Bank of Russia

Uncertain technical change, Bayesian learning and the Green Paradox

WORKING PAPER SERIES

No. 133 / August 2024

A.A. Reentovich

Alexander Reentovich

Bank of Russia, Research and Forecasting Department

E-mail: ReentovichAA@cbr.ru

Bank of Russia Working Paper Series is anonymously refereed by members of the Bank of Russia Research Advisory Board and external reviewers.

The author is grateful to Andrey Sinyakov for drawing author's attention to the topic. The author also thanks Ksenia Yudaeva, Konstantin Styrin, Mark Sidorovsky, the anonymous referee and the participants of two Bank of Russia research seminars for valuable comments and suggestions. The author thanks Alexander Tsemakhman for preparing the Russian text.

Cover image: Shutterstock/FOTODOM

Address: 12 Neglinnaya street, Moscow, 107016
Tel.: +7 495 771-91-00, +7 495 621-64-65 (fax)
Website: www.cbr.ru

Uncertain technological change, Bayesian learning and the Green Paradox

Alexander Reentovich

Bank of Russia. ReentovichAA@cbr.ru

Abstract

The phenomenon of the 'Green Paradox' has been widely discussed in the climate economics literature for the last 15 years. The term refers to a situation in which a well-intended climate policy leads to adverse results, such as a rise in greenhouse gas emissions. The emergence of the paradox is usually attributed to the exhaustibility of fossil fuels: firms extracting these resources seek to equalise the present value of resource rents in each time period, so, anticipating future tax increases, they increase current production. On the other hand, if the productivity growth in the green sectors is driven by learning-by-doing, the Green Paradox does not emerge. In this paper, I show that the Green Paradox may arise as a consequence of an ex ante optimal policy in the absence of exhaustible resources if technological change in the clean sector is subject to uncertainty. That is, if the true speed of technological progress is an unknown parameter, economic agents have to form their expectations regarding future technological development with the help of Bayes' rule and make their decisions accordingly. If the market expects (a priori) the demand for dirty capital to shrink more rapidly due to technological progress than policymakers do, the latter must cut carbon taxes or even subsidize investment in dirty capital to avoid underinvestment in this type of capital and, consequently, the underproduction of energy in the present. If the flow of subsidies is persistent enough, CO_2 emissions may rise.

Keywords: Climate change, Bayesian learning, Green Paradox, Learning-by-doing, heterogeneous beliefs.

JEL classification: D81, D83, D84, Q54, Q55, Q58.

1 Introduction

Global warming is a serious threat to humankind. On the other hand, it is literally a textbook example of a negative externality – a problem thought to be resolvable by appropriate Pigouvian taxes. However, considerable doubts have been raised regarding the efficiency of such measures. As Sinn (2008; 2012) has argued, a well-intended policy may lead to results that are opposite to what is planned. That is, carbon emissions may rise in response to policies implemented to reduce them, at least temporarily. This situation is called the Green Paradox (GP).

The initial argument in favor of the GP drew upon the work of Hotelling (1931): if there is an exhaustible resource, it is optimal for resource extracting firms to equalize the present value of the resource rent in all periods. Given that the demand for such resource is downward sloping in each period, an increase in future sales taxes[1] stimulates the redistribution of extraction from the future, after tax hike, to the present (the future price will rise and the current price will fall until the present values of the resource rents equalize). The entire stock of the resource will eventually be extracted if each unit of it still yields positive rents. Thus, carbon taxes do not prevent climate change but make it happen earlier!

In the 2010s, this simple model was generalized and many simplified assumptions were dropped. Still, the discussion remained centered on fossil fuels extraction (Jensen, Mohlin, et al. 2015). This present paper shows that the GP can arise even in an economy without exhaustible resources. The necessary condition is the presence of durable dirty capital goods such as fossil-fuel-based power stations.[2]

Additionally, the term 'Green Paradox' has become synonymous with the 'unintended' consequences of climate policies(Jensen, Mohlin, et al. 2015, p. 246). This paper adopts even more general approach: any climate policy (of a benevolent government) is deemed paradoxical if it raises carbon emissions

[1] High taxes are assumed to be postponed for political reasons. However, the same effect arises if taxes are introduced immediately but are expected to increase at a rate greater than the interest.

[2] In fact, Smulders et al. (2012) have already demonstrated the possibility of the GP in the abscence of resource scarcity. However, they propose a mechanism quite different from the one presented here and they do not address the question of the optimality of the paradoxical policy assuming an ad hoc tax increase.

(whether intentionally or not), at least in the short run, compared to some benchmark, e.g. business-as-usual (competitive) scenario (BAU). Specifically, I show that an ex ante optimal policy may lead to GP.

Several authors have argued that green technological progress (GTP) in the form of learning-by-doing in renewable energy sector counteracts or alleviates the effects of the GP (Nachtigall and Rübbelke 2016; Hannesson 2018). The main reason is that green energy producers increase their current output if they expect their future productivity to rise due to learning-by-doing effect. Then, the current energy price falls, which makes extracting fossil fuels today less attractive[3]. Besides, technological progress in the green sector makes a backstop technology (which renders all the use of fossil fuels unceconomical) available sooner, so more carbon will be left in the ground. However, no existing studies look at GTP as an uncertain process by its very nature. On the contrary, in this paper, I propose a model in which green technological change is a stochastic process subject to parametric uncertainty which is dynamically resolved via Bayesian learning. I show that in this setting, GTP may be a cause of the GP itself. That is, if the non-renewable sector expects GTP to be fast and the state expects it to be slow, the government should reduce the carbon tax or even turn it into carbon subsidy to avoid underinvestment in dirty capital and, therefore, energy underproduction today. Thus, an optimal climate policy stimulates an increase in emissions.

The modelling of GTP this way is inspired by recent contributions to the literature on economic growth, such as the works of Mirman et al. 2016 or Fu and Le Riche 2021. These papers acknowledge the uncertainty regarding the future development of technology and embed Bayesian learning into different types of growth models.[4]

The notion of Bayesian learning under uncertainty is not new to climate economics – see, for example, the literature review contained in (Lemoine and Rudik 2017). However, most attention is paid to uncertainty regarding climate sensitivity (Kelly and Kolstad 1999; Hwang et al. 2017; Rudik et al.

---

[3]If green firms do not fully endogenize learning-by-doing (e.g. when its speed depends on the aggregate stock of clean capital, and not on stock owned by a separate firm), green subsidies become necessary for Pareto-optimal outcome.

[4]Interestingly, Fu and Le Riche (2021) are the first to consider a decentralized equilibrium under Bayesian learning though they do not allow agents to have different priors.

2018) or future policy measures (Dalby et al. 2018). Although there are a number of climate-related publications, such as those of Jensen and Traeger (2014) or García-León (2016), which take growth or technological uncertainty into account, these papers all concentrate on modelling separate agents: either a policymaker (Social Planner) or an investor[5]. Contrarily, the proposed model considers the interaction of multiple agents with heterogeneous prior beliefs. Thus, this paper falls into the literature on heterogeneous beliefs in climate-economic models (Bréchet et al. 2014; Kiseleva 2016; Nutz and Stebegg 2022), though these papers neither consider uncertainty regarding the development of green technologies nor treat agents as 'Bayesian statisticians'.

The contribution of this paper is therefore as follows: first, I consider an economy in which production with clean capital is characterized by a technology subject to stochastic technological change and parametric uncertainty, which agents resolve using Bayes' rule; second, I consider the interaction of several agents with heterogeneous priors; third, I show that in such circumstances, the GP is a possible outcome of the first-best policy even in the absence of imperfect competition, information asymmetry, or resource exhaustibility. It turns out in the course of the study that even a weak GP is possible only under extreme values of the model parameters. Nevertheless, policymakers should take heterogeneity of beliefs into account, because the optimal carbon tax is not equal to the social costs of carbon (SCC) unless the expectations of firms and the government coincide.

The rest of the paper is organized as follows: Section 2 outlines the basic model and lists the main theoretical findings; Section 3 introduces the computational techniques used to solve the basic model and its extensions and presents the results of numerical experiments; and Section 4 concludes the paper and discusses future work and policy implications.

[5]In the work of Karp and J. Zhang (2001) the policymaker interacts with a number of firms but the latter do not learn and optimize only intratemporally.

2 Basic model

Consider a partial equilibrium model with the following structure: suppose the economy consists of a policymaker, an energy producing firm and two firms producing different kinds of capital: a clean (ecologically friendly) one and a dirty (ecologically unfriendly) one (the variables corresponding to the former and to the latter are denoted by superscripts c and d, respectively). All firms are assumed to interact via perfectly competitive markets. The economy exists for an infinite number of periods, so each agent faces a dynamic optimization problem in discrete time.

