Mathematical Finance

WILEY

MATHEMATICAL FINANCE An International Journal of Mathematics, Statistics and Financial Economics

ORIGINAL ARTICLE OPEN ACCESS

Random Carbon Tax Policy and Investment Into Emission Abatement Technologies

Katia Colaneri[1] | Rüdiger Frey[2] | Verena Köck[2]

[1]Department of Economics and Finance, University of Rome Tor Vergata, Rome, Italy | [2]Institute for Statistics and Mathematics, Vienna University of Economics and Business, Vienna, Austria

Correspondence: Rüdiger Frey (ruediger.frey@wu.ac.at)

Received: 27 May 2024 | **Revised:** 27 March 2026 | **Accepted:** 22 April 2026

Keywords: carbon taxes | climate policy uncertainty | emission abatement | stochastic control | stochastic differential games

ABSTRACT

We analyze the problem of a profit-maximizing electricity producer, subject to carbon taxes, who decides on investments into CO_2 abatement technologies. We assume that the carbon tax policy is random and that the investment in the abatement technology is divisible, irreversible, and subject to transaction costs. Two frameworks for randomness in taxes are considered. First, we assume a precise probabilistic model for the tax process, namely a pure jump Markov process (so-called tax-risk). Second, we analyze the case of a producer who is uncertainty-averse with respect to the tax evolution and who uses a differential game as conceptual tool to decide on optimal production and investment. We provide a rigorous mathematical treatment of both settings, including the analysis of the associated nonlinear PDEs. Numerical methods are employed to investigate the optimal investment strategies. We find that in the tax-risk case, investment in abatement technologies is generally lower than in a benchmark scenario with deterministic taxation. Nevertheless, factors such as production technology, investment divisibility, tax rebates, and credibility of the tax policy introduce interesting twists. In contrast, the uncertainty-averse framework may lead to increased investment as uncertainty rises.

1 | Introduction

Carbon taxes and trading of emission certificates are key policy tools for reducing carbon pollution and hence for mitigating climate change. Academic contributions in this field from an environmental economics perspective have mainly focused on *optimal* tax schemes or optimal carbon prices for an efficient emission reduction, see, for instance, the seminal contributions by Nordhaus (1993), Nordhaus (2019), Golosov et al. (2014), and Acemoglu et al. (2012). More recently, this problem has been addressed within the literature on continuous-time stochastic control by, for example, Aid and Biagini (2023), Aid et al. (2025), or Carmona et al. (2021) (these papers are discussed in Section 1.1). While the design of an optimal tax scheme or carbon price is a very relevant research question, in reality, emission tax policy is affected by many unpredictable factors such as changes in political sentiment and election results, lobbying by industry groups, or developments in international climate policy. Therefore, future tax rates are random and long-term emission tax schemes announced by governments are not fully credible from the viewpoint of carbon-emitting producers. This is a prime example of the so-called *climate policy uncertainty*. In environmental economics, it is often argued that policy uncertainty has a negative impact on investments in carbon abatement technology. For instance, the British newspaper The Economist (2023) writes the following:

> Political polarisation [regarding the relevance of climate change] means bigger flip-flops when power changes hands: imagine France under the wind-farm-

> loathing Marine Le Pen. Everywhere, making climate policy less predictable makes it harder for investors to plan for the long term, as they must.

From a policy perspective, the International Energy Agency report (Yang, 2008) investigates how abrupt shifts in carbon prices, often resulting from changes in carbon tax policy, affect investment incentives for low-carbon power-generation technologies. The report concludes that "the greater the level of policy uncertainty, the less effective climate change policies will be at incentivising investment in low-emitting technologies." Empirical support for this view is provided by Basaglia et al. (2025), who construct newspaper-based indices of climate policy uncertainty and use regression analysis to show that higher uncertainty significantly depresses investment in carbon-intensive sectors.

The academic literature on climate policy uncertainty and the adoption of green technologies frequently employs real options models. Notably, Fuss et al. (2008), Blyth et al. (2007), and Yang et al. (2008), the latter two closely linked to the IEA report Yang (2008), investigate how stochastic carbon pricing affects investment decisions in carbon capture and storage. These models consider the following two sources of uncertainty: (i) market-driven fluctuations in electricity and carbon prices, and (ii) policy uncertainty. Policy uncertainty is modeled as a government announcement at a predetermined future time $\bar{t}$, where it is revealed, with known probabilities, whether the carbon pricing scheme will be implemented. Investment is assumed to be immediate, indivisible, and irreversible. Under these conditions, the optimal strategy is to delay investment until the policy decision at $\bar{t}$ is revealed, underscoring the deterrent effect of policy uncertainty on investment. In contrast, Hagspiel et al. (2025) present evidence that unannounced subsidy withdrawals can, in certain cases, lead to higher investment in green technologies than announced ones. This suggests that policy uncertainty may, under specific circumstances, encourage rather than hinder the adoption of green technologies.

The present paper adds to this strand of literature in various ways. We analyze the behavior of a profit-maximizing electricity producer who is subject to emission taxes and has the option to invest in emission abatement technologies. We employ a continuous-time framework that accommodates a broad range of models for randomness in carbon tax policy. Unlike much of the real options literature, we consider divisible investments such as the incremental installation of solar panels, and we assume that the producer chooses the rate at which she invests. Investment is irreversible and subject to transaction costs, which prevent instantaneous adjustments in the investment level. Our framework captures stylized forms of emission abatement, including retrofitting existing gas-fired plants with carbon capture and storage or filter technologies, and investing in new green technologies with low marginal production costs. These cases are discussed in detail in both the theoretical analysis and the numerical simulations.

We explore two distinct approaches to modeling randomness in carbon taxes. First, we assume a precise probabilistic model for the tax process, namely a pure jump Markov process. In decision-theoretic terms, this corresponds to the paradigm of risk, so that we refer to this situation as *tax-risk*. In that case, the producer is confronted with a stochastic control problem with the investment rate as control variable. Second, we consider a producer who is uncertainty-averse with respect to future tax rates and we use a max–min criterion to determine her optimal production and investment strategy. A possible interpretation of the max–min criterion is that of a game between the producer and a fictitious adversary (called nature). The objective of the producer remains that of maximizing expected profits, whereas nature chooses a tax process to minimize the profits of the producer. This setting is referred to as *tax uncertainty*.

We provide a rigorous mathematical analysis of the producer's optimization problem in both settings. In the tax-risk scenario, we characterize the value function as the unique viscosity solution of the corresponding HJB equation, applying general results from Pham (1998), and we establish conditions for the existence of classical solutions. In the tax uncertainty case, the problem can be analyzed as a stochastic differential game; we prove the existence of an equilibrium and characterize the game's value via a classical solution to the Bellman–Isaacs equation. As explicit solutions are available only in exceptional cases, we perform numerical experiments to study the producer's investment behavior.

In our numerical experiments on tax-risk, we consider two models of tax dynamics. In the first, the government may increase taxes at a random future time, for instance to comply with international climate agreements. In the second, elevated taxes may be reversed, for example, if a government with a "brown" agenda replaces a "green" one. We compare the producer's investment decisions in these stochastic settings to a benchmark case with deterministic taxes. Our results indicate that tax uncertainty generally reduces the firm's willingness to invest in abatement technologies, supporting the view that carbon tax volatility can undermine climate policy. However, there are some new interesting twists. Specifically, producers may invest preemptively, before a tax increase is enacted, to hedge against potential future tax burdens. This hedging behavior is not observed in real options models such as Fuss et al. (2008), where tax uncertainty leads to investment delays until policy decisions are clarified. We also examine the effect of an emission-independent tax rebate and find that it incentivizes abatement investment. Finally, our findings highlight the critical role of investor expectations: a credible tax policy—where firms believe that announced tax increases will be implemented and sustained—is significantly more effective than a noncredible one. Next, we analyze optimal investment within a stochastic differential game framework for an uncertainty-averse producer. Notably, the results are reversed in this setting: Greater uncertainty enhances investment in carbon abatement technology, thus benefiting society. Furthermore, a rebate typically reduces investment. These findings highlight the critical role of the modeling paradigm in determining how climate policy uncertainty influences abatement investment.

The remainder of the paper is structured as follows. Section 1.1 reviews some relevant literature; Section 2 introduces the model and the electricity producer's optimization problem; in Section 3, we present specific examples of electricity generation and emission abatement technologies; Section 4 analyzes the control problem under tax-risk; tax uncertainty and the resulting differential game are addressed in Section 5.1; Section 6 reports numerical results and their economic interpretation; and Section 7 concludes.

1.1 | Discussion of Some Related Contributions

We continue with a brief discussion of related contributions. Within the framework of the stochastic control literature on tax- and carbon pricing schemes, Aid and Biagini (2023) study an optimal dynamic carbon emission regulation for a set of firms, in the presence of a regulator who may choose dynamically the emission allowances to each firm. The problem is formulated as a Stackelberg game between the regulator and the firms in a jump diffusion setup with linear quadratic costs. This formulation allows for a closed-form expression of the optimal dynamic allocation policies. Aid et al. (2025) investigate the optimal regulatory incentives that trigger the development of green electricity production in a monopoly and in a duopoly setup. The regulator wishes to encourage green investments to limit carbon emissions, while simultaneously reducing the intermittency of the total energy production. Their main result is a characterization of the regulatory contract that naturally includes interesting agreements like rebate. Carmona et al. (2021) analyze mean field control and mean field game models of electricity producers who interact via electricity spot markets and who can decide on the composition of their energy mix (brown or green) in the presence of a carbon tax. Producers have to balance the cost of intermittency and the amount of carbon tax they pay. The paper analyzes competitive (Nash equilibrium) and cooperative (social optimum) solutions to this problem via systems of forward–backward SDEs. It also includes a study of a Stackelberg game between a tax-setting regulator and the mean field of producers.

The impact of tax policy uncertainty on firm level and aggregate investment is studied in the seminal paper Hassett and Metcalf (1999) in a general real options model that is not directly related to investment in green technology. Lavigne and Tankov (2026), Dumitrescu et al. (2024), and Flora and Tankov (2023) have done interesting theoretical research on the implication of randomness in climate policy more generally. Lavigne and Tankov (2026) consider a mean-field game model for a large financial market where firms determine their dynamic emission strategies under climate transition risk in the presence of green and neutral investors. They show among others that uncertainty about future climate policies leads to overall higher emissions in equilibrium. In a similar spirit, Dumitrescu et al. (2024) study the impact of transition scenario uncertainty on the pace of decarbonization and on output prices in the electricity industry. Flora and Tankov (2023) analyze how uncertainty in climate transition pathways affects investment in assets that may become stranded. Empirical studies on the impact of carbon taxes on green technology adaption include Aghion et al. (2016) who study in particular the effects of taxes and fuel prices on investment in technological innovation using data for the automobile sector, and Martinsson et al. (2024) where data on CO_2 emissions from Swedish manufacturing sector are used to estimate the impact of carbon pricing on firm-level emission intensities.

2 | The Optimization Problem of the Electricity Producer

Throughout, the paper we fix a horizon date T and a probability space $(\Omega, \mathcal{F}, \mathbf{P})$ with filtration $\mathbb{F} = (\mathcal{F}_t)_{0\le t\le T}$ representing the information flow. In the sequel, all processes are assumed to be $\mathbb{F}$-adapted and expectations are taken with respect to the probability measure $\mathbf{P}$.

We consider a profit-maximizing electricity producer subject to carbon taxes, with the option to invest in CO_2 emission abatement technology. Let τ_t denote the tax per unit of emission and X_t the value of abatement investment at time t. The producer determines the electricity output over time and controls the investment process $X = (X_t)_{0\le t\le T}$ through her investment in abatement technology.

To capture the randomness in electricity and fuel prices, as well as in the productivity of electricity generation technologies, we introduce an exogenous d-dimensional factor process $Y = (Y_t)_{0\le t\le T}$. We assume that Y is a d-dimensional diffusion process,

$$dY_t = \beta(t, Y_t)\,dt + \alpha(t, Y_t)\,dB_t, \quad Y_0 = y \in \mathbb{R}^d, \tag{1}$$

where B is a d-dimensional Brownian motion, and where the drift $\beta(t, y) \in \mathbb{R}^d$ and the dispersion $\alpha(t, y) \in \mathbb{R}^{d\times d}$, for $(t, y) \in [0, T] \times \mathbb{R}^d$, satisfy standard conditions for existence and uniqueness of the SDE (1). Moreover, we denote the generator of Y by $\mathcal{L}^Y$, which reads as as follows:

$$\mathcal{L}^Y f(y) = \sum_{i=1}^{d} f_{y_i}(y)\beta_i(t, y) + \frac{1}{2}\sum_{i,j=1}^{d} f_{y_i y_j}(y)\mathfrak{S}_{ij}(t, y),$$

where $\mathfrak{S}(t, y) = \alpha(t, y)\cdot\alpha^\top(t, y)$. We leave the specification of Y fairly general at this stage. In this way, our analytical results remain valid regardless of the specific model used. Moreover, a specific model for the dynamics of Y would not simplify the mathematical analysis.

2.1 | Instantaneous Electricity Production

We assume that the electricity market is perfectly competitive so that the producer acts as a price taker, that is she takes the price $p_t = p(Y_t)$ of one unit of electricity as given and adjusts the quantity produced in order to maximize instantaneous profits. This situation might arise in the context of a merit order system, where the electricity spot price is determined by the short run marginal production cost of the power plant that is on the margin of the electricity production system. For a given investment value x, tax rate τ and value y of the factor process, we denote the cost of producing q units of electricity by $C(q, x, y, \tau)$. Hence the instantaneous profit is given by the following:

$$\Pi(q, x, y, \tau) = p(y)q - C(q, x, y, \tau) + \nu_0(q)\tau\,. \tag{2}$$

The term $\nu_0(q)\tau$ models a *tax rebate* that depends on the amount q of energy produced and on the current tax rate, but not on the actual emissions of the producer. Tax rebates of this form penalize (reward) producers with high (low) emissions compared to the industry average and are part of many proposals for carbon taxes.

The producer chooses the output q to maximize her instantaneous profit and we denote the maximal profit by the following:

$$\Pi^*(x, y, \tau) = \max_{q\in[0,q^{\max}]} \Pi(q, x, y, \tau), \quad (3)$$

where the constant $q^{\max} > 0$ denotes the maximum capacity of the production technology. In the numeric examples, we also consider the situation where the quantity to be produced is fixed to some $\bar{q}$, for instance, since the producer has entered into long-term delivery contracts. In that case, one obviously has $\Pi^*(x, y, \tau) = \Pi(\bar{q}, x, y, \tau)$.

The next assumption gives conditions ensuring that the function Π^* is well-defined and enjoys certain regularity properties.

Assumption 2.1.

i. There are functions $C_0, C_1 : [0, q^{\max}] \times \mathbb{R} \times \mathbb{R}^d \to [0, \infty)$ such that

$$C(q, x, y, \tau) = C_0(q, x, y) + C_1(q, x, y)\tau\,;$$

moreover, C_0 and C_1 are increasing, strictly convex and C^1 in q and C_1 is bounded.

ii. C_0 and C_1 are Lipschitz continuous in x, y uniformly in $q \in [0, q^{\max}]$.

iii. The function ν_0 is differentiable, increasing, and concave on $[0, q^{\max}]$.

iv. The function $p : \mathbb{R}^d \to \mathbb{R}$ is Lipschitz continuous in y.

In economic terms, the function C_1 measures the emissions from producing q units of electricity given the investment level x and the value of the factor process; these are then multiplied with τ to give the instantaneous carbon tax payments; the function C_0 gives the emission-independent production cost. Specific examples are discussed in Section 3 below.

