Ellerman & Wing 2003, Jotzo & Pezzey 2007, Newell & Pizer 2008, Doda 2016). For example, Burtraw, Holt, Palmer & Shobe (2020) note that a policy that conditions the current quota allocation on previous prices increases welfare relative to a standard cap or tax. The failure of such conditioning is less harmful in a smart cap because of its automatic adjustment to the price of certificates. We still recommend explicit conditioning on observables in order to permit tailoring the smart cap to those cost shocks that are less well observed.

The closest real-world implementation of a self-adjusting cap is the recently enacted market stability reserve in the EU ETS, which addresses the prevailing oversupply of allowances and cancels banked permits in a rather complicated fashion. We refer to

¹Gerlagh & Heijmans (2020) discuss a mechanism that achieves almost the first best for a particular type of stock pollutant – one in which stock-related damages arise only in the final period. However, for most pollutants, including climate change, damages in a period depend on the stock in that period. This mechanism therefore has limited applicability.

²The intertemporal exchange rate is the number of permits in period t that can be exchanged for one permit in period $t + 1$. Pizer and Prest also consider a climate application. Their assumption that marginal flow damages (and, thus, the SCC) are independent of the stock of atmospheric carbon assumes away the interaction between stock pollutants and technology shocks that turns a plausibly small but non-zero marginal damage slope crucial for climate policy.

Perino (2018), Perino & Willner (2016), Kollenberg & Taschini (2016), Fell (2016), and Silbye & Birch-Sørensen (2019) for detailed discussion and critical assessments. Perino, Ritz & van Benthem (2020) and Jarke & Perino (2017) show how interacting climate policies sometimes reinforce and other times offset each other.

Our focus is on achieving efficient regulation of a stock pollutant when there is asymmetric information and shocks are persistent. As an important byproduct, we obtain a simple and intuitive criterion for ranking the standard tax and quota, two second-best policies that do not overcome the problem of asymmetric information. Weitzman (1974) provided the criterion for ranking these two policies for flow pollutants; a number of papers have extended his results to stock pollutants.³ We provide a much simpler and more intuitive criterion, showing that (for a stock pollutant) the choice between a standard tax and cap depends on the relative slopes of marginal abatement costs and the smart tax – not the SCC. Thus, we obtain an exact and very simple analog between the tax-quota ranking criteria for a flow versus a stock pollutant.

2 Smart Tax and Smart Cap

Regulators usually set policy without knowing firms' abatement cost. This asymmetry of information arises both because firms have genuinely private information and because they make emissions decisions more frequently than regulators revise policy. In the latter case, firms condition their decisions on information that is publicly available, but unknown when the regulator sets the policy. For example, the 2008 recession reduced firms' incentives to emit, contributing to the low permit prices in the European carbon trading system. We emphasize the asymmetry arising from private information; the second source of asymmetry can be eliminated by announcing future state-contingent policies that depend on future public information.

The smart cap uses the market to aggregate information and to implement the first-best emission allocation. We use the one-period (or flow pollution) model to review a smart tax and introduce the smart cap. Here (by assumption), uncertainty affects

³These papers include Hoel & Karp (2001), Newell & Pizer (2003), and Karp & Zhang (2005). Fischer & Springborn (2011), and Heutel (2012) use stochastic general equilibrium frameworks to compare tax versus quantity regulation, emphasizing the effect of business cycles. Our companion paper Karp & Traeger (2018) discusses further implications for taxes versus quantities when these are the only feasible policy options. It also relates our findings to Weitzman's (1974) and Stavins's (1996) (static) insights on the role of correlated shocks. Stavins (2020) reviews tax and quantity regulation in theory and practice.

abatement costs, but not social damages. We then modify the one-period model to provide intuition for the dynamic setting. Next, we consider the stability of competitive equilibria, and we examine the effect of market power.

2.1 A Static Model

A representative firm's marginal benefit $MB(E|\theta)$ from emissions E (equal to its marginal abatement cost)⁴ depends on the random variable θ . The marginal social damage of emissions is $MD(E)$. We assume that $MD(E)$ and $MB(E|\theta)$ are positive, continuously differentiable, and that benefits of emissions are concave: $MB_E(E|\theta) < 0$. We define θ so that a larger realization increases the marginal benefit of emissions, thereby raising abatement costs; for example, a large θ represents lower than expected green technological progress, or higher than expected economic growth and demand for fossil fuels.

A regulator distributes Q tradable emission certificates, and announces a "redemption function", $q(p)$, where p is the endogenous market price of a certificate. One certificate allows the firm to emit $q(p)$ units, so the cost to the firm of one unit of emissions is $p^E = \frac{p}{q(p)}$. The endogenous "smart cap" is $Qq(p)$, equal to the number of units of emissions. Firms choose their level of emissions, E , and they trade emission certificates at price p .

The optimal emission price equals the marginal damage of emissions:

$$p^E = MD(E).$$

A *smart tax* (a function of E), equal to the right side of this equation, implements the first best level of emissions. Using the market clearing condition for emission certificates, $E = Qq(p)$, we obtain an implicit formula for the optimal *redemption function* $q(p)$

$$\frac{p}{q(p)} = MD(Q \cdot q(p)). \quad (1)$$

Totally differentiating this equation gives the slope of the redemption function

$$q'(p) = \frac{1}{MD'(E)E + MD(E)} \quad \text{where } E = Qq. \quad (2)$$

⁴Abatement A is the difference between business as usual emissions E^{BAU} and actual emissions E . Abatement costs are $C(A|\theta) = C(E^{BAU} - E|\theta) = B(E^{BAU}|\theta) - B(E|\theta)$. Deriving this equation w.r.t. emissions implies $MC(A|\theta) = MB(E|\theta)$.

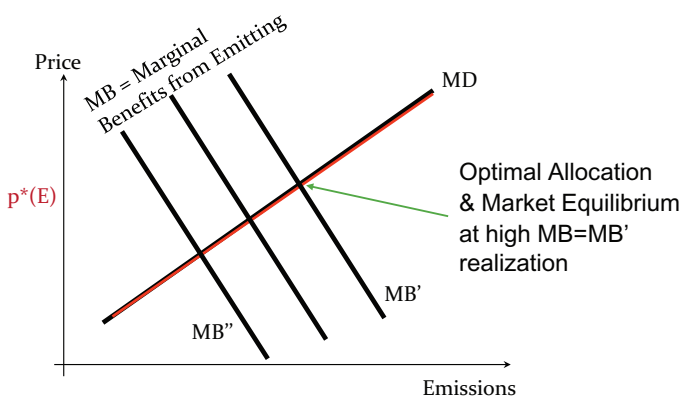


Figure 1: Static setting. The optimal carbon tax (smart tax, red) equals the marginal damage (MD) curve. The three downward sloping lines represent different realizations of the technology shock, each of which results in a different marginal benefit (MB) curve. The green arrow identifies the optimal allocation under the high marginal benefits. Under a smart tax (or a smart cap), this optimal allocation is also the market equilibrium where firms equate the marginal benefits from emissions with their private cost of emitting another unit.

If marginal damages are flat ($MD'(E) \approx 0$), the redemption function $q(p)$ is approximately linear in the certificate price, and the optimal smart tax is approximately constant.⁵ In this situation, a hybrid cap with a price ceiling and floor might be difficult to implement because the regulator would have to buy or sell many certificates to defend the floor or ceiling. The smart cap, in contrast, responds smoothly to shocks.

The redemption function's slope varies inversely with marginal damages. With low marginal damages, a higher certificate price generates a large increase in allowable emissions. With high marginal damages, the same price increase leads to a smaller increase in emissions. With strictly convex damages, higher emissions are increasingly costly to society. The term $MD'(E)$ decreases the redemption function's slope, leading to a smaller expansion of a smart cap following an increase in the certificate price.

In the classic prices versus quantities setting (Weitzman 1974) marginal damages are linear in emissions $MD(E) = a + bE$. The red line in Figure 1 illustrates the smart tax, coinciding with the MD -curve. Equation (1) becomes a quadratic equation with positive (because $q \geq 0$) root and results in the smart cap

$$q(p) = \frac{1}{2Qb} \left(-a + \sqrt{a^2 + 4Qbp} \right),$$

⁵More generally, the redemption function $q(p)$ is linear in the certificate price if and only if the damage function satisfies $a_0 + aE + c \ln E$, implying marginal damages $MD(E) = a + \frac{c}{E}$. Damages are increasing and concave for $c > 0$. Damages are convex but start out falling and only increase for $E > \frac{c}{a}$ for $c > 0$.

with slope

$$q'(p) = \frac{1}{a + 2bE} \quad \text{where } E = Qq(p).$$

The slope of the marginal damage function, $MD'(E) = b$, is a measure of the damage convexity. Following the discussion above (or using L'Hospital's Rule) we find a linear redemption function $q(p) \rightarrow \frac{p}{a}$ as $b \rightarrow 0$. The smart cap becomes inversely proportional to the (then constant) marginal damages a . Another interesting case arises as we rotate the marginal damage curve counterclockwise around some point (E^*, p^*) , increasing b and making offsetting changes in a . As $b \rightarrow \infty$ the smart cap approaches the constant cap E^* .

We obtain the (optimal) equilibrium emission from the representative firm's optimality conditions. This firm sets its marginal benefit from emissions equal to the emission price $MB(E, \theta) = p^E \Leftrightarrow MB(Qq(p), \theta) = \frac{p}{q(p)}$. If marginal benefit is linear, $MB(E|\theta) = \theta - fE$, the firm's optimality condition is ⁶

$$\theta - fE = \frac{p}{q(p)} \quad \Rightarrow \quad E = q(p)Q = \frac{\theta - a}{b + f}.$$

The redemption function $q(p)$ depends only on the certificate price; but the equilibrium price, and thus the equilibrium value of the redemption function, depends on the realization of the technology shock. The equilibrium cap is directly proportional to the net benefit $\theta - a$ of the first unit of emissions, and inversely proportional to the sum of the slopes of marginal costs and damages.

The representative firm formulation, where marginal benefits of emissions depend on the aggregate technology shock, conceals an important advantage of a smart cap over a smart tax. Individual firms have little knowledge of the mitigation technologies in other sectors, and thus are unable to predict aggregate emission levels. They therefore cannot predict the equilibrium value of the end-of-period smart tax. In contrast, the cap and trade market *aggregates information* governing technology and expected emission levels across firms. This market helps firms to form the correct price expectations, *enabling them to take optimal emissions decisions*.

We continue to use the representative firm model for exposition. However, because of the importance of firm heterogeneity in motivating the smart cap, we briefly consider the model with a continuum of firms, with mass normalized to unity. Firm

⁶The firm's first order condition implies $fQq^2 - \theta q + p = 0$. The smart cap satisfies $bQq^2 + aq - p = 0$. For $p > 0$, these two equations imply $qQ = (\theta - a)/(f + b)$.

i 's technology realization in state of the world s is $\theta_s(i)$, with marginal benefits from emissions $MB_i(E_i|\theta_s(i))$. Facing the emission price p^E (the same for every firm), the firm's optimal emission level is $E_i(\theta_s(i); p^E) = MB_i^{-1}(p^E|\theta_s(i))$. The aggregate emissions level is $E(\theta_s; p^E) = \int_i E_i(\theta_s(i); p^E) di = \int_i MB_i^{-1}(p^E|\theta_s(i)) di$. Then, we can define the representative firm's (aggregate) marginal benefit curve $MB(p^E|\theta_s(i))$ by solving (for p^E) the implicit equation $E(\theta_s; p^E) = \bar{E}$ for all relevant emission \bar{E} and shock θ_s levels. The solution is usually at least locally well-defined, and in our leading example of linear-quadratic benefits it implies globally the linear-quadratic representative firm model.⁷

2.2 Dynamic Insights

Climate change is a dynamic problem. As emissions accumulate in the atmosphere, marginal damages likely increase. Here, optimal policy depends on the shadow cost of the pollution stock, called the social cost of carbon (SCC) in the climate setting.

With persistent technology shocks, today's innovation affects future abatement costs, altering future emissions levels. Consequently, today's technology shock affects the future marginal damages arising from today's emissions. Thus, the SCC depends on both today's realization of the technology shock and on current emissions, E . We write the SCC as $SCC(E|\theta)$, a function of emissions, conditional on the shock realization. As with marginal damages, we assume that $SCC(E|\theta)$ is continuously differentiable in both arguments. Here, to explain the basic insight as simply as possible, we take the function $SCC(E|\theta)$ as exogenous; Section 3 derives this function from primitives.

We denote the smart tax as $SCC^*(E)$, and obtain its formula using the optimality condition

$$MB(E|\theta) = SCC(E|\theta) \quad \forall \theta.$$

We denote the optimal emissions level as a function of the shock by $E^*(\theta)$, and its inverse by $E^{*-1}(E)$. The smart tax is a function of emissions but not the shock

$$SCC^*(E) \equiv SCC(E|E^{*-1}(E)) \quad \forall E \in \{E|\exists \theta \text{ s.th. } E = E^*(\theta)\}.$$

⁷Here, $E(\theta_s; p^E) = \int_i \frac{\theta_s(i) - p^E}{f_i} di = \int_i \frac{\theta_s(i)}{f_i} di - p^E \int_i \frac{1}{f_i} di$. Defining $f \equiv \left(\int_i \frac{1}{f_i} di\right)^{-1}$ and $\theta_s \equiv f \int_i \frac{\theta_s(i)}{f_i} di$ we obtain $E(\theta_s; p^E) = \frac{\theta_s - p^E}{f}$ and $MB(E|\theta_s) = \theta_s - f E$.

By construction, a market equilibrium that satisfies $MB(E|\theta) = SCC^*(E)$ delivers the first best emissions level at the optimal carbon price.

Proposition 1 compares the slope of the smart tax and the slope of the SCC for an arbitrary realization of the shock. Both of these slopes are evaluated at $E^*(\theta)$. The cases correspond to the relative responsiveness of marginal benefits versus marginal damages (SCC) to a realization of the shock. We introduce the notation MB_θ for $\frac{\partial MB(E|\theta)}{\partial \theta}$, E given, and similarly SCC_θ for $\frac{\partial SCC(E|\theta)}{\partial \theta}$.

Proposition 1 *The slope of the smart tax satisfies*

$$SCC^*_E = \frac{MB_\theta}{MB_\theta - SCC_\theta} SCC_E + \frac{-SCC_\theta}{MB_\theta - SCC_\theta} MB_E, \quad (3)$$

where all functions are evaluated for the same shock realizations θ and emission levels $E^*(\theta)$. Assuming $MB_E < 0$ and $SCC_E, MB_\theta > 0$ we find

- (i) $0 < SCC_\theta < MB_\theta \Rightarrow SCC^*_E > SCC_E$
- (ii) $SCC_\theta = MB_\theta \Rightarrow SCC^*_E = +\infty$
- (iii) $MB_\theta < SCC_\theta \Rightarrow SCC^*_E < 0 (< SCC_E)$
- (iv) $SCC_\theta = 0 \Rightarrow SCC^*_E = SCC_E$ (as in the static setting)
- (v) $SCC_\theta < 0 \Rightarrow SCC^*_E < SCC_E$ (SCC^*_E can be negative).

Figure 2 illustrates the proposition. The solid curves labeled MB and MD show the marginal benefits and the social cost from emitting, given the expected technology level θ . If the realization of θ equals its expected value, the intersection of these curves identifies the optimal emission level. The dashed curves correspond to a lower realization of θ , implying cheaper than expected abatement, e.g., due to an unexpected innovation in green technology. The figure assumes that the shock also reduces future abatement costs, thereby reducing future emissions. Under the assumption of convex damages, the lower future emissions reduce the future marginal damages associated with today's emissions. Thus, the technology shock causes the SCC to shift down to the dashed curve.

In Figure 2, the optimal allocation for the low realization of θ lies to the lower left of the expected allocation. This graphical feature corresponds to case (i) of Proposition 1, where the marginal benefits respond more strongly to the technology shock than do

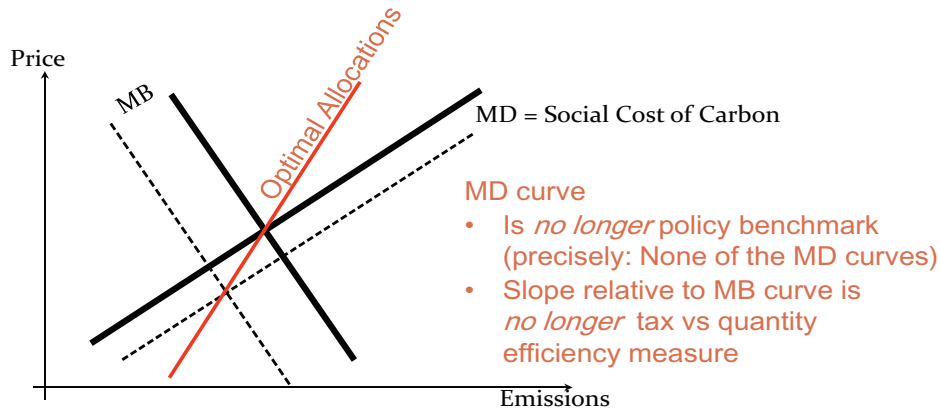


Figure 2: Dynamic setting. The optimal carbon tax (smart tax, red) as a function of the emissions level. The MD curve depicts the marginal damages from emissions, here the social cost of carbon. Black solid lines depict the expected MD and MB curves. Dashed lines depict the case of a better than expected technological innovation. As in the static setting, the innovation shifts down the marginal benefits from emissions curve (abatement cost). In contrast to the static setting, the technological innovation now also shifts down the MD curve: better technology in the future reduces future emissions and, thereby, reduces the marginal damage caused by today’s emissions. The smart tax no longer coincides with (any of) the MD curve.