Firm i ($i \in \{c, d\}$) chooses level of investment I_t^i to adjust existing capital stock K_{t-1}^i. The firms bear convex adjustment costs given by $(I_t^i)^{\alpha_i}, \alpha_i \in \{2, 4, 6, ...\}$. The law of motion for K_t^i is given by:

$$K_t^i = (1 - \delta_i)K_{t-1}^i + I_t^i, \tag{2.1}$$

where $\delta_i \in (0, 1)$. Both firms maximize the discounted flow of expected profit.[6]

The instantaneous welfare of the benevolent policymaker is given by[7]:

$$\begin{aligned} W_t &= U(Y_t) - F(I_t^d, I_t^c) - G(D_t) \\ &= \frac{Y_t^{1-\eta}}{1-\eta} - [(I_t^d)^{\alpha_d} + (I_t^c)^{\alpha_c}] - \kappa_d \frac{D_t^{1+\chi}}{1+\chi} \end{aligned} \tag{2.2}$$

where $\eta, \chi > 0$, Y_t is the amount of energy produced, $[(I_t^d)^{\alpha_d} + (I_t^c)^{\alpha_c}]$ shows the total adjustment costs, and D_t denotes the stock of pollution. The market time discount factor is denoted by by $\beta = (1 + r)^{-1} \in (0, 1)$ where r is the interest rate. Consumers' demand for energy is represented by the condition $p_t = U'(Y_t) = Y_t^{-\eta}$.

In each time period, the stock of pollution D_t increases proportionally to currently working stock of dirty capital, while the current stock of CO_2

[6] Differentiating capital into two categories is becoming popular in climate economics literature (Jin and Z. X. Zhang 2019; Baldwin et al. 2019).

[7] This is essentially the same function as that used in (Gronwald et al. 2017): the social utility of energy minus the total production costs minus the social costs of pollution. However, I consider a somewhat more general case which allows the final two terms to be nonlinear (strictly convex).

partially dissipates into the atmosphere:

$$D_t = (1-\delta)D_{t-1} + \gamma K_t^d. \tag{2.3}$$

where $\delta \in (0,1), \gamma > 0$.

Both capital stocks are rented for one period (the rental prices are denoted by p_t^c, p_t^d) to the energy producer. The state may levy sales taxes (τ_t^c, τ_t^d) on both of them. The energy producer combines the two types of capital using a stochastic technology characterized by gradually increasing productivity of the green capital[8,9]:

$$Y_t = A_y \left[\nu(A_t K_t^c)^\xi + (1-\nu)(K_t^d)^\xi\right]^{1/\xi}. \tag{2.4}$$

where $A_y > 0, \nu, \xi \in (0,1)$ and

$$A_t = \begin{cases} A_{t-1} + \lambda^m A_0, & \text{with probability } \theta, \\ A_{t-1}, & \text{with probability } 1-\theta, \end{cases} \tag{2.5}$$

where $m \leq t$ is the number of times productivity has risen before t (inclusively), $\theta \in (0,1), \lambda \in (0,1)$[10]. Note that

$$A_t \underset{t\to\infty}{\longrightarrow} \frac{A_0}{1-\lambda} \quad \text{a.s.} \tag{2.6}$$

[8] One may think of the energy production in my model in terms of (Baldwin et al. 2019): suppose energy (or electricity) Y_t is produced according to the CES technology $Y_t = A_y\left[\nu(H_t)^\xi + (1-\nu)(\Gamma_t^{ED})^\xi\right]^{1/\xi}$ where $H_t \equiv A_t K_t^c$ and Γ_t^{ED} refer to 'renewable production capacity' ('human capital in the form of knowledge, as well as the infrastructure itself') and 'dirty production capacity', respectively. In turn, dirty energy production requires combining dirty capital with fossil fuels O_t in a fixed proportion: $\Gamma_t^{ED} = \min\{K_t^d, \zeta O_t\}, \zeta > 0$. However, extraction is costless and unbounded from above so $\Gamma_t^{ED} \equiv K_t^d$.

[9] Note also that a somewhat nonstandard timing is used for convenience of presentation: I assume that new investment is installed immediately, so the capital of the current period is used in current production and current emissions increase proportionally to current dirty capital (and hence dirty investment). Under the standard timing of stocks, current investment would influence emissions only after two periods:

$$\begin{aligned} &D_{t+1} = (1-\delta)D_t + \gamma K_t^d, \\ &K_{t+1}^d = (1-\delta_d)K_t^d + I_t^d, \\ \Rightarrow \quad &D_{t+1} = (1-\delta)[(1-\delta)D_{t-1} + \gamma K_{t-1}^d] + \gamma[(1-\delta_d)K_{t-1}^d + I_{t-1}^d]. \end{aligned}$$

This feature does not entail any logical inconsistency because the investment good is not produced with capital as it is in most general equilibrium models; it could also easily be dropped. (Compare the Control Model from (Cyert and DeGroot 1974)).

[10] As in (Hannesson 2018), I model GTP as an exogenous process because 'green subsidies' are not in the focus of the present paper.

where 'a.s.' means 'almost sure convergence'. In what follows the term 'new' steady state refers to the asymptotic steady state with $A_t \equiv A_0/(1-\lambda)$.

Finally, GTP is characterized by parametric uncertainty, and agents must update their prior beliefs regarding the law behind GTP using Bayes' rule. That is, θ is unknown, however, it is known that $\theta \in \{\theta_H, \theta_L\}, \theta_H > \theta_L$,[11] and agent i has a prior belief in the form:

$$\theta = \begin{cases} \theta_H, & \text{with probability } q_0^i, \\ \theta_L, & \text{with probability } 1 - q_0^i. \end{cases} \tag{2.7}$$

The dynamics of agent i's belief is described by Bayes' rule[12]:

$$q_t^i = \begin{cases} \frac{\theta_H q_{t-1}^i}{\theta_H q_{t-1}^i + \theta_L (1-q_{t-1}^i)}, & A_t > A_{t-1}, \\ \frac{(1-\theta_H) q_{t-1}^i}{(1-\theta_H) q_{t-1}^i + (1-\theta_L)(1-q_{t-1}^i)}, & A_t = A_{t-1}. \end{cases} \tag{2.8}$$

The key feature of this model is that agents may have different prior beliefs. Thus, there are three types of posterior probabilities (and corresponding expectation operators) in the economy: q_t^p, q_t^c, q_t^d ($\mathbb{E}_t^p, \mathbb{E}_t^c, \mathbb{E}_t^d$) for the policymaker, the clean capital producer and the dirty capital producer, respectively.[13] It seems reasonable to assume different priors, since green technologies are completly novel and agents do not have a common reliable source of information regarding their future development.

Assume the following timing: before $t = 0$ the economy rests in the deterministic steady state with $A_t \equiv A_0$; at time 0, it becomes known that GTP will begin from period 1, so in period 0 the agents make their decisions according to their prior beliefs; starting from period 1, the agents observe the evolution of technology, update their beliefs, and make their decisions accordingly. The following proposition summarizes the first-best policy for the benevolent policymaker to implement. Analytical characterizations of the centralized and decentrazalized solutions (as well as proofs of all the

[11] In what follows expressions like 'the economy is in j-regime' or 'the state of the world is j' mean that the actual $\theta = \theta_j$.

[12] Assume that the technology shock and the update of the agents' beliefs happen in the beginning of the period, prior to any trade. At is assumed to be directly observable by everybody.

[13] The energy producer faces in fact static problem so this agent's expectations do not matter.

propositions stated below) are presented in Appendix A.

Proposition 1 *The optimal policy (OP) in the economy is given by:*

$$\hat{\tau}_t^d = \underbrace{\kappa_d \gamma \left[D_t^\chi + \beta(1-\delta)\mathbb{E}_t^p\left(D_{t+1}^\chi + \beta(1-\delta)D_{t+2}^\chi + ...\right)\right]}_{SCC} + \underbrace{\beta(1-\delta_d)\alpha_d\left[\mathbb{E}_t^d\left((I_{t+1}^d)^{\alpha_d-1}\right) - \mathbb{E}_t^p\left((I_{t+1}^d)^{\alpha_d-1}\right)\right]}_{SCDE^d} \tag{2.9}$$

and

$$\hat{\tau}_t^c = \underbrace{\beta(1-\delta_c)\alpha_c\left[\mathbb{E}_t^c\left((I_{t+1}^c)^{\alpha_c-1}\right) - \mathbb{E}_t^p\left((I_{t+1}^c)^{\alpha_c-1}\right)\right]}_{SCDE^c}. \tag{2.10}$$

That is, the optimal carbon tax is the sum of two components: the social costs of carbon (SCC) and the social costs of discrepancy in expectations (the SCDE). Unlike many existing models, the optimal carbon tax $\hat{\tau}_t^d$ may be negative if the SCDE are negative and greater in absolute value than SCC. This is the case if the managers of the dirty capital producing firm expect the demand for their output (and, hence, their own demand for investment) to shrink more rapidly due to the technological progress than the state expects. K^d then becomes sub-optimally low, and the state should subsidize dirty investment. On the contrary, if the dirty capital industry is more sceptical regarding the speed of GTP than the state, the SCDE become positive, and the state should restrict the production of K^d more agressively than under coinciding expectations (parametric certainty being a special case of this).