Under Assumption 2.1(i) and (iii), there is a unique optimal instantaneous energy output $q^* \in [0, q^{\max}]$ for every (x, y, τ). Taking derivatives with respect to q gives the first-order condition

$$p(y) - \partial_q C_0(q, x, y) - (\partial_q C_1(q, x, y) - \partial_q \nu_0(q))\tau = 0, \quad (4)$$

which has at most one solution due to the strict convexity of the cost functions. If we consider moreover the boundary cases $q = 0$ and $q = q^{\max}$, we get the following:

$$q^* = \begin{cases} 0, & \text{if } p(y) - \partial_q C_0(0, x, y) - (\partial_q C_1(0, x, y) - \partial_q \nu_0(0))\tau < 0; \\ q^{\max}, & \text{if } p(y) - \partial_q C_0(q^{\max}, x, y) - (\partial_q C_1(q^{\max}, x, y) - \partial_q \nu_0(q^{\max}))\tau > 0; \\ \text{the solution of (4)}, & \text{else.} \end{cases} \quad (5)$$

2.2 | Optimal Investment

We assume that the investment value X has dynamics as follows:

$$X_t = X_0 + \int_0^t \gamma_s ds - \int_0^t \delta X_s ds + \sigma W_t, \quad t \le T. \quad (6)$$

Here $W = (W_t)_{t\ge 0}$ is a Brownian motion, independent of the Brownian motion B in the dynamics of Y, $0 \le \delta < 1$ is the depreciation rate and the term σW_t models exogenous fluctuations of the investment value, due, for example, to random replacement costs. The process $\boldsymbol{\gamma} = (\gamma_t)_{0\le t\le T}$ represents the rate at which the producer invests into abatement technology. We assume that the investment is *irreversible*, that is we introduce the constraint that $\gamma_t \ge 0$ for all t. Moreover, we assume that the investment is subject to proportional *transaction* costs given by $\kappa\gamma^2$. These costs penalize a rapid build-up of abatement technology. By $\mathcal{A}$, we denote the set of all *admissible* investment strategies, that is the set of all adapted nonnegative càdlàg processes $\boldsymbol{\gamma}$ with $\mathbb{E}\left[\int_0^T \gamma_t^2\, dt\right] < \infty$.

The producer uses a cash account D to finance her investments and to invest the profits from selling electricity. We assume that there is a constant interest rate $r \ge 0$ that applies to borrowing and lending. Hence D has the dynamics as follows:

$$dD_t = \big(rD_t + \Pi^*(X_t, \tau_t, Y_t) - (\gamma_t + \kappa\gamma_t^2)\big)dt, \quad D_0 = 0. \quad (7)$$

We interpret the horizon date T as lifetime of the electricity production technology, and we model the residual value of the investment by a function $h(X_T)$, which is nonnegative, increasing and continuous and whose form will depend on the type of abatement technology.

The goal of the electricity producer is to maximize $\mathbb{E}\left[e^{-rT}(D_T + h(X_T))\right]$, the expected discounted value of her terminal cash position and of the residual investment. Next we show that the initial value of the cash account does not affect the investment decision of the producer. In fact, using Equation (7), we get that

$$e^{-r(T-t)}D_T = D_t - \int_t^T re^{-r(s-t)}D_s ds + \int_t^T e^{-r(s-t)}(rD_s + \Pi^*(X_s, Y_s, \tau_s) - \gamma_s - \kappa\gamma_s^2)ds.$$

Hence,

$$\mathbb{E}\left[e^{-rT}(D_T + h(X_T))\right] = D_0 + \mathbb{E}\left[\int_0^T e^{-rs}\big(\Pi^*(X_s, \tau_s, Y_s) - \gamma_s - \kappa\gamma_s^2\big)ds + e^{-rT}h(X_T)\right], \quad (8)$$

and the goal of the producer amounts to maximizing the second term in Equation (8).

In this paper, we study the optimization problem of the producer in two settings that differ with respect to the modeling of randomness in carbon taxes. In Section 4, we analyze the case of tax-risk, where the tax process τ follows a precise probabilistic model, namely a pure jump process with given jump intensity and jump size distribution. In Section 5.1, we consider an alternative approach corresponding to the case of tax uncertainty. There the electricity producer considers a set of possible future tax scenarios and uses a max–min criterion to determine her production and investment strategy.

3 | Production Technologies: Examples

We now discuss two specific examples for the production and emission abatement technology that will be used in the numerical experiments.

3.1 | The Filter Technology

In this example, we assume that the producer is using a brown technology such as coal-fired power plants but is able to reduce CO_2 pollution by investing in a carbon capture and storage or filter technology. We let ζ be the input good, that is, the amount of raw material (coal or gas), which is needed to produce electricity. We suppose that one unit of raw material has a cost of $\bar{c}(y)$ dollars and that for each unit of raw material used in the production process, the amount of emitted CO_2 is e_0. If filters are installed, emissions per unit of raw material are reduced by $e(x)$. The emission reduction depends clearly on the quality and the number of filters, and hence on the investment level x. Given an investment level x, total emissions for ζ units of raw material are thus given by $\zeta(e_0 - e(x))$.

We denote by $P(\zeta)$ the amount of electricity that can be produced using ζ units of raw material, for a continuous increasing and concave function P with $P(0) = 0$. Denote by $Q(\cdot)$ the inverse function of P. Then, to produce the amount q of electricity, the producer needs $\zeta = Q(q)$ units of raw material and hence the incurred cost (production cost and taxes) is given by the following:

$$C(q, x, y, \tau) = Q(q)(\bar{c}(y) + (e_0 - e(x))\tau). \tag{9}$$

Note that in this example, the functions C_0 and C_1 from Assumption 2.1 are given by $C_0(q, x, y) = Q(q)\bar{c}(y)$ and $C_1(q, x, y) = Q(q)(e_0 - e(x))$. Recall that we interpret the horizon date T as lifetime of the brown power plant. It makes sense to assume that the residual value of the filters installed is zero once the power plant is no longer in operation, so that for the filter technology, we take $h(X_T) = 0$.

3.2 | Two Technologies

Next we consider a situation where the energy producer has the option to replace a brown technology, such as coal or gas power plants, by a green technology, such as wind or solar energy. We denote by ζ_b the amount of input material for the brown technology and suppose that one unit of input material costs $c_b(y)$ dollars and leads to e_b tons of CO_2. Let $P_b(\zeta)$ be the amount of electricity that can be produced with ζ units of raw material and assume that P_b is increasing and concave and $P_b(0) = 0$.

The input material to produce green, on the other hand, has a price of zero (for instance wind or sun) and for simplicity, we assume that green technology does not emit CO_2. We associate the investment level x with the amount of green production facilities (solar panels or wind turbines) installed and we denote by $P_g(x)$ be the maximum amount of electricity that can be produced with the green technology for a given investment level x. We assume that the maintenance cost $c_g(x)$ of the green technology only depends on the investment x. Denote by $Q_b(\cdot)$ the inverse function of P_b. Then the total cost for producing q units of energy is

$$C(q, x, y, \tau) = \begin{cases} c_g(x) & \text{if } q - P_g(x) \le 0\,; \\ c_g(x) + (c_b(y) + e_b\tau)Q_b\big(q - P_g(x)\big) & \text{if } q - P_g(x) > 0\,; \end{cases} \tag{10}$$

equivalently, $C(q, x, y, \tau) = c_g(x) + (c_b(y) + e_b\tau)Q_b((q - P_g(x)^+)$. In this example, we get that

$$C_0(x, y, q) = c_g(x) + c_b(y)Q_b\Big(\big(q - P_g(x)\big)^+\Big), \tag{11}$$

$$C_1(x, y, q) = e_bQ_b\Big(\big(q - P_g(x)\big)^+\Big). \tag{12}$$

These functions satisfy the regularity conditions stated in Assumption 2.1(i)–(iii) if c_g and P_g are Lipschitz, Q_b is C^1, increasing, strictly convex and $(Q_b)'(0) = 0$, with the exception that C_0 is strictly convex in q only for $q \ge P_g(x)$. This is however sufficient for the existence of a unique optimal electricity output q^*, which is given by Equation (5). We will also assume that c_g and P_g are increasing in the investment level x to make the model reasonable from an economic viewpoint.

In the numerical experiments, we take $Q_b(q) = aq^{3/2}$, which fits the conditions above. Moreover, we use a function P_g of the following form:

$$P_g(x) = p_g[(x - \bar{x})^+]^\alpha, \quad \alpha \in (0, 1),$$

for some productivity parameter $p_g \in (0, 1)$, where $\bar{x}$ represents initial expenses such as land acquisition and infrastructure development for connecting to the grid that the electricity company must bear when building a green power plant. This example shows that our model may account for threshold effects, even when the investment occurs continuously in time.

To simplify the exposition, we concentrate on the case where the maximum production level $q^{\max}$ is independent of x. In the two-technology example, this bound can be interpreted as the maximum amount of electricity that could be absorbed by the grid. However, it might also make sense to consider the case where $q^{\max}$ depends on the amount of green technology installed and hence on the investment level, that is, to model the maximum capacity as a function $\bar{q}(x)$. This choice brings additional technicalities in some of the theoretical results. We refer to the comment on maximum capacity expansion in Appendix B.1 for further discussion.

4 | Tax Risk and Stochastic Control

In this section, we analyze the case of tax-risk, where the tax process $\boldsymbol{\tau}$ follows a fully specified probabilistic model, namely a Markovian pure jump process with given jump intensity and jump size distribution. From now on, we use the notation $\mathbb{E}_t[\cdot]$ to indicate the conditional expectation given $X_t = x, Y_t = y, \tau_t = \tau$. The reward function of the optimization problem (8) is thus given by

$$J(t,x,y,\tau,\boldsymbol{\gamma}) = \mathbb{E}_t\left[\int_t^T e^{-r(s-t)}\big(\Pi^*(X_s,\tau_s,Y_s) - \gamma_s - \kappa\gamma_s^2\big)ds + e^{-r(T-t)}h(X_T)\right], \tag{13}$$

and we denote by $V(t,x,y,\tau) = \sup\{J(t,x,y,\tau,\boldsymbol{\gamma}) : \boldsymbol{\gamma} \in \mathcal{A}\}$ the corresponding value function. The main goal of this section is to characterize V as viscosity solution of a certain Hamilton–Jacobi–Bellman (HJB) equation and to give criteria for the existence of a classical solution.

4.1 | The Tax Process

In the tax-risk case, we work with the following mathematical model for the tax process. Let $N(dt,dz)$ be a homogeneous Poisson random measure with intensity measure $m(dz)dt$, where $m(dz)$ is a finite measure on a compact set $\mathcal{Z} \subset \mathbb{R}$, that is, $m(\mathcal{Z}) = M < \infty$. Then the tax process satisfies the following SDE:

$$\tau_t = \tau_0 + \int_0^t \int_{\mathcal{Z}} \Gamma(s, Y_{s-}, \tau_{s-}, z)\, N(ds,dz), \quad t \in [0,T], \tag{14}$$

for a function $\Gamma(t,y,\tau,z) : [0,T]\times\mathbb{R}^d\times\mathbb{R}\times\mathcal{Z} \to \mathbb{R}$ such that Equation (14) has a unique solution, which is then automatically a Markov process. (A set of sufficient conditions is listed in Assumption 4.2-(iii)). In the sequel, to avoid technicalities, we assume that there exists $\tau^{\max} > 0$ such that $\tau_t \in [0, \tau^{\max}]$ for all $t \in [0,T]$. This translates into the following conditions on the function Γ:

$$\sup_{z\in\mathcal{Z}} \Gamma(t,y,\tau,z) \le \tau_{\max} - \tau \text{ and } \inf_{z\in\mathcal{Z}} \Gamma(t,y,\tau,z) \ge -\tau, \tag{15}$$

for all $t \in [0,T]$, $y \in \mathbb{R}^d$, $\tau \in [0,\tau^{\max}]$. We denote by $\mathcal{L}^\tau$ the generator of τ, for given t and value of the factor process y, that is

$$\mathcal{L}^\tau f(t,y,\tau) = \int_{\mathcal{Z}} \big\{f(t,y,\tau+\Gamma(t,y,\tau,z)) - f(t,y,\tau)\big\}\, m(dz). \tag{16}$$

The process τ features totally inaccessible jump times, implying that tax rate changes are "surprises" for the producer. This constitutes a stylized model, as in practice tax changes may be announced randomly but are in effect implemented at fixed future dates to minimize disruption to firms. However, when the delay between announcement and implementation is short, this simplification has limited impact on the qualitative features of the optimal investment strategy due to transaction costs. In particular, firms cannot defer investment decisions until a tax increase is announced, as rapidly expanding abatement capacity within the short interval between announcement and implementation is prohibitively costly.

Example 4.1. A basic example is given by a two-state Markov chain with states $\tau^1 < \tau^2$ and transition rate matrix $G = (g_{ij})_{i,j\in\{1,2\}}$, where $g_{ii} = -g_{ij}$ for $j \ne i$. The generator of τ is then given by

$$\mathcal{L}^\tau f(\tau) = \sum_{j=1}^{2} \mathbf{1}_{\{\tau=\tau^j\}} \sum_{i=1}^{2} g_{ji} f(\tau^i).$$

Here, $\tau_t = \tau^1$ may represent a low tax regime under a government with limited environmental focus, while $\tau_t = \tau^2$ corresponds to a higher tax consistent with a government with a green policy agenda; transitions between tax levels then reflect shifts in political leadership. A Markov switching model is a natural choice to represent tax uncertainty and it was also considered in the seminal paper Hassett and Metcalf (1999). In Section 6, we consider two special Markov switching models for the tax evolution, namely the *tax increase* and the *tax reversal*. In the tax increase case, we assume that $\tau_0 = 0$ and that the process jumps upward to $\tau^2 > 0$ at a random time. This is a stylized model for the situation where a government plans to rise carbon taxes in order to comply with international climate agreements but where the exact timing of the tax rise depends on random political factors. In the numerical experiments, we moreover assume that the high tax value is an absorbing state. In the case of a tax reversal, the tax is initially set to a high value ($\tau_0 = \tau^2$), but it jumps downward to τ^1 at a random time, due to lobbying efforts or a change in government composition. In this case, we assume that after the tax reversal, the tax rate may increase again to τ^2.

A two-state Markov chain can be embedded into the family of pure jump processes of the form (14) with the following identification: We let $\mathcal{Z} = \{0,1\}$, $m^\tau(dz) = g_{12}\,\delta_{\{0\}}(dz) + g_{21}\,\delta_{\{1\}}(dz)$ and put

$$\Gamma(\tau,0) = [\tau^2-\tau]^+ - [\tau^1-\tau]^+ \text{ and } \Gamma(\tau,1) = [\tau-\tau^2]^+ - [\tau-\tau^1]^+\,.$$

In this way, Γ is bounded and Lipschitz in τ (see Assumption 4.2 below).

4.2 | Properties of the Value Function

In the sequel, we use the notation L_f for the Lipschitz constant of a function f. We need the following set of conditions for our analysis.

Assumption 4.2.

i. Function $h(x)$ is Lipschitz in x.
ii. Functions β, α are continuous and globally Lipschitz.
iii. Function $\Gamma(t,y,\tau,z)$ is continuous in t,y,τ,z, Lipschitz in y,τ, for all $t \in [0,T]$, and for all $z \in \mathcal{Z}$ and satisfies Equation (15).

Lemma 4.3. *Let V be the value function of the problem (13). Suppose that Assumptions 2.1 and 4.2 hold. Then,*

i. *Π^* is Lipschitz continuous in (x, y, τ);*

ii. *V is Lipschitz in x, uniformly in t, τ, y, with Lipschitz constant*

$$L_V = \frac{L_{\Pi^*}(1 - e^{-(r+\delta)T})}{r+\delta} + L_h. \tag{17}$$

iii. *Suppose moreover that Π^* and h are increasing in x, then V is increasing in x as-well.*

Proof. We begin with statement (i). To prove this, we will use Assumption 2.1. By direct computations, we get that

$$|\Pi^*(x^1, y^1, \tau^1) - \Pi^*(x^2, y^2, \tau^2)| = \left| \max_{q\in[0,q^{\max}]} \Pi(x^1, y^1, \tau^1, q) - \max_{q\in[0,q^{\max}]} \Pi(x^2, y^2, \tau^2, q) \right| \tag{18}$$

$$\leq \max_{q\in[0,q^{\max}]} \left|\Pi(x^1, y^1, \tau^1, q) - \Pi(x^2, y^2, \tau^2, q)\right|. \tag{19}$$

Moreover, we get from the definition of Π that

$$\begin{aligned}\left|\Pi(x^1, y^1, \tau^1, q) - \Pi(x^2, y^2, \tau^2, q)\right| &\leq |p(y^1) - p(y^2)|q + |C_0(q, x^1, y^1) - C_0(q, x^2, y^2)| \\ &\quad + |\tau^1 C_1(q, x^1, y^1) - \tau^2 C_1(q, x^2, y^2)| + |\tau^1 - \tau^2|\nu_0(q) \\ &\leq q^{\max} L_p |y^1 - y^2| + L_{C_0}(|y^1 - y^2| + |x^1 - x^2|) + \tau^{\max} L_{C_1}(|y^1 - y^2| + |x^1 - x^2|) \\ &\quad + (||C_1||_\infty + ||\nu_0||_\infty)|\tau^1 - \tau^2|,\end{aligned}$$

so that Π is Lipschitz continuous in (x, y, τ) uniformly in $q \in [0, q^{\max}]$. Here $||\cdot||_\infty$ represents the supremum norm.