(expected future) marginal damages. As a consequence, the slope of the smart tax is positive and larger than the slope of marginal damages (SCC).

Figure 3 represents cases (ii) and (iii) of Proposition 1. If the technological innovation shifts the MD curve and the MB curve by the same amount ($MB_{\theta} = SCC_{\theta}$), then the smart tax is vertical (left graph). In this case, cap and trade is optimal regardless of the relative slopes of the MB and the MD curves. If the technological innovation shifts the MD curve even more that it shifts the MB curve ($MB_{\theta} < SCC_{\theta}$), then the slope of the smart tax is negative. In this case, it is optimal to emit more under a lower tax despite the better abatement technology, because the climate change problem has become substantially less bad.

The fourth case in Proposition 1 is analogous to the static case. Here, the SCC curve does not respond to the technological innovation, and it directly gives the smart tax. The fifth case describes the scenario where a shock increases abatement costs but reduces marginal damages. Here the slope of the smart tax is smaller than the slope of the MD curve (and possibly negative).

Our quantitative analysis in the next sections identifies case (i) of Proposition 1 and Figure 2 as the most likely (or at least “base”) scenario in the case of climate

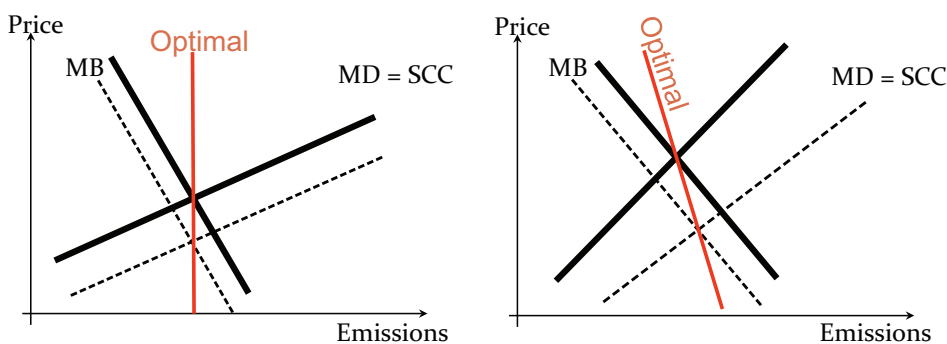


Figure 3: Dynamic setting, analogous to Figure 2. The optimal carbon tax (smart tax, red) as a function of the emissions level. If the MD curves shifts down as much as the MB curve, the smart tax is vertical and a standard cap and trade system is first best (left). If the MD curves shifts more than the MB curve, the smart tax falls with emissions (right). Here, a low emission price signals sufficiently large falls in future abatement costs that it is optimal to increase current emissions.

change. Here, the SCC 's response to the technology shock is smaller than that of the MB curve and the slope of the smart tax is positive, finite, and steeper than that of the SCC curve.

We replace MD with SCC^* in equation (2) to relate the slopes of the smart cap and the smart tax.

$$q'(p) = \frac{1}{SCC^*_E E + SCC^*} \Leftrightarrow SCC^*_E = \frac{p^E}{E} \left(\frac{1}{\varepsilon_{q,p}} - 1 \right) \quad (4)$$

The right side uses the definition of the redemption function's elasticity w.r.t. the certificate price, $\varepsilon_{q,p}(p) = \frac{dq}{dp} \frac{p}{q}$. For cases (i) and (iv) of Proposition 1, the left side of equivalence (4) implies that the smart cap expands with the certificate price. Case (ii) implies a standard cap, one that does not respond to the price. Case (iii) implies a negatively sloped smart tax, and case (v) implies a potentially negatively sloped smart tax. The left side of equivalence (4) shows that a downward sloping smart tax is consistent with a smart cap that increases in the certificate price, provided that the optimal carbon price is sufficiently high. The right side of the equivalence (4) rephrases these cases using the redemption function's elasticity w.r.t. certificate price. A positively sloped inelastic redemption function, $\varepsilon_{q,p}(p) \in (0, 1)$, corresponds to a positively sloped smart tax. Both a positively sloped elastic redemption function and a negatively sloped redemption function correspond to a negatively sloped smart tax.

2.3 Stability and Market Power

This section discusses stability of the competitive equilibrium and market power. Here we assume existence of the smart tax and cap, a subject taken up in section 3. We also assume that the smart cap is defined (only) on the domain of prices for which there exists some technology realization such that the smart cap induces a socially optimal allocation; and we assume $MB_\theta \neq SCC_\theta$. We exclude the case of a vertical smart cap in order to obtain a continuously differentiable redemption function.

Under the Walrasian auctioneer, a market equilibrium (p^*, E^*) is locally stable if excess demand is strictly positive below p^* and strictly negative above p^* in some neighborhood of the equilibrium. It is globally stable if this condition on excess demands holds for arbitrary deviations of the price.

Proposition 2 (i) *A **smart tax** decentralizes the socially optimal emissions level as a locally stable competitive equilibrium if and only if*

- (a) $SCC_E^*(E(p^E)) \geq 0$, or
- (b) $SCC_E^*(E(p^E)) < MB_E(E(p^E); \theta)$,

*evaluated at the equilibrium.*⁸ *It is globally stable if everywhere one of conditions (a) or (b) is met.*

(ii) *A **smart cap** decentralizes the socially optimal emissions level as a locally stable competitive equilibrium if and only if*

$$\frac{MB_E E^* + MB}{SCC_E^* E^* + SCC^*} < 1. \quad (5)$$

at the equilibrium emission level E^ . If $MB(E, \theta)$ is strictly monotonic in the technology level θ (for any given emission level E), then, a smart cap satisfying condition (5) at every social optimum is globally stable if $\epsilon_{q,p}(p) \neq 1$ for all p .*

Under a positively sloped smart tax, both equilibria are stable. If the smart tax is negatively sloped, it is stable if and only if it is steeper (more negative) than the marginal benefit curve. For our dynamic linear quadratic model in section 3 these conditions are always met. In case (b), the smart cap is stable if and only if it has a negative slope.

⁸We are ignoring the case where $SCC_E^*(E(p^E)) = MB_E(E(p^E); \theta)$ at the equilibrium point but $SCC_E^*(E(p^E)) < MB_E(E(p^E); \theta)$ holds everywhere in the neighborhood. This case is also stable by our definition.

We now consider market power. If the smart cap/tax applies to a single sector, firms might be able to exercise market power. We consider the extreme case of a monopsony facing a smart cap as the simplest illustration of the implications of market power.⁹

For the next proposition, we define $\epsilon_{E,p^E} = \frac{dE}{dp^E} \frac{p^E}{E}$, the elasticity of emissions in tons of carbon w.r.t. the emissions price; this elasticity is distinct from the elasticity of the redemption function $\epsilon_{q,p}$ defined earlier.

Proposition 3 (i) *Consider a monopsonist that faces a smart cap $q(p)$. If this monopsony problem is concave with an interior solution, the monopsonist's marginal benefits from emissions satisfy*

$$MB(E) = \frac{p^E}{\epsilon_{q,p}} = p^E \frac{1 + \epsilon_{E,p^E}}{\epsilon_{E,p^E}}. \quad (6)$$

(ii) *Facing the smart cap designed to support the optimal (interior) outcome under competition, a monopsonist increases profits by emitting (weakly) less than the optimal (competitive) amount if $SCC_E^* \geq 0$, and by emitting more than the optimal amount if $SCC_E^* < 0$.*

(iii) *The policy maker can induce the monopsony to emit at the optimal level using a smart cap $q^m(p)$ that solves the following ODE*

$$q^{m'}(p) = \frac{1}{SCC^*(Qq^m(p))} > 0, \quad (7)$$

provided that this ODE has a solution that generates a strictly concave optimization problem. The monopsony's second order condition is locally satisfied in any equilibrium if and only if $MB_E(E^(\theta); \theta) < SCC_E^*(E^*(\theta))$ for all θ .*

Equation (6) is a familiar result. It states that a monopsony chooses the optimal level of an input, here emissions, by setting the marginal benefit of the input equal to the marginal outlay. If the smart cap was designed for a competitive market, the monopsony emits *less* than the competitive level if and only if the slope of the smart tax is positive. In this case, the slope of the smart cap is also positive (equation 4).

We can compare the optimal redemption functions under monopsony and competition by comparing the ODEs that the two function satisfy, equations (4) and (7). We illustrate this procedure using the case $SCC_E^* \geq 0$. Here, the redemption function designed to counter market power is steeper than the redemption function under

⁹We do not consider a monopsony who faces a *smart tax* because this case is standard: the smart tax is the inverse supply of emissions.

competition. For a given change in the price of certificates, the smart cap expands and contracts more strongly to offset market power. Why? Here the monopsony exercises market power by reducing emissions in order to lower the certificate price. The regulator discourages this behavior by making it more expensive, in terms of reduced emissions, for the monopsony to achieve a given reduction in the certificate price.

For example, suppose that the range of optimal emissions is $[E^L, E^H]$, with corresponding domain of the competitive certificate price $[p^L, p^H]$. If we choose the boundary condition for the monopsony redemption function to satisfy $q^{-1}(\frac{E^L}{Q}) \leq p^L$, then the redemption function under monopsony lies below and is flatter, in the (E, p) plane, than the redemption function under competition. For every realization of the shock, the monopsony emits at the optimal level but pays a lower price.

3 The Dynamic Model

This section uses a dynamic version of Weitzman's (1974) familiar static linear-quadratic model. The full-information *SCC* increases with emissions, but the smart tax might either increase or decrease in emissions. The smart tax implements the full-information (first best) level of emissions as a unique stable competitive equilibrium. The stability condition under the smart cap is more delicate. The smart tax provides an extremely simple way of expressing the welfare ranking of the standard tax and quota, one that exactly parallels Weitzman's ranking for the static model. We also examine certificate trading across periods and quantify the smart cap and smart tax.

3.1 Model and Analytic Results

We measure the pollution stock S_t at the beginning of period t by its deviation from the zero-cost level (e.g., the pre-industrial level of GHG). The stock of pollution at the end of the period is

$$S_{t+1} = \delta S_t + E_t,$$

where the parameter δ , $0 < \delta \leq 1$, measures the pollutant's persistence.

At the beginning of period t , the policy maker and all firms know the value of the random variable θ_{t-1} . Firms, but not the policy maker, then observe the innovation $\varepsilon_t \sim iid(0, \sigma^2)$. We move straight to the model of the representative firm because

the market for certificates aggregates the individual shocks as described at the end of section 2.1. Thus, we have the equation of motion

$$\theta_t = \rho\theta_{t-1} + \varepsilon_t,$$

with shock (or technology) persistence $0 < \rho \leq 1$.¹⁰ The realization of ε_t alters the marginal benefit of emissions, via a change in technology affecting emissions intensity, or a change in economic activity affecting emissions demand.

Only a fraction, $0 < \alpha \leq 1$, of this innovation is embodied in the current period, so firms in period t operate with technology level

$$\hat{\theta}_t = \rho\theta_{t-1} + \alpha\varepsilon_t.$$

We can interpret α as a share of firms adopting the new technology in the current period, as in the literature on technology diffusion (Rogers 1995). More generally, a higher α represents a quicker response of firms to the shock.

The benefits from emissions (or abatement costs) depend linearly on the technology level,

$$B(E_t, \hat{\theta}_t, t) = (h_t + \hat{\theta}_t) E_t - \frac{f}{2} E_t^2,$$

so the marginal benefits of emissions are $\frac{\partial B}{\partial E} = h_t + \hat{\theta}_t - fE_t$, with $f > 0$. Hereafter, we assume that the marginal benefits at zero emissions, $h_t + \hat{\theta}_t$, and the full-information (first best) level of emissions are both positive with probability one. Flow damages are quadratic in the pollution stock

$$D(S_t) = \frac{b}{2} S_t^2,$$

with $b > 0$. The policy maker with discount factor $0 < \beta < 1$ maximizes

$$\mathbb{E}_t \sum_{s=t}^{\infty} \beta^{s-t} \left(\left(h_s + \hat{\theta}_s - \frac{1}{2} f E_s \right) E_s - \frac{1}{2} b S_s^2 \right).$$

The policy maker is aware that future optimal emission policies depend on the future realizations of the state variables.

To derive the smart tax and cap, we first solve the full information optimum, where the social cost of carbon, SCC , is a linear function of the stock of carbon and

¹⁰For $\rho = 0$ the social cost of carbon is independent of θ_{t-1} , and the separation implies similar results as in the static model. We ignore the empirically less relevant case $\rho < 0$.

the technology level (see equation 21 in Appendix C). An additional unit of emissions produces an additional unit of the next-period stock, so the derivative of the SCC w.r.t. current emissions is the same as the derivative w.r.t. next-period stock. This derivative is positive because damages are convex. The full information SCC depends on the innovation ε_t . The smart tax, SCC^* , in contrast, is independent of this innovation.

Proposition 4 (i) *The smart tax is*

$$SCC_t^* = A_0 S_t + A_1 \theta_{t-1} + \gamma E_t + a_t. \quad (8)$$

(ii) *The smart tax's emissions' slope, γ , can take either sign. There exists $\alpha^* \in (0, \beta)$ such that for $\alpha > \alpha^*$*

$$\gamma = \frac{\partial SCC_t^*}{\partial E_t} > \frac{\partial SCC_t}{\partial E_t} > 0, \quad (9)$$

and for $\alpha < \alpha^*$

$$\gamma = \frac{\partial SCC_t^*}{\partial E_t} < 0 \text{ and } \frac{\partial SCC_t}{\partial E_t} > 0. \quad (10)$$

For $\alpha = \alpha^*$, the slope of the smart tax is infinite, and a conventional cap and trade achieves the first best emission allocation. As α passes through α^* (from below), the slope of the smart tax switches from $-\infty$ to $+\infty$, and for $\alpha > \alpha^*$ the slope of the smart tax decreases continuously in α .

(iii) *The smart tax supports the optimal level of emissions as a globally stable competitive equilibrium for all $\alpha \in (0, 1]$, i.e., for both positive and negative γ .*

Proposition 4 shows that our dynamic model can produce cases *i* – *iii* of Proposition 1.¹¹ For $\alpha \approx 1$, a positive shock ε causes a larger increase in the marginal benefit of emissions than in the social cost of carbon ($MB_\varepsilon > SCC_\varepsilon$) and the smart tax is steeper than the SCC (case *i* in Proposition 1). In this case, a positive shock increases the optimal emissions level. For α small, a shock has little effect on the present period's marginal benefit of emissions, but a non-negligible effect on the SCC . In this case, the smart tax has a negative slope (case *iii* in Proposition 1). Here, a positive shock lowers the optimal level of emissions. Finally, if $\alpha = \alpha^*$, the shock equally affects marginal

¹¹The proof gives the formula for α^* and for the functions A_0 , A_1 , γ and a_t ; a_t depends on time because of the trend h_t . The functions A_0 and A_1 , like γ , switch signs at α^* .

benefits and damages from emissions, and a conventional cap and trade-system achieves the first best emission allocation (case *ii* in Proposition 1).

The proposition also shows that a negatively sloped smart tax requires $\alpha < 1$; here the technological innovation is observed but not fully implemented in the current period. Only then does the market price carry sufficient information to affect future marginal damages from emissions more than present marginal benefits. Regardless of the value of α , the smart tax implements the social optimum as a stable competitive equilibrium.

We now consider the smart cap. To simplify notation, we define $\hat{A}_t \equiv A_0 S_t + A_1 \theta_{t-1} + a_t$, thereby collecting all of the time-dependent variables in the formula for the smart tax, apart from the current emissions level, E_t . With this definition, the smart tax is $SCC_t^* = \hat{A}_t + \gamma E_t$. We follow the same logic as in Section 2.1. The firm's price of a unit of emissions is $p_t^E = \frac{p_t}{q_t(p_t)}$. We construct the smart cap so that it implements the optimal level of emissions, i.e., we set $\frac{p_t}{q_t(p_t)} = \hat{A}_t + \gamma E_t$. The subscript on q_t serves as a reminder that the redemption function depends on time via the function \hat{A}_t . We have

Proposition 5 *If $\gamma \geq 0$, the redemption function*

$$q_t^+(p_t) = \frac{1}{2\gamma Q} \left(-\hat{A}_t + \sqrt{\hat{A}_t^2 + 4\gamma Q p_t} \right)$$

implements the first-best emission level as a stable competitive equilibrium. This redemption function increases in the price of certificates.