The same logic applies to the market for clean capital. That is, the state would impose a tax on clean capital, not a subsidy, if the clean industry were too optimistic regarding GTP (from the policymaker's perspective).

Note that the result of Proposition 1 depends neither on the nature of the stochasticity in the model nor on the specific priors which the agents might have. The sufficient condition for non-zero the SCDE is heterogeneity of expectations.

The following Proposition summarizes two main properties of the SCDE relevant for further discussion:

Proposition 2 *For any index $i \in \{c, d\}$:*

1.

$$sign(SCDE_t^i) = sign([q_0^i - q_0^p][(I_{t+1}^i)^{\alpha_i - 1}|_{A_{t+1}>A_t} - (I_{t+1}^i)^{\alpha_i - 1}|_{A_{t+1}=A_t}]); \tag{2.11}$$

2.

$$SCDE_t^i \underset{t\to\infty}{\longrightarrow} 0 \quad a.s. \tag{2.12}$$

Proposition 2 provides insight into how the optimal tax trajectory might look like. Initially, if dirty capital producers expect rapid increase in the productivity of clean capital, and these expectations are not shared by the state, the latter undertaxes (or even subsidizes) the production of dirty capital to prevent the production of energy from falling below the optimal level. However, as time passes the discrepancy of expectations lowers and the motive of reducing carbon emissions becomes dominant for the policymaker, which then eventually introduces a positive carbon tax (though for any finite t, it may still be below SCC by a small margin).[14]

I can now discuss under which conditions the GP might occur in the economy described. To start with, I introduce formal definitions of various degrees of the GP in the spirit of (Gronwald et al. 2017). Define the strict climate policy as follows: $\tau^d = SCC, \tau^c = 0$. That is, under this regime, the government neglects the SCDE completely. It seems sensible to discuss this option because the heterogeneity of expectations is not currently in the mainstream of climate policy debates. Let P be a certain policy (the optimal policy (OP), business-as-usual (BAU) or the strict policy (S)) and X_t^P be the value of variable X_t under this policy.

Definition 1 *A semi-weak (or the weakest) GP occurs when $D_t^{OP} > D_t^S$ for some $t \geq 0$.*

[14]Note that part 2 of the Proposition 2 is a manifestation of the well-known result that sooner or later Bayesian agents learn the true state of the world irrespective of their priors (Cyert and DeGroot 1974; Demers 1991). However, in the present model the transition to this state of 'full information' (or 'rational expectations') is of special interest as it goes hand in hand with the climate transition whose properties (like the emergence of the GP) this study is devoted to.

Definition 2 *Policy P invokes a weak GP if $D_t^P > D_t^{BAU}$ for some $t \geq 0$.*

Definition 3 *Policy P invokes a strong GP if:*

$$\frac{1}{1+\chi}\sum_{t=0}^{\infty}\beta^t(D_t^P)^{1+\chi} > \frac{1}{1+\chi}\sum_{t=0}^{\infty}\beta^t(D_t^{BAU})^{1+\chi}. \tag{2.13}$$

Definition 4 *Policy P invokes an extreme GP if*

$$\sum_{t=0}^{\infty}\beta^t W_t^P < \sum_{t=0}^{\infty}\beta^t W_t^{BAU}. \tag{2.14}$$

Note that I introduce a novel definition of a 'semi-weak' or ('the weakest') GP which corresponds to the case in which carbon emissions are higher under the OP than under the strict policy (but may still be lower than under BAU). This situation occurs when $SCDE^d$ are negative, at least in certain periods, so the discounted flow of carbon taxes is lower under OP than under the strict policy.

However, under the optimal taxation scheme, the economy may face a 'classic' weak GP if the initial $SCDE^d$ are negative, persistent, and so great in absolute value that the carbon subsidies in the initial periods outweigh the discounted carbon taxes in the later periods in the dirty capital producing firm's Euler equation (see (A.9) in the Appendix). The level of pollution in the initial periods is thus higher under the OP than under BAU. If the negative $SCDE^d$ are persistent enough, a strong GP may arise under the OP, i.e., the discounted sum of the disutility from pollution under this policy may be higher than under BAU.

It may seem that neither a strong nor extreme GP is possible under the optimal policy because, by construction, the benevolent government will not undertake any action which reduces the discounted flow of welfare. However, note that the definitions of the strong and the extreme GP involve ex post pollution and welfare, respectively. If the policymaker's prior belief is severely flawed (e.g., $q_0^p = 10^{-10}$ whereas $\theta = \theta_H$) and the firms' beliefs are not, then the ex ante optimal policy may actually lead to results worse than those obtained under BAU. Besides, the first-best policy is constructed to be optimal on average but not for each and every realization of $\{A_t\}_{t=0}^{\infty}$. Thus far, it remains unclear if such dramatic fall in welfare may occur under reasonable assumptions.

To answer this question the next Section employs numerical methods to assess the probability of observing various degrees of the GP in the economy described under different values of model parameters.

3 Numerical experiments

3.1 Solution method and calibration

To solve the model, that is, to find an approximation of the policy function that maps current state of the economy into the agents' controls, I resort to a variation of the Euler-residual minimization approach (Maliar et al. 2021). In essence, this requires the training of a neural network so as to minimize the mean sum of the squared residuals of the Euler equations:

$$\Xi_n = \frac{1}{n}\sum_{i=1}^{n}\sum_{j=1}^{l} v_j R_j^2(s_i, s_i'(\phi)), \tag{3.1}$$

where n is the batch size, l is the number of Euler equations, R_j^2 is the square of the j-th Euler residual, v_j is the corresponding weight (a hyper-parameter of the algorithm), s and s' are the states of the current and the following period, respectively, and ϕ is the vector of the parameters of the neural network[15]. For each i, the loss function is calculated as follows: each $s_t = (K_{t-1}^d, K_{t-1}^c, D_{t-1}, A_t, q_t^c, q_t^d, q_t^p)$ is randomly drawn from the ergodic set[16] and fed into the network, which returns the current controls (I_t^c, I_t^d and Φ_t – the Social Planner's Lagrange multiplier on emissions[17]) used to calculate current endogenous states K_t^d, K_t^c, D_t and static variables; after that, two ver-

[15]To be precise, I use the Euler residuals normalized by respective prices.

[16]In the context of this paper I am interested in the transition from one steady state to another so capital stocks and damages are sampled uniformly from the interval between their old and new steady state values; to get A_t samples I first sample m from a discrete uniform distribution on $[0, 200]$ and then calculate $A_t = A_0(\lambda^{m+1} - 1)/(\lambda - 1)$; the posterior probabilities are all sampled from $U[0, 1]$.

[17]Technically, I use the same network for both the centralized and decentralized solutions but in the latter case the output of the net corresponding to Φ_t is not used anywhere.

sions of the exogenous state $(A_{t+1}, q^c_{t+1}, q^d_{t+1}, q^p_{t+1})$ transition are calculated: with a rise in A and with no rise in A; for each of them next period controls and endogenous variables are calculated and weighted by the corresponding probabilities to find the relevant Bayesian expectations in the Euler equations.

As for the neural network architecture, I use a 3-layer perceptron with sigmoid activation. Its approximation capacity appears to be enough for the simplistic settings of this paper. I train the network using the Adam algorithm, making 50,000 gradient steps with an initial learning rate of 0.003 (decreasing 3 times every 20,000 steps) and batch size $n = 128$. Goodness-of-fit is checked by inspecting the mean trajectories of the simulated capital stocks produced (whether they approximate new steady state levels calculated analytically). This procedure (with appropriate modifications) is used to solve all the models presented further in this Section.