Next we establish (ii). By (i), we know that Π^* is Lipschitz continuous in (x, y, τ) with Lipschitz constant L_{Π^*}. Next we let X^1 and X^2 be the solutions of Equation (6) with the initial conditions $x^1 \neq x^2$, respectively, that is $X^i_t = x^i + \int_0^t (\gamma_s - \delta X_s)\, ds + \sigma W_t$, for $i = 1, 2$. Then we get that $X^1_t - X^2_t = (x^1 - x^2)e^{-\delta t}$ for $t \geq 0$ and

$$|V(t, x^1, y, \tau) - V(t, x^2, y, \tau)| \tag{20}$$

$$\leq \sup_{\gamma\in\mathcal{A}} \mathbb{E}_t \left[\int_t^T \left|\Pi^*(X^1_s, Y_s, \tau_s) - \Pi^*(X^2_s, Y_s, \tau_s)\right| e^{-r(s-t)} ds + e^{-r(T-t)} \left|h(X^1_T) - h(X^2_T)\right| \right] \tag{21}$$

$$\leq \sup_{\gamma\in\mathcal{A}} \mathbb{E}_t \left[\int_t^T L_{\Pi^*} |X^1_s - X^2_s| e^{-r(s-t)} ds + e^{-r(T-t)} L_h |X^1_T - X^2_T| \right] \tag{22}$$

$$\leq \int_t^T L_{\Pi^*} |x^1 - x^2| e^{-(r+\delta)(s-t)} ds + e^{-(r+\delta)(T-t)} L_h |x^1 - x^2| \tag{23}$$

$$= |x^1 - x^2| \left(\frac{L_{\Pi^*}(1 - e^{-(r+\delta)(T-t)})}{r+\delta} + e^{-(r+\delta)(T-t)} L_h \right). \tag{24}$$

This shows that V is Lipschitz in x, uniformly in t, τ, y with uniform Lipschitz constant L_V.

For (iii) note that if h and Π^* are increasing in x, the reward function (13) of the problem is increasing in x, which carries over to the value function V by definition. □

Remark 4.4 (Comments and extensions). 1. The argument in the proof of Lemma 4.3-(i) can be applied to establish the regularity of the function Π^* in cases where investment may expand the maximum capacity. Such an extension is particularly relevant in the context of the two-technology example.

2. It is possible to show that if Π^* and h are concave in x, then V is also concave in x. This situation arises, for instance, if $\Pi(q, x, y, \tau)$ is concave in x and if q^* is a fixed quantity.

The mathematical details of these two extensions are discussed in Appendix B.1.

4.3 | Viscosity Solutions

For mathematical reasons, we first assume that the set of admissible controls is bounded and we denote by $\mathcal{A}^{\bar\gamma} \subset \mathcal{A}$ the set of all adapted càdlàg processes $\boldsymbol{\gamma}$ with $0 \leq \gamma_t \leq \bar\gamma$ for all t. Let,

$$V^{\bar\gamma}(t, x, y, \tau) := \sup_{\gamma\in\mathcal{A}^{\bar\gamma}} \mathbb{E}_t \left[\int_t^T \left(\Pi^*(X_s, Y_s, \tau_s) - \gamma_s - \kappa\gamma_s^2\right) e^{-r(s-t)} ds + e^{-r(T-t)} h(X_T) \right]. \tag{25}$$

As a first step, we show that $V^{\bar\gamma}$ is the unique viscosity solution of the HJB equation

$$\begin{aligned} &v_t(t, x, y, \tau) + \Pi^*(x, y, \tau) + \mathcal{L}^\tau v(t, x, y, \tau) + \mathcal{L}^Y v(t, x, y, \tau) + \frac{\sigma^2}{2} v_{xx}(t, x, y, \tau) \\ &+ \sup_{0\leq\gamma\leq\bar\gamma} \left\{ v_x(t, x, y, \tau)(\gamma - \delta x) - (\gamma + \kappa\gamma^2) \right\} = -rv(t, x, y, \tau) \end{aligned} \tag{26}$$

with the terminal condition $v(T, x, y, \tau) = h(x)$.

Proposition 4.5. *The function $V^{\bar\gamma}$ is Lipschitz in (x, y, τ) and Hölder in t and the unique viscosity solution of Equation (26). Moreover, a comparison principle holds for that equation.*

Proof. It is easy to check that for the problem (25), the hypotheses (2.1)-(2.5) of Pham (1998) are satisfied. Then, the result follows from Pham (1998, Theorem 3.1). Note that in Pham (1998), it is assumed that the controls take values in a compact set, so that the results of that paper apply only to the case where $\boldsymbol{\gamma} \in \mathcal{A}^{\bar\gamma}$. □

Next we want to prove that $V^{\bar\gamma}$ is independent of $\bar\gamma$ for sufficiently large values of $\bar\gamma$. For this, we use that, in view of Lemma 4.3, the value function $V^{\bar\gamma}$ is Lipschitz with Lipschitz constant L_V as in Equation (17); in particular, the Lipschitz constant may be taken independent of $\bar\gamma$.

Proposition 4.6. *Consider constants $\bar{\gamma}^1 < \bar{\gamma}^2$ such that $\frac{(L_V-1)^+}{2k} < \bar{\gamma}^1$. Then $V^{\bar{\gamma}^1} = V^{\bar{\gamma}^2}$.*

Proof. For $i = 1, 2$, denote by $V^{\bar{\gamma}^i}(t, x, y, \tau)$ the value function of the optimization problem with strategies in $\mathcal{A}^{\bar{\gamma}^i}$. Since $\bar{\gamma}^1 < \bar{\gamma}^2$, it is immediate that $\mathcal{A}^{\bar{\gamma}^1} \subset \mathcal{A}^{\bar{\gamma}^2}$ and hence $V^{\bar{\gamma}^2}(t, x, y, \tau) \geq V^{\bar{\gamma}^1}(t, x, y, \tau)$. To establish the opposite inequality, that is, $V^{\bar{\gamma}^1}(t, x, y, \tau) \geq V^{\bar{\gamma}^2}(t, x, y, \tau)$, we prove that $V^{\bar{\gamma}^1}$ is a viscosity supersolution of the HJB equation (26) with $\bar{\gamma} = \bar{\gamma}^2$.

Fix some point (t_0, x_0, y_0, τ_0). Since $V^{\bar{\gamma}^1}$ is a viscosity solution (and hence in particular a supersolution) of Equation (26) with $\bar{\gamma} = \bar{\gamma}^1$, for every smooth function $\phi(t, x, y, \tau)$ such that

$$\phi(t, x, y, \tau) \leq V^{\bar{\gamma}^1}(t, x, y, \tau) \text{ for all } (t, x, y, \tau) \text{ and } \phi(t_0, x_0, y_0, \tau_0) = V^{\bar{\gamma}^1}(t_0, x_0, y_0, \tau_0) \tag{27}$$

it holds that

$$-\Big(\phi_t(t_0, x_0, y_0, \tau_0) + \Pi^*(x_0, y_0, \tau_0) + \mathcal{L}^\tau V^{\bar{\gamma}^1}(t, x, y, \tau) + \mathcal{L}^Y \phi(t, x, y, \tau) + \frac{\sigma^2}{2}\phi_{xx}(t, x, y, \tau)$$
$$+ \sup_{0 \leq \gamma \leq \bar{\gamma}^1} \{\phi_x(t, x, y, \tau)(\gamma - \delta x) - (\gamma + \kappa\gamma^2)\} - rV^{\bar{\gamma}^1}(t, x, y, \tau)\Big) \geq 0 \tag{28}$$

It follows from Equation (27) that $\phi_x(t_0, x_0, y_0, \tau_0) \leq L_V$ with ϕ_x being the partial derivative of ϕ with respect to x. Note that the supremum in Equation (28) is attained at $\gamma^* = \frac{(\phi_x-1)^+}{2k} \leq \frac{(L_V-1)^+}{2k} < \bar{\gamma}^1$. Hence we can replace $\bar{\gamma}^1$ with $\bar{\gamma}^2$ in Equation (28) without changing the supremum. This implies that $V^{\bar{\gamma}^1}$ is a supersolution of Equation (26) with $\bar{\gamma} = \bar{\gamma}^2$ and completes the proof. □

Summarizing, we have the following result.

Theorem 4.7. *Fix $\bar{\gamma} > \frac{(L_V-1)^+}{2k}$. Then the value function V of the optimization problem (13) equals $V^{\bar{\gamma}}$. It follows that V is Lipschitz in (x, y, τ), Hölder in t and the unique viscosity solution of the HJB equation (26).*

Proof. In the sequel, we show that $V(t, x, y, \tau) = V^{\bar{\gamma}}(t, x, y, \tau)$ for $\bar{\gamma} > \frac{(L_V-1)^+}{2k}$, hence V is the unique viscosity solution of the HJB equation (26) and inherits the regularity properties of $V^{\bar{\gamma}}$ from Proposition 4.5.

In view of Proposition 4.6, it remains to show that $V(t, x, y, \tau) = \lim_{m\to\infty} V^m(t, x, y, \tau)$ (V^m is the solution of Equation 26 with $\bar{\gamma} = m$.) The inequality $V(t, x, y, \tau) \geq \lim_{m\to\infty} V^m(t, x, y, \tau)$ is clear, since $\mathcal{A}^m \subset \mathcal{A}$. For the converse inequality, we observe that for all $\boldsymbol{\gamma} \in \mathcal{A}$, there is a sequence of strategies $\boldsymbol{\gamma}^m \in \mathcal{A}^m$ such that $\lim_{m\to\infty} \sup_{0\leq t\leq T} |\gamma_t^m - \gamma_t| = 0$ $\mathbf{P}$-a.s. To show this, it is sufficient to take $\gamma_t^m = \gamma_t \wedge m$. Moreover, it is easily seen that the reward function is continuous with respect to $\boldsymbol{\gamma}$ so that we have the convergence $J(t, x, y, \tau, \boldsymbol{\gamma}^m) \to J(t, x, y, \tau, \boldsymbol{\gamma})$. Now we choose $\varepsilon > 0$ and a strategy $\boldsymbol{\gamma}^\varepsilon \in \mathcal{A}$ such that $J(t, x, y, \tau, \boldsymbol{\gamma}^\varepsilon) \geq V(t, x, y, \tau) - \varepsilon/2$. Let $\{\boldsymbol{\gamma}^{m,\varepsilon}\}_{m\in\mathbb{N}}$ such that

$$\lim_{m\to\infty} \sup_{0\leq t\leq T} |\gamma_t^{m,\varepsilon} - \gamma_t^\varepsilon| = 0 \ \ \mathbf{P}\text{-a.s.},$$

hence $J(t, x, y, \tau, \boldsymbol{\gamma}^{m,\varepsilon}) \to J(t, x, y, \tau, \boldsymbol{\gamma}^\varepsilon)$. Then, there is $m^*(\varepsilon) \in \mathbb{N}$ such that for all $m > m^*(\varepsilon)$, it holds that $V^m(t, x, y, \tau) \geq J(t, x, y, \tau, \boldsymbol{\gamma}^m) \geq V(t, x, y, \tau) - \varepsilon$. Since ε is arbitrary, we get the result. □

4.4 | Classical Solution

In this paragraph, we discuss conditions ensuring that the value function is a classical solution of the HJB equation (26).

Theorem 4.8. *Fix $\bar{\gamma} > \frac{(L_V-1)^+}{2\kappa}$. Assume that $\sigma > 0$ and that there is some $\bar{\alpha} > 0$ such that for all $\xi \in \mathbb{R}^d$, $\xi^\top \mathfrak{S}(t, y)\xi > \bar{\alpha}||\xi||^2$. Then the value function $V(t, x, y, \tau)$ is the unique classical solution of the HJB equation (26).*

The proof of this result is given in Appendix B.

Corollary 4.9. *Under the assumptions of Theorem 4.8, the optimal strategy satisfies $\gamma_t^* = \gamma^*(t, X_t, Y_t, \tau_t) = \frac{(V_x(t,X_t,Y_t,\tau_t)-1)^+}{2\kappa}$ for every $t \in [0, T]$.*

Proof. Since the function V is a classical solution of Equation (26), we get that $\gamma^*(t, X_t, Y_t, \tau_t)$ is the optimal strategy by verifying first- and second-order conditions. □

In Appendix B.3, we discuss an example which shows that the assumption $\sigma > 0$, that is, strict ellipticity of the generator of the controlled process X, plays a crucial role in obtaining the classical-solution characterization discussed in this section. This example illustrates in particular that, when the assumption is not satisfied, the value function may be a strict viscosity solution of the HJB equation.

5 | Tax Policy Uncertainty and Stochastic Differential Games

Climate policy variables such as emission tax rates are the result of unpredictable political processes. Consequently, specifying a reliable probabilistic model for the evolution of future tax rates is challenging. In this section, we therefore analyze optimal abatement investment under the assumption that the producer is *uncertainty-averse* with respect to future taxes. In that case, the decision on the optimal instantaneous electricity output q_t cannot be separated from the optimal investment decision, and we denote by $\mathcal{Q}$ the set of admissible production plans $\boldsymbol{q} = (q_t)_{0\leq t\leq T}$. We assume that the producer's preferences account for uncertainty in future tax paths. Specifically, she considers a set $\mathcal{S}^\tau$ of possible tax scenarios and selects her production plan $\boldsymbol{q} \in \mathcal{Q}$ and investment strategy $\boldsymbol{\gamma} \in \mathcal{A}$ by solving the following max–min problem:

$$\max_{\boldsymbol{q}\in\mathcal{Q}, \boldsymbol{\gamma}\in\mathcal{A}} \min_{\tau\in\mathcal{S}^\tau} J(\boldsymbol{\tau}, \boldsymbol{\gamma}, \boldsymbol{q}). \tag{29}$$

Here the payoff function J is given by

$$J(\tau, \boldsymbol{\gamma}, \boldsymbol{q}) = \mathbb{E}\left[\int_0^T (\Pi(q_s, X_s, Y_s, \tau_s) - \gamma_s - \kappa\gamma_s^2)e^{-rs}ds + h(X_T)e^{-rT}\right] + \nu_1\rho(\tau), \tag{30}$$

$\nu_1 > 0$ is a given constant and the nonnegative functional ρ *penalizes* the tax processes $\boldsymbol{\tau} \in \mathcal{S}^\tau$ according to their deviation from some reference tax plan $\bar{\tau}$ (precise definitions of all objects in Equations 29–30 are given in Section 5.1 below).

The criterion (29)–(30) may be interpreted as a sequential game between the producer and a hypothetical adversary ("nature"). In this game, the producer first selects a production plan $\boldsymbol{q}$ and an investment strategy $\boldsymbol{\gamma}$. In response, nature chooses a tax process $\boldsymbol{\tau}$ that minimizes the producer's penalized expected profits, so that $\boldsymbol{\tau}$ represents a worst-case tax scenario. To determine an optimal strategy, the producer anticipates nature's reaction and selects $\boldsymbol{\gamma}$ and $\boldsymbol{q}$ to maximize the minimized expected profits. Nature's choice is penalized by the function $\rho(\boldsymbol{\tau})$. Tax processes with large $\rho(\boldsymbol{\tau})$ are thus disadvantaged in the worst-case scenario determination, making them less relevant to the producer's decision. This explains why the term $\nu_1\rho(\boldsymbol{\tau})$ is added rather than subtracted in Equation (30). The parameter $\nu_1 > 0$ represents the producer's uncertainty aversion or sensitivity to tax uncertainty. When ν_1 is near zero, the producer considers all possible $\boldsymbol{\tau} \in S^\tau$, indicating high uncertainty about future tax rates. Conversely, a large ν_1 leads to the near exclusion of tax processes that deviate significantly from a reference tax plan, implying low perceived tax uncertainty. Note that in Equation (30), the producer evaluates expected profits, meaning that under Equations (29)–(30), she is uncertainty-averse regarding the evolution of taxes but remains risk-neutral concerning standard investment and price risks. This distinction aligns with the differing nature of these risks.