If $\gamma < 0$ the redemption function

$$q_t^- = \frac{1}{-2\gamma Q} \left(\hat{A}_t + \sqrt{\hat{A}_t^2 + 4\gamma Q p} \right)$$

implements the first-best emissions level as a stable competitive equilibrium on the domain $p_t \in [0, \frac{\hat{A}_t^2}{-4\gamma Q}]$, with the range of emissions $E_t \in [\frac{\hat{A}_t}{-2\gamma}, \frac{\hat{A}_t}{-\gamma}]$. Here, the redemption function decreases in the certificate price.

Consistent with Proposition 2 for the general case, the smart cap is stable and increases with the certificate price when the smart tax increases. Proposition 5 gives the precise form of the optimal smart cap. For $\gamma < 0$, where the smart tax decreases, the set of emissions levels supported by the smart cap equals the set where $SCC^* > 0$. This is the range satisfying the local stability inequality (5).

The market response to either the smart tax or the smart cap enables the regulator to recover one piece of hidden information. Optimal policy depends on the persistence of this shock, measured by ρ , and the speed at which firms respond to the shock, measured by α . We can interpret the shock as being related either to technology or to the business cycle. A technology shock tends to be genuinely asymmetric information. Firms' response to the policy reveals the hidden technology shock. The macro shock is unknown when the regulator announces the smart tax or smart cap, but observed by both firms and the regulator during the compliance period. Therefore, the policy can be conditioned on the macro shock. If the macro shock is iid, we do not need to modify the model presented above. More plausibly, if expectations of the macro shock depend on information such as current and lagged macro conditions, those variables become part of the information set. The full information SCC and both the smart tax and the smart cap then depend on those variables, but the structure of the policy does not change.

A famous result, due to Weitzman (1974), states that in the linear-quadratic model with additive shocks and a flow pollutant, the standard tax welfare-dominates the standard quota if and only if the slope of marginal damages is less than the slope of marginal abatement costs (equal to the slope of marginal benefit of emissions). The literature reviewed in Footnote 3 studies the more complicated welfare comparison between the standard tax and quota for a stock pollutant. The smart tax provides a novel and intuitive link between the models with flow and stock pollutants. Section 2.1 notes that the marginal damage function coincides with the smart tax for a flow pollutant. Thus, for a flow pollutant we can restate Weitzman's result as "For a flow pollutant, taxes welfare-dominate quotas if and only if the slope of the smart tax is less than the slope of marginal abatement costs". The same comparison holds with stock pollutants:

Proposition 6 *With stock pollutants, (standard) taxes welfare-dominate (standard) quotas if and only if the slope of the smart tax, γ , is less than the slope of the marginal abatement cost, f .*

3.2 Inter-period trading and optimality

At least eight cap and trade programs, including California's Low Emission Vehicle Program, the EPA's SO₂ and NO_x programs, and the EU's Emissions Trading Scheme,

allow intertemporal banking of permits (Holland & Moore 2013). Intertemporal trading can smooth carbon price fluctuations triggered by technology or growth shocks, potentially increasing welfare. A smart cap does not require intertemporal trading, because emissions respond to shocks optimally by construction. Intertemporal trading may nevertheless be relevant if the commitment phase is long or if the institutional framework does not permit conditioning the smart cap on macroeconomic indices.

We assume that the regulator allows interperiod trading within a *commitment phase* lasting T periods. For example, each period might last for one year and a commitment phase may be a decade ($T = 10$). Firms can trade certificates across the T periods; the certificate cap, Q , applies to the entire *commitment phase*. Later commitment phases might have a different number of periods, but we assume that: policy is set optimally in the future; the horizon is infinite; and the certificates of the current phase cannot be used in later phases. In the model above firms do not have to make intertemporal decisions. Keeping with this setting, we assume the existence of a risk neutral arbitrageur whose actions make firms indifferent about the timing of emissions and certificate purchases.

To avoid the need for double-subscripts, we consider the case of the first T -period commitment phase, with the initial period set at $t = 1$. The smart cap's redemption function in period t , $q_t(p_t)$, determines the exchange ratio q_t between certificates and carbon dioxide emitted in period t . For E_t emissions in period t , the representative firm has to deliver $\frac{E_t}{q_t}$ certificates at the end of the commitment phase. Market clearing requires $\sum_{t=1}^T \frac{E_t}{q_t} = Q$, where Q is the total number of certificates for this commitment phase. Focusing on the fundamental issues of inter-period trading, we assume that innovations are immediately adopted: $\alpha = 1 \Rightarrow \hat{\theta}_t = \theta_t$.

Proposition 7 *There exists a sequence of redemption functions $q_t^{p_1, \dots, p_{t-1}}(p_t)$, $t \in \{1, \dots, T\}$, and an allocation of certificates $Q(p_1, \dots, p_{T-1})$ supporting the optimal emissions trajectory as a decentralized intertemporal equilibrium.*

The aggregate number of certificates for this commitment phase depends on the sequence of certificate prices within that phase. If the aggregate number Q was fixed, the certificates remaining at the beginning of period T would be stochastic. However, a given redemption function q_T achieves the first best allocation only for a specific number of certificates. Therefore, the aggregate number of certificates has to depend on the earlier prices in order to guarantee that the number of certificates remaining in

period T , together with the redemption function in that period, support the optimal emissions level.

Under banking and borrowing – in both a standard and a smart cap – intertemporal arbitrage implies that the price of certificates has to rise at the rate of interest. An emissions price growing at the rate of interest is generally not optimal, which is an issue for standard emission trading schemes. Proposition 7’s period-dependence of the redemption function, i.e., the exchange ratio between emissions and certificates, decouples the price increase of emissions from the intertemporal arbitrage condition to achieve first best.¹²

Proposition 7 conditions the redemption functions on the certificate price in earlier periods. Such conditioning allows the mechanism to incorporate the carbon stock fluctuations resulting from the sequence of technology shocks over the course of a commitment phase. Over a fairly short commitment phase, e.g. a decade, the stock of carbon is likely to vary much less than the technology variable. Then, it seems reasonable to neglect a conditioning of the redemption functions on the earlier period’s price realizations.

3.3 Quantification

We use our results to study global climate change. As the introduction notes, many countries are either planning to use or currently using taxes or cap and trade systems to reduce their CO₂ emissions. We quantify the smart tax and cap for the case of global cooperation.

Output, Abatement, and Emissions. Global world output in 2020 is 130 trillion USD using purchasing power parity weights (IMF 2020). We use Nordhaus & Sztorc’s (2013) DICE model to estimate the 2020 marginal abatement cost slope as $f = 2.5 * 10^{-9} \frac{USD}{tCO_2^2}$. Much of our analysis depends only on the slopes of the marginal abatement cost and marginal damage curves. The absolute levels of the social cost of carbon also depends on $h = 101 \frac{USD}{tCO_2}$, the intercept of marginal abatement costs. We assume that this intercept falls exogenously by 1% per year.¹³ This calibration implies a business

¹²In a standard cap and trade system with deterministic technological change (Kling & Rubin 1997) or uncertainty about abatement costs (Yates & Cronshaw 2001) it is not optimal that the emissions prices grows at the consumption discount factor. The stock pollutant creates additional reasons for the optimal expected marginal abatement cost to vary over time. In these circumstances, the literature suggests using certificate discount factors. Our redemption function $q_t(p_t)$ already decouples certificate prices from absolute emissions and incorporates such discount factors.

¹³These values derive from the optimized DICE 2013 run for the year 2020. We set the expected

as usual emission level of $E^{BAU} = 40 \text{ GtCO}_2$, implying that we abate a few percent of BAU emissions in 2020.

Technology diffusion. We obtain an estimate (or guesstimate) of the technology diffusion parameter α by regressing US CO₂ emissions in 1995-2010 against (stocks and flows of) green patents on. We assume a 5 year commitment periods and $\rho = 1$, i.e., no decay of innovations (patents). We restrict attention to “major” green patents, those registered in all three major patent offices, United States, Europe, and Japan. We summarize details in Appendix A.¹⁴ Our preferred estimate lies slightly above $\alpha \approx 1/4$; about one quarter of the long-run impact of the innovation shocks occur within the current commitment phase. Other relevant innovations, which are not being patented, might be adopted faster leading to a somewhat higher overall adoption share α . We present results for $\alpha \in \{0.25, 0.5, 1\}$.

Climate. We use the model of transient climate response to cumulative emissions (TCRE) to calibrate climate dynamics. Recent climate modeling shows that average global atmospheric temperature can be well-approximated as a linear function of cumulative historic emissions. The consensus report IPCC (2013) states that the proportionality factor between cumulative emissions and temperature, TCRE, is likely in the range between 0.8°C and 2.5°C for each 1000 GtC (10^{12} tons of carbon). We use the mid-value TCRE, $1.65 * 10^{-15} \frac{^\circ\text{C}}{\text{GtC}}$.¹⁵ Our state variable, S_t , is cumulative historic emissions, which are proportional to temperature; the persistence factor is $\delta = 1$.

We briefly comment on the intuition of the TCRE model. In the actual climate system, most carbon dioxide emissions are eventually removed from the atmosphere, but each emission unit has a cumulative impact on temperature over time through its greenhouse effect. Scientific models of climate change find that the removal of carbon from the atmosphere and the delayed warming response to an increase in carbon *concentrations* approximately cancel each other, making cumulative historic emissions a good proxy for temperature.

value of the technology shock in the present period to zero, thereby making the calibration results independent of α and ρ .

¹⁴Our preferred estimate is $\alpha = 0.28$. Extending the time series of our estimation further back to 1990 renders the time series nonstationary and delivers the slightly higher estimate $\alpha = 0.34$. Going back to 1985, the coefficients lose significance and the coefficient on new patents has the wrong sign. Controlling for oil prices and allowing for a break in the quadratic trend does not regain significance, but reasonable coefficient estimates that imply $\alpha = 0.27$. Controlling for oil prices and introducing a flexible break point in the trend does not affect the estimates of the shorter time series.

¹⁵The TCRE is usually expressed w.r.t. tons of carbon (C), which is how we cite it here. However, our other values follow the convention expressing the SCC in USD per ton of carbon dioxide (CO_2).

Damages. DICE assumes no damages at the pre-industrial temperature level and global damages of approximately 1% of world output at a $2^\circ C$ warming (Nordhaus & Sztorc 2013). Our *baseline* calibration of the damage function uses this assumptions, producing $b_{base} = 1.3 * 10^{-13} \frac{USD}{tCO_2^2}$. We also introduce a “*concerned*” scenario that assumes today’s damage from global warming is zero, but a $3^\circ C$ warming causes a loss of 5% of world output. This scenario implies a more convex damage function with $b_{concerned} = 6.6 * 10^{-13} \frac{USD}{tCO_2^2}$. We can also interpret this scenario as reflecting concern about tipping points.

Expected optimal SCC. We test our calibration by calculating the implied optimal carbon tax under the expected technology realization. At the optimal emission allocation, the smart tax equals the SCC by construction. For an annual rate of pure time preference (rptp) of 1.5% ($\beta = 0.985$) we obtain an optimal carbon tax of $26 \frac{USD}{tCO_2}$. This tax is a little higher than in DICE, which has recently been discovered to exaggerate the temperature delay in warming (a feature we avoid by using the TCRE model). Reducing the rptp to 0.5% ($\beta = 0.995$), the median response of Drupp, Freeman, Groom & Nesje’s (2018) expert survey, approximately doubles this tax ($55 \frac{USD}{tCO_2}$). These values suggest that the model calibration is reasonable. The corresponding optimal emission levels are $E^{opt} = 29 GtCO_2$ for $\beta = 0.985$, and $E^{opt} = 18 GtCO_2$ for $\beta = 0.995$. Under the 1.5% rptp, the concerned scenario using the more convex damage function increases the tax only mildly to $30 \frac{USD}{tCO_2}$.¹⁶

Results Base Calibration. Figure 4 presents the smart tax and cap assuming a five-year commitment period and immediate adoption of the new innovation ($\alpha = 1$). The left panel graphs the smart tax as well as the SCC and the marginal benefits from emissions under the expected technology realization. By construction, all the lines intersect at the expected price and emission levels. For other realizations of technology, the equilibrium moves along the smart tax. We observe that (i) the smart tax is substantially steeper than the SCC curve and (ii) the (absolute of the) MB-curve’s slope is greater than the slope of the smart tax. By Proposition 6, taxes are preferred over quantities in this baseline scenario with $\alpha = 1$.

The smart cap shown on the right of Figure 4 eliminates the welfare loss of a tax. To make it easy to compare the smart tax and the smart cap, we depict the overall (global) cap in $GtCO_2$. We set the number of certificates, Q , equal to the optimal emission level

¹⁶The concerned scenario reduces current damages but increases the damages resulting from a higher level of global warming. The optimal expected carbon tax assumes optimal future mitigation policy and, thus, temperature is unlikely to reach high levels.




Figure 4: Calibrated Smart Tax and Cap. Left: SCC and marginal benefits of emissions (MB) under the expected technology realization as well as smart tax (independent of technology realization). Right: Smart Cap, which is the redemption function times expected emissions.

under the expected technology realization. With this choice, the certificate price under the expected technology realization coincides with the smart tax of $26 \frac{USD}{tCO_2}$. Greener than expected technological progress, causing a downward shift in the demand for emissions (the MB curve), leads to a lower certificate price and a contraction of the smart cap. Similarly, less green technological progress increases the certificate price and expands the smart cap. The redemption function's graph is identical to that of the smart cap once we change the scale on the vertical axis from aggregate emissions to the emission level per certificate.

Figure 5 varies the speed of firms' technology adoption, with the solid graphs replicating those of Figure 4, where $\alpha = 1$. The dashed graph uses our preferred estimate $\alpha \approx 0.25$, where only one quarter of firms adopt the new technology innovations within the 5-year commitment period. The reduced speed of adoption substantially increases the slope of the smart tax and flattens the slope of the smart cap, which graphs emissions over price rather than price over emissions. We note that the slope of the dashed smart tax exceeds that of the MB-curve (depicted in Figure 4); thus, by Proposition 6 quantities dominate taxes for $\alpha = 0.25$. The dash-dotted line of $\alpha = 0.5$, assuming that half of the firms adopt the new innovation during the 5-year commitment period. This value represents that less fundamental non-patented innovations might also be adopted more quickly, increasing α . For $\alpha = 0.5$, the smart tax and the MB-curve have almost the same slope; here, the welfare difference between a tax and a standard cap is close to zero.

Concerned Scenario. Figure 6 presents the results for the *concerned* scenario,




Figure 5: Variations of speed of technology adoption. Immediate full adoption ($\alpha = 1$), half of firms adopt during a 5-year commitment period ($\alpha = .5$), one quarter of firms adopts during a 5-year committeemen period ($\alpha = .25$). Left: Smart tax. Depicted SCC assumes expected technology realization. Right: Smart Cap.

where damages are more convex (initially lower and then higher). The smart tax and the SCC under the expected technology realization are higher than in the baseline. They increase substantially faster for lower than expected technological progress, because the resulting higher future emissions increase damages more strongly with more convex damages. Similarly, higher than expected green progress reduces the equilibrium prices more strongly; here, a reduction in the future CO₂ stock implies a stronger reduction of future damages than in the baseline. The smart cap, a function of the certificate price, shows the same qualitative features as a function of the certificate price. Reducing the speed of technology adoption, α , rotates the smart tax graph counter-clockwise (making it steeper) and the smart cap graph clockwise (making it less steep). For $\alpha \lesssim 0.4$, the smart tax and cap have negative slopes (see as well right graph in Figure 7).

In particular, our estimate $\alpha \approx 0.25$ implies a decreasing smart tax and cap. Here, a higher than expected green technological progress not only lowers the cost of abatement, but also reduces long-term damages sufficiently that it is optimal to respond with both a price reduction and an emissions increase (moving down on the smart tax graph and up on the smart cap graph). The planner knows that most of the improved technology will be adopted in the next period, lowering future emissions and the marginal damage associated with current emissions. Similarly, less green progress increases the emissions price and, given the damage convexity, urges us to cut more emissions.




Figure 6: Concerned scenario. Smart tax (left) and smart cap (right) assuming a higher damage convexity. The different curves correspond to immediate full adoption ($\alpha = 1$), half of firms adopting ($\alpha = .5$), and one quarter of firms adopting during a 5-year commitment period ($\alpha = .25$).