The trained network is used to simulate the trajectories of (K^d_t, K^c_t, D_t, W_t) under different policies for each of 10000 pre-simulated 1000-periods long trajectories of $\{A_t\}$. To estimate the probability of observing a weak GP in period t, I calculate the share of D^{BAU}_t being smaller than D^{OP}_t. To estimate the probability of a strong (extreme) GP I calculate the needed sum for each trajectory and then average them. The expected pollution at period t is calculated as the average D_t across the trajectories.

Next, Table 3.1 lists the values of model parameters which are used in all experiments if nothing else is stated explicitly. As this paper is more concerned with qualitative results than quantitative results, the values for calibration do not correspond to any real country or region but are set according to the following considerations: I assume both capital stocks to be symmetric in all aspects except initial productivity[18] in order to concentrate on the effects of GTP exclusively; I take the instantaneous disutility of pollution to be linear following (Gronwald et al. 2017)[19]; I also assume that priors

[18] That is, they have the same depreciation rate of 0.025 (Baldwin et al. 2019, δ^D) and are close substitutes (hence ξ, ν), but the productivity of green capital is initially lower by a factor of 2. The rate of depreciation of capital coincides with the rate of dissipation of carbon dioxide in the atmosphere (δ), taken from `https://euanmearns.com/the-half-life-of-co2-in-earths-atmosphere-part-1/`

[19] To make the SCC constant so that the strict policy simulation becomes easier: Equation (2.9) implies that $SCC \equiv \gamma\kappa_d/(1-\beta(1-\delta))$ if $\chi = 0$.

Table 3.1: Baseline calibration

Parameter	Value
η	1.4
χ	0
κ_d	0.01
β	0.99
λ	0.9
θ_L	0.2
θ_H	0.8
A_0	0.5
δ	0.025
γ	0.01
$\alpha_c = \alpha_d$	2
$\delta_c = \delta_d$	0.025
ξ	0.99
ν	0.5
A_y	1
q_0^p	0.01
q_0^d	0.99
q_0^c	0.99

q_0^i take 'extreme' values to make the effect in question more pronounced. The elasticity of the marginal utility of consumption η is assumed to be equal to the OECD average (Evans 2005)[20]. The interest rate is taken approximately equal to 1% ($\beta = 0.99$), which roughly corresponds to the average Fed funds rate for the years since 2008. θ_L, θ_H are chosen so that there is a noticeable difference between the regimes, and λ is chosen so that the productivity of green capital in the new stationary state is significantly higher than that of brown capital. The simulations are carried out assuming that the true state of the world is L.

[20] I assume that energy is the only product consumed by the households. Actually, an extensive meta-analysis by Labandeira et al. (2017) suggests that the short-run elasticity of electricity demand ($-1/\eta$) is as low in absolute value as 0.1 (though the value of $-1/1.4 \approx -0.7$ falls into the intervals reported by several surveys, according to (Labandeira et al. 2017, Table 1, Electricity, Energy); moreover, it is close to the estimates for Switzerland (Filippini 2011) and Norway (Nesbakken 1999)). However, setting $\eta = 10$ results in both capital stocks being lower in the new steady state than in the initial one (also under an alternative calibration in accordance with (Baldwin et al. 2019): $\xi = 0.46, \nu = 0.32, \delta_c = 0.04$) which seems counterintuitive. Green capital stock does rise in the new steady state under $\eta = 10$ if λ is assumed to be lower than 0.9, but even in this case the difference between the 'old' and 'new' K^c is quite small (not to mention that decreasing λ reduces the final 'ceiling' on the productivity of green capital).

3.2 Basic model under baseline calibration

The numerical simulations show that the weak GP effect is not present in the basic model under the optimal policy and the baseline calibration[21]. However, in accordance with the theory, the greater the dirty firm's belief in rapid GTP is, the less pollution is observed under BAU (see Fig. 3.1). Moreover, given $q_0^d > q_0^p$, the weakest GP is present in the model, as indicated by Fig. 3.2.

The absence of a weak GP can be explained by the high speed of the Bayesian learning: the agents' posterior beliefs become near to one another too qui so the low carbon taxes (or carbon subsidies) are not persistent enough to outweigh high taxes in the dirty firm's Euler equation. Besides, the SCDE might not be great enough in absolute value to compensate for the high SCC in Eq. (2.9). A natural question thus arises: might the GP be observed if the learning rate is lower, or if firms disregard future taxes to a large extent, or if the policymaker is less concerned with climate change? Thus, in the following subsections, I consider the implications of lower κ_d, lower β, and higher conservatism in the probability judgement for the GP.

3.3 Basic model under lower β

Suppose now that the discount factor is lower (the interest is higher). In this case, dirty firms are more concerned with near-term subsidies than with distant taxes, whereas the public (and the benevolent policymaker) cares less about the future stock of pollution, so the SCC are lower in each period. Both of these factors promote the emergence of the GP which is vindicated by the results presented in Fig. 3.3. Moreover, the probability of a strong GP is 1 for both $\beta = 0.95$ and $\beta = 0.9$ (i.e., $r \approx 0.053$ and $r \approx 0.11$, respectively) and the extreme GP occurs under $\beta = 0.9$ with probability 0.13.

[21]In fact, the calculations show that it is present in the very initial periods, but it is not persistent and disappears if the neural network used to approximate the policy function is trained with Euler residuals normalized by the initial steady state values of the relevant Lagrange multipliers. This is why I regard the initial result as spurious.

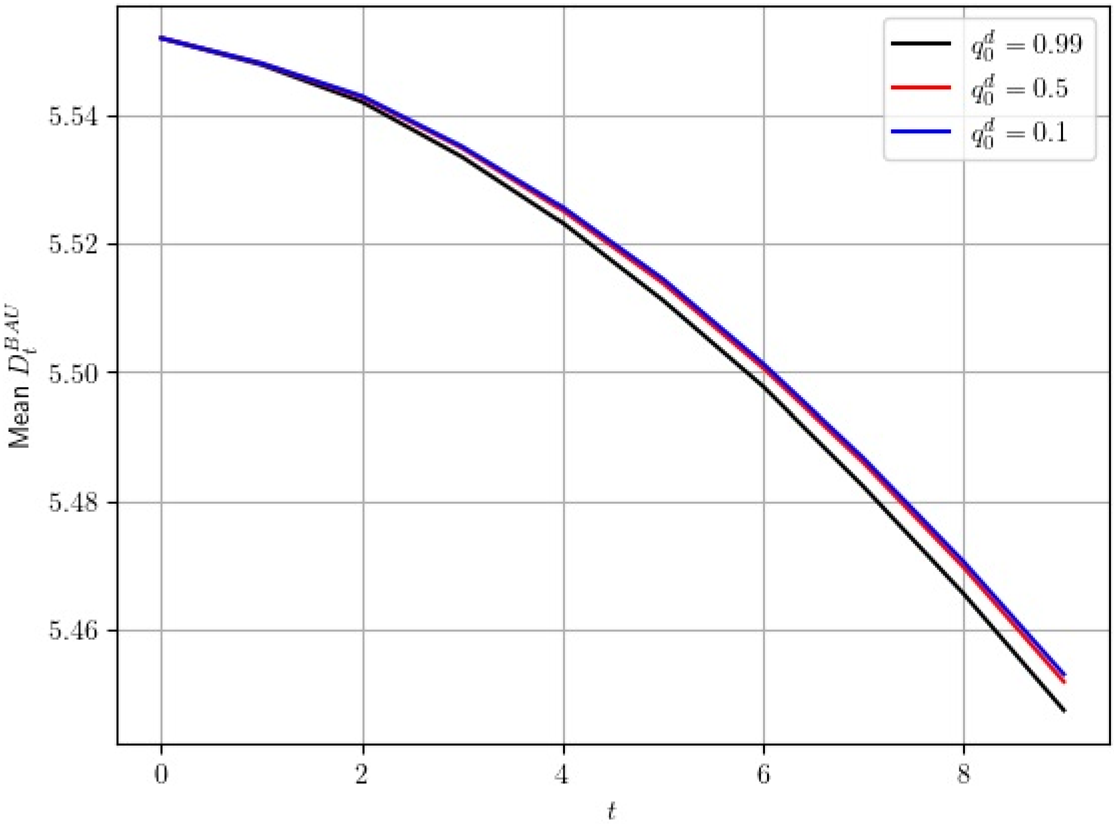

Figure 3.1: Basic model: expected pollution under BAU for different values of q_0^d. Source: Author's calculations.