The mathematical analysis of the optimization problem (29) builds on the interpretation of Equation (29) as a game. More precisely, we consider a stochastic differential game between the producer and the opponent with payoff J as in Equation (30), where the goal of the producer is to maximize J over $\boldsymbol{\gamma}$ and $\boldsymbol{q}$, whereas nature wants to minimize J by her choice of $\boldsymbol{\tau}$, and we identify the equilibrium strategies $(\boldsymbol{\gamma}^*, \boldsymbol{q}^*)$ and $\boldsymbol{\tau}^*$. Standard min–max theory then implies that $\boldsymbol{\gamma}^*$ and $\boldsymbol{q}^*$ solve the optimization problem (29) and that $\boldsymbol{\tau}^*$ constitutes the corresponding worst-case tax scenario. *Comments.* A decision-theoretic foundation for preference functionals akin to Equation (29) is provided by Maccheroni et al. (2006); an alternative approach to model ambiguity aversion is discussed by Klibanoff et al. (2005). Optimization problems similar to Equation (29) and stochastic differential games have been previously employed to model the decision-making of uncertainty-averse investors in mathematical finance. Notable applications include option hedging under volatility uncertainty, as discussed in Herrmann et al. (2017).

We stress that the game between the producer and nature serves as a conceptual tool to solve the optimization problem (29) for the uncertainty-averse producer and that it does not correspond to a "real" game. In particular, the fictitious adversary should not be interpreted as a regulator or government entity. A government's objective function, aimed at maximizing social welfare, would typically consider factors such as emissions, energy production, and tax revenue—none of which are included in the reward function (30). For an example of a regulator's objective function incorporating these factors, we refer to Carmona et al. (2021, eq. 18).

5.1 | The Differential Game

We now describe the game between the producer and her fictitious opponent in detail. The dynamics of the factor process Y and of the stochastic investment X are given by Equations (1) and (6), respectively. We consider the profit function $\Pi(q, x, y, \tau)$ introduced in Equation (2), where the cost function satisfies the Assumption 2.1.

We assume that the set $\mathcal{S}^\tau$ consists of all adapted tax processes with values in a band around some deterministic tax plan $\bar{\tau} : [0, T] \mapsto [0, \infty)$, which can be interpreted as the producer's prediction of the future tax evolution or as the future carbon tax rate officially announced by the government at $t = 0$. Given functions $\tau^{\min}, \tau^{\max} : [0, T] \to [0, +\infty)$ with $\tau^{\min}(t) \le \bar{\tau}(t) \le \tau^{\max}(t)$ for every $t \in [0, T]$, we thus define $\mathcal{S}^\tau$ as the set of all adapted processes $\boldsymbol{\tau} = (\tau_t)_{0\le t\le T}$ such that $\tau^{\min}(t) \le \tau_t \le \tau^{\max}(t)$ for all $t \in [0, T]$. For $0 \le t \le T$, we introduce the penalty function as follows:

$$\boldsymbol{\tau} \mapsto \rho_t(\boldsymbol{\tau}) = \mathbb{E}_t\left[\int_t^T e^{-r(s-t)}(\tau_s - \bar{\tau}(s))^2\, ds\right], \tag{31}$$

and we let $\rho := \rho_0$. Next, we denote by $\mathcal{Q}$ the set of all adapted production processes $\boldsymbol{q} = (q_t)_{0\le t\le T}$ taking values in $[0, q^{\max}]$, for some $q^{\max} > 0$ that represents the maximum capacity of production. Finally, recall that $\mathcal{A}$ denotes the set of admissible investment strategies, that is, the set of all adapted processes $\boldsymbol{\gamma} = (\gamma_t)_{0\le t\le T}$ with values in $[0, +\infty)$ and $\mathbb{E}\left[\int_0^T \gamma_t dt\right] < \infty$.

The special form of Equation (31) permits us to move the penalization inside the expectation in Equation (30). Given an uncertainty aversion parameter $\nu_1 > 0$, the payoff function of the differential game between the producer and the fictitious adversary is thus given by the following:

$$J(t, x, y, \boldsymbol{\tau}, \boldsymbol{\gamma}, \boldsymbol{q}) = \mathbb{E}_t\Big[h(X_T)e^{-r(T-t)} \tag{32}$$

$$+ \int_t^T \left(\Pi(q_s, X_s, Y_s, \tau_s) - \gamma_s - \kappa\gamma_s^2 + \nu_1(\tau_s - \bar{\tau}(s))^2\right)e^{-r(s-t)}ds\Big]. \tag{33}$$

In the game, the goal of the producer is to maximize Equation (32) over admissible investment rates $\boldsymbol{\gamma} \in \mathcal{A}$ and production plans $\boldsymbol{q} \in \mathbf{Q}$; the opponent, on the other hand, chooses the tax process $\boldsymbol{\tau} \in \mathcal{S}^\tau$ in order to minimize Equation (32). A pair of strategies $(\boldsymbol{\gamma}^*, \boldsymbol{q}^*)$ (for the producer) and $\boldsymbol{\tau}^*$ (for the opponent) an *equilibrium* for the game if for any $\boldsymbol{\tau} \in \mathcal{S}^\tau$, $\boldsymbol{\gamma} \in \mathcal{A}$, $\boldsymbol{q} \in \mathcal{Q}$,

$$J(0, X_0, Y_0, \boldsymbol{\tau}^*, \boldsymbol{\gamma}, \boldsymbol{q}) \le J(0, X_0, Y_0, \boldsymbol{\tau}^*, \boldsymbol{\gamma}^*, \boldsymbol{q}^*) \le J(0, X_0, Y_0, \boldsymbol{\tau}, \boldsymbol{\gamma}^*, \boldsymbol{q}^*). \tag{34}$$

We then call $u(t, x, y) := J(t, x, y, \boldsymbol{\tau}^*, \boldsymbol{\gamma}^*, \boldsymbol{q}^*)$ the *value* of the game. In the sequel, we show that under certain regularity conditions, the game (32) has equilibrium strategies $(\boldsymbol{\gamma}^*, \boldsymbol{q}^*)$ and $\boldsymbol{\tau}^*$ in feedback form, which implies that the value of the game is well-defined. Moreover, it follows from standard min–max theory that $\boldsymbol{\gamma}^*$ and $\boldsymbol{q}^*$ solve the optimization problem (29) and that $\boldsymbol{\tau}^*$

constitutes the corresponding worst-case tax scenario, see, for instance, Rockafellar (1970, Lemma 36.2).

Our analysis of the game between the producer and nature is based on Friedman (1972). Subsequent contributions to the mathematical theory of stochastic differential games include Fleming and Souganidis (1989), and more recently, Possamaï et al. (2020).

5.2 | Characterization of Equilibrium Strategies

In the following, we characterize the game's value and equilibrium strategies. In stochastic differential games, this is typically achieved using the Bellman–Isaacs equation. However, since in our model, the tax value τ chosen by the opponent affects only the instantaneous profit, the Bellman–Isaacs equation can be reduced to a standard HJB equation.

We define the function

$$g(q,\tau;x,y)=\Pi(q,x,y,\tau)+\nu_1(\tau-\bar{\tau}(t))^2\,, \tag{35}$$

and recall that $\Pi(q,x,y,\tau)=p(y)q-(C_0(q,x,y)+C_1(q,x,y)\tau)+\nu_0(q)\tau$. In Lemma 5.2, we show that for every fixed (x,y), the function g admits a unique saddle point (q^*,τ^*). Hence we may define functions $\hat{q}(x,y)$ and $\hat{\tau}(x,y)$ that map (x,y) to the associated saddle point of g. Denote by

$$\begin{aligned} G(x,y):&=g(\hat{q}(x,y),\hat{\tau}(x,y);x,y)=\max_q\min_\tau g(q,\tau;x,y)\\ &=\min_\tau\max_q g(q,\tau;x,y) \end{aligned} \tag{36}$$

the corresponding saddle value, where the maximum is taken over $q\in[0,q^{\max}]$ and the minimum over $\tau\in[\tau^{\min},\tau^{\max}]$. In the next result, we show that the equilibrium strategy and the value of the game can be characterized in terms of an HJB equation with running reward given by the function G.

Proposition 5.1. *Suppose that for fixed (x,y), the function g has a saddle point $(\hat{q}(x,y),\hat{\tau}(x,y))$ and that the PDE*

$$\begin{aligned} 0=\ & u_t(t,x,y)+G(x,y)+\mathcal{L}^Y u(t,x,y)+\frac{\sigma^2}{2}u_{xx}(t,x,y)\\ &+\sup_{\gamma\ge 0}\left((\gamma-\delta x)u_x(t,x,y)-\gamma-\kappa\gamma^2\right)-ru(t,x,y) \end{aligned}$$

with the final condition $u(T,x,y)=h(x)$ has a classical solution. Let $\hat{\gamma}(t,x,y)=(u_x(t,x,y)-1)^+/(2\kappa)$. Then u is the value function of the game and the strategies $\boldsymbol{q}^*=(\hat{q}(X_t,Y_t))_{0\le t\le T}$, $\boldsymbol{\gamma}^*=(\hat{\gamma}(t,X_t,Y_t))_{0\le t\le T}$, and $\boldsymbol{\tau}^*=(\hat{\tau}(X_t,Y_t))_{0\le t\le T}$ are equilibrium strategies for the game.

Proof. This proposition can be established via classical verification arguments. Suppose that the opponent uses the strategy $\boldsymbol{\tau}^*$ and denote by X the solution of the SDE

$$dX_t=(\hat{\gamma}(t,X_t,Y_t)-\delta X_t)dt+\sigma dW_t\,.$$

Since $(\hat{q}(x,y),\hat{\tau}(x,y))$ is a saddle point of g, we have $G(x,y)=\sup_{q\in[0,q^{\max}]}g(q,\hat{\tau}(x,y);x,y)$, and we may rewrite the PDE (37) in the following form:

$$u_t(t,x,y)+\mathcal{L}^Y u(t,x,y)+\frac{\sigma^2}{2}u_{xx}(t,x,y)-\delta x u_x(t,x,y) \tag{38}$$

$$+\sup_{q\in[0,q^{\max}]}g(q,\hat{\tau}(x,y);x,y)+\sup_{\gamma\ge 0}\left\{\gamma u_x(t,x,y)-\gamma-\kappa\gamma^2\right\}=ru(t,x,y)\,. \tag{39}$$

Moreover $u(T,x,y)=h(x)$, so that this is the HJB equation for the control problem

$$\max_{\boldsymbol{q}\in\mathcal{Q},\boldsymbol{\gamma}\in\mathcal{A}}\mathbb{E}_t\left[\int_t^T\left(g(q_s,\hat{\tau}(X_s,Y_s);X_s,Y_s)-\gamma_s-\kappa\gamma_s^2\right)e^{-r(s-t)}ds+e^{-r(T-t)}h(X_T)\right]. \tag{40}$$

A standard verification result for stochastic control problems such as Theorem 3.5.2 in Pham (2009) now shows that u is the value function for the control problem (40) and that $\boldsymbol{q}^*$ and $\boldsymbol{\gamma}^*$ are an optimal strategy in Equation (40). A similar argument shows that $\boldsymbol{\tau}^*$ is optimal against $\boldsymbol{q}^*$ and $\boldsymbol{\gamma}^*$, which completes the proof. □

Next we verify that the assumptions of Proposition 5.1 are satisfied. We begin with the existence of a unique saddle point for g. We omit the arguments x,y to ease the notation. For fixed $q\in[0,q^{\max}]$, the function $\tau\mapsto g(q,\tau)$ is strictly convex and has a unique minimum on $[\tau^{\min},\tau^{\max}]$ which we denote by $\tau(q)$. Similarly, the function $q\mapsto g(q,\tau)$ is strictly concave and has a unique maximum $q(\tau)$ on $[0,q^{\max}]$. A saddle point (q^*,τ^*) of g is characterized by the equations

$$\tau^*=\tau(q^*)\quad \text{ and } q^*=q(\tau^*). \tag{41}$$

We use first-order conditions to identify $\tau(q)$ and $q(\tau)$. It holds that

$$\tau(q)=\left\{\bar{\tau}+\frac{1}{2\nu_1}\left(C_1(q)-\nu_0(q)\right)\right\}\vee\tau^{\min}\wedge\tau^{\max} \tag{42}$$

The optimal instantaneous production $q(\tau)$ is determined as in Section 2. In particular, the first-order condition characterizing $q(\tau)$ is $p-\partial_q C_0(q)-(\partial_q C_1(q)-\partial_q\nu_0(q))\tau=0$, and $q(\tau)$ is therefore given by Equation (5). The existence of a unique solution to Equation (41) is established in the next lemma, whose proof is given in Appendix C.

Lemma 5.2. *Suppose that the cost function C satisfies Assumption 2.1 and that the functions C_0, C_1, and ν_0 are moreover C^2 in q. Then, for every fixed (x,y), the function $g(q,\tau;x,y)$ has a unique saddle point $(q^*,\tau^*)=:(\hat{q}(x,y),\hat{\tau}(x,y))$.*

The next theorem summarizes the mathematical analysis of the stochastic differential game.

Theorem 5.3. *Suppose that $\sigma^2>0$, that the generator $\mathcal{L}^Y$ is strictly elliptic, that the cost function satisfies Assumption 2.1 and that C_0 and C_1 are C^2 in q. Then the PDE (37) has a unique classical solution u, which is the value function of the game. Moreover, the strategies $\boldsymbol{q}^*$, $\boldsymbol{\gamma}^*$, and $\boldsymbol{\tau}^*$ from Proposition 5.1 are equilibrium strategies for the game.*

Proof. In view of Proposition 5.1, we need to show the existence of a classical solution to the PDE (37). For this, we first show that the function G from Equation (36) is Lipschitz in (x, y). The definition of G implies that $|G(x', y') - G(x, y)| \le \sup_{(q,\tau)\in B} |g(q, \tau; x', y') - g(q, \tau; x, y)|$, where $B = [0, q^{\max}] \times [\tau^{\min}, \tau^{\max}]$. Now,

$$|g(q,\tau;x',y') - g(q,\tau;x,y)| \le q^{\max}|p(y') - p(y)| + |C_0(q,x',y') - C_0(q,x,y)|$$

$$+ \tau^{\max}|C_1(q,x',y') - C_1(q,x,y)| \le C|(x',y') - (x,y)|,$$

where the last inequality follows from the Lipschitz conditions in Assumption 2.1. Existence and uniqueness of a classical solution to Equation (37) now follow by similar arguments as in Section 4.4. In fact, the analysis of Equation (37) is even simpler than the analysis of the HJB equation in Section 4.4, since there are no jump terms in the equation. □

5.3 | Properties of the Optimal Tax Rate and Production

We now discuss the properties of the saddle point $(\tau^*, q^*) = (\hat{\tau}(x), \hat{q}(x))$, omitting the dependence on y for clarity. We first consider the case where the rebate $\nu_0(q)$ is zero. In that setting, the producer's profit $\Pi(q, x, \tau)$ is decreasing in τ. Since nature selects the worst-case tax scenario τ^* in order to minimize profits, it follows that $\tau^* = \tau(q^*)$ is larger than $\bar{\tau}$, consistent with Equation (42). This, in turn, incentivizes higher investment compared to the reference tax scenario $\bar{\tau}$. Consequently, increased uncertainty can be socially beneficial, as confirmed numerically in Figure 5 below. This reveals a key difference between a producer who is averse to tax uncertainty and a scenario where tax dynamics are assumed to be known.

If $\nu_0(q) > 0$, we obtain from Equation (42) that $\tau^* \ge \bar{\tau}$ if and only if $\nu_0(q^*) \le C_1(q^*, x)$, that is if the tax paid on emissions is higher than the tax repayment due to the rebate. In particular, with rebate and for full abatement, that is for $C_1 \equiv 0$, we have $\tau^*(q) < \bar{\tau}$.

These properties are reported in Figure 1 which depicts the saddle point $(\hat{\tau}(x), \hat{q}(x))$ for a cost function associated with the two-production-technology model (see Section 3.2 for the model and Section 6.2 for parameter specifications). We impose the capacity constraints $q^{\min} = 5$ and $q^{\max} = 10$, where the lower bound represents contractual obligations ensuring a minimum energy supply. The tax rate τ varies within $[0.5, 1.5]$, with the most plausible value set at $\bar{\tau} = 1$. The rebate function is given by $\nu_0(q) = e_b Q(\alpha q)$ for different α values, while deviations from $\bar{\tau}$ incur a penalty $\nu_1(\tau - \bar{\tau})^2$. The left panels represent the case of high uncertainty, modeled by a small penalty on deviations from $\bar{\tau}$ ($\nu_1 = 1$), while the right panels correspond to low uncertainty, where deviations are strongly penalized ($\nu_1 = 20$). Each panel considers scenarios without rebate ($\alpha = 0$, blue dashed line) and with rebate ($\alpha = 0.5$, solid red line).