Given the novelty of the finding, it merits discussing another variation of the intuition. Lower than expected green progress is bad news for both firms and the environment. In the baseline scenario (or, here, for $\alpha \geq 0.4$), the optimal policy uses the environment to smooth shocks to the firms; if abatement turns out to be very expensive, we allow firms to emit more. However, if damages are sufficiently convex and α is low, the future environmental damage implied by the lack of green progress is too costly to tolerate such smoothing at the expense of the environment. Instead of using costs to the environment as a substitute for costs to the firms, the policy maker now treats them as complements. Under bad news we increase the unit price *and* cut emissions. Conversely, under good news, we lower the price and permit firms to emit more.

Reduction of Time Preference. The left graph in Figure 7 reduces the rate of pure time preference from an annual 1.5% to 0.5% in the base scenario. The implications are qualitatively similar to those observed in the previous variation with more convex damages. Here, the policy maker pays more attention to future damages. As a result, the SCC under the expected realization increases substantially and the smart tax rotates counter-clockwise for any speed of technology adoption. An adoption share of $\alpha = 0.5$ during the 5-year commitment period makes the smart tax vertical and the smart cap horizontal (not shown). Under these assumptions, the smart cap corresponds with the classical cap and the ordinary cap and trade system reaches first best.

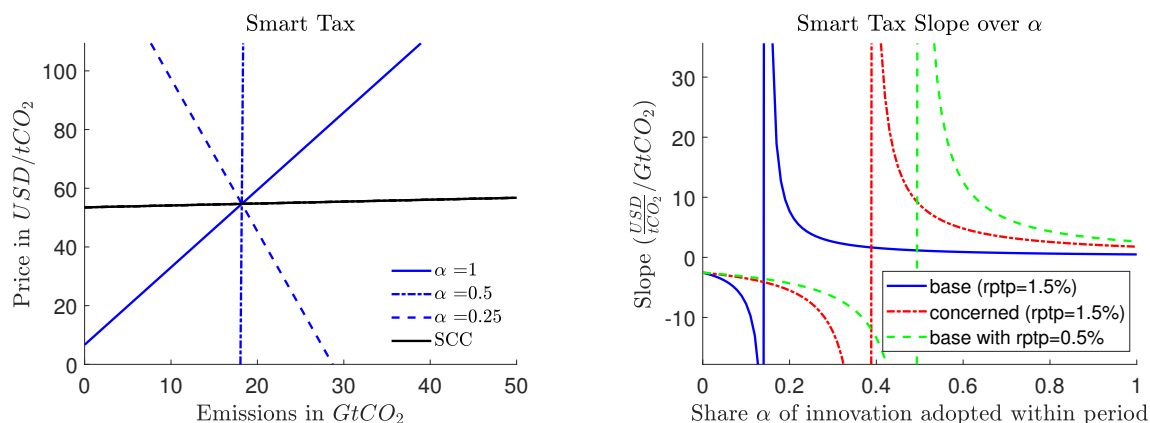


Figure 7: Left: Smart tax under the baseline calibration but with a reduced rate of pure time preference (0.5% instead of 1.5%). The share α of immediate adopters varies in three discrete steps from one to one quarter. Right: Slope of the smart tax. The share α of immediate technology adopters varies continuously along the horizontal axis and the three curves correspond to the different scenarios.

Slope over Adoption Share. The right panel of Figure 7 plots the slopes of the smart tax over the “speed of adoption”, i.e., the share of firms that adopt α within the 5-year commitment period. Starting from the right, we observe that the smart tax is most sensitive to emissions under the reduced discount rate and more sensitive in the concerned scenario than in the baseline. This difference in slope (sensitivity) increases as we reduce the share α . The vertical lines identify the values of α at which the slope of the smart cap flips sign: $\alpha \approx 0.5$ for the rate of pure time preference of 0.5% (green dashed), α just below 0.4% for the concerned scenario (red dash-dotted), and in the base scenario the adoption share within the commitment period would have to fall all the way to $\alpha = 0.14$ (half or our preferred estimate) to turn a standard cap first best.

4 Practical Implementation of a Smart Cap

This section discusses the practical implementation of the smart cap and some easily implemented compromises to improve efficiency in pre-existing cap and trade systems. In the real world, (i) business cycles have a major impact on emissions and certificate prices, (ii) information is revealed continuously over the course of a commitment phase and certificates are traded continuously and (iii) political institutions tend to favor simplicity and minimal change. While the smart cap can help with point (i), we repeat that it is better to deal with this issue by explicitly conditioning the (smart or standard) cap on GDP or alternative business cycle indicators. Thus, this section is

mostly concerned with points (ii) and (iii). That said, much of this discussion also applies to cost shocks generated by business cycles or other sources of price shocks.

Section 3.2 explains how a sequence of announced redemption functions can achieve or improve efficiency when shocks and trading occur repeatedly during a commitment period. For example, we can choose annual (or monthly) compliance periods with annual (or monthly) redemption functions. Proposition 7 would motivate a dense sequence of redemption functions that respond directly to preceding price realizations. As we noted, this conditioning enables the smart cap to respond to the small fluctuations of the CO₂ *stock* during a commitment period, but these are unimportant during a five-year period. Thus, we recommend an annual redemption function using a weighted average of carbon prices over the course of the year. Certificates will be traded throughout the year and beyond as certificate delivery is usually required only a few months after the end of the period.¹⁷ The annual redemption functions would be announced for a 5 year commitment period, changing primarily to reflect expected technological progress and economic growth.

The smart cap's certificates are not in units of CO₂. Similar trading is already common in fishery regulation. The smart cap's certificates correspond to individual fishing quotas, giving their owner claim to a share of the total allowable catch (here: emission level). In regulating fisheries, the regulator sets the total allowable catch period by period. In the smart cap, the total emission level is determined endogenously to address the asymmetric information problem. Both markets trade shares of a pie of varying size. The firm's burden in forming expectations about future emission prices is higher than in the standard cap, but it is present in both. A specialized arbitrageur can reduce this burden by selling claims in units of CO₂ to the firms. We note that expectation formation under the current market stability reserve in the EU ETS is also very complicated; there, estimates over future carbon prices diverge substantially (Perino 2018, Silbye & Birch-Sørensen 2019).¹⁸

If a market for flexible certificates is not politically acceptable, there are more conservative approaches that can incorporate much of the smart cap's efficiency gain while

¹⁷As in the standard cap, non-compliance is subject to monitoring and fines. Given the slightly more sophisticated market clearing conditions, one could consider convex fines, increasing non-linearly in the certificate gap. In case market clearing fails, one could permit firms to submit next period certificates, which can be discounted resembling a fine.

¹⁸Most emissions trading systems with banking and borrowing, including the EU ETS, are unclear about the terminal conditions that are important in determining the certificate price. This ambiguity complicates firms' long-term planning. Because the smart cap does not require banking, it reduces this source of uncertainty.

keeping certificates labeled in units of CO₂. In a first simplified alternative, the regulator can make the current period's cap a function of last period's closing (or average) price. The redemption function would expand or contract aggregate emissions with a short delay. Including banking and/or borrowing – perhaps with some discounting of previous period's certificates – could help to incorporate future adjustments into the present period's expectations and actions. In a second alternative, the regulator can use existing auction systems to implement a simple analogue of the smart cap's redemption function during each auction. Auctions in the EU ETS and other cap and trade systems require firms to submit demand functions, permitting the auctioneer to determine the equilibrium price and settle the certificate distribution. Such an auction can readily match demand with the redemption function or smart tax rather than a fixed amount of certificates. Here, the redemption function essentially becomes an offer curve. A third alternative, using an even smaller change to existing markets, makes the number of certificates auctioned at a given date depend on the price of the previous auction(s). Then, auctions would provide fixed quantities. However, expectations would already respond immediately to the (slightly) lagged quantity response; e.g., auctions in the ETS usually take place every two weeks. We emphasize that such offer curves or delayed response functions should rely on the smart tax rather than the marginal damages or the SCC curve. The disadvantage of this approach is that previously sold certificates do not respond to price signals, requiring that new auctions respond more strongly and possibly limiting their leverage. However, any of these approaches can provide substantial efficiency improvements while keeping the system as close as possible to existing forms of the standard cap. We also note that all of these suggestions are quantity-based regulation, and thus can be implemented and changed by simple majority in the EU ETS; in contrast, a carbon tax (price instrument) requires unanimous approval.

Policy groups and lobbyists strongly influence policy. Environmental organizations and citizen groups that favor strong climate policy have little tolerance for the low, and much lower-than-expected, carbon prices that emerged in many carbon trading systems during the past decade. Firms are afraid that unforeseen shocks can cause certificate prices to increase steeply above expected levels. Here, the smart cap provides a natural compromise – at least under the assumption of an increasing redemption function. If abatement turns out to be cheaper, the policy instrument automatically ramps up reduction efforts and prevents the price from falling too much. If the certificate price

threatens to “go through the roof” in a standard cap, the smart cap would expand, lowering the pressure on firms. Hence, a smart cap is not only a more efficient policy instrument, but it also promises to ease political compromises across different interest groups. Outside of the economic discipline, cap and trade systems, including hybrid system, are criticized for eliminating the moral incentive to reduce emissions: if an environmentally conscientious actor felt compelled to reduce emissions for non-pecuniary reasons, cost-minimizing actors would crowd out those reductions and eliminate the non-pecuniary incentive.¹⁹ In contrast, the smart cap rewards morally-motivated emissions reductions by increasing the emissions price and reducing aggregate emissions.

5 Discussion

We introduce a new cap and trade system that efficiently controls stock pollution under asymmetric information about firms’ abatement technologies. The prime application is the mitigation of climate change. To date, cap and trade-systems are the main market-based approach for the regulation of greenhouse gas emissions. Recent years have exposed major inefficiencies in standard cap and trade systems’ response to cost shocks. The smart cap’s ability to endogenously respond to shocks by optimally relaxing or tightening the cap reduces the cost of greenhouse gas mitigation. The smart cap is also more efficient than a standard tax, while building on existing institutions and maintaining the political advantages of established cap and trade systems.

Climate change is a stock pollution problem. The literature has repeatedly pointed out that the SCC as a function of emissions is relatively flat. In many settings, the SCC curve is the stock analogue of the marginal damage curve for a flow pollutant. Therefore, building on Weitzman’s (1974) reasoning for a flow pollutant, a common conjecture is that optimal emission prices should respond relatively little to shocks, whereas CO₂ emissions should be very responsive. If this conjecture were correct, the optimal smart cap would be very elastic. It would be similar to a standard tax, which would unambiguously dominate the standard cap and trade. We explain why this intuitive argument is wrong for the case of climate change. The SCC’s slope w.r.t. emissions does not represent the optimal equilibrium price change. A shock to the abatement technology has a persistent impact on the emission flow, thereby

¹⁹Jarke & Perino (2017) explain that under incomplete coverage of the cap, inter-sectoral leakage can translate individual effort into an overall emission reduction. Yet, they also show that it can lead to an overall emissions increase under different circumstances.

changing the future emission stock and future marginal damages. As a result, lower than expected green technological progress implies not only higher current marginal abatement costs for firms, but also a higher SCC. Permitting firms to increase emissions under an unfavorable realization of the technology shock also comes at an increased cost to the environment. As a result, optimal emissions (the smart cap) should be less elastic than the slope of the SCC suggests.

We present a simple general model as well as a linear-quadratic quantitative dynamic integrated assessment model of climate change. We show that the share of firms adopting technological innovations within a given commitment period crucially determines the optimal responsiveness of emissions to changes in the market price of certificates. If a smaller share of firms adopts new innovations immediately, then observed shocks in the certificate market have stronger persistent implications for future adopted technology, emissions, and for social costs. A smaller share of adoption within a commitment period flattens the smart cap and steepens the optimal price response to emissions. We call that price response to emissions the smart tax. It is the optimal non-linear emissions tax, i.e. it is society's inverse supply function for emissions. We show that this smart tax, rather than the SCC, permits extending Weitzman's (1974) intuitive reasoning from a flow pollutant to a stock pollutant. A standard tax dominates a standard cap and trade if and only if marginal abatement costs are steeper than the smart tax curve.

In most circumstances, including our baseline calibration, a policy maker uses emission levels to smooth the cost shocks to firms, resulting in an upward sloping relation between emission price and emission level, i.e., upward sloping smart tax and cap. We characterize stability of equilibria and show that, whenever the smart tax is upwards sloping, a monopsony would use the market power to reduce emissions. However, the smart tax and cap might slope downwards if, e.g., damages are more convex or we reduce the pure rate of time preference from 1.5% to a recent (median) expert suggestion of 0.5%. Then, an unfavorable technology realization increases the SCC sufficiently that emissions have to be reduced even under a higher marginal abatement cost; the optimal price-emission relationship turns negative.

Current mitigation levels are rarely optimal. However, the smart cap enables society to abate more at a lower cost. In addition, a smart cap allows the regulator to balance carbon price and emission targets. Several sectors argue that the risk of a high carbon price hurts their economic competitiveness. Consumers and environmental interest

groups fear that a low carbon price leaves cheap abatement options on the table. The smart cap equips the regulator with a compromise: if abatement turns out to be cheap, we abate more, but if it turns out to be expensive, we abate less. Thus, we believe that the smart cap is not only more efficient as a policy instrument, but also helpful in reaching a compromise across different lobbies. The smart cap's ability to endogenously contract also addresses a criticism that non-economists frequently raise against the classical cap. If an individual or a firm reduces their emissions out of a moral obligation, these reduction would be perfectly crowded out by other emission sources. With a positively sloped smart cap, those actions would still be rewarded, even if not to the full extent.

The time horizon of setting and revising caps or taxes is historically long, around 5-10 years. Often these adjustments track international negotiations, which have proven even more inert. On these time horizons uncertainties about green technological progress, economic growth, and global convergence become even more important. To date, climate negotiations have focused entirely on quantity targets. Our paper's insights also emphasize the relevance of a "smart cap" negotiation, i.e., policy makers should agree to do more if mitigation turns out cheaper than expected, and less if it turns out more expensive. This approach to negotiations is more efficient, and likely politically more palatable at the same time.

A large literature discusses distributional impact and political economy aspects of pollution regulation. In many aspects, the smart cap is a combination of a standard cap and trade system and a tax. As with a standard cap, many of the arguments favoring the auctioning of certificates as compared to their grandfathering also apply to the smart cap. We leave a detailed discussion of these interesting and important aspects to future research.

References

- Boleslavsky, R. & Kelly, D. L. (2014), 'Dynamic regulation design without payments: The importance of timing', *Journal of Public Economics* **120**, 169–180.
- British Petroleum (2020), 'Statistical review of world energy'. Retrieved 10/18/2020.
<https://www.bp.com/content/dam/bp/business-sites/en/global/corporate/xlsx/energy-economics/statistical-review/bp-stats-review-2020-all-data.xlsx>

- Burtraw, D., Holt, C., Palmer, K. & Shobe, W. (2020), ‘Quantities and prices: Price-responsive allocane supply in environmental markets’, *RFF Working Paper 20-17* .
- Dasgupta, P., Hammond, P. & Maskin, E. (1980), ‘On imperfect information and optimal pollution control’, *The Review of Economic Studies* **47**(5), 857–860.
- Doda, B. (2016), ‘How to price carbon in good times... and bad!’, *Climate change, 2016-01* **7**(1), 135–144.
- Drupp, M. A., Freeman, M. C., Groom, B. & Nesje, F. (2018), ‘Discounting disentangled’, *American Economic Journal: Economic Policy* **10**(4), 109–134.
- Ellerman, A. D. & Wing, I. S. (2003), ‘Absolute versus intensity-based emission caps’, *Climate Policy* **3**, S7–S20.
- European Patent Office (2020), ‘PATSTAT’. Retrieved 04/05/2020.
<https://www.epo.org/searching-for-patents/business/patstat.html>
- Fell, H. (2016), ‘Comparing policies to confront permit over-allocation’, *Journal of Environmental Economics and Management* **80**, 53–68.
- Fell, H., Burtraw, D., Morgenstern, R. D. & Palmer, K. L. (2012), ‘Soft and hard price collars in a cap-and-trade system: A comparative analysis’, *Journal of Environmental Economics and Management* **64**, 183–198.
- Fell, H. & Morgenstern, R. D. (2010), ‘Alternative approaches to cost containment in a cap-and-trade system’, *Environmental Resource Economics* **47**, 275–297.
- Fischer, C. & Springborn, M. (2011), ‘Emissions targets and the real business cycle: Intensity targets versus caps or taxes’, *Journal of Environmental Economics and Management* **62**, 352–366.
- Gerlagh, R. & Heijmans, R. (2020), ‘Regulating stock externalities’, *Unpublished* .
- Grüll, G. & Taschini, L. (2011), ‘Cap-and-trade properties under different hybrid scheme designs’, *Journal of Environmental Economics and Management* **61**(1), 107–118.