3.4 Basic model under lower κ_d

Suppose now that the relative weight of the disutility of pollution in social welfare (κ_d) is lower. In this case, the population (and the benevolent policymaker) cares less about climate change, the SCC are lower in each period, so the carbon taxes are lower and the carbon subsidies (if any) are higher. Fig. 3.4 proves that lowering κ_d actually leads to a weak GP. However, the probability of a strong/extreme GP is always 0 due to the low persistency of excessive pollution under OP.

3.5 Basic model under lower θ_H, θ_L

Consider the case in which the probability of an increase in A_t in both regimes is lower than in the baseline calibration by the same factor, i.e., assume now that $\theta_H = 0.6$, $\theta_L = 0.15$. Thus, it becomes harder for agents to distinguish the two regimes because A_t changes more rarely under them. It can be clearly

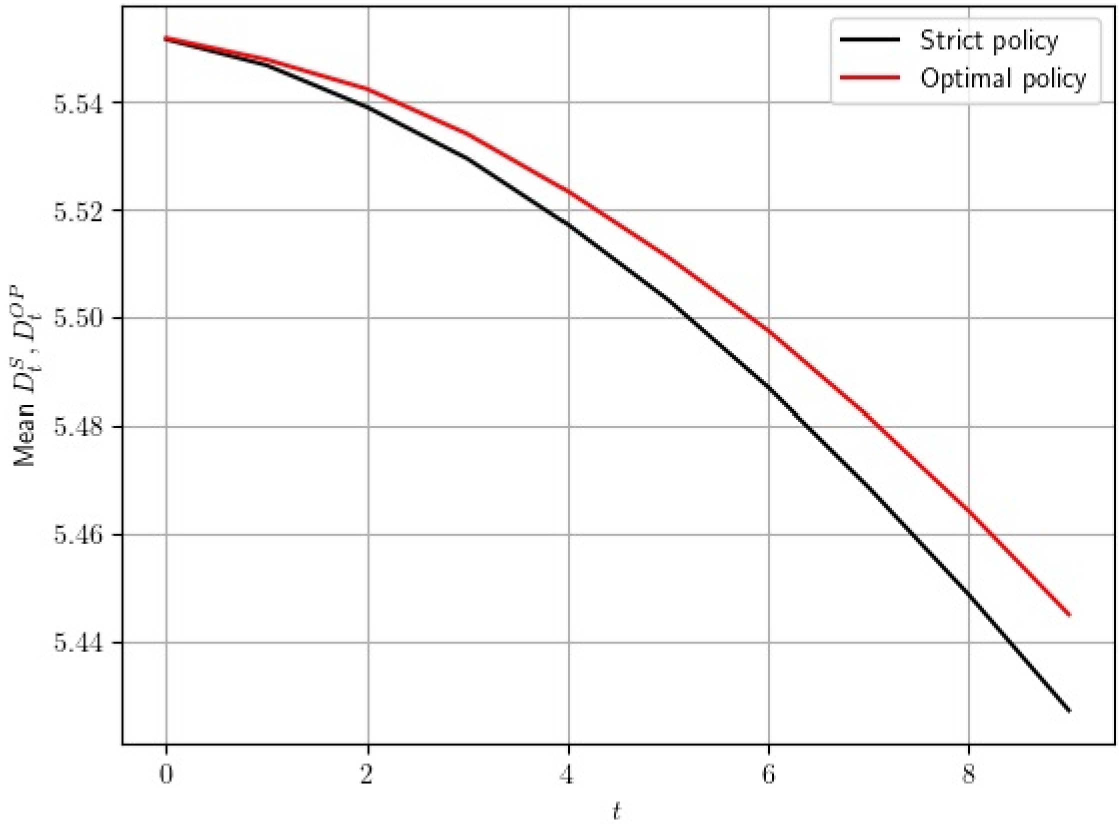

Figure 3.2: Basic model: expected pollution under optimal and strict policies. Source: Author's calculations.

seen from Fig. 3.5 that the probability of a weak GP is non-negligible under this setting. However, the excessive pollution under OP is not persistent enough to generate a strong/extreme GP and its probability is again 0.

3.6 Model with conservative agents

Assume that the agents are bounded-rational[22] and reluctant to change their views so they update their beliefs according to the following rule:

$$q_t^i = \begin{cases} (1-\omega)\frac{\theta_H q_{t-1}^i}{\theta_H q_{t-1}^i + \theta_L (1-q_{t-1}^i)} + \omega q_{t-1}^i, & A_t > A_{t-1}, \\ (1-\omega)\frac{(1-\theta_H) q_{t-1}^i}{(1-\theta_H) q_{t-1}^i + (1-\theta_L)(1-q_{t-1}^i)} + \omega q_{t-1}^i, & A_t = A_{t-1}. \end{cases} \tag{3.2}$$

where $\omega \in (0,1)$. For simplicity, assume that ω is common for all agents. Fig. 3.6 shows the probability that a weak GP is observed in each period

[22] For an overview of behavioral patterns relevant for climate economics see (Gsottbauer and van den Bergh 2011).

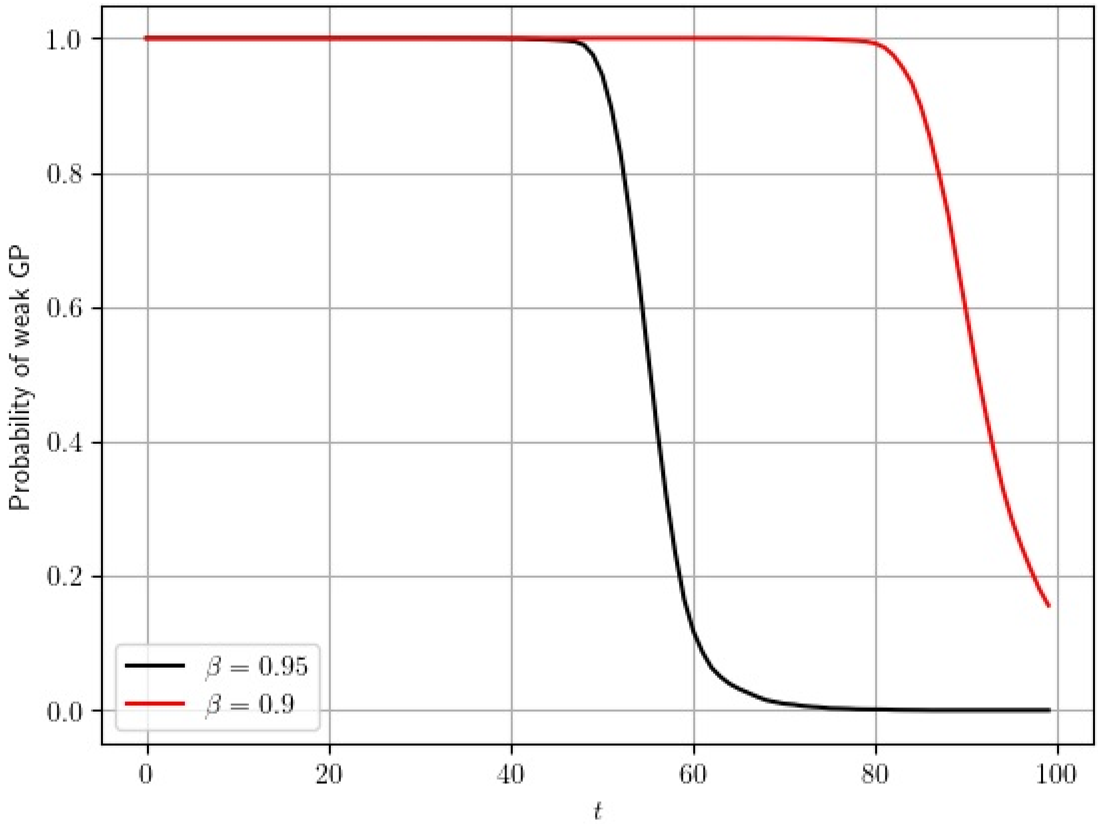

Figure 3.3: Basic model: the probability of weak GP for different β. Source: Author's calculations.

for different values of ω. As conjectured above, strong conservatism in probability judgment leads to a weak GP. Moreover, the persistency of the GP effect increases in the degree of anchoring (ω). However, the probability of a strong/extreme GP appears to be zero, i.e., the high level of pollution in the initial periods does not outweigh the low level in subsequent periods.

3.7 Model with noisy technological progress

Finally, the results of Kelly and Kolstad (1999) suggest that adding noise to the process of technological change obscures the technology trend so that learning (and the closure of the gap in expectations) takes longer. One might thus conjecture that this will lead to the same results as in the case of conservative agents.