The graphs show that the optimal production function $\hat{q}(x)$ is increasing in x as higher investment reduces tax payments, thereby lowering marginal costs. Moreover, the rebate stimulates production, as evidenced by the solid red line positioned above the dashed blue line. Notably, $\hat{q}(x)$ exhibits robustness to tax uncertainty, displaying similar behavior across both panels in the upper row. In contrast, the optimal tax rate $\hat{\tau}(x)$ is more sensitive to ν_1. For small ν_1 (high uncertainty), $\hat{\tau}(x)$ spans the entire interval $[\tau^{\min}, \tau^{\max}]$, with constraints $\tau \le \tau^{\min} = 0.5$ and $\tau \ge \tau^{\max} = 1.5$ being binding. For large ν_1 (low uncertainty), these constraints are inactive, and $\hat{\tau}(x)$ remains close to $\bar{\tau}$, consistent with Equation (42). Furthermore, $\hat{\tau}(x)$ decreases with x, aligning with economic intuition: Higher investment reduces emissions and, consequently, carbon tax revenues, leading to lower tax rates under rebate and penalization mechanisms.

6 | Numerical Experiments

In this section, we report the results of numerical experiments that study the impact of transaction cost, production technology, market structure, and randomness in the tax system on the investment strategy and the optimal electricity output of the producer. Throughout, we use the deep-learning algorithm proposed in Frey and Köck (2022) to compute the value function and the optimal investment rate. We refer to Appendix A for the details on the numerical methodology.

In Section 6.1, we present results in the context of the filter technology from Section 3.1, in Section 6.2, we discuss results for the two technologies from Section 3.2. In both cases, we work under tax-risk and assume that the tax process follows the two-state Markov chain introduced in Example 4.1 with $\tau^1 = 0$ and $\tau^2 > 0$ and transition intensity matrix G. In the tax increase case, we assume $\tau_0 = \tau^1$, $g_{12} = 0.25$, and $g_{21} = 0$, meaning that τ^2 is an absorbing state. In the tax reversal case $\tau_0 = \tau^2$ and $g_{12} = g_{21} = 0.25$. In Section 6.3, we finally discuss examples for the stochastic differential game in the context of the two technologies.

6.1 | Filter Technology and Tax-Risk

We now discuss results of numerical experiments for the filter technology. We use the following parameters: $\delta = 0.05$, $\sigma = 0.05$, $r = 0.02$, the time horizon is $T = 15$ years and $h(x) = 0$, which is in line with the fact that filters lose their value at the end of the lifetime of the underlying power plant. The transaction costs parameter is usually set to $\varkappa = 0.5$. The high tax value is set to $\tau^2 = 0.2$. We work with a *cost function* of the form (9), where the cost of one unit of raw material is constant and equal to $\bar{c}$, the quantity of raw material is specified as $Q(q) = aq^{\frac{3}{2}}$, and where the abatement function is given by the following:

$$e(x) = \left(e_1 x - \frac{e_1^2}{4e_0}x^2\right)\mathbf{1}_{\{x\le 2e_0/e_1\}} + e_0\mathbf{1}_{\{x>2e_0/e_1\}} \qquad (43)$$

In the numerical experiments, we use the parameter values $a = 1.25$, $\bar{c} = 1$, $e_0 = 1.5$, $e_1 = 0.5$. Note that for the chosen parameters, the abatement cost (43) is globally nondecreasing in x, concave and differentiable and that the maximum abatement level is e_0.

These parameter values are chosen to obtain a qualitatively reasonable behavior of the production function. They are, however,

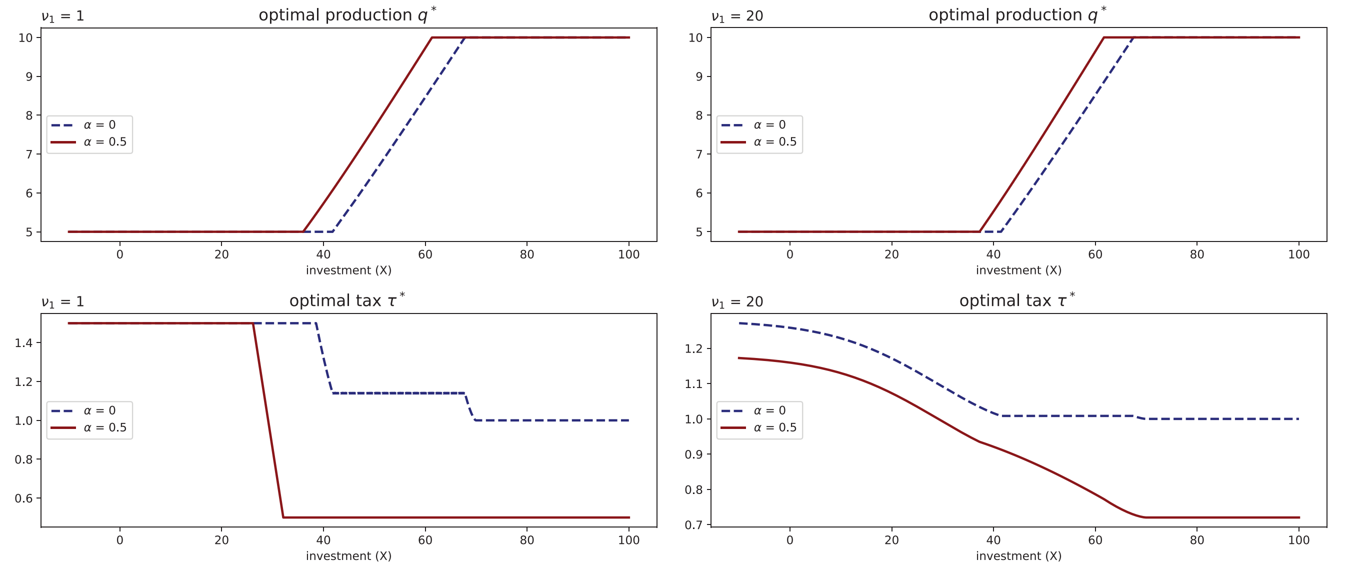

FIGURE 1 | Saddle point $(\hat{\tau}(x)$ (top) and $\hat{q}(x))$ (bottom) for the cost function from the two technologies example, with $\bar{\tau} \equiv 1$. Left: small ν_1, right: large ν_1. The value $\alpha = 0$ corresponds to no rebate, $\alpha = 0.5$ to a positive rebate. Note that the left and the right panel on the bottom use a different scale. [Color figure can be viewed at wileyonlinelibrary.com]

not calibrated to a real production or abatement technology, for the following reasons. First, real abatement cost functions and production technology characteristics are highly firm- and sector-specific, and comprehensive, publicly available data capturing these aspects are scarce. Moreover, in many cases, abatement technologies are still evolving, making it difficult to define stable, reliable cost structures. However, given the outcome of several robustness checks, we are confident that the qualitative insights from our analysis remain informative nonetheless.

6.1.1 | Fixed Electricity Output

In this section, we assume that electricity production is fixed and equal to $q^{\max} = 4$, for instance, since the producer has entered into long-term delivery contracts. In that case, the investment decision of the electricity producer is not affected by the rebate or by the electricity price, so that the only risk that is relevant for the investment decision of the producer is the randomness in taxes.

In Figure 2, we plot single trajectories of the cumulative investment for the tax increase (left panel) and the tax reversal (right panel), for $\varkappa = 0.5$ and, for comparison purposes, for $\varkappa = 0.2$ (low transaction costs). We make the following observations. First, investments are larger for low transaction costs. Moreover, the investment level decreases as time approaches the horizon date T. This is due to the fact that γ_t^* is equal to zero for t close to T, since in that case, the tax savings generated by new investment over the remaining lifetime of the power plant are too small to warrant the expenditure. Finally, the producer reacts to changes in the tax regimes. Indeed, when a change in the tax rate occurs the trajectory of the investment process suddenly exhibits a change in the slope (i.e., a kink), which, intuitively, corresponds to a jump in the investment rate. In particular, in the tax increase scenario, the investment rate γ_t jumps upward as the tax rate switches from τ^1 to τ^2. Interestingly, the producer starts to invest already at $t = 0$, even if the tax rate is equal to zero for small t. In this way, she *hedges* against an anticipated tax increase. In fact, due to transaction costs, it would be too costly to wait until the upward jump in taxes actually occurs and to invest only thereafter. This hedging behavior distinguishes our model from the real options literature such as Fuss et al. (2008), where it is optimal to wait if and when a regulator acts and to invest only afterwards. In the tax reversal case, investments start at a high rate due to the high taxation of emissions. As soon as taxes switch to τ^1, the producer reduces or even stops her investment so that X_t decreases due to depreciation.

Next we present numerical results on the distribution of the emissions. For comparison purposes, we also consider a deterministic *benchmark tax scenario* $\bar{\tau}(t)$, which is computed as follows: In the tax increase case, we assume that $\bar{\tau}(t)$ is linear increasing that is $\bar{\tau}(t) = bt$; in the tax reversal case, we assume that the reverence tax rate is constant, $\bar{\tau}(t) = \bar{\tau}$. The parameters b and $\bar{\tau}$ are set so that the expected average tax rate is identical in the benchmark scenario and in the case with random taxes. The benchmark scenario in the case of a tax increase corresponds to the situation where a government incrementally raises the carbon tax in a series of small deterministic steps. A real-world example of such a scheme is the carbon tax policy implemented by the Singaporean government, as discussed in Tan and Tan (2022).

In Table 1, we report the values of several statistics of the emission distribution (5% quantile, mean and 95% quantile) at two evaluation dates, after 10 and 15 years, for $\varkappa = 0.5$, in the two tax scenarios and in the benchmark case. The values suggest that the benchmark tax regime leads to emission levels that are lower than the mean emissions under random tax rates. This confirms the intuition that randomness in future tax rates reduces investments into carbon abatement technologies, so that a deterministic tax policy would be beneficial for stipulating emission reduction.

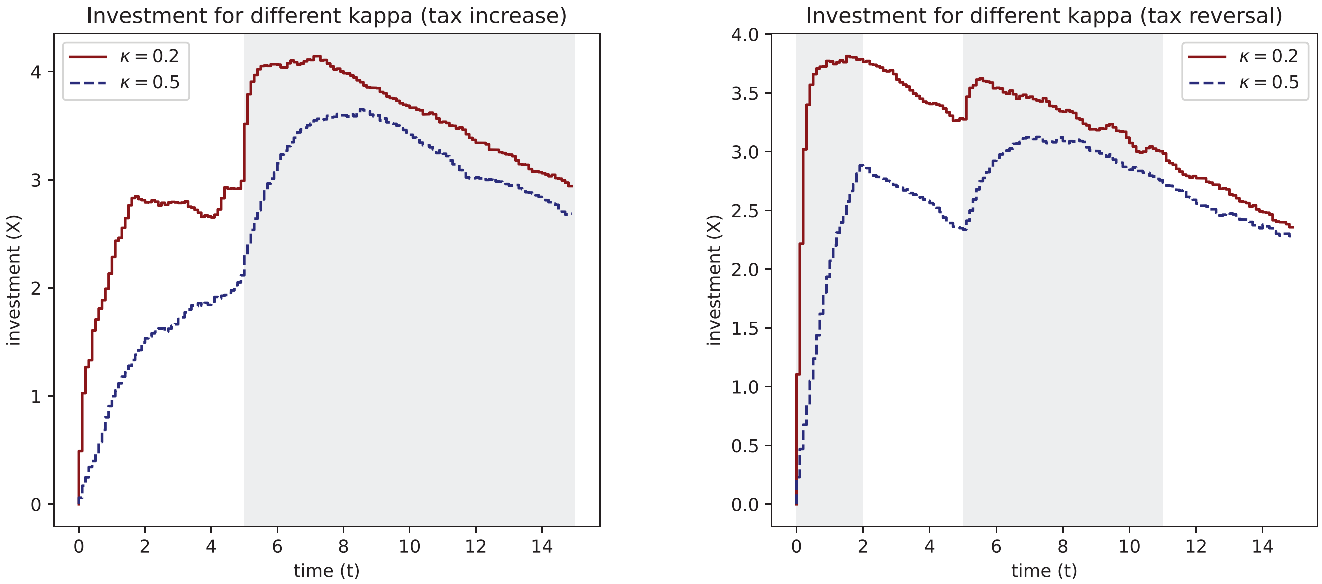

FIGURE 2 | Single trajectory of cumulative investment. Left: tax increase, right: tax reversal. Low transaction costs are represented by solid line, high transaction costs by the dashed line. The gray (white) shaded areas correspond to time periods with high tax rate (low tax rate), respectively. [Color figure can be viewed at wileyonlinelibrary.com]

TABLE 1 | Statistics of the emissions distribution under both stochastic tax regimes for the filter technology example.

Statistics of emission distribution								
	***t* = 10**				***t* = 15**			
	5%	**Mean**	**95%**	**Bench**	**5%**	**Mean**	**95%**	**Bench**
Tax increase	3.27	5.27	7.82	5.07	5.21	7.31	11.34	6.38
Tax reversal	4.31	5.11	7.43	4.18	6.74	7.83	10.78	6.52

TABLE 2 | Statistics of the emissions distribution for the random tax increase: wrong belief versus correct belief on the switching intensity (filter technology).

Statistics of emission distribution, different beliefs						
	***t* = 10**			***t* = 15**		
	5%	**Mean**	**95%**	**5%**	**Mean**	**95%**
Wrong belief	3.72	7.96	15.03	5.54	10.20	21.41
Correct belief	3.33	5.27	7.82	5.21	7.31	11.34

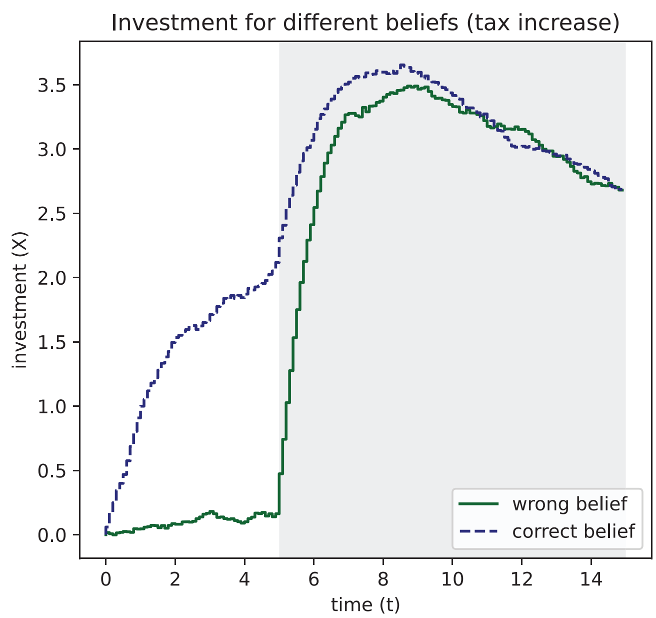

FIGURE 3 | Example trajectory of total investment for the tax increase case for two investors with different beliefs about switching intensity (filter technology). [Color figure can be viewed at wileyonlinelibrary.com]

6.1.2 | Credibility of Tax Policy

Next we investigate the impact of the credibility of an announced carbon tax policy on the investment decisions of producers, and consequently, on the effectiveness of the policy. To shorten exposition, we focus on the case of a random tax increase (the results for the tax reversal are qualitatively similar). We compare a producer who does not believe in the announced future tax increase, and as a result, operates with a very low switching intensity ($g_{12} = 0.05$), to a producer with the correct switching intensity of $g_{12} = 0.25$. Table 2 presents the quantiles and average emissions for both producers. The results show significantly higher emissions in the case of incorrect beliefs.