- Hepburn, C. (2006), ‘Regulation by Prices, Quantities, or Both: A Review of Instrument Choice’, *Oxford Review of Economic Policy* **22**(2), 226–247.
- Heutel, G. (2012), ‘How should environmental policy respond to business cycles? optimal policy under persistent productivity shocks’, *Review of Economic Dynamics* **15**, 244–264.
- Hoel, M. & Karp, L. (2001), ‘Taxes and quotas for a stock pollutant with multiplicative uncertainty’, *Journal of Public Economics* **82**, 91–114.
- Holland, S. & Moore, M. (2013), ‘Market design in cap and trade programs: permit validity and compliance timing’, *Journal of Environmental Economics and Management* **66**(3), 671–87.
- IMF (2020), ‘World Economic Outlook Database. International Monetary Fund. October 2020’.
- IPCC (2013), *Climate Change 2013: The Physical Science Basis*, Cambridge University Press, Cambridge.
- Jarke, J. & Perino, G. (2017), ‘Do renewable energy policies reduce carbon emissions? On caps and inter-industry leakage’, *Journal of Environmental Economics and Management* **84**, 102–124.
- Jotzo, F. & Pezzey, J. C. (2007), ‘Optimal intensity targets for greenhouse gas emissions trading under uncertainty’, *Environmental and Resource Economics* **38**, 259–284.
- Karp, L. & Traeger, C. (2018), ‘Prices versus quantities reassessed’, *CESifo Working Paper* (7331).
- Karp, L. & Zhang, J. (2005), ‘Regulation of stock externalities with correlated abatement costs’, *Environmental and Resource Economics* **32**, 273–299.
- Kling, C. & Rubin, J. (1997), ‘Bankable permits for the control of environmental pollution’, *Journal of Public Economics* **64**, 101–115.
- Kollenberg, S. & Taschini, L. (2016), ‘Emission trading systems with cap adjustments’, *Journal of Environmental Economics and Management* **80**, 20–36.

- Kwerel, E. (1977), ‘To tell the truth: Imperfect information and optimal pollution control’, *The Review of Economic Studies* **44**(3), 595–601.
- Newell, R. G. & Pizer, W. A. (2003), ‘Regulating stock externalities under uncertainty’, *Journal of Environmental Economics and Management* **45**, 416–432.
- Newell, R. G. & Pizer, W. A. (2008), ‘Indexed regulation’, *Journal of Environmental Economics and Management* **56**, 221–233.
- Newell, R., Pizer, W. & Zhang, J. (2005), ‘Managing permit markets to stabilize prices’, *Environmental and Resource Economics* **31**, 133–157.
- Nordhaus, W. & Sztorc, P. (2013), *DICE2013R: Introduction and User’s Manual*, dice-model.net.
- Perino, G. (2018), ‘New EU ETS Phase 4 rules temporarily puncture waterbed’, *Nature Climate Change* **8**, 260–271.
- Perino, G., Ritz, R. & van Benthem, A. (2020), ‘Overlapping climate policies’, *NBER Working Paper 25643*.
- Perino, G. & Willner, M. (2016), ‘Procrastinating reform: The impact of the market stability reserve on the eu ets’, *Journal of Environmental Economics and Management* **80**, 37–52.
- Pizer, W. A. (2002), ‘Combining price and quantity controls to mitigate global climate change’, *Journal of Public Economics* **85**, 409–434.
- 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.
- Requate, T. & Unold, W. (2001), ‘Pollution control by options trading’, *Economic Letters* **73**, 353–358.
- Roberts, M. J. & Spence, M. (1976), ‘Effluent charges and licences under uncertainty’, *Journal of Public Economics* **5**, 193–208.
- Rogers, E. (1995), *Diffusion Innovations*, Free Press.

- Silbye, F. & Birch-Sørensen, P. (2019), ‘National climate policies and the european emissions trading system’, *Nordic Economic Policy Review* **1**, 63–106.
- Stavins, R. (1996), ‘Correlated uncertainty and policy instrument choice’, *Journal of Environmental Economics and Management* **30**(2), 218–232.
- Stavins, R. N. (2020), ‘The future of u.s. carbon-pricing policy’, *Environmental and Energy Policy and the Economy* **1**, 8–64.
- Weitzman, M. L. (1974), ‘Prices versus quantities’, *Review of Economic Studies* **41**, 477–491.
- Weitzman, M. L. (1978), ‘Optimal rewards for economic regulation’, *American Economic Review* **68**(4), 683–691.
- World Bank (2014), ‘CO2 emissions (in kt)’. Retrieved 04/05/2020.
<https://data.worldbank.org/indicator/EN.ATM.CO2E.KT?end=2016&start=1960&view=chart>
- World Bank (2019), ‘GDP (Current LCU)’. Retrieved 04/08/2020.
<https://data.worldbank.org/indicator/NY.GDP.MKTP.CN?end=2011&start=1960>
- Yates, A. & Cronshaw, M. (2001), ‘Pollution permit markets with intertemporal trading and asymmetric information’, *Journal of Environmental Economics and Management* **42**, 104–118.

Appendix

A Calibration Details for Technology Adoption Share α

Method and Estimation Equation. The present section derives a crude estimate of our parameter α , the adoption share of an innovations during a commitment period. As in section 3.3, we assume commitment periods of 5 years and that technology is fully persistent. We use green patents as a proxy for our technology stock. New patents over the course of a commitment period represent the innovations (expected trend plus shock). To identify α , we regress business as usual emissions on patents in the period preceding serious CO₂ regulation. Business as usual emissions in our model are

$$E_t = \frac{1}{f} (h_t + \theta_{t-1} + \alpha \epsilon_t).$$

We assume that the adoption share α_t applies equally to the deviation from the trend as to the expected patent trend, i.e., $h_t = \bar{h}_{t-1} + \alpha \Delta_t$ and $\bar{h}_t = \bar{h}_{t-1} + \Delta_t$. Our identification assumes that the technology stock is linear in the amount of green patents $\bar{h}_{t-1} + \theta_{t-1} = \gamma_1 + \gamma_2 P_{t-1}$ with $\gamma_1 \in \mathbb{R}$ and $\gamma_2 \in \mathbb{R}_+$. Then, our estimation equation becomes

$$E_t = \frac{\gamma_2}{f} (P_{t-1} + \alpha(P_t - P_{t-1})) + \gamma_3 \quad (11)$$

with $\gamma_3 \in \mathbb{R}$. Translated into an empirical model controlling for a time trend and GDP, and using an annual time step, the estimation takes the form

$$\begin{aligned} \sum_{\tau=t-4}^t Emissions_{\tau} &= \mu_0 + \mu_1 Trend_t + \mu_2 \sum_{\tau=t-4}^t Gdp_{\tau} \\ &+ \mu_3 \sum_{\tau=t-4}^t New\ green\ patents_{\tau} + \mu_4 \sum_{\tau=0}^{t-5} New\ green\ patents_{\tau} + \sum_{\tau=t-4}^t \zeta_{\tau}. \end{aligned}$$

If follows from equation (11) that we obtain the adoption share α as the ratio of the regression coefficients on new patents and the stock of patents before the beginning of the commitment period, $\frac{\mu_4}{\mu_3} = \frac{\gamma_2 \alpha}{\gamma_2} = \alpha$. Our base scenario controls only for *GDP* and uses a quadratic trend, our more sophisticated estimation scenario also controls for oil prices and allows for a break-point in the trend (more below). We assume that the

estimation error ζ_T is i.i.d. and $\zeta_t \sim N(0, \sigma^2)$. By construction, the observations have serially correlated errors for the adjacent 5 years. Therefore, we use the Newey-West estimator with five lags, which is robust to the serial correlation between errors up to five periods.

Data. We use the PATSTAT database of the (European Patent Office 2020). We use green patents as identified by the classification code Y02E.²⁰ As is common in the literature, we only consider major patents, i.e., patents that are registered at all of the three major patent offices, the United States, Europe, and Japan. We use the earliest filing date of the patent as the time of innovation. We obtain data on emissions and other controls variables from World Bank (2014, 2019) and British Petroleum (2020). Even if short, we consider the period from 1995 to 2010 for our preferred estimates. After 2010, implementations of the Kyoto protocol would interfere with our assumption of business as usual emissions. Before 1995, detrending the variables becomes difficult. In particular, the 1980s have seen a strong change in the growth of green patents and starting before 1995 we are not able to render the time series stationary. Table 1 gives an overview of the data for our preferred estimation period, 1995 to 2010.

²⁰We use a patent if the Y02E classification is part of their classification codes.

Table 1: Descriptive Statistics

	Mean	Std. Dev.
CO ₂ emissions ¹ (Yearly)	5452829	275448.5
Early period ²	5329343	265499.1
Late period ³	5622621	193621.5
Total green patents ⁵ (Yearly)	5083.632	3329.472
Early period	2776.273	1115.06
Late period	8256.25	2619.059
Gross Domestic Product ⁴ (Yearly)	1.08×10^{13}	2.88×10^{12}
Early period	8.67×10^{12}	1.54×10^{12}
Late period	1.36×10^{13}	1.28×10^{12}
Oil Price ⁶ (Yearly)	37.63731	25.74642
Early period	19.77602	4.594336
Late period	62.19659	21.99409

¹ Measured in kilotons of carbon dioxide mass per year.

²The early period runs from 1995 to 2002.

³The late period runs from 2002 to 2010.

⁴Measured in US dollars (at current prices). For a given year.

⁵Green patents registered at all the leading patent offices (in the United States, Europe, and Japan) in a year.

⁶ Measured in US dollars (at nominal prices). Average over a year.

Results. Table 2 presents the result of our base estimation, which only controls for GDP and uses a quadratic time trend $Trend_t = \beta_0 t + \beta_1 t^2$. We find an estimate of $\alpha = 0.28$ for our preferred estimation period. Going back to 1990, our time series are no longer stationary but we still obtain significant regression coefficients. The estimate increases to $\alpha = 0.34$. Once we enter the 1980s, one of our coefficients loses statistical significance and the sign is off (market in red). Table 3 tries to get a handle on (some of the) issue controlling for the somewhat volatile oil prices and allowing for a break-point in the quadratic trend.²¹ The table leaves the estimate of $\alpha = 0.28$ in our preferred period unchanged. Starting in 1990, it slightly reduces the estimate to $\alpha = 0.34$. The coefficient on new patents in the time series starting in 1985 remains statistically insignificant, but it changes both sign and magnitude to a more

²¹The regression includes a linear-quadratic spline with one break-point

$$\beta_1 \times Time_t + \beta_2 \times Time_t^2 + \beta_3 \times Time_t \times Dummy_t + \beta_4 \times (Time_t^2) \times Dummy_t$$

where the dummy-variable specifies the breaking point, which we choose to minimize the sum of squared residuals in the regression.

reasonable values, delivering $\alpha = 0.27$. Based on these results, we take $\alpha \approx 0.25$ as our low estimate of the share of adoption during a 5 year period, and we also present the results for a somewhat larger $\alpha = 0.5$. These values bound our estimates. We pick a somewhat high upper value because non-patented innovation might be implemented somewhat faster, increasing the overall adoption share of green innovation relative to our green-patents-based adoption share.

Table 2: Base Estimate of Adoption Share α

	1995-2010	1990-2010	1985-2010
$\sum_{\tau=0}^{t-5} \text{New green patents}_{\tau}$			
<i>Coefficient</i> ¹	-175.02	-221.49	-195.93
<i>Standard Error</i> ²	9.17	8.19	17.26
<i>P – Value</i>	0.00***	0.00***	0.00***
$\sum_{\tau=t-4}^t \text{New green patents}_{\tau}$			
<i>Coefficient</i> ³	-48.35	-74.16	29.57
<i>Standard Error</i> ²	19.14	21.26	31.04
<i>P – Value</i>	0.03**	0.00***	0.35
	0.28	0.34	-0.15

Significance levels are * for 0.1, ** for 0.05 and *** for 0.01

¹ Read the unit of measurement for the coefficient as "Kilotons of Carbon Dioxide Mass per New Green Patent".

² Read the unit of measurement for the standard error as "Kilotons of Carbon Dioxide Mass".

³ Read the unit of measurement for the coefficient as "Kilotons of Carbon Dioxide Mass per New Green Patent in Stock of Green Patents".

Table 3: Estimate of Adoption Share α Controlling Additionally for Oil Prices and Allowing for a Break-Point in the Quadratic Trend

	1995-2010	1990-2010	1985-2010
$\sum_{\tau=0}^{t-5} \text{New green patents}_{\tau}$			
<i>Coefficient</i> ¹	-175.74	-185.66	-89.46
<i>Standard Error</i> ²	26.13	15.45	24.44
<i>P – Value</i>	0.00***	0.00***	0.00***
$\sum_{\tau=t-4}^t \text{New green patents}_{\tau}$			
<i>Coefficient</i> ³	-48.48	-63.45	-26.20
<i>Standard Error</i> ²	20.57	23.11	22.38
<i>P – Value</i>	0.04***	0.02**	0.06
α -coefficient	0.28	0.34	0.27

Significance levels are * for 0.1, ** for 0.05 and *** for 0.01

¹ Read the unit of measurement for the coefficient as "Kilotons of Carbon Dioxide Mass per New Green Patent".

² Read the unit of measurement for the standard error as "Kilotons of Carbon Dioxide Mass".

³ Read the unit of measurement for the coefficient as "Kilotons of Carbon Dioxide Mass per New Green Patent in Stock of Green Patents".

B Proofs for Section 2

Proof of Proposition 1. For all shock levels θ the first best emission level satisfies

$$\begin{aligned}
 SCC(E, \theta) &= MB(E, \theta) \\
 \Rightarrow SCC_E(E, \theta)dE + SCC_\theta(E, \theta)d\theta &= MB_E(E, \theta)dE + MB_\theta(E, \theta)d\theta \quad (12) \\
 \Rightarrow \frac{dE}{d\theta} &= \frac{MB_\theta(E, \theta) - SCC_\theta(E, \theta)}{SCC_E(E, \theta) - MB_E(E, \theta)}
 \end{aligned}$$

and therefore

$$\begin{aligned}
 SCC^*_E(E) &= SCC_E(E, \theta) + SCC_\theta(E, \theta) \frac{d\theta}{d\epsilon} \\
 &= SCC_E(E, \theta) + SCC_\theta(E, \theta) \frac{SCC_E(E, \theta) - MB_E(E, \theta)}{MB_\theta(E, \theta) - SCC_\theta(E, \theta)} \quad (13) \\
 &= \frac{SCC_E(MB_\theta - SCC_\theta) + SCC_\theta(SCC_E - MB_E)}{MB_\theta - SCC_\theta} \\
 &= \frac{MB_\theta}{MB_\theta - SCC_\theta} SCC_E + \frac{-SCC_\theta}{MB_\theta - SCC_\theta} MB_E. \quad (14)
 \end{aligned}$$

We assume $SCC_E > MB_E < 0$, and $MB_\theta > 0$. We characterize the five cases of Proposition 1 one by one. i) If $0 < SCC_\theta < MB_\theta$ then both numerator and denominator of the fraction in equation (13) are positive and the second summand in the equation is adding a positive amount to SCC_E so that $SCC^*_E > SCC_E$. ii) If $SCC_\theta = MB_\theta$ then the right hand side of equation (13) goes to infinity indicating $SCC^*_E = +\infty$. More precisely, before dividing by the term converging to zero, equation (12) cannot be satisfied for $SCC_\theta = MB_\theta$ in general, but only holds for a particular emission level, which is the optimal cap. iii) If $MB_\theta < SCC_\theta$ then the denominator of the two fractions in equation (14) is negative. Thus, the sign of SCC^*_E is negative if and only if $SCC_\theta MB_E - MB_\theta SCC_E < 0 \Leftrightarrow \frac{SCC_\theta}{MB_\theta} > \frac{SCC_E}{MB_E}$, which is satisfied because by assumption $\frac{SCC_\theta}{MB_\theta} > 1 > \frac{SCC_E}{MB_E}$. iv) If $SCC_\theta = 0$ then $SCC^*_E = SCC_E$. v) If $SCC_\theta < 0$ then equation (14) is a standard weighted mean between SCC_E and $MB(E)$ and, hence, smaller than SCC_E . ■

Proof of Proposition 2. Local stability is equivalent to a negative slope of excess demand evaluated at the equilibrium.