To test the above hypothesis, suppose now that the technology can be decomposed into two separately unobservable components, trend A_t^T and

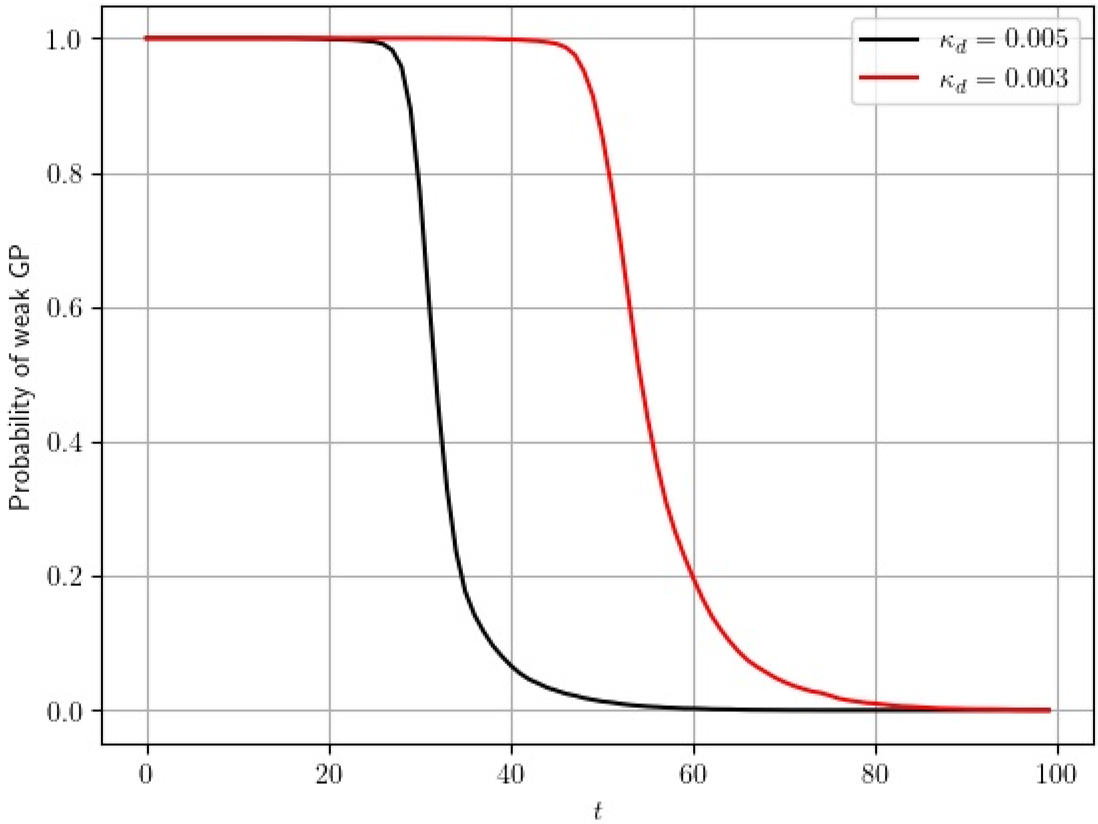

Figure 3.4: Basic model: the probability of weak GP for different κ_d. Source: Author's calculations.

cycle A_t^C:

$$A_t = A_t^T A_t^C, \quad A_t^C = \exp(\varepsilon_t), \quad \varepsilon_t \sim \mathcal{N}\left(-\frac{\sigma^2}{2}, \sigma^2\right). \tag{3.3}$$

Here the evolution of the trend is described by the standartd equation (2.5) whereas the cycle component is just a log-normal random noise with zero autocorrelation, unity mean, and standard deviation equal to $\sqrt{\exp(\sigma^2) - 1}$.

Now, each q_t^i is updated according to the rule

$$q_t^i = \frac{f_t(A_t|H) q_{t-1}^i}{f_t(A_t|H) q_{t-1}^i + f_t(A_t|L)(1 - q_{t-1}^i)} \tag{3.4}$$

where $f_t(A_t|j)$ is the probability density of $\log A_t$ given that the economy is in j-regime, $j \in \{H, L\}$, estimated at the observed value of A_t. In turn,

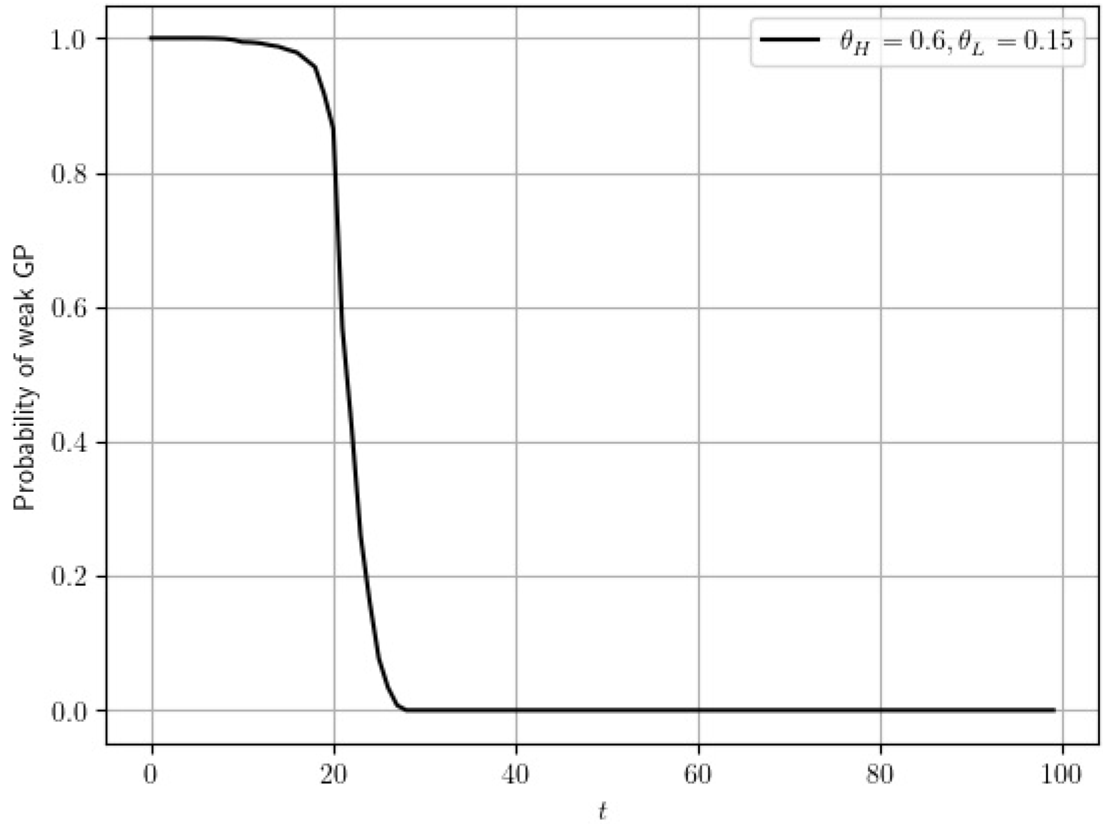

Figure 3.5: Basic model: the probability of weak GP under lower θ_H, θ_L. Source: Author's calculations.

according to the law of total probability:

$$f_t(x|j) = \sum_{k=0}^{t} \binom{t}{k} \theta_j^k (1-\theta_j)^{t-k} \mathcal{N}(x|\log\mu_k - \sigma^2/2, \sigma^2), \tag{3.5}$$

$$\mu_k = \frac{A_0(\lambda^{k+1}-1)}{\lambda - 1} \tag{3.6}$$

where $\mathcal{N}(x|a, b^2)$ is the probability density function of the normal distribution with mean a and standard deviation b.

Here, when calculating the expected values in the Euler equations, I resort to a simple Monte-Carlo method: first, I include t in the vector of states[23]; then, for a given t for each regime (H, L), I draw 100 next period $\log A_{t+1}$ from the distribution (3.5); for each of them, I calculate the necessary variables for the following period and average them (to find the expected values conditional on regime $\mathbb{E}[X_{t+1}|j], j \in \{H, L\}$). Then, I find the weighted average of these

[23]In this subsection the distribution of A_{t+1} depends not only on A_t but also on t itself. There are also other minor modifications to the solution code which I omit here due to their purely technical nature.