To explain these differences, in Figure 3, we compare the trajectory of cumulative investment of both producers for a given tax path. We see that the producer with $g_{12} = 0.05$ does not hedge against the potential rise in tax rates but increases investment only after the tax increase. These observations are highly relevant

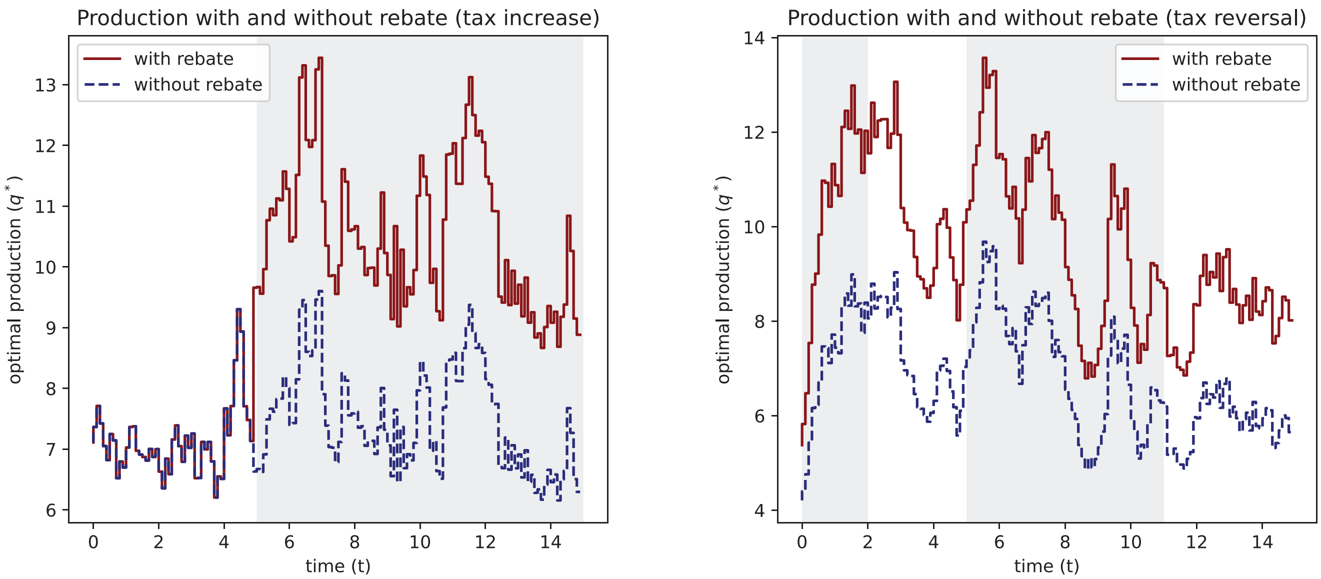

FIGURE 4 | Trajectory of the optimal production q^*, with rebate (solid red line) and without rebate (dashed blue line). Left: tax increase. Right: tax reversal (filter technology). [Color figure can be viewed at wileyonlinelibrary.com]

from a policy perspective as they show that a carbon tax policy must be credible to substantially speed up the transition to greener energy production.

6.1.3 | Stochastic Price and Endogenous Electricity Output

Now we enrich the setup assuming that the selling price of electricity is random and given by $p_t = \exp(Y_t)$, where the process Y is the solution of the one-dimensional SDE

$$dY_t = \theta(\mu - Y_t)\,dt + \alpha\,dB_t, \quad Y_0 = \ln(p_0),$$

for a one-dimensional Brownian motion $B = (B_t)_{t\geq 0}$ that is independent of W. We fix $\mu = \ln(5)$, $\theta = 1$, $\alpha = 0.1$, and $p_0 = 5$. The dynamics of X and τ are as in Section 6.1.1. The producer optimizes the instantaneous electricity production $q_t^* = q^*(X_t, Y_t, \tau_t)$. We also consider a tax rebate, which is modeled by the function $\nu_0(q) = \frac{1}{2}Q(q)e_0$, with $Q(q) = aq^{\frac{3}{2}}$, that is the tax payments of the producer are fully refunded when half of the emissions are abated.

Modeling the price process as the exponential of an Ornstein–Uhlenbeck process is a well-established approach. The mean-reversion component effectively captures stochastic price fluctuations around an average value. This model has been widely used in the literature, including Geman and Roncoroni (2006), Fuss et al. (2008), and Benth et al. (2012). Energy prices are also subject to temporary price spikes caused by events such as power plant outages or capacity shortages, which are not captured by our model; however, since these effects are not persistent, they do not have a significant impact on long-term planning of the firm investments. Similarly, we found that parameters for the process Y have no significant bearing on the qualitative results regarding the optimal strategy.

Figure 4 plots trajectories of the optimal production for the random tax increase (left panel) and for the random tax reversal (right panel). We compare the cases with rebate (solid lines) with that of no-rebate ($\nu_0(q) \equiv 0$, dashed lines). The plots are obtained for the same selected price trajectory. We see that optimal production q^* reacts to three different factors: (i) there are instantaneous jumps occurring at tax switches; (ii) between two consecutive tax jumps production fluctuates in reaction to price changes; (iii) with tax rebate production is both larger and more volatile than for $\nu_0(q) \equiv 0$.

The implications for optimal investment are presented in Table 3. We compare outcomes with and without a rebate for both tax scenarios. The benchmark tax scenarios are as in Section 6.1.1. The results show that investment levels are consistently higher when a rebate is applied, and that uncertainty in future tax rates tends to discourage investment. In particular, benchmark investment levels are always higher than the expected investment under stochastic tax policies, and in most cases exceed the 95% quantile of the corresponding investment distribution.

Remark 6.1. In practice, an increase in taxation may depress electricity supply by raising production costs, potentially inducing a discrete upward shift in equilibrium electricity prices. Formally, this suggests that the price process p depends on the prevailing tax rate τ. This dependence is not modeled explicitly here, as we do not adopt an equilibrium framework. Nevertheless, its qualitative implications for optimal investment closely parallel those of a tax rebate. Indeed, a rebate operates as a transfer proportional to τ, so that, following a tax increase, producers receive higher net income per unit of electricity sold. In light of the results in this section, we therefore conjecture that, in equilibrium, investment in abatement technologies—particularly the filter technology considered here—would exceed that observed in a setting with exogenous prices.

TABLE 3 | Statistics of the investment distribution under stochastic tax regimes with endogenous production (filter technology).

Statistics of investment distribution									
		$t = 10$				**$t = 15$**			
Tax scenario	**Rebate**	**5%**	**Mean**	**95%**	**Benchmark**	**5%**	**Mean**	**95%**	**Benchmark**
Increase	Yes	4.70	4.76	4.82	5.19	4.13	4.28	4.43	4.84
	No	4.32	4.40	4.54	5.03	3.53	3.71	3.90	4.40
Reversal	Yes	3.51	4.29	4.65	4.70	2.86	3.73	4.14	3.82
	No	3.18	4.01	4.40	4.44	2.60	3.24	3.60	3.54

TABLE 4 | Statistics of the investment distribution for the random tax increase (two technologies).

Statistics of the investment distribution								
	$t = 10$				**$t = 15$**			
	5%	**Mean**	**95%**	**Benchmark**	**5%**	**Mean**	**95%**	**Benchmark**
Rebate	45.18	57.35	64.16	58.87	45.28	58.47	62.65	62.19
No-rebate	40.74	43.48	45.19	43.36	39.14	41.48	43.08	41.79

6.2 | Two Technologies With Tax-Risk

This section discusses numerical experiments based on the setup described in Example 3.2, where a producer can invest in both green and brown production technologies. In these experiments, the cost function takes the following form:

$$C(q, x, y, \tau) = (c_b + e_b\tau)Q_b\left(\left(q - P_g(x)\right)^+\right), \tag{44}$$

where $c_b = 1, e_b = 1, Q_b(q) = q^{3/2}$, and $P_g(x) = p_g(x - \bar{x})^+$. Here, $\bar{x} = 20$ represents an initial expenditure required before green investment can generate electricity, such as land acquisition for solar farms or infrastructure costs for connecting a solar park to the grid. We set the productivity parameter to $p_g = 0.2$ and the maximum production capacity to $q^{\max} = 10$. Additionally, we fix the following parameter values: $T = 15$ years, $h(x) = (0.7x)^+$, $\delta = 0.02$, $\sigma = 0.2$, $r = 0.04$, and $\varkappa = 0.5$, and we set the high tax value $\tau^2 = 1$. Moreover, we assume that the selling price of electricity remains constant at $p = 2.1$, while the output level q is endogenous. Furthermore, we assume a rebate of the form $\nu_0(q)\tau = e_b Q_b(\alpha q)\tau$, where $Q_b(q) = q^{3/2}$ and α takes different values. Specifically, $\alpha = 0$ represents the case without a rebate, while $\alpha > 0$ corresponds to a scenario where the rebate is positive and exceeds tax payments once the producer generates more than a fraction $1 - \alpha$ of the total output using green technology.

We found that the numerical results for tax increase and tax reversal are qualitatively similar. To shorten the exposition, we therefore concentrate on the case of the tax increase. In Table 4, we compare the average investments, the 5% and 95% quantiles of the investment distribution, with and without rebates, after 10 and 15 years. Additionally, we present results for the benchmark tax scenarios from Section 6.1.1. We make the following observations. First, the statistics of the investment distribution with rebates are consistently higher than those without rebates, that is the introduction of a rebate scheme enhances investment. Indeed, since the rebate is proportional to q and τ, when taxes are high, the producer gains considerably from a higher level of green investment, because green electricity production increases due to the higher rebate. Second, the two-technology case further confirms that the benchmark scenarios result in higher investment levels compared to the average investment under stochastic tax rates.

6.3 | Two Technologies With Tax Uncertainty

Finally, we present numerical results for the uncertainty-averse investor from Section 5.1. In that case, tax rates emerge endogenously as the equilibrium of a stochastic differential game. The analysis follows the setup of Example 3.2 (two-technology case) and employs the same parameters as Section 6.2, except that now $T = 10$. Tax rates are constrained to the interval $[0.5, 1.5]$, with the most plausible rate set as $\bar{\tau} \equiv 1$. The tax rebate is defined as $\nu_0(q)\tau = e_b Q(\alpha q)\tau$ for $\alpha \in \{0, 0.5\}$, while deviations from $\bar{\tau}$ are penalized by $\nu_1(\tau - \bar{\tau})^2$, with $\nu_1 \in \{1, 20\}$. The equilibrium output $\hat{q}(x)$ and tax rate $\hat{\tau}(x)$ for this setup are analyzed in Section 5.3, specifically in Figure 1.

In Figure 5, we plot the average investment $\mathbb{E}[X_t]$ under different values for rebate and penalization. The left panel corresponds to the case of high uncertainty ($\nu_1 = 1$), the right panel corresponds to the case of low uncertainty ($\nu_1 = 20$).

The figure shows that the results from the tax-risk paradigm are reversed. Specifically, the average investment under high uncertainty is significantly greater than under low uncertainty, implying that high uncertainty can be beneficial from a societal perspective. This occurs because the equilibrium tax rate for $\nu_1 = 1$ is higher than that for $\nu_1 = 20$ due to weaker penalization,

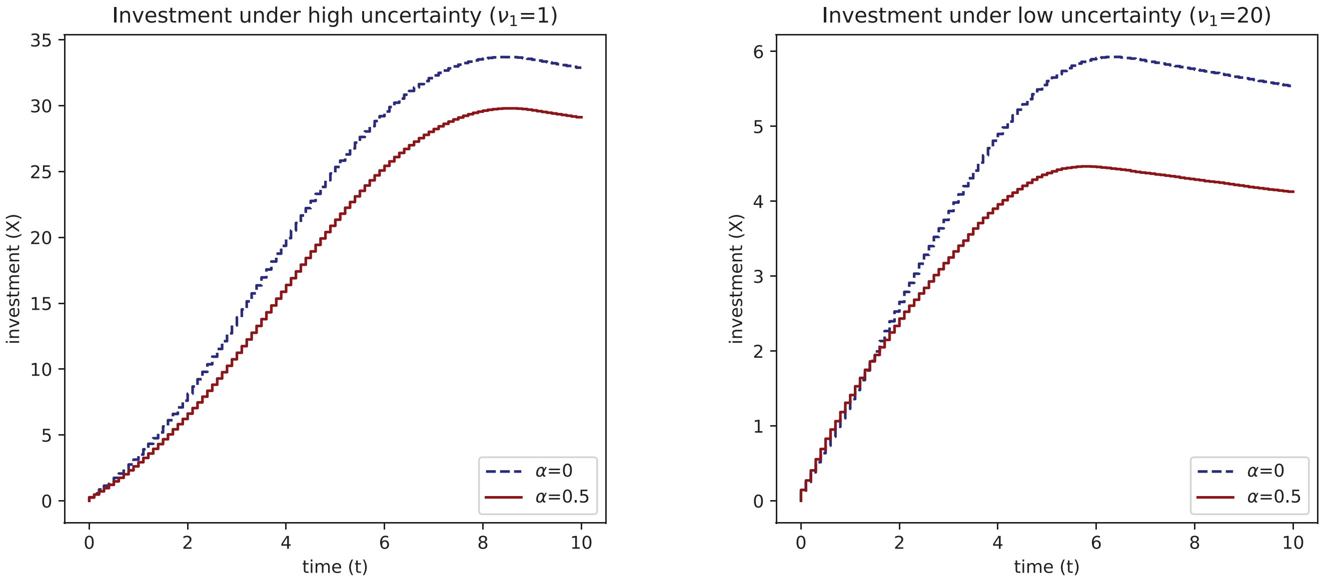

FIGURE 5 | Average investment $\mathbb{E}[I_t]$ under tax uncertainty for different values of the rebate. Left: high uncertainty ($\nu_1 = 1$), right: low uncertainty ($\nu_1 = 20$). [Color figure can be viewed at wileyonlinelibrary.com]

see Figure 1; a higher tax in return rate stimulates increased investment. Furthermore, under tax uncertainty, the introduction of a rebate reduces investment, whereas in the tax-risk scenario, a rebate led to an increase in investment. This difference arises because the rebate lowers the equilibrium tax rate in the strategic interaction between the producer and the opponent (see again Figure 1).

7 | Conclusion

In this paper, we investigated the impact of randomness in carbon tax policies on the investment behavior of a stylized, profit-maximizing electricity producer who is subject to carbon taxation and can invest in technologies for the abatement of CO_2 emissions. Extending existing literature, we considered a framework in which investments in abatement technologies are divisible, irreversible, and subject to transaction costs. We explored two approaches to model the randomness in taxation. In the first one, we assumed a well-specified probabilistic model for the tax process, leading to a stochastic control problem for the producer's investment strategy. In the second approach, we modeled the producer as uncertainty-averse, which gives rise to a stochastic differential game with a fictitious adversary. We provided a rigorous mathematical analysis of both models, and of the corresponding nonlinear P(I)DE. Furthermore, numerical methods were employed to investigate the quantitative characteristics of the optimal investment strategies.

Our numerical analysis indicates that, under tax-risk, the producer tends to be less willing to invest in abatement technologies compared to a benchmark scenario with a deterministic tax scheme. The policy implications of our findings are the following: a carbon tax scheme that is initially mild and defers increases to random future dates may delay necessary investments in green technologies. Conversely, a policy that is too stringent may generate strong political pressure to revert to lower taxes, which would be counterproductive for reducing carbon emissions. Moreover, we found that a tax policy must be credible (i.e., producers must be confident that an announced tax increase will actually be implemented) to have a sizable impact on the transition to low-carbon energy production. These results reinforce the widely held view that randomness in carbon taxation is generally detrimental to climate policy.

Interestingly, our results reveal that this conclusion is reversed under tax uncertainty. In scenarios characterized by high uncertainty, the producer invests more than in low-uncertainty settings where tax levels are nearly deterministic. From a societal perspective, this suggests that increased uncertainty can be beneficial. This is an interesting observation, which shows that the paradigm used to model the decision making process of the producer is a crucial determinant for the impact of randomness in climate policy. While it is beyond the scope of this paper to determine which paradigm, risk, or uncertainty, more accurately reflects real-world investor behavior, it is interesting to investigate the difference further. Intuitively, we believe that the recommendations from the tax-risk model are more directly relevant for climate policy design.

Acknowledgments

The work of Katia Colaneri was partially supported by the European Union–Next Generation EU—Project PRIN 2022 [2022BEMMLZ—CUP E53D23005660006] with the title Stochastic control and games and the role of information. Katia Colaneri is member of Gruppo Nazionale per l'Analisi Matematica, la Probabilitàe le loro Applicazioni (GNAMPA) of Istituto Nazionale di Alta Matematica (INdAM).

Open Access funding provided by Wirtschaftsuniversitat Wien.

Conflicts of Interest

The authors declare no conflicts of interest.

Data Availability Statement

The paper uses only simulations and no real data. The Python source code and the simulation data are available from the authors upon reasonable request.

References

Acemoglu, D., P. Aghion, L. Bursztyn, and D. Hemous. 2012. "The Environment and Directed Technical Change." *American Economic Review* 102, no. 1: 131–166.