(i) The inverse supply of emissions is $p^{E, \text{supply}} = SCC^*(E)$ and the inverse demand is $p^{E, \text{demand}} = MB(E; \theta)$. If $SCC^*_E(E(p^E)) = 0$, the slope of excess demand equals the

slope of the demand for emissions, which is negative by concavity of marginal benefits. In the case where $SCC_E^*(E(p^E))$ is nonzero, denoting $X(p^E; \theta)$ as excess demand, the slope of excess demand equals²²

$$\frac{\partial X(p^E; \theta)}{\partial p^E} = \frac{SCC_E^*(E(p^E)) - MB_E(E(p^E); \theta)}{MB_E(E(p^E); \theta) SCC_E^*(E(p^E))}. \quad (15)$$

Local stability is equivalent to the right side of (15) being negative. By concavity of benefits, $MB_E(E(p^E); \theta) < 0$ is negative.

Sufficiency of (a) and (b): If $SCC_E^*(E(p^E)) > 0$ the slope of excess demand is negative and we conclude that local stability holds. As explained above, local stability also holds in the case where $SCC_E^*(E(p^E)) = 0$, because there the inverse supply function is vertical. Thus, case (a) implies local stability. If $SCC_E^*(E(p^E)) < 0$ the slope of excess demand is negative and, thus, case (b) also implies local stability. If either (a) or (b) holds everywhere, then excess demand is positive for prices below the equilibrium prices and negative for higher prices, implying global stability and uniqueness.

Necessity of (a) and (b): Local stability implies a downward sloping excess demand. This condition is met if $SCC_E^*(E(p^E)) = 0$, where the inverse supply is vertical. For $SCC_E^*(E(p^E)) \neq 0$, local stability implies that the right side of equation (15) is negative, which is equivalent to

$$\begin{aligned} & SCC_E^*(E(p^E)) > MB_E(E(p^E); \theta) \text{ and } SCC_E^*(E(p^E)) > 0 \\ \text{or} \\ & SCC_E^*(E(p^E)) < MB_E(E(p^E); \theta) \text{ and } SCC_E^*(E(p^E)) < 0 \\ \Leftrightarrow \\ & SCC_E^*(E(p^E)) > 0 \quad \text{or} \quad SCC_E^*(E(p^E)) < MB_E(E(p^E); \theta), \end{aligned}$$

which implies cases (a) or (b).

(ii) We first establish continuous differentiability of the smart cap's redemption function $q(p)$. By equation (3), continuous differentiability of MB and SCC together with our assumption $MB_\theta \neq SCC_\theta$, implies continuous differentiability of SCC^* . Given continuous differentiability (and thus continuity) of SCC^* , equation (4) implies continuous differentiability of the redemption function.

²²The slope of demand is one over the slope of inverse demand and analogously for supply, so that $\frac{\partial X(p^E; \theta)}{\partial p^E} = \frac{1}{MB_E(E(p^E); \theta)} - \frac{1}{SCC_E^*(E(p^E))}$.

Under the smart cap, the supply of certificates is fixed at Q , and thus independent of the certificate price. We first prove the statement regarding *local stability*. Suppressing the argument θ , the competitive inverse demand for certificates, Q^c , satisfies $p = MB(Q^c q(p)) q(p)$, so the slope of industry demand satisfies

$$\begin{aligned} dp &= MB_E(Q^c q(p)) [q(p) dQ^c + Q^c q'(p) dp] q(p) + MB(Q^c q(p)) q'(p) dp \\ \Rightarrow [1 - MB_E(Q^c q(p)) Q^c q'(p) - MB(Q^c q(p)) q'(p)] dp & \\ &= MB_E(Q^c q(p)) [q(p)]^2 dQ^c \\ \Rightarrow \frac{dQ^c}{dp} &= \frac{1 - [MB_E Q^c q(p) + MB] q'(p)}{[q(p)]^2 MB_E}. \end{aligned} \quad (16)$$

The denominator of this expression is negative, so the slope of industry demand, and thus the slope of excess demand, is negative if and only if the numerator is positive. At the equilibrium price $p = p^*$, the equilibrium requirement that supply equals demand implies $Q = Q^c$. By construction, the equilibrium price supports the optimal level of emissions, E^* , so in equilibrium $Q^c q(p^*) = E^*$. Using this equality in equation (16) yields

$$\frac{dQ^c}{dp} < 0 \Leftrightarrow 1 - [MB_E E^* + MB] q'(p) > 0.$$

Using the left side of equation 4 we obtain the condition for local stability

$$1 > \frac{MB_E E^* + MB}{SCC_E^* E + SCC^*}$$

stated in the proposition as equation (5).

We proceed to show *global stability* of the smart cap under the assumption that $\epsilon_{q,p}(p) \neq 1$ for all p . This assumption implies that $\frac{p}{q(p)}$ is strictly monotonic as

$$\frac{d}{dp} \frac{p}{q(p)} = \frac{q(p) - pq'(p)}{q(p)^2} = \frac{q(p) \left(1 - \frac{p}{q} q'(p)\right)}{q(p)^2} = \frac{1 - \epsilon_{q,p}}{q(p)} \neq 0$$

and we established above continuous differentiability of the smart cap (implying continuity also of $\epsilon_{q,p}$). As a result, the relation between p and $p^E = \frac{p}{q(p)}$ is one to one. The assumption that condition (5) holds implies that the slope of excess demand is negative in the neighborhood of any socially optimal equilibrium. Thus, it suffices to show that excess demand does not *cross* zero again outside of a socially optimal equilibrium. By equation (16) continuous differentiability of redemption function and marginal benefits

implies continuity of excess demand. Thus, it suffices to show that excess demand does not equal zero outside of a socially optimal equilibrium (because crossing implies an intersection).

Assume that some certificate price \bar{p} implies a zero excess demand for some technology realization $\bar{\theta}$. Market clearing implies that $MB(\bar{E}; \bar{\theta}) = \frac{\bar{p}}{q(\bar{p})}$ for some emissions level \bar{E} . We have assumed that the smart cap is defined (only) on the domain of prices for which there exists a technology shock such that the price induces a social optimum. We denote by θ^* the technology level that induces the social optimum for the certificate price \bar{p} . Market clearing in the social optimum under technology realization θ^* implies that $MB(E^*; \theta^*) = \frac{\bar{p}}{q(\bar{p})}$ for some emissions level E^* . Moreover, because the relation between p and $p^E = \frac{p}{q(p)}$ is one to one and because of market clearing under both \bar{p} and $\bar{\theta}$ and under \bar{p} and θ^* we find that $\bar{E} = q(\bar{p})Q^c = E^*$. It follows that

$$MB(q(\bar{p})Q^c; \bar{\theta}) = MB(\bar{E}; \bar{\theta}) = \frac{\bar{p}}{q(\bar{p})} = MB(E^*; \theta^*) = MB(q(\bar{p})Q^c; \theta^*).$$

Because $MB(E; \theta)$ is strictly monotonic in θ it follows from $MB(q(\bar{p})Q^c; \bar{\theta}) = MB(q(\bar{p})Q^c; \theta^*)$ that $\bar{\theta} = \theta^*$. Therefore, the arbitrary market clearing equilibrium with zero excess demand we started with, characterized by \bar{p} and $\bar{\theta}$, is a social optimum; the only zeros of excess demand are the socially optimal market equilibria where excess demand slopes correctly by our assumption that condition (5) holds. ■

Proof of Proposition 3. We use a general redemption function, $q(p)$. To distinguish this function from the smart cap for the competitive industry, we denote the latter as $q^c(p)$ in this appendix. Given that $Q^m = Q$, the monopsony chooses p :

$$\max_{p \geq 0} [B(Qq(p); \theta) - Qp].$$

The first order condition at an interior point is

$$MB(Qq(p); \theta) q' = 1 \quad \Rightarrow \quad MB(Qq(p); \theta) = \frac{1}{q'} = \frac{\frac{p}{q}}{q' \frac{p}{q}} = \frac{p^E}{\varepsilon_{q,p}}, \quad (17)$$

implying the first equality in equation (6). To establish the second equality, we differentiate the definition $p^E = \frac{p}{q}$ and simplify, to obtain $\frac{dp^E}{dp} = \frac{q - pq'}{q^2}$. The equilibrium

condition $qQ = E$ implies

$$\begin{aligned} Qq' &= \frac{dE}{dp^E} \frac{dp^E}{dp} \Rightarrow q' \frac{p}{q} = \frac{dE}{dp^E} \frac{dp^E}{dp} p^E \frac{1}{Q} \Rightarrow \\ \varepsilon_{q,p} &= \frac{dE}{dp^E} \frac{p^E}{E} \frac{E}{Q} \frac{dp^E}{dp} = \varepsilon_{E,p^E} \left(q \frac{q - pq'}{q^2} \right) = \varepsilon_{E,p^E} \left(1 - \frac{pq'}{q} \right) \\ &= \varepsilon_{E,p^E} (1 - \varepsilon_{q,p}) \Rightarrow \varepsilon_{q,p} = \frac{\varepsilon_{E,p^E}}{(1 + \varepsilon_{E,p^E})}. \end{aligned}$$

(ii) Now we examine the case where $q(p) = q^c(p)$, the smart cap for the competitive industry. Using equation (17), we find that the monopsony's marginal increase in profits resulting from a change in certificate price is $MB(Qq^c(p); \theta) \frac{\varepsilon_{q^c,p}}{p^E} Q - Q$. Evaluating this derivative at the competitive equilibrium (where $MB = p^E$), the monopsony's marginal benefits of a change in the price equal $(\varepsilon_{q^c,p^c} - 1)Q$, where the superindex c denotes the price in the competitive equilibrium. Thus, the monopsony benefits from raising the certificate price above the competitive level if and only if $\varepsilon_{q^c,p^c} \geq 1$.

By definition of the elasticity, ε_{q^c,p^c} and $q'(p)$ have the same sign. Therefore, the monopsony wants to raise the certificate price above the competitive level if and only if it wants to increase emissions above the competitive (socially optimal) level. The inequality $\varepsilon_{q^c,p^c} \geq 1$ (using equation (4)) is equivalent to $SCC_E^* \leq 0$. Given that this relations is "if and only if", the monopsony benefits from reducing emissions below the competitive level if and only $SCC_E^* \geq 0$.²³ These actions strictly increase or reduce emissions if the inequalities are strict and the competitive equilibrium is interior.

(iii) Here we consider the case where the redemption function induces the monopsony to emit at the socially optimal level. The monopsony's first order condition, equation (17), implies $q' = \frac{1}{MB(Qq(p); \theta)}$ and the condition for social optimality requires $MB(E; \theta) = SCC^*(E)$. Thus, the optimizing monopolist will emit at the socially optimal level if $MB(Qq(p); \theta) = \frac{1}{q'} = SCC^*(E)$. Consequently, a smart cap satisfying equation (7) induces the social optimum if the monopsony's first order conditions correctly characterize the optimum.

²³If $\varepsilon_{q^c,p^c} < 0$ then: the monopsony's marginal benefit of a higher price is negative; $SCC_E^* < 0$; and in addition $q' < 0$. In this case, the monopsony wants to lower the price below the competitive level as a means of increasing emissions. Here also we have $SCC_E^* < 0$ and the monopsony wants to increase emissions above the competitive level.

The monopsony's second order condition is

$$MB(Qq^m; \theta) \frac{d^2 q^m(p)}{dp^2} + MB_E(Qq^m; \theta) Q \left(\frac{dq^m(p)}{dp} \right)^2 \leq 0 \quad (18)$$

We now show that this equation is locally satisfied in the neighborhood of an equilibrium if and only if $MB_E(E; \theta) < SCC_E^*$. Differentiating equation (7) gives

$$\frac{d^2 q^m(p)}{dp^2} = -\frac{SCC_E^* \cdot Q \frac{dq^m(p)}{dp}}{(SCC^*)^2} = -\frac{SCC_E^* \cdot Q}{(SCC^*)^3}$$

Substituting this expression into equation (18) gives

$$\left(MB_E - MB \frac{SCC_E^*}{SCC^*} \right) \frac{Q}{(SCC^*)^2} = (MB_E - SCC_E^*) \frac{Q}{(SCC^*)^2} \quad (19)$$

where the last equality uses $MB(Qq(p); \theta) = \frac{1}{q'} = SCC^*(E)$. The right side of equation (19) is negative if and only if $MB_E(E; \theta) < SCC_E^*$. There may exist a family of solutions to equation (7) because the boundary condition is unspecified. ■

C Proofs for Section 3

We first obtain the solution to the full-information problem where the regulator observes ε_t and chooses emissions directly. This solution provides the full-information $SCC_t(S_{t-1}, E_t, \theta_{t-1}, \varepsilon_t)$ and the full-information value function. We use $SCC_t(\cdot)$ to construct the smart tax, Proposition 4, which we then use to construct the smart cap, Proposition 5.

We choose a unit of time to be one year, and we use the parameter ϕ to denote the number of years in a compliance period. This formulation enables us to change the length of a period without recalibrating the model. We assume that (i) emissions are constant during a period (equal to ϕ years), (ii) the change in the stock occurs at the end of the period, and (iii) the payoff flow is not discounted within a period. Thus, for example, $S_{t+1} = \delta S_t + E_t \phi$, where E_t equals annual emissions during period t . The parameter δ depends on ϕ . With the annual persistence level $\hat{\delta}$, we set $\delta = \hat{\delta}^\phi$. We make an analogous adjustment in the discount factor, β , and the serial correlation parameters, ρ . The shock evolves according to $\theta_t = \rho \theta_{t-1} + \varepsilon_t$. We measure the pollution stock in gigatons (Gt) of CO_2 and payoffs in giga dollars (G\$). The marginal benefit

of emissions (equal to the marginal abatement cost) is measured in $\frac{G\$}{GtCO_2} = \frac{\$}{tCO_2}$.

The flow payoff during a period lasting ϕ years is

$$\left((h_t + \rho\theta_{t-1} + \alpha\varepsilon_t) x_t - \frac{1}{2} f x_t^2 - \frac{1}{2} b S_t^2 \right) \phi.$$

We use $j_t(S_t, \theta_{t-1}, \varepsilon_t)$ to denote the value function, i.e. the value of the program under full information and optimal emissions. The full-information optimization problem is

$$j_t(S_t, \theta_{t-1}, \varepsilon_t) = \max_{\{x_s\}_{s=t}^{\infty}} \mathbb{E}_{\{\varepsilon_j\}_{j=t+1}^{\infty}} \sum_{s=0}^{\infty} \beta^s \left((h_t + \rho\theta_{t-1} + \alpha\varepsilon_t) x_{t+s} - \frac{1}{2} f x_{t+s}^2 - \frac{1}{2} b S_{t+s}^2 \right) \phi. \quad (20)$$

The value function, $j_t(S_t, \theta_{t-1}, \varepsilon_t)$, depends on the current state variable, $(S_t, \theta_{t-1}, \varepsilon_t)$, and on calendar time via the time dependence of h_t . The value function is quadratic in the state variable. The social cost of carbon is defined as the expectation, over the next period shock, of the present value of the cost of beginning the next period with one additional unit of the pollution stock: $SCC_t = \beta \mathbb{E}_{\varepsilon_{t+1}} \frac{\partial j_t(S_{t+1}, \theta_{t+1}, \varepsilon_{t+1})}{\partial S_{t+1}}$. We have

Lemma 1 *The social cost of carbon and the optimal emissions rule are linear functions of the information state, S_t , θ_{t-1} , and ε_t . The coefficients of these linear functions are constant; the intercepts depend on the trajectory of the exogenous demand shifter, $\{h_{t+s}\}_{s=0}^{\infty}$.*

(i) *The social cost of carbon is*

$$SCC_t = \beta (-v_{1,t+1} + \lambda(\delta S_t + \phi E_t) + \mu(\rho\theta_{t-1} + \varepsilon_t)). \quad (21)$$

Using the definition

$$\varpi \equiv f \left(1 - \beta\delta^2 - \beta \frac{b}{f} \phi^2 \right), \quad (22)$$

the constant coefficients are

$$\lambda = \frac{1}{2\beta\phi} \left(-\varpi + \sqrt{\varpi^2 + 4\beta\phi^2 b f} \right) > 0 \quad (23)$$

and

$$\mu = \rho\beta\delta\phi \frac{\lambda}{f(1 - \rho\beta\delta) + \beta\phi\lambda}; \quad 1 > \mu > 0. \quad (24)$$

The time-varying intercept, $v_{1,t}$, is the solution to the difference equation

$$v_{1,t} = \beta \frac{\delta (f\nu_{1,t+1} - \lambda\phi h_t)}{f + \beta\lambda\phi} \Rightarrow v_{1,t} = -\frac{\beta\delta\phi\lambda}{f + \beta\phi\lambda} \sum_{j=0}^{\infty} \left(\frac{\beta\delta f}{f + \beta\phi\lambda} \right)^j h_{t+j}. \quad (25)$$

A sufficient condition for existence of $v_{1,t}$ is that the elements of the sequence $\{h_{t+j}\}_{j=0}^{\infty}$ are finite. We assume that $h_t = \eta^t h_0$ falls at a constant rate ($\eta < 1$). Then

$$v_{1,t} = -\frac{\lambda\beta\delta h_t\phi}{f(1 - \beta\eta\delta) + \beta\phi\lambda} < 0 \text{ (given } h > 0\text{)}. \quad (26)$$

(ii) The full information optimal emissions rule is

$$E_t = Z_{0t} + H\varepsilon_t + z_1 S_t + z_2 \theta_{t-1} \quad (27)$$

with

$$\begin{aligned} Z_{0t} &= \frac{h_t + \beta v_{1,t+1}}{(f + \beta\phi\lambda)}, \quad H = \frac{\alpha - \beta\mu}{f + \beta\phi\lambda} \\ z_1 &= \frac{-\beta\lambda\delta}{f + \beta\phi\lambda} < 0 \text{ and } z_2 = \rho \frac{1 - \beta\mu}{f + \beta\phi\lambda} > 0 \end{aligned} \quad (28)$$

Comments on the Lemma. The slope, w.r.t. the pollution stock, of the SCC is $\beta\lambda$; the coefficient $\beta\mu$ shifts the intercept of the SCC in response to the technology level that the next period inherits, θ_t . The optimal level of emissions is a decreasing function of the current stock of pollution and an increasing function of the lagged technology level, θ_{t-1} . An increase in θ_{t-1} increases the current marginal benefit of an additional unit of emissions; i.e., it increases the marginal cost of abatement, thereby increasing the demand for emissions.