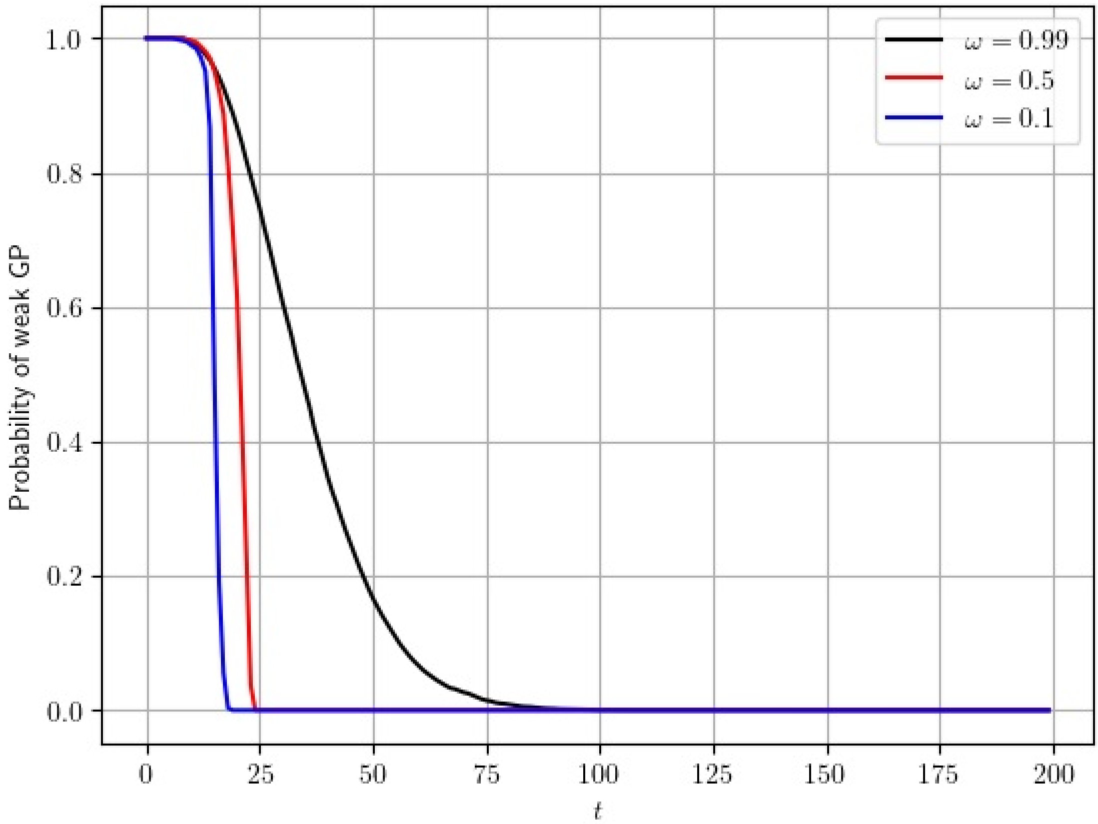

Figure 3.6: Model with conservative agents: probability of weak GP for different ω. Source: Author's calculations.

averages, according to the law of total expectation: $\mathbb{E}_t^i X_{t+1} = q_t^i \mathbb{E}[X_{t+1}|H] + (1 - q_t^i)\mathbb{E}[X_{t+1}|L]$.

Surprisingly, adding noise to the technological change[24] leads to zero probability of a weak GP. To grasp the intuition behind this result (alas, this is not a completely rigorous proof), consider the following Proposition (which may be quite interesting in and of itself):

Proposition 3 *Ceteris paribus, if $\xi + \eta > 1$, the one-period-ahead expected price for dirty capital at time t is higher in the model with noisy technological change than in the basic model.*

This result is in line with the finding of Hartman (1972) that more uncertainty (in the form of mean-preserving spread) regarding the future prices for a firm's output leads to higher expected marginal revenues and higher current investment[25]. However, Proposition 3 demonstrates that the same logic

[24] I tested values of 0.5, 0.75 and 1 for the standard deviation of A_t^C.

[25] In fact, both these results are driven by Jensen's inequality.

applies if the competitor's productivity undergoes mean-preserving spread, not the price for the firm's own output.

The symmetric statement for p^c is not true in general, since this price is not concave or convex for all values of A. However, assuming that capital stocks are equal to their initial steady state values and considering $A \leq 5$, it can clearly be seen that p^c and p^d are strictly concave and convex in A, respectively (see Fig. 3.7). So, it seems likely that, under a more volatile A, the expected price of clean (dirty) capital is lower (higher), at least for small t.

Therefore, initially both competitive firms and the Social Planner have a higher (lower) propensity to invest in dirty (clean) capital when more uncertainty is present. However, the Social Planner increases K^d to a lesser extent when it is concerned with increased pollution. This is why the desire of the policymaker to subsidize the production of dirty capital is lower in the noisy GTP model. Thus, rising uncertainty in the productivity of green capital counteracts the GP.

4 Discussion and conclusion

In this paper, I have shown that the Green Paradox may arise in an economy with clean and dirty durable capital goods (even without exhaustible resources) for the following reasons. If the environmentally friendly sector is characterized by technological change of unknown speed, economic agents must form their expectations regarding future technological development with the help of Bayes' rule and act accordingly. If the dirty capital producers expect (a priori) the demand for their output to shrink more rapidly due to technological progress than the state expects, the latter must cut carbon taxes or even subsidize investment in dirty capital to avoid underproduction today. If the flow of subsidies is persistent enough, CO_2 emissions may rise.

The mechanism of the Green Paradox presented in this paper is rather non-standard, and, more importantly, unlike many existing models, the Green Paradox effect under the best policy is always optimal and desirable (at least

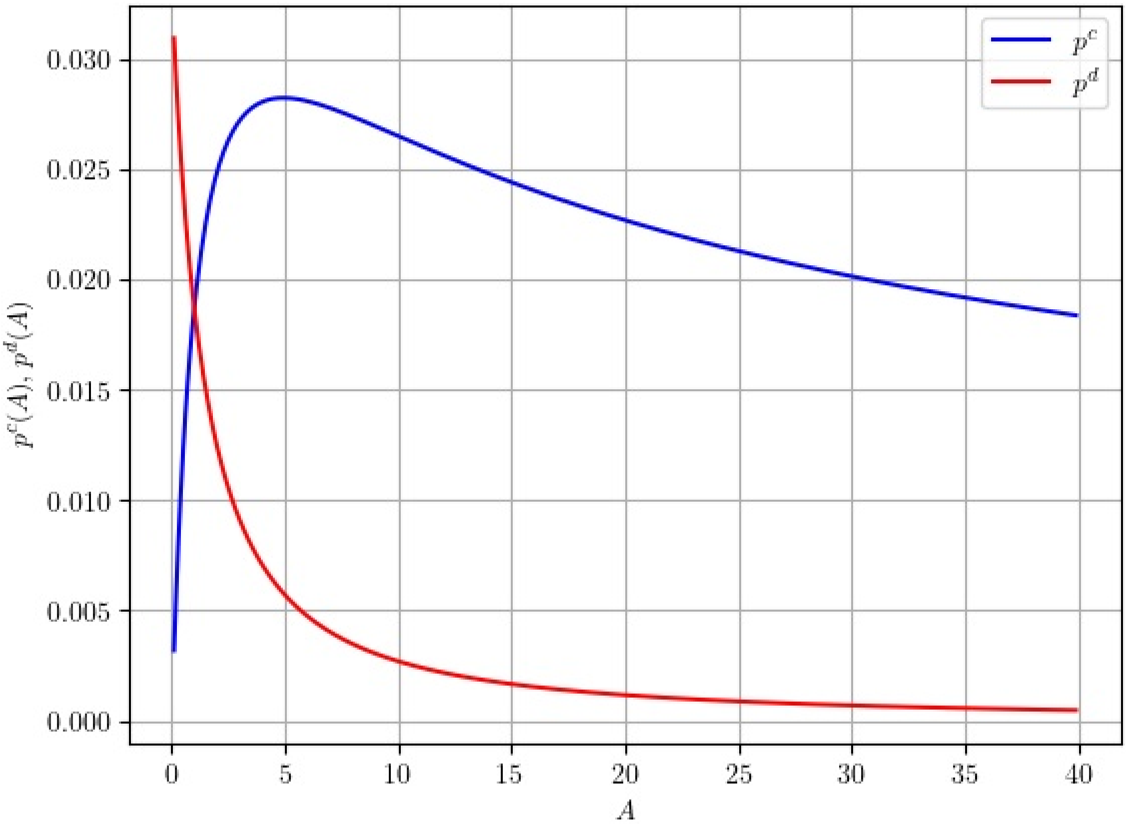

Figure 3.7: Model with noisy technological progress: prices of capital stocks for different A under initial K^d, K^c.

ex ante) from the standpoint of social welfare. The implementation of such optimal (but paradoxical) policies would be challenging.