Aghion, P., A. Dechezleprêtre, D. Hemous, R. Martin, and J. Van Reenen. 2016. "Carbon Taxes, Path Dependency, and Directed Technical Change: Evidence From the Auto Industry." *Journal of Political Economy* 124, no. 1: 1–51.

Aid, R., A. Kemper, and N. Touzi. 2025. "A Principal-Agent Model for Optimal Incentives in Renewable Investments." *International Journal of Theoretical and Applied Finance* 28: 01n02.

Aid, R., and S. Biagini. 2023. "Optimal Dynamic Regulation of Carbon Emissions Market." *Mathematical Finance* 33, no. 1: 80–115.

Basaglia, P., C. Berestycki, S. Carattini, A. Dechezleprêtre, and T. Kruse. 2025. "Climate Policy Uncertainty and Firms' and Investors' Behavior." Working Paper No. 11782. CESifo.

Beck, C., S. Becker, P. Cheridito, A. Jentzen, and A. Neufeld. 2021. "Deep Splitting Method for Parabolic PDEs." *SIAM Journal on Scientific Computing* 43, no. 5: A3135–A3154.

Benth, F., R. Kiesel, and A. Nazarova. 2012. "A Critical Empirical Study of Three Electricity Spot Price Models." *Energy Economics* 34, no. 5: 1589–1616.

Blyth, W., R. Bradley, D. Bunn, C. Clarke, T. Wilson, and M. Yang. 2007. "Investment Risks Under Uncertain Climate Change Policy." *Energy Policy* 35, no. 11: 5766–5773.

Carmona, R., G. Dayanikli, and M. Laurière. 2021. "Mean Field Models to Regulate Carbon Emissions in Electricity Production." *Dynamic Games and Applications* 12: 897–928.

Dumitrescu, R., M. Leutscher, and P. Tankov. 2024. "Energy Transition Under Scenario Uncertainty: A Mean-Field Game of Stopping With Common Noise." *Mathematics and Financial Economics* 18: 233–274.

Fleming, W., and M. Soner. 2006. *Controlled Markov Processes and Viscosity Solutions*. Vol. 25. Springer Science & Business Media.

Fleming, W., and P. Souganidis. 1989. "On the Existence of Value Function of Two-Player, Zero-Sum Stochastic Differential Games." *Indiana University Mathematics Journal* 38, no. 2: 293–314.

Flora, M., and P. Tankov. 2023. "Green Investment and Asset Stranding Under Transition Scenario Uncertainty." *Energy Economics* 124: 106773.

Frey, R., and V. Köck. 2022. "Deep Neural Network Algorithms for Parabolic PIDEs and Applications in Insurance and Finance." *Computation* 10, no. 11: 201.

Friedman, A. 1972. "Stochastic Differential Games." *Journal of Differential Equations* 11, no. 1: 79–108.

Fuss, S., J. Szolgayova, and M. Obersteiner. 2008. "Investment Under Market and Climate Policy Uncertainty." *Applied Energy* 85: 708–721.

Geman, H., and A. Roncoroni. 2006. "Understanding the Fine Structure of Electricity Prices." *Journal of Business* 79, no. 3: 1225–1261.

Golosov, M., J. Hassler, P. Krusell, and A. Tsyvinski. 2014. "Optimal Taxes on Fossil Fuel in General Equilibrium." *Econometrica* 82, no. 1: 41–88.

Hagspiel, V., P. M. Kort, and X. Wen. 2025. "Green Technology Investment: Announced vs. Unannounced Subsidy Retraction." *Journal of Economic Dynamics and Control* 170: 105030.

Hassett, K. A., and G. Metcalf. 1999. "Investment With Uncertain Tax Policy: Does Random Tax Policy Discourage Investment." *Economic Journal* 109, no. 457: 372–393.

Herrmann, S., J. Muhle-Karbe, and F. Seifried. 2017. "Hedging With Small Uncertainty Aversion." *Finance and Stochastics* 21: 1–64.

Klibanoff, P., M. Marinacci, and S. Mukerji. 2005. "A Smooth Model of Decision Making Under Ambiguity." *Econometrica* 73, no. 6: 1849–1892.

Ladyženskaja, O. A., V. Solonnikov, and N. Ural'ceva. 1988. *Linear and Quasi-Linear Equations of Parabolic Type*. Vol. 23. American Mathematical Society.

Lavigne, P., and P. Tankov. 2026. "Decarbonization of Financial Markets: A Mean-Field Game Approach", arXiv preprint, https://arxiv.org/abs/2301.09163.

Maccheroni, F., M. Marinacci, and A. Rustichini. 2006. "Ambiguity Aversion, Robustness, and the Variational Representation of Preferences." *Econometrica* 74, no. 6: 1447–1498.

Martinsson, G., L. Sajtos, P. Strömberg, and C. Thomann. 2024. "The Effect of Carbon Pricing on Firm Emissions: Evidence From the Swedish CO2 Tax." *Review of Financial Studies* 37: 1848–1886.

Nordhaus, W. 1993. "Optimal Greenhouse-Gas Reductions and Tax Policy in the DICE Model." *American Economic Review* 83, no. 3: 313–317.

Nordhaus, W. 2019. "Climate Change: The Ultimate Challenge for Economics." *American Economic Review* 109, no. 6: 1991–2014.

Pham, H. 1998. "Optimal Stopping of Controlled Jump Diffusion Processes: A Viscosity Solution Approach." *Journal of Mathematical Systems, Estimation and Control* 8: 1–27.

Pham, H. 2009. *Continuous-Time Stochastic Control and Optimization With Financial Applications*. Vol. 61. Springer Science & Business Media.

Possamaï, D., N. Touzi, and J. Zhang. 2020. "Zero-Sum Path-Dependent Stochastic Differential Games in Weak Formulation." *Annals of Applied Probability* 30, no. 3: 1415–1457.

Rockafellar, R. T. 1970. *Convex Analysis*. Princeton University Press.

Tan, J., and X. Tan. 2022. "Putting a Price on Carbon: Hard Targets to Keep Singapore Going." *ScholarBank@NUS Repository*, 1–11.

The Economist. 2023. "How to Deal With the Global Climate Backlash." October 14.

Yang, M. 2008. "Climate Policy Uncertainty and Investment Risk." Discussion Paper. International Energy Association.

Yang, M., W. Blyth, R. Bradley, D. Bunn, C. Clarke, and T. Wilson. 2008. "Evaluating the Power Investment Options With Uncertainty in Climate Policy." *Energy Economics* 30: 1933–1950.

Appendix A: Details on the Numerical Methodology

For the numerical experiments in Section 6, we implemented the deep splitting method that was proposed by Beck et al. (2021) and extended to partial integro-differential equations (PIDEs) by Frey and Köck (2022). This approach uses deep neural networks to approximate the solution of a PIDE together with the gradients. Hence, we are able to compute the value function for the considered stochastic control problems and determine investments in green technology accordingly. In this section, we present the basic idea of the algorithm. We consider a PIDE of the following form:

$$\begin{cases} u_t(t,\psi) + \mathcal{L}u(t,\psi) = f\left(t,\psi,u(t,\psi),\partial_\psi u(t,\psi)\right) & \text{on } [0,T)\times\mathbb{R}^n\,, \\ u(T,\psi) = g(\psi) & \text{on } \mathbb{R}^n\,. \end{cases} \tag{A1}$$

Here $n = d+2$, $\psi = (x,y,\tau) \in \mathbb{R}^n$, $\partial_\psi u$ is the gradient of u with respect to the space variable, u_t the derivative with respect to the time variable, and

$$\mathcal{L}u(t,\psi) := b(t,\psi)\cdot\partial_\psi u(t,\psi) + \frac{1}{2}\sum_{i,j=1}^{n}(\Sigma\Sigma^\top)_{ij}(t,\psi)u_{\psi_i\psi_j}(t_{\psi_i\psi_j},\psi)$$

$$+ \int_{\mathbb{R}} u(t,\psi+\widetilde{\Gamma}(t,\psi,z)) - u(t,\psi)\,m(dz)\,,$$

where $b(t,\psi) = (-\delta x, \alpha(t,y), 0)^\top \in \mathbb{R}^n$, $\Sigma(t,\psi)\in\mathbb{R}^{n\times n}$ has components $\Sigma_{1,1}(\psi) = \sigma$, $\Sigma_{i,j}(t,\psi) = \alpha_{i-1,j-1}(t,y)$ for $i,j = 2\ldots,n-1$ and all other components equal to zero, and $\widetilde{\Gamma}(t,\psi,z)\in\mathbb{R}^n = e_n\Gamma(t,\psi,z)$, where e_n is the nth standard vector in $\mathbb{R}^n$. Next, we consider an auxiliary process, denoted as Ψ, whose dynamics correspond to the generator $\mathcal{L}$,

$$\Psi_t = \Psi_0 + \int_0^t b(\Psi_s)\,ds + \int_0^t \Sigma(\Psi_s)\,d\widetilde{W}_s + \int_0^t\int_{\mathbb{R}^d}\widetilde{\Gamma}(s,\Psi_{s-},z)N(ds,dz)\,. \tag{A2}$$

Specifically, in our context, $\Psi_t = (X_t^0, Y_t, \tau_t)$, where X^0 is the uncontrolled version of the process X, that is, for $\gamma_t = 0$. The first step of the considered numerical algorithm is to divide the time horizon into N equidistant grid points $0 = t_0 < t_1 < \ldots < t_N = T$, where each interval is $\Delta t := 1/N$. Then, we discretize the process Ψ using a method such as the Euler–Maruyama scheme along the given time grid. This discretization yields approximations for Ψ_{t_i} at each time step t_i. We denote these approximation points as $\widehat{\Psi}_{t_i}$. For the solution u, we consider a BSDE representation. Given that the PIDE admits a classical solution, we can apply It ô's formula and express the solution as follows:

$$u(t_i,\Psi_{t_i}) = u(t_{i+1},\Psi_{t_{i+1}}) - \int_{t_i}^{t_{i+1}} f(s,\Psi_s,u(s,\Psi_s),\partial_\psi u(s,\Psi_s))\,ds - \int_{t_i}^{t_{i+1}}\Sigma(s,\Psi_s)^\top\partial_\psi u(s,\Psi_s)\,d\widetilde{W}_s \tag{A3}$$

$$-\int_{t_i}^{t_{i+1}}\int_{\mathbb{R}^d} u(s,\Psi_s+\widetilde{\Gamma}(s,\Psi_{s-},z)) - u(s,\Psi_{s-})\,(N(ds,dz)-m(dz)ds)\,, \tag{A4}$$

Both the integral with respect to the Brownian motion and the integral with respect to the compensated jump measure are martingales (assuming sufficient regularity of u). Taking conditional expectations leads to the following:

$$u(t_i,\Psi_{t_i}) = \mathbb{E}\left[u(t_{i+1},\Psi_{t_{i+1}}) - \int_{t_i}^{t_{i+1}} f(s,\Psi_s,u(s,\Psi_s),\partial_\psi u(s,\Psi_s))\,ds\middle|\Psi_{t_i}\right]. \tag{A5}$$

The discretization allows us to approximate the integral term in the conditional expectation by $f(t_{i+1},\widehat{\Psi}_{t_{i+1}},u(s,\widehat{\Psi}_{t_{i+1}}),\partial_\psi u(t_{i+1},\widehat{\Psi}_{t_{i+1}}))\Delta t$. Using the L^2-minimality of conditional expectations, we represent $u(t_i,\Psi_{t_i})$ as the unique solution of the minimization problem over all C^1 functions

$$\min_{U\in C^1}\mathbb{E}_{t_i}\left[\left(U - u(t_{i+1},\widehat{\Psi}_{t_{i+1}}) + f(t_i,\widehat{\Psi}_{t_{i+1}},u(t_{i+1},\widehat{\Psi}_{t_{i+1}}),\partial_\psi u(t_{i+1},\widehat{\Psi}_{t_{i+1}}))\Delta t\right)^2\right]$$

This minimization problem serves as a loss function for deep neural networks in the deep splitting algorithm, and the algorithm can be summarized as follows.

Deep splitting algorithm. Fix a class $\mathcal{N}$ of C^1 functions $\mathcal{U} : \mathbb{R}^d \to \mathbb{R}$ that are given in terms of neural networks with fixed structure. Then the algorithm proceeds by backward induction as follows.

1. Let $\widehat{\mathcal{U}}_N = g$.
2. For $i = N-1,\ldots,1,0$, choose $\widehat{\mathcal{U}}_i$ as minimizer of the loss function $L_i : \mathcal{N}\to\mathbb{R}$,

$$\mathcal{U}\mapsto\mathbb{E}\left[\left|\widehat{\mathcal{U}}_{i+1}(\widehat{\Psi}_{t_{i+1}}) - \mathcal{U}(\widehat{\Psi}_{t_i}) - \Delta t\, f\left(t_i,\widehat{\Psi}_{t_{i+1}},\widehat{\mathcal{U}}_{i+1}(\widehat{\Psi}_{t_{i+1}}),D_\psi\widehat{\mathcal{U}}_{i+1}(\widehat{\Psi}_{t_{i+1}})\right)\right|^2\right]. \tag{A6}$$

Specifically, to address the numerical solution of this problem in our case studies, we generate simulations of trajectories for the processes Ψ. These simulations were carried out over the time interval [0,15], discretized into 150 equally spaced time points (that is $N = 150$ intervals). The inherent nonlinearity of this problem is represented by the following function:

$$f(t, x, y, \tau, u_\psi) = \Pi^*(\psi) + \frac{\left((\partial_{\psi_1} u - 1)^+\right)^2}{4\kappa} = \Pi^*(x, y, \tau) + \frac{\left((\partial_x u - 1)^+\right)^2}{4\kappa}.$$

We use deep neural networks with two hidden layers, each containing 40 nodes. In total, each experiment involves 150 networks. The neural networks are initialized with random values using the Xavier initialization scheme. We employ mini-batch optimization with a mini-batch size of $M = 10,000$, incorporating batch normalization. The training process spans 10,000 epochs, and the loss function is minimized through the Adam optimizer. The learning rate starts at 0.01 and with a decay of 0.1 every 4000 steps. The activation function for the hidden layers is the sigmoid function, while the output layer uses the identity function.

An advantage of this methodology is flexibility. The approach allows for effortless dimensionality adjustments in the state process Ψ or modifications of its dynamics, with the only necessary adaptation being the Euler–Maruyama scheme for Ψ. For further details on deep splitting algorithms for general nonlinear PIDEs, we refer to Frey and Köck (2022).

Appendix B: Some Discussions and Proofs for the Tax-Risk Setting

In this section, we present various technical results that are related to the characterization of the value function as classical solution of the HJB equation (26).

B.1 | Comments and Extensions of Lemma 4.3

Here we discuss a few comments on possible extensions of the result of Lemma 4.3, as anticipated in Remark 4.4.

1. *Maximum capacity expansion*. In some examples, it may make sense to assume that investment can expand maximum capacity. In such case, a similar argument as in the proof of Lemma 4.3-(i) can be used to get regularity of the function Π^*. How to do that is briefly outlined next. If the maximum capacity depends on the investment level, that is, $q^{\max}(x)$, for some Lipschitz continuous, increasing and bounded function, the above arguments can be extended. Indeed,

$$\left|\Pi^*(x^1, y^1, \tau^1) - \Pi^*(x^2, y^2, \tau^2)\right| \tag{B1}$$

$$= \left|\max_{q\in[0,q^{\max}(x^1)]} \Pi(x^1, y^1, \tau^1, q) - \max_{q\in[0,q^{\max}(x^2)]} \Pi(x^2, y^2, \tau^2 q)\right| \tag{B2}$$

$$\leq \left|\max_{q\in[0,q^{\max}(x^1)]} \Pi(x^1, y^1, \tau^1, q) - \max_{q\in[0,q^{\max}(x^1)]} \Pi(x^2, y^2, \tau^2, q)\right| \tag{B3}$$

$$+ \left|\max_{q\in[0,q^{\max}(x^2)]} \Pi(x^2, y^2, \tau^2, q) - \max_{q\in[0,q^{\max}(x^1)]} \Pi(x^2, y^2, \tau^2, q)\right|. \tag{B4}$$

In the last expression, the first term is estimated exactly as in the proof of Lemma 4.3-(i). In the second term, Lipschitzianity in x is proved using Lipschitzianity of the function $q^{\max}(x)$.