The optimal level of emissions might increase or decrease with the current shock, ε_t . A positive technology shock (for example) raises both the current demand for emissions and (with $\rho > 0$) future demand. For α close to 1 the increase in current demand is large relative to the increase in future demand. Here, $H > 0$: optimal emissions increase with the technology shock. In contrast, for small α , a positive shock raises the current demand for emissions relatively little compared to the rise in future demand. The anticipation of the higher future demand causes the regulator to reduce current emissions in response to a positive shock: $H < 0$.

Note that μ does not depend on α . For the knife-edge case $\alpha - \beta\phi\mu = 0$, the optimal level of emissions is independent of the current technology shock, although it still depends on S_t and θ_{t-1} . For this knife-edge case, the feedback quota under

asymmetric information is first-best.

Sketch of Proof of Lemma 1. The proof is straightforward but tedious, so we only sketch the steps here, relegating the details to Referees' Appendix D. We refer to the value function conditional on the current technology shock, $j_t(S_t, \theta_{t-1}, \varepsilon_t)$ as the ex post value function, and its expectation $J_t(S_t, \theta_{t-1}) \equiv \mathbb{E}_{\varepsilon_t} j_t(S_t, \theta_{t-1}, \varepsilon_t)$ as the ex ante-value function. It is well known that for the linear-quadratic control problem with additive errors, $j_t(S_t, \theta_{t-1}, \varepsilon_t)$ is a linear-quadratic function. Therefore, $J(S_t, \theta_t)$ is also linear-quadratic.

The dynamic programming equation under full information is

$$\begin{aligned}
 j_t(S_t, \theta_{t-1}, \varepsilon_t) = & \\
 \max_E \left[\left((h_t + \rho\theta_{t-1} + \alpha\varepsilon_t) E - \frac{1}{2}fE^2 - \frac{1}{2}bS_t^2 \right) \phi \right. & \\
 \left. + \beta \mathbb{E}_{\varepsilon_{t+1}} j_{t+1}(S_{t+1}, \theta_t, \varepsilon_{t+1}) \right] = & \quad (29) \\
 \max_E \left[\left((h_t + \rho\theta_{t-1} + \alpha\varepsilon_t) E - \frac{1}{2}fE^2 - \frac{1}{2}bS_t^2 \right) \phi \right. & \\
 \left. + \beta J_{t+1}(S_{t+1}, \theta_t) \right]. &
 \end{aligned}$$

We replace $J_{t+1}(S_{t+1}, \theta_t)$ with a linear-quadratic trial solution in the last line of equation 29 and perform the optimization, writing the full information decision rule as a function of $(S_t, \theta_{t-1}, \varepsilon_t)$ and of the parameters of the trial solution. The SCC is simply $\beta \frac{\partial J_{t+1}(S_{t+1}, \theta_t)}{\partial S_{t+1}}$. This derivative depends on $(S_t, \theta_{t-1}, \varepsilon_t)$ and on the emissions level, E .

Because $J_{t+1}(S_{t+1}, \theta_t)$ is a quadratic function, $\frac{\partial J_{t+1}(S_{t+1}, \theta_t)}{\partial S_{t+1}}$ is a linear function of the state, (S_{t+1}, θ_t) , as shown in equation 21, and the optimal decision rule is a linear function of the state, as shown in equation 27. Performing the optimization in equation 29 we obtain the optimal decision rule in terms of the coefficients of the ex ante value function. We then substitute this decision rule into the dynamic programming equation and take expectations, replacing $\mathbb{E}_{\varepsilon_t} j_t(S_t, \theta_{t-1}, \varepsilon_t)$, with $J_t(S_t, \theta_{t-1})$. We then equate coefficients of terms that are independent of (S_{t+1}, θ_t) , linear in (S_{t+1}, θ_t) and quadratic in (S_{t+1}, θ_t) , thereby obtaining the formulae for the coefficients (equations 22, 23, 24 and 25) ■

Proof of Proposition 4. Part i. For $H \neq 0$, we invert the full-information emissions

rule, equation 27, to write the shock as a function of the optimal level of emissions

$$\varepsilon_t = \frac{E_t - (Z_{0t} + z_1 S_t + z_2 \theta_{t-1})}{H}. \quad (30)$$

We want to obtain the formulae for the coefficients of the smart tax, a linear function of the observables at time t , S_t , θ_{t-1} , and E_t :

$$A_0 S_t + A_1 \theta_{t-1} + \gamma E_t + \alpha_t.$$

If firms set their marginal benefit of emissions (per year) equal to this tax, the equilibrium condition is

$$h_t + \rho \theta_{t-1} + \alpha \varepsilon_t - f E_t = A_0 S_t + A_1 \theta_{t-1} + \gamma E_t + \alpha_t. \quad (31)$$

By construction, $SCC^* = SCC$ evaluated at the optimal level of emissions. Using equation 30 to eliminate ε_t from the left side of equation 31, and collecting terms, we write the left side of equation 31 as

$$-\alpha \frac{z_1}{H} S_t + \left(\rho - \alpha \frac{z_2}{H} \right) \theta_{t-1} + \left(\frac{\alpha}{H} - f \right) E_t + h_t - \alpha \frac{Z_{0t}}{H} \quad (32)$$

Substituting the definitions of z_1 , z_2 , and Z_{0t} in equation 28 into expression 32 and then equating coefficients with the right side of equation 31 produces the formula for the coefficients of the smart tax:

$$\begin{aligned} A_0 &= \frac{\alpha \beta \lambda \delta}{\alpha - \beta \mu} \text{ and } A_1 = -\rho \beta \mu \frac{1 - \alpha}{\alpha - \beta \mu} \\ \gamma &= \beta \frac{\alpha \lambda \phi + \mu f}{\alpha - \beta \mu} \text{ and } a_t = -\beta \frac{\mu h_t + \alpha \nu_{1,t+1}}{\alpha - \beta \mu} \end{aligned} \quad (33)$$

contained in Proposition 4.

For $H = 0$ the optimal emissions level is independent of the current shock, ε_t . As $\alpha \rightarrow \beta \mu$ the smart tax becomes steeper. In the limit

Part ii. Using the inequalities in equations 23 and 24, the numerators of A_0 and γ are strictly positive, and for $\rho > 0$ the numerator of A_1 is weakly negative. Therefore the signs of these coefficients depend on the sign of the denominator, $\alpha - \beta \mu$. By equation 24, μ is independent of α , so there exists $\alpha^* = \beta \mu$ with (for $\rho > 0$) $0 < \alpha^* < \beta < 1$. For $\alpha < \alpha^*$ $A_0 < 0$, $\gamma < 0$ and $A_1 > 0$. These inequalities are reversed for $\alpha^* < \alpha$. At

$\alpha = \beta\mu$ the smart tax is a vertical line, i.e. a correspondence, not a function.

Part iii The argument for this claim parallels the proof of Proposition 2, so we do not repeat it. ■

Proof of Proposition 5. Using the definition of \hat{A}_t we write the smart tax as $SCC_t^* = \hat{A}_t + \gamma E_t$. The firm's price of a unit of emissions is $p^E = \frac{p}{q_t(p)}$. We define the smart cap using $\frac{p}{q_t(p)} = \hat{A}_t + \gamma E$. Inserting the market clearing condition, $E = Qq_t(p)$, gives

$$\frac{p}{q_t(p)} = \hat{A}_t + \gamma Qq_t(p). \quad (34)$$

The assumption that the first best level of emissions is non-negative implies $q_t(p) \geq 0$. Simplifying equation 34, gives the quadratic equation $\hat{A}_t q_t(p) + \gamma Q q_t^2(p) - p = 0$. The two roots are

$$q^+ = \frac{1}{2\gamma Q} \left(-\hat{A}_t + \sqrt{\hat{A}_t^2 + 4\gamma Q p} \right),$$

$$q^- = \frac{1}{2\gamma Q} \left(-\hat{A}_t - \sqrt{\hat{A}_t^2 + 4\gamma Q p} \right).$$

For $\gamma > 0$ and $p \geq 0$, the correct root is q^+ . We can exclude the other root because $q^- \leq 0$; however, emissions are always positive. The slope of the smart cap is

$$\frac{dq^+}{dp} = \frac{1}{\sqrt{\hat{A}_t^2 + 4\gamma Q p}} > 0.$$

From equation 5 in Proposition 2, we know that the smart cap supports the socially optimal level of emissions as a stable competitive equilibrium when $\gamma = SCC_E^* > 0$. This smart cap is an increasing function of the certificate price.

Proposition 4 shows that the smart tax is stable even for $\gamma < 0$. Therefore, condition (b) in Part i of Proposition 2 holds when $\gamma < 0$. In this case, from Part ii of Proposition 2, the smart cap supports the optimal outcome as a stable competitive equilibrium if and only if the slope of the smart cap is negative (i.e. where the denominator of equation equation 5 is negative).

Thus, for $\gamma < 0$ the smart cap equals the negative root, given by q^- ; the derivative of this root w.r.t. the certificate price is negative. For $\gamma < 0$, this root is defined only for $p \leq -\frac{\hat{A}_t^2}{4\gamma Q}$. At this supremum, the redemption function is $q^- = -\frac{\hat{A}_t}{2\gamma Q}$, and emissions equal $E = -\frac{\hat{A}_t}{2\gamma}$. At the infimum price, $p = 0$, the value of the negative

root is $q^- = -\frac{\hat{A}_t}{\gamma Q}$; here, emissions equal $-\frac{\hat{A}_t}{\gamma}$. In summary, q^- is a decreasing function of p defined for $p \in [0, \frac{\hat{A}_t^2}{-4\gamma Q}]$. Over this domain, the smart tax is positive and the marginal outlay curve corresponding to that inverse supply function is negative. The corresponding range of emissions is $E \in (-\frac{\hat{A}_t}{2\gamma}, -\frac{\hat{A}_t}{\gamma})$. ■

Proof of Proposition 6. Karp and Traeger (2019) show that taxes welfare-dominate quotas if and only if $\frac{\lambda}{f} < \frac{1}{\beta} - \frac{2\mu}{\alpha}$, i.e. if and only if

$$\lambda < f \left(\frac{1}{\beta} - \frac{2\mu}{\alpha} \right) = \frac{f}{\alpha\beta} (\alpha - 2\beta\mu). \quad (35)$$

The proof of Proposition 4 provides the formula for the slope of the smart tax, $\gamma = \frac{\beta\alpha\lambda + \mu f}{\alpha - \beta\mu}$. From Lemma 1 we have $\lambda > 0$ and (for $\rho > 0$, as we assume) $\mu > 0$. Therefore, if $\alpha - \beta\mu < 0$ the slope of the smart tax is negative, and thus less than f , which is positive. In this case, $\alpha - 2\beta\mu < 0$ so inequality 35 is never satisfied. Thus, for $\alpha - \beta\mu < 0$ the slope of the smart tax is less than f and quotas dominate taxes.

For $\alpha - \beta\mu > 0$, where $\gamma > 0$, we have

$$\begin{aligned} \gamma < f &\Leftrightarrow \beta \frac{\alpha\lambda + \mu f}{\alpha - \beta\mu} - f = \alpha \frac{\beta}{\alpha - \beta\mu} \lambda + \left(\beta \frac{\mu}{\alpha - \beta\mu} - 1 \right) f < 0 \Leftrightarrow \\ \lambda &< \frac{\alpha - \beta\mu}{\alpha\beta} \left(1 - \beta \frac{\mu}{\alpha - \beta\mu} \right) f = \frac{f}{\alpha\beta} (\alpha - 2\beta\mu), \end{aligned}$$

reproducing inequality 35. ■

Proof of Proposition 7. Permitting trade of certificates across periods introduces two changes to the original smart cap construction; (i) we have a joint constraint on the certificate number and (ii) we have to respect the arbitrage opportunity for certificates. Our proof proceeds in 4 steps. In step 1 we establish two observations repeatedly used in the proof. The first observation shows that we can rescale the certificate price without changing the physical emission allocation. The second observation shows that there is a unique relation between the technology realization and the certificate price. Step 2 constructs redemption functions that satisfy a no-arbitrage condition for the two-period case. Step 3 uses the two-period case as the basis for an inductive proof extending the construction to an arbitrary number of periods. Steps 2 and 3 construct certificate allocations in each period that achieve optimal emissions in all periods and states of the world. Step 4 defines the aggregate certificate supply and shows that the constructed market allocations indeed form an equilibrium.

Notation: We denote the periods of a given trading phase by $t = 1, \dots, T$. We let $q_t^*(\cdot)$ and Q_t^* denote a period t redemption function and certificate allocation that yield an optimal emission level in the standard smart cap, where emission certificates are valid only in a given period.

Step 1: Two observations.

Observation 1: We can rescale a (one-period) smart cap by $\lambda > 0$ as $\lambda * p_t$, $\lambda * q_t$ and $\frac{Q_t}{\lambda}$ without changing the physical allocation and price of emissions. As a result, by rescaling the number of certificates (and the redemption function accordingly), we can set the certificate price to any desired level while maintaining optimality. This observation will help us in satisfying the no-arbitrage equation; for all periods $t \in \{1, \dots, T - 1\}$ and in every state of the world (θ_{t-1}, S_t) no-arbitrage requires

$$p_t(\theta_t) = \beta \mathbb{E}_t[p_{t+1} | \theta_t] \quad \forall \theta_t, \quad (36)$$

Expectations are taken w.r.t. θ_{t+1} conditional on θ_t , i.e., they are expectations w.r.t. ϵ_{t+1} . We note that the conditional next-period expectations also depend on S_{t+1} , but given a particular mechanism, S_{t+1} is a direct consequence of the technology realization in θ_t that we already conditioned upon.

Observation 2: Any redemption function achieving the first best allocation has to imply that the certificate price p_t responds in a strictly monotonic manner to the realization of the technology level θ_t . From equations (21) and (24), the full information SCC and thus the optimal emissions price, p_t^E , is a strictly increasing function of the shock. Differentiating both sides of the definition $p_t^E = \frac{p_t}{q_t(p_t)}$ yields

$$p_t'(\theta_t) \frac{q_t(p_t(\theta_t)) - p_t(\theta_t) q_t'(p_t(\theta_t))}{q_t(p_t(\theta_t))^2} = \frac{dp_t^E(\theta_t)}{d\theta_t} \neq 0. \quad (37)$$

Given that emissions are strictly positive by assumption ($q_t > 0$), a necessary condition for satisfying inequality (37) is $p_t'(\theta_t) \neq 0$.

Step 2: Two-period construction \mathcal{E} inductive basis.

We start with the first period redemption function $q_1^*(\cdot)$. We *assume* that the equilibrium allocation of certificates in the first period is Q_1^* , ensuring an optimal emission allocation. Each realization of the technology shock implies the optimal certificate price $p_1(\theta_1)$. By Observation 2, this relation is strictly monotonic and, thus, invertible. We denote the inverse relationship by the function f_1 so that $\theta_1 = f_1(p_1)$. The realization of θ_1 also affects the emissions in period 1 and, thus, S_2 . As a result, period 2 starts

out with the state variables $\theta_1 = f_1(p_1)$ and $S_2(\theta_1) = S_2(f_1(p_1))$. We denote the latter relationship by $S_2 = g_1(p_1)$.