2. *Concavity of the value function*. If Π^* and h are concave in x, then, it can be proved that V is also concave x. Before going to the proof of this result, we highlight that an example where Π^* is concave arises, for instance, if $\Pi(x, y, \tau, q)$ is concave in x and q^* is a fixed quantity. Indeed, the function $\Pi^*(x, y, \tau)$ is the result of an optimization and hence not an input variable of our model. This implies in particular that we cannot simply impose concavity, but we need to verify it, and, in general, even if Π is concave, the supremum over q may not be so. To establish concavity of the value function, one can follow the steps below. We let for simplicity $t = 0$. Consider $X_0^1, X_0^2 > 0$ and strategies $\boldsymbol{\gamma}^1, \boldsymbol{\gamma}^2 \in \mathcal{A}$. Denote by X^j, $j = 1, 2$, the investment process with initial value X_0^j and strategy $\boldsymbol{\gamma}^j$ and let for $\lambda \in [0, 1]$, $0 \leq t \leq T$, $\bar{X}_t = \lambda X_t^1 + (1 - \lambda) X_t^2$. Then it is easily seen that

$$d\bar{X}_t = \lambda\gamma_t^1 + (1 - \lambda)\gamma_t^2 - \delta\bar{X}_t dt + \sigma dW_t$$

so that $\bar{X}$ is the investment process corresponding to the strategy $\bar{\boldsymbol{\gamma}} = \lambda\boldsymbol{\gamma}^1 + (1 - \lambda)\boldsymbol{\gamma}^2$ with initial value $\bar{X}_0$ (Here we use that the dynamics of X are linear). Concavity of π^* and h now imply the following:

$$J(0, \bar{X}_0, y, \tau, \bar{\boldsymbol{\gamma}}) \geq \lambda J(0, X_0^1, y, \tau, \boldsymbol{\gamma}^1) + (1 - \lambda) J(0, X_0^2, y, \tau, \boldsymbol{\gamma}^2).$$

Concavity of V follows from this inequality, if we choose γ^j as an ε-optimal strategy for the problem with initial value X_0^j.

B.2 | Proof of Theorem 4.8

From Proposition 4.7, the function $V(t,x,y,\tau)$ is Lipschitz continuous in (x,y), Hölder in t and the unique viscosity solution of the PIDE

$$v_t(t,x,y,\tau) + \Pi^*(x,y,\tau) + \int_{\mathcal{Z}} v(t,x,y,\tau+\Gamma(t,y,\tau,z))m(dz) \tag{B5}$$

$$+\sum_{i=1}^{d}\beta_i(t,y)v_{y_i}(t,x,y,\tau) + \frac{\sigma^2}{2}v_{xx}(t,x,y,\tau) + \frac{1}{2}\sum_{i,j=1}^{d}\mathfrak{S}_{ij}(t,y)v_{y_iy_j}(t,x,y,\tau) \tag{B6}$$

$$+\sup_{0\le\gamma\le\bar{\gamma}}(\gamma(v_x(t,x,y,\tau)-1)-\kappa\gamma^2) - \delta x v_x(t,x,y,\tau) = (r+m(\mathcal{Z}))v(t,x,y,\tau), \tag{B7}$$

with the terminal condition $v(T,x,y,\tau)=h(x)$. For fixed τ, we define the function $f^\tau(t,x,y) := \int_{\mathcal{Z}} V(t,x,y,\tau+\Gamma(t,y,\tau,z))m(dz) + \Pi^*(x,y,\tau)$. Then for every fixed τ, $V^\tau(t,x,y) := V(t,x,y,\tau)$ is a viscosity solution of the equation

$$u_t(t,x,y) + \sum_{i=1}^{d}\beta_i(t,y)u_{y_i}(t,x,y) + \frac{\sigma^2}{2}u_{xx}(t,x,y) + \frac{1}{2}\sum_{i,j=1}^{d}\mathfrak{S}_{ij}(t,y)v_{y_iy_j}(t,x,y,\tau) \tag{B8}$$

$$+\sup_{0\le\gamma\le\bar{\gamma}}(\gamma(u_x(t,x,y)-1)-\kappa\gamma^2) - \delta x u_x(t,x,y) + f^\tau(t,x,y) = Ru(t,x,y), \tag{B9}$$

with $u(T,x,y)=h(x)$ and $R = r + m(\mathcal{Z})$. Note that this is a quasilinear parabolic PDE since there are no nonlocal terms and since for all $p\in\mathbb{R}$,

$$\sup_{0\le\gamma\le\bar{\gamma}}\{p\gamma-\gamma-\kappa\gamma^2\} = \begin{cases} 0 & \text{if } p<1 \\ \dfrac{[(p-1)^+]^2}{4\kappa} & \text{if } 1\le p\le 2\kappa+1 \\ \kappa\bar{\gamma}^2 & \text{if } p>2\kappa+1 \end{cases}$$

Our goal is to show that this PDE has a classical solution, which coincides with V^τ. We proceed in several steps.

Step 1. Fix $K>0$ and define the set $Q_K=[0,T]\times B_K$, where $B_K=\{\mathbf{x}\in\mathbb{R}^{d+1} : \|\mathbf{x}\|^2\le K^2\}$, and let $\mathcal{G}_K=\{T\}\times B_K\cup[0,T)\times S_K$ where $S_K=\{\mathbf{x}\in\mathbb{R}^{d+1} : \|\mathbf{x}\|^2=K^2\}$. Consider the terminal boundary value problem consisting of the PDE (B9) and the boundary condition $u=V^\tau$ on $\mathcal{G}_K$. We now use Theorem 6.4 in Ladyženskaja et al. (1988, Ch. 5) to show that this terminal boundary value problem has a classical solution that is moreover smooth on the interior of Q_K. For this, we formulate Equation (B9) as a parabolic equation in divergence form. We define for $y=(y_1,\dots,y_d)$, $p_2=(p_{2,1},\dots,p_{2,d})$ the functions

$$A(t,x,y,u,p_1,p_2) = \sum_{i=1}^{d}\beta_i(t,y)p_{2,i} + \sup_{0\le\gamma\le\bar{\gamma}}\{\gamma(p_1-1)-\kappa\gamma^2\} - \delta x p_1 - Ru + f^\tau(t,x,y) \tag{B10}$$

$$a(t,x,y,u,p_1,p_2) = A(t,x,y,u,p_1,p_2) + \sum_{i,j=1}^{d}\partial y_i\mathfrak{S}_{ij}(t,y)p_{2,j} \tag{B11}$$

$$a_1(t,x,y,u,p_1,p_2) = \frac{\sigma^2}{2}p_1 \tag{B12}$$

$$a_{2,i}(t,x,y,u,p_1,p_2) = \frac{1}{2}\sum_{j=1}^{d}\mathfrak{S}_{ij}(t,y)p_{2,j}, \quad i=1,\dots p \tag{B13}$$

Then Equation (B9) can be written in divergence form as in eq. (6.1) of Ladyženskaja et al. (1988, Chapter 5):

$$\partial_t u + \partial_x a_1(t,x,y,u,u_x,u_y) + \sum_{i=1}^{d}\partial y_i a_{2,i}(t,x,y,u,u_x,u_y) - a(t,x,y,u,u_x,u_y) = 0. \tag{B14}$$

Note that the signs differ from those in Ladyženskaja et al. (1988) since we are dealing with a *terminal* value condition.

Next we show that the assumptions of Theorem 6.4 in Ladyženskaja et al. (1988, Ch. 5) are satisfied on the domain Q_K. Note first that the set S_K is the boundary of the $d+1$-dimensional circle so it is smooth and hence satisfies condition (A) (see Ladyženskaja et al. 1988, 9). Moreover,

$$A(t,x,y,u,0,0)u = -(r+m(\mathcal{Z}))\,u^2 + f^\tau(t,x,y)\,u \geq -b_1u^2 - b_2$$

for $b_1, b_2 \geq 0$, since the functions $f^\tau(t,x,y)$ are bounded on Q_K. To see the latter recall that $f^\tau(t,x,y) := \int_{\mathcal{Z}} V(t,x,y,\tau+\Gamma(t,y,\tau,z))m(dz) + \Pi^*(x,\tau,y)$, and $\Pi^*(x,\tau,y)$ and V are bounded on the bounded set Q_K, respectively, $Q_K \times [0,\tau^{\max}]$. Hence the inequality holds. That guarantees that condition (a) of Theorem 6.4 (Ladyženskaja et al. 1988, Ch. 5) holds. Conditions (3.1), (3.2), (3.3), and (3.4) (Ladyženskaja et al. 1988, Ch. 5) are immediate. In particular, the condition $\sigma^2 > 0$ and the strict ellipticity of $\mathfrak{S}(t,y)$ ensure that the crucial condition (3.1) holds. Finally, since $V(t,x,y,\tau)$ is a Lipschitz viscosity solution of the HJB equation, the boundary condition is Lipschitz, which in particular implies condition (c) of Theorem 6.4 (Ladyženskaja et al. 1988, Ch. 5). By applying Theorem 6.4 in Ladyženskaja et al. (1988, Ch. 5), we thus get that in any interior subdomain Q_K, the HJB equation has a classical solution $U^\tau(t,x,y)$, which coincides with $V^\tau(t,x,y)$ on the boundary $\mathcal{G}_K$.

Step 2. Next we show that $U^\tau(t,x,y) = V(t,x,y,\tau)$ in the interior of Q_K for every K, which allows to conclude that $V(t,x,y,\tau)$ is smooth in the interior of Q_K. To prove this, we apply the comparison principle given by Fleming and Soner (2006, Corollary 8.1, Ch. 5). Note that inequality (7.1) on page 218 of the book is implied by in particular by Lipschitzianity of the functions α, β, Γ in y. Then we obtain that $U^\tau(t,x,y) = V(t,x,y,\tau)$ on Q_K.

Since K was arbitrary, we finally get that V is smooth everywhere. Hence V is also a classical solution of the HJB equation (26), which concludes the proof.

B.3 | An Example With Strict Viscosity Solution

In the following section, we present an example illustrating that, in general, the value function may be nonsmooth and hence a strict viscosity solution of the HJB equation. Specifically, we examine the cost function associated with the filter technology, assuming a fixed electricity price $\bar{p}$ and a fixed production quantity $\bar{q}$. To present this example with minimal technical difficulties, we make certain assumptions. We set r and δ to zero, take the residual value as $h(X_T) = 0$, and assume deterministic tax rate equal to $\bar{\tau} > 0$. Additionally, we adopt the abatement technology $e(x) = (1-x)^+$ and assume no external variations in the investment level ($\sigma = 0$). This assumption is crucial for our example, since for $\sigma > 0$, the HJB equation has a classical solution by Theorem 4.8 Section 4.4.

In this setting, $X_t = X_0 + \int_0^t \gamma_s ds$, and the value function is given by the following:

$$V(t,x) = \sup_{\gamma\in\mathcal{A}} \mathbb{E}_t\left[\int_t^T \left(\bar{p}\bar{q} - \bar{q}(\bar{c} + (1-X_s)^+\bar{\tau}) - \gamma_s - \kappa\gamma_s^2\right)ds\right] =: (\bar{p}\bar{q} - \bar{q}\bar{c})(T-t) + \tilde{V}(t,x) \tag{B15}$$

where

$$\tilde{V}(t,x) = \sup_{\gamma\in\mathcal{A}} \mathbb{E}_t\left[\int_t^T \bar{q}(1-X_s)^+\bar{\tau} - \gamma_s - \kappa\gamma_s^2\,ds\right]. \tag{B16}$$

In the sequel, we concentrate on $\tilde{V}$ and w.l.o.g. we take $\bar{q} = 1$. Note first that for $x \geq 1$, the optimal strategy is $\boldsymbol{\gamma}^* = 0$, since choosing $\gamma_s > 0$ is costly but generates no additional reduction in emissions. Therefore, $\tilde{V}(t,x) = 0$ for $x \geq 1$. Below we show that

$$\tilde{V}(t,x) \leq -(1-x)(1 \wedge (T-t)\bar{\tau}), \quad x \leq 1. \tag{B17}$$

Let $\tilde{V}_{x^-}(t,1)$ be the left derivative of $\tilde{V}(t,\cdot)$ at $x=1$. It follows that

$$\tilde{V}_{x^-}(t,1) = \lim_{h\to 0^+} \frac{1}{(-h)}(\tilde{V}(t,1-h) - \tilde{V}(t,1)) \geq (1 \wedge (T-t)\bar{\tau}\Big).$$

Hence $\tilde{V}(t,\cdot)$ has a kink at $x=1$, and from Equation (B16), we get that V is a strict viscosity solution of the HJB equation.

Now we turn to the inequality (B17). Obviously,

$$\tilde{V}(t,x) \leq w(t,x) := \sup_{\gamma\in\mathcal{A}} \mathbb{E}_t\left[\int_t^T -(1-X_s)^+\bar{\tau} - \gamma_s\,ds\right]. \tag{B18}$$

Since in Equation (B18) transaction costs are zero, the producer can push x instantaneously to any level $x' > x$, incurring a cost of size $x'-x$. It follows that for $x<1$, the "limiting optimal strategy" in Equation (68) is to push the investment level to 1 immediately at t, provided the resulting tax savings $\bar{\tau}(1-x)(T-t)$ exceed the cost $(1-x)$, and to choose $\boldsymbol{\gamma} \equiv 0$ otherwise. This gives

$$u(t,x) = \begin{cases} -(1-x) & \text{if } \bar{\tau}(T-t) > 1\,, \\ -\bar{\tau}(1-x)(T-t) & \text{if } \bar{\tau}(T-t) \leq 1\,, \end{cases}$$

that is, $u(t,x) = -(1-x)(1 \wedge (T-t)\bar{\tau})$, $x \le 1$, which implies (B17).

Appendix C: Differential Game

C.1 | Proof of Lemma 5.2

Define the compact and convex set $B := [0, q^{\max}] \times [\tau^{\min}, \tau^{\max}]$ and the function $F : B \to B$ by $F(q,\tau) = (q(\tau), \tau(q))'$. Note that $q(\tau)$ and $\tau(q)$ and hence F are continuous on B (since $\partial_q C_0$ and $\partial_q C_1$ are strictly increasing and since $\nu_1 > 0$). By Equation (41), (q^*, τ^*) is a saddle point of g if and only if it is a fixed point of F on B. The existence of a fixed point of F follows immediately from Brouwers fixed point theorem, which establishes the existence of a saddle point of g.

For uniqueness, note that the pair (q^*, τ^*) is a saddle point if and only if q^* satisfies the fixed point relation $q^* = q(\tau(q^*))$ and if $\tau^* = \tau(q^*)$. Define the mapping $\varphi : [0, q^{\max}] \to \mathbb{R}$ with

$$\varphi(q) := p - \partial_q C_0(q) - (\partial_q C_1(q) - \partial_q \nu_0(q))\tau(q).$$

By the first-order condition characterizing $q(\tau)$, a solution $q^* \in [0, q^{\max}]$ is a solution of the equation $q^* = q(\tau(q^*))$ if one of the following three conditions hold (i) $\varphi(q^*) = 0$; (ii) $\varphi(0) < 0$, in which case $q^* = 0$; (iii) $\varphi(q^{\max}) > 0$, in which case $q^* = q^{\max}$. Below, we show that φ is strictly decreasing. It follows that there is at most one $q^* \in [0, q^{\max}]$ that fulfills (i), (ii), or (iii) and hence at most one saddle point.

To show that φ is strictly decreasing, we first compute the derivative of φ for those values of q with $\tau(q) \in (\tau^{\min}, \tau^{\max})$. We get the following:

$$\begin{aligned}\partial_q \varphi(q) &= -\partial_q^2 C_0 - (\partial_q^2 C_1(q) - \partial_q^2 \nu_0(q))\tau(q) - (\partial_q C_1(q) - \nu_0'(q))\partial_q \tau(q) \\ &= -\partial_q^2 C_0 - (\partial_q^2 C_1(q) - \partial_q^2 \nu_0(q))\tau(q) - \frac{1}{2\nu_1}(\partial_q C_1(q) - \partial_q \nu_0(q))^2,\end{aligned}$$

which is negative due to the assumptions on C_0, C_1, and ν_0. For values of q where the constraints on τ bind, we have $\partial_q \tau(q) = 0$ and $\partial_q \varphi(q) = -\partial_q^2 C_0 - (\partial_q^2 C_1(q) - \partial_q^2 \nu_0(q))\tau(q) < 0$. It follows that φ is absolutely continuous with strictly negative derivative and hence strictly decreasing.