For each realization of θ_1 we have to satisfy the no-arbitrage equation (36) between periods 1 and 2. By Observation 1, we can rescale the period 2 certificate price p_2 to any desired level by rescaling $q_2^*(\cdot)$ and Q_2^* accordingly. Thus, for every $\theta_1 = f_1(p_1)$ and $S_2 = g_1(p_1)$ we can define $q_2^{p_1}(\cdot)$ and $Q_2(p_1)$, such that the equilibrium price expectation in period 2 satisfies the no-arbitrage equation (36) between periods 1 and 2. Again, at this point we merely *assume* that the equilibrium certificate allocation will indeed be $Q_2(p_1)$ and return to this assumption in step 4.

Summary and basis of the inductive proof. For $t = 1$ we have established (i) the existence of redemption functions $q_t^{p_1, \dots, p_{t-1}} = q_1$ and $q_{t+1}^{p_1, \dots, p_t} = q_2^{p_1}$ satisfying the no-arbitrage condition (36) under the assumption that equilibrium certificate allocations are $Q_t(p_1, \dots, p_{t-1}) = Q_1^*$ and $Q_{t+1}(p_1, \dots, p_t) = Q_2(p_1)$, and (ii) the existence of functions $f_t^{p_1, \dots, p_{t-1}}(p_t)$ and $g_t^{p_1, \dots, p_{t-1}}(p_t)$ such that $\theta_t = f_t^{p_1, \dots, p_{t-1}}(p_t)$ and $S_{t+1} = g_t^{p_1, \dots, p_{t-1}}(p_t)$. These functions are conditioned on the first period stock variables, which we suppress, as well as the certificate price realizations of the preceding trading periods.

Step 3: Inductive step – extension to arbitrary number of periods.

We now construct a period $t + 2$ redemption function satisfying the no-arbitrage condition (36) between periods $t + 1$ and $t + 2$ as well as functions $f_{t+1}^{p_1, \dots, p_t}$ and $g_{t+1}^{p_1, \dots, p_t}$ expressing the states as a realization of past shocks. Given the stock levels θ_t and S_{t+1} , we employ Observation 2 just as we did for the first period to derive functions f_{t+1}^* and g_{t+1}^* such that $\theta_{t+1} = f_{t+1}^*(p_{t+1})$ and $S_{t+2} = g_{t+1}^*(p_{t+1})$. Incorporating stock dependence, we condition these functions explicitly on the states θ_t and S_{t+1} at the beginning of period $t + 1$, which we can express as functions of the historic price realizations by the induction hypothesis. Thus, we obtain the desired functions $f_{t+1}^{p_1, \dots, p_t}$ and $g_{t+1}^{p_1, \dots, p_t}$ needed for our induction step.²⁴

We (continue to) assume given certificate allocations in period $t+2$, proving they are part of an intertemporal equilibrium in step 4. Given θ_{t+1} and S_{t+2} we can rescale the period $t+2$ certificate price p_{t+2} to any desired level by rescaling $q_{t+2}^*(\cdot)$ and Q_{t+2}^* accordingly (Observation 1). Using this observation, we rescale (for each realization of p_{t+1} and corresponding realizations of $\theta_{t+1} = f_{t+1}^{p_1, \dots, p_t}(p_{t+1})$ and $S_{t+2} = g_{t+1}^{p_1, \dots, p_t}(p_{t+1})$) the

²⁴In detail, making the conditionality of the functions f_{t+1}^* explicit we have $\theta_{t+1} = f_{t+1}^*(p_{t+1}; \theta_t, S_{t+1})$ and define $f_{t+1}^{p_1, \dots, p_t}(p_{t+1}) = f_{t+1}^*(p_{t+1}; f_t^{p_1, \dots, p_{t-1}}(p_t), g_t^{p_1, \dots, p_{t-1}}(p_t))$ and similarly for $g_{t+1}^{p_1, \dots, p_t}$.

certificate price p_{t+2} such that $p_{t+1}(\theta_{t+1}) = \beta \mathbb{E}_{t+1}[p_{t+2}|\theta_{t+1}]$, thereby satisfying the no-arbitrage equation (36) between periods $t+1$ and $t+2$.²⁵ Recognizing that the relation is conditional on the states θ_{t+1} and S_{t+2} , which we expressed as functions of the historic price shocks, we write the resulting smart cap function as $q_{t+2}^{p_1, \dots, p_{t+1}}(p_{t+2})$ and define the corresponding certificate levels for the period $t+2$ allocation as $Q_{t+2}(p_1, \dots, p_{t+1})$. Thereby we completed the inductive step. We are done in period $T-1$ as there is no further arbitrage opportunity in period T .

Step 4: Certificate allocation.

We now return to the point that we have only one certificate constraint for all periods in the trading phase. We let the aggregate emission certificate level supplied to the market be

$$Q(p_1, \dots, p_T) = \sum_{t=1}^T Q_t(p_1, \dots, p_{t-1}), \quad (38)$$

which the decision maker announces at the beginning of the first period. By construction, the certificate allocation using $Q_t(p_1, \dots, p_{t-1})$ certificates in period t is feasible. Again by construction, such a certificate allocation assigning the use of $Q_t(p_1, \dots, p_{t-1})$ certificates to period t implies optimal emission levels in all periods and every state of the world. We are left to show that this certificate allocation is also a market equilibrium.

By construction of the individual redemption functions, for any price sequence p_1, \dots, p_T , firms request $Q_t(p_1, \dots, p_{t-1})$ certificates in period t at the corresponding price p_t in every state of the world (θ_t, S_t) . By our rescaling of the individual redemption functions, the no-arbitrage equation (36) is satisfied in every period and state of world. By definition of the certificate supply (38) the certificate market clears. ■

²⁵As remarked in Observation 1, conditioning on θ_{t+1} simultaneously conditions on S_{t+2} . As remarked in Observation 2, p_{t+1} and θ_{t+1} are strictly monotonic transformations and we can exchange p_{t+1} and θ_{t+1} as the conditioning variables.

D Detailed Steps Solving the Linear-Quadratic Model

Details for Proof of Lemma 1. We define

$$Y_t = \begin{pmatrix} S_t \\ \theta_{t-1} \end{pmatrix}, Q = \begin{pmatrix} -b & 0 \\ 0 & 0 \end{pmatrix}, A = \begin{pmatrix} \delta & 0 \\ 0 & \rho \end{pmatrix}, \quad (39)$$

$$W = \begin{pmatrix} 0 & \rho \end{pmatrix}, B = \begin{pmatrix} \phi \\ 0 \end{pmatrix}, D = \begin{pmatrix} 0 \\ 1 \end{pmatrix}.$$

With these definitions, we write the equation of motion as

$$Y_{t+1} = AY_t + BE_t + D\varepsilon_t$$

and the period payoff as

$$\left((h_t + \alpha\varepsilon_t) E - \frac{1}{2}fE^2 + \frac{1}{2}Y_t'QY_t + WY_tE \right) \phi.$$

Our trial solution for the ex ante value functions is

$$J_t(S_t, \theta_{t-1}) = V_{0t} + V_{1t}'Y_t + \frac{1}{2}Y_t'V_2Y_t \quad (40)$$

with the scalar V_{0t} and

$$V_{1t} = \begin{pmatrix} v_{1t} \\ v_{2t} \end{pmatrix} \text{ and } V_2 = - \begin{pmatrix} \lambda & \mu \\ \mu & \tau \end{pmatrix}. \quad (41)$$

With this notation, we rewrite the right side of the dynamic programming equation 29 as

$$\begin{aligned} & \max_E \left[(h_t + \alpha\varepsilon_t) E - \frac{1}{2}fE^2 + \frac{1}{2}Y_t'QY_t + WY_tE \right] \phi + \\ & \beta(V_{0,t+1} + V_{1,t+1}'(AY_t + BE + D\varepsilon_t)) \\ & + \frac{1}{2}(AY_t + BE + D\varepsilon_t)'V_2(AY_t + BE + D\varepsilon_t)]. \end{aligned} \quad (42)$$

The marginal benefit of an additional unit of emissions in each of the next ϕ years is

$$(h_t + \alpha\varepsilon_t - fE + WY_t) \phi = (h_t - fE + \rho\theta_{t-1} + \alpha\varepsilon_t) \phi.$$

An additional unit of annual emissions increases the next-period stock by ϕ units. The present discounted value of the future stream of marginal cost arising from an additional unit of emissions in each of the next ϕ years is the negative of

$$\begin{aligned} & \beta (V'_{1,t+1}B + B'V_2 (AY_t + D\varepsilon_t) + B'V_2BE) = \\ & -\beta\phi (-v_{1,t+1} + \mu(\rho\theta_{t-1} + \varepsilon_t) + \lambda(\delta S_t + \phi E)). \end{aligned}$$

Therefore, the present discounted value of the future stream of marginal cost arising from a single additional unit of emissions, the SCC_t , is equals

$$SCC_t = \beta (-v_{1,t+1} + \mu(\rho\theta_{t-1} + \varepsilon_t) + \lambda(\delta S_t + \phi E)),$$

as in equation 21.

The first order condition equates the marginal benefit from one additional unit of carbon to its marginal cost:

$$(h_t - fE + \rho\theta_{t-1} + \alpha\varepsilon_t) = \beta (-v_{1,t+1} + \mu(\rho\theta_{t-1} + \varepsilon_t) + \lambda(\delta S_t + \phi E))$$

Solving for E gives the optimal emissions rule

$$E = Z_{0t} + H\varepsilon_t + ZY_t$$

with the definitions

$$\begin{aligned} Z_{0t} &= \frac{h_t + \beta v_{1,t+1}}{f + \beta\lambda\phi}, \quad H = \frac{\alpha - \beta\mu}{f + \beta\lambda\phi} \\ Z &= \begin{pmatrix} z_1 & z_2 \end{pmatrix} = \begin{pmatrix} \frac{-\beta\lambda\delta}{f + \beta\lambda\phi} & \rho \frac{1 - \beta\mu}{f + \beta\lambda\phi} \end{pmatrix}, \end{aligned} \tag{43}$$

as in equations 27 and 28. Below we confirm the second order condition for maximization, $f + \beta\lambda\phi > 0$.

Next we obtain the formulae for the coefficients of the linear-quadratic function $V_t(Y_t)$. We first use the decision rule 27 to eliminate E from the right side of the DPE,

equation 29. The maximized DPE is

$$\begin{aligned}
& j_t(S_t, \theta_{t-1}, \varepsilon_t) = \\
& [(h_t + \alpha\varepsilon_t)(Z_{0t} + H\varepsilon_t + ZY_t) - \frac{1}{2}f(Z_{0t} + H\varepsilon_t + ZY_t)'(Z_{0t} + H\varepsilon_t + ZY_t) \\
& + \frac{1}{2}Y_t'QY_t + WY_t(Z_{0t} + H\varepsilon_t + ZY_t)]\phi \\
& + \beta[V_{0,t+1} + V_{1,t+1}'(AY_t + B(Z_{0t} + H\varepsilon_t + ZY_t) + D\varepsilon_t) \\
& + \frac{1}{2}(AY_t + B(Z_{0t} + H\varepsilon_t + ZY_t) + D\varepsilon_t)'V_2(AY_t + B(Z_{0t} + H\varepsilon_t + ZY_t) + D\varepsilon_t)].
\end{aligned}$$

Now we take expectations w.r.t. ε_t of both sides of (the maximized) DPE, using $\mathbb{E}\varepsilon = 0$, $\mathbb{E}\varepsilon^2 = \sigma^2$ and the definition $J_t(S_t, \theta_{t-1}) \equiv \mathbb{E}_{\varepsilon_t} j_t(S_t, \theta_{t-1}, \varepsilon_t)$. Collecting terms, the result is

$$\begin{aligned}
& V_{0t} + V_{1t}'Y_t + \frac{1}{2}Y_t'V_2Y_t = \\
& (h_tZ_{0t} - \frac{1}{2}fZ_{0t}^2)\phi + \beta(V_{0t+1} + V_{1,t+1}'BZ_{0t} + \frac{1}{2}Z_{0t}^2B'V_2B) \\
& + \frac{1}{2}((2\alpha H - fH^2)\phi + \beta((D + BH)'V_2(D + BH)))\sigma^2 \tag{44} \\
& + ((h_tZ - fZ_{0t}Z + Z_{0t}W)\phi + \beta(V_{1,t+1}'(A + BZ) + Z_{0t}B'V_2(A + BZ)))Y_t + \\
& \frac{1}{2}Y_t'((Q - fZ'Z + W'Z + Z'W)\phi + \beta(A + BZ)'V_2(A + BZ))Y_t
\end{aligned}$$

To obtain the formulae for the parameters of V_2 we equate coefficients of the terms that are quadratic in Y_t , resulting in

$$V_2 = ((Q - fZ'Z + W'Z + Z'W)\phi + \beta(A + BZ)'V_2(A + BZ)).$$

Using the definition of Z in equation 43 and of V_2 in equation 41 and then carrying out the matrix multiplication produces

$$- \begin{pmatrix} \lambda & \mu \\ \mu & \tau \end{pmatrix} = \begin{pmatrix} K_1 & K_2 \\ K_2 & K_3 \end{pmatrix} \tag{45}$$

with the definitions

$$\begin{aligned} K_1 &= -\frac{1}{f+\beta\lambda\phi} (f\beta\lambda\delta^2 + b\beta\lambda\phi^2 + bf\phi) \\ K_2 &= -\beta\delta\frac{\rho}{f+\beta\lambda\phi} (f\mu + \lambda\phi) \\ K_3 &= -\frac{\rho^2}{f+\beta\lambda\phi} (-\phi\beta^2\mu^2 + \lambda\tau\phi\beta^2 + 2\phi\beta\mu + f\tau\beta - \phi). \end{aligned}$$

Comparing the 1,1 elements on both sides gives the relation

$$\lambda = \frac{1}{f + \beta\lambda\phi} (f\beta\lambda\delta^2 + b\beta\lambda\phi^2 + bf\phi).$$

This equation has a positive and a negative root. Using the definition of ϖ in equation 22 gives the positive root in equation 23. For large S_t the value function must be negative; therefore it must be the case that $-\lambda < 0$. Therefore, $f + \beta\lambda\phi > 0$, establishing the second order condition for optimality.

Comparing the 1,2 elements on the left and right side of equation 45 implies

$$\mu = \beta\delta\frac{\rho}{f + \beta\lambda\phi} (f\mu + \lambda\phi).$$

The solution to this equation produces equation 24. We now establish the inequality $1 > \mu > 0$. Because $\rho\beta\delta < 1$ and $\lambda > 0$, both the numerator and the denominator of μ are positive; thus, $\mu > 0$. We also have

$$\mu < 1 \Leftrightarrow \rho\beta\delta\phi\lambda < f + \beta\phi\lambda - \rho\beta\delta f \Leftrightarrow \lambda\beta\phi(\rho\delta - 1) < f(1 - \rho\beta\delta).$$

The last equality holds, because the left side is negative and the right side is positive.

To obtain the formulae for the parameters of V_{1t} we equate coefficients of Y_t on the two sides of equation 44 to obtain

$$V'_{1t} = (h_t Z - f Z_{0t} Z + Z_{0t} W) \phi + \beta (V'_{1t+1} (A + BZ) + Z_{0t} B' V_2 (A + BZ)).$$

We require only the first element of this vector. Equating the 1,1 elements on both sides, we obtain the difference equation 25, repeated here

$$\nu_{1t} = \beta \frac{\delta (f\nu_{1,t+1} - \lambda\phi h_t)}{f + \beta\lambda\phi}.$$

By inspection $\frac{\beta\delta f}{f+\beta\lambda\phi} < 1$. Therefore, a sufficient condition for the solution ν_{1t} to exist, i.e. for the infinite sum in equation 25 to be bounded, is that h_t is bounded for all t . Then we obtain

$$v_{1,t} = \beta \frac{\delta (f\nu_{1,t+1} - \lambda\phi h_t)}{f + \beta\lambda\phi} \Rightarrow v_{1,t} = -\frac{\beta\delta\phi\lambda}{f + \beta\phi\lambda} \sum_{j=0}^{\infty} \left(\frac{\beta\delta f}{f + \beta\phi\lambda} \right)^j h_{t+j}. \quad (46)$$

Our quantitative analysis assumes that $h_{t+j} = \eta^j h_t$ falls at a constant rate ($\eta < 1$). Then

$$\begin{aligned} v_{1,t} &= -\frac{\beta\delta\phi\lambda}{f + \beta\phi\lambda} h_t \sum_{j=0}^{\infty} \left(\eta \frac{\beta\delta f}{f + \beta\phi\lambda} \right)^j \\ &= -\frac{\beta\delta\phi\lambda}{f + \beta\phi\lambda} h_t \frac{1}{1 - \frac{\eta\beta\delta f}{f + \beta\phi\lambda}} \\ &= -\frac{\beta\delta\phi\lambda h_t}{f + \beta\phi\lambda - \eta\beta\delta f} \\ &= -\frac{\lambda\beta\delta h\phi}{f(1 - \beta\delta\eta) + \beta\phi\lambda} < 0 \end{aligned}$$

given $h > 0$. ■