However, the numerical analysis has shown a hopeful picture: even the weak Green Paradox (temporary increase in emissions) under the optimal policy should be considered a limiting case: it is possible with an extremely low weight for the disutility of carbon in the social welfare function, a high interest rate, a low speed of technological progress (under any regime), or conservative updates of posterior beliefs. The first condition is not met in practice, as even under the baseline calibration, the emissions in the new steady state under the optimal policy amount to 35% of the emissions under business as usual, whereas real-world politicians usually advocate a completely carbon-free economy. Interest rates, in turn, were quite low in the developed countries during the 2010s, and it is unclear if their current surge will last long. In addition, recent studies in the field of cognitive psychology show that the concept of conservatism is not empirically relevant in probability judgement (Woike et al. 2023).

Moreover, the GP requires a specific relationship between prior beliefs, namely $q_0^d > q_0^p$, i.e., the future of green technologies must seem brighter to dirty capital producers than to policymakers. If this inequality has the opposite sign, the optimal carbon tax exceeds the social costs of carbon. It remains a question for future empirical studies which values these parameters take in the real world. Research in this direction is quite relevant, because (and this is the main policy implication of this paper) the optimal carbon tax differs from the social costs of carbon unless $q_0^d = q_0^p$. Moreover, designing quantitative parameters of optimal climate policies requires introducing heterogeneous beliefs in models with a more elaborate structure than the one presented here has.

Finally, there are several areas for future theoretical studies. First, in this paper, I treat technological change in the green sector as a purely exogenous process. However, in many papers (e.g. Nachtigall and Rübbelke 2016; Baldwin et al. 2019), it is assumed to depend on the level of production or capital in the clean sector. Introducing such a feedback mechanism (while keeping the assumptions of parametric uncertainty and stochasticity) would not only make the model more realistic but also allow agents to experiment, i.e., to change their controls to make the signal more informative and to learn actively. In addition, it should not be forgot that fossil fuel extraction is costly and bounded from above in the real world. The implications of these features under Bayesian learning and heterogeneous priors remain to be studied.

Moreover, it would be interesting to study whether green subsidies can be used instead of brown subsidies when the Green Paradox arises according to the mechanism described above. In other words, the question is under what conditions (if any) the benevolent government might decide to increase the stock of clean capital to compensate for the low amount of dirty capital in order to avoid the underproduction of energy. Comparing the relative merits of different policy instruments (such as when only one of them is available to the policymaker) in the context of the model above remains a subject for further research.

Next, the only source of (parametric) uncertainty considered in the present paper is the change in the productivity of green capital. It remains to be studied how the uncertainties regarding the parameters of climate sensitivity, the political environment, production conditions in the dirty energy sector as well as in other industries, etc., influence the climate transition under Bayesian learning and the heterogeneity of agents' prior beliefs.

Another set of policy-relevant research questions concerns the signals that influence the formation of expectations:[26]: what are these signals, can the policymaker control them, might the manipulation of these signals be a (partial) substitute for carbon taxes or green subsidies,[27] and can the expectations gap be closed at the beginning of the green transition to prevent currently favored strict policies from becoming suboptimal? Answering these questions is likely to have a large impact on how we think of managing climate change.

References

Baldwin, Elizabeth, Yongyang Cai, and Karlygash Kuralbayeva (2019). *To Build or not to Build? Capital Stocks and Climate Policy*. OxCarre Working Papers 204, Oxford Centre for the Analysis of Resource Rich Economies, University of Oxford.

Bréchet, Thierry, Carmen Camacho, and Vladimir M. Veliov (2014). "Model predictive control, the economy, and the issue of global warming". In: *Annals of Operations Research* 220.1, pp. 25–48.

Cyert, Richard M. and Morris H. DeGroot (1974). "Rational Expectations and Bayesian Analysis". In: *Journal of Political Economy* 82.3, pp. 521–536.

Dalby, Peder A.O. et al. (2018). "Green investment under policy uncertainty and Bayesian learning". In: *Energy* 161, pp. 1262–1281.

[26] In this paper, the only signal regarding the future development of technology is the current state of technology itself. One may ask if, say, any scientific forecasts regarding the trajectory of $\{A_t\}$ influence the evolution of beliefs.

[27] According to Demers (1991), 'Bayesian' firms reduce the current level of irreversible investment if they are waiting for an informative signal regarding future market conditions. Thus, a policymaker may face a number of non-obvious trade-offs. For example, where should a given amount of budget money be allocated: to green subsidies or to a research project which will someday clarify the uncertainty which dirty capital producers face now?

Demers, Michel (1991). “Investment under Uncertainty, Irreversibility and the Arrival of Information Over Time”. In: *Review of Economic Studies* 58.2, pp. 333–350.

Evans, David (2005). “The Elasticity of Marginal Utility of Consumption: Estimates for 20 OECD Countries”. In: *Fiscal Studies* 26.2, pp. 197–224.

Filippini, Massimo (2011). “Short- and long-run time-of-use price elasticities in Swiss residential electricity demand”. In: *Energy Policy* 39.10, pp. 5811–5817.

Fu, Wentao and Antoine Le Riche (2021). “Endogenous growth model with Bayesian learning and technology selection”. In: *Mathematical Social Sciences* 114, pp. 58–71.

García-León, David (2016). *Adapting to Climate Change: an Analysis under Uncertainty.* Working Papers 2016.10, Fondazione Eni Enrico Mattei.

Gronwald, Marc, Ngo Van Long, and Luise Roepke (2017). *Three Degrees of Green Paradox: The Weak, The Strong, and the Extreme Green Paradox.* Cahiers de recherche 02-2017, Centre interuniversitaire de recherche en économie quantitative, CIREQ.

Gsottbauer, Elisabeth and Jeroen C. J. M. van den Bergh (2011). “Environmental Policy Theory Given Bounded Rationality and Other-regarding Preferences”. In: *Environmental and Resource Economics* 49, pp. 263–304.

Hannesson, Rögnvaldur (2018). *The Green Paradox and learning by doing.* Discussion Papers 2018/17, Norwegian School of Economics, Department of Business and Management Science.

Hartman, Richard (1972). “The effects of price and cost uncertainty on investment.” In: *Journal of Economic Theory* 5.2, pp. 258–266.

Hotelling, Harold (1931). “The economics of exhaustible resources”. In: *Journal of Political Economy* 39.2, pp. 137–175.

Hwang, In Chang, Frédéric Reynès, and Richard S.J. Tol (2017). “The effect of learning on climate policy under fat-tailed risk”. In: *Resource and Energy Economics* 48, pp. 1–18.

Jensen, Svenn, Kristina Mohlin, et al. (2015). "An Introduction to the Green Paradox: The Unintended Consequences of Climate Policies". In: *Review of Environmental Economics and Policy* 9.2, pp. 246–265.

Jensen, Svenn and Christian P. Traeger (2014). "Optimal climate change mitigation under long-term growth uncertainty: Stochastic integrated assessment and analytic findings". In: *European Economic Review* 69, pp. 104–125.

Jin, Wei and Zhong Xiang Zhang (2019). *Capital Accumulation, Green Paradox, and Stranded Assets: An Endogenous Growth Perspective.* Working Papers 2018.33, Fondazione Eni Enrico Mattei.

Kamihigashi, Takashi (2005). "Necessity of the transversality condition for stochastic models with bounded or CRRA utility". In: *Journal of Economic Dynamics and Control* 29.8, pp. 1313–1329.

Karp, Larry and Jiangfeng Zhang (2001). *Bayesian Learning and the Regulation of Greenhouse Gas Emissions.* UC Berkeley: Department of Agricultural and Resource Economics. URL: https://escholarship.org/uc/item/2fr0783c.

Kelly, David and Charles Kolstad (1999). "Bayesian learning, pollution, and growth". In: *Journal of Economic Dynamics and Control* 23.4, pp. 491–518.

Kiseleva, Tatiana (2016). "Heterogeneous Beliefs and Climate Catastrophes". In: *Environmental and Resource Economics* 65.3, pp. 599–622.

Labandeira, Xavier, José M. Labeagac, and Xiral López-Otero (2017). "A meta-analysis on the price elasticity of energy demand". In: *Energy Policy* 102, pp. 549–568.

Lemoine, Derek and Ivan Rudik (2017). "Managing Climate Change Under Uncertainty: Recursive Integrated Assessment at an Inflection Point". In: *Annual Review of Resource Economics* 9, pp. 117–142.

Maliar, Lilia, Serguei Maliar, and Pablo Winant (2021). "Deep learning for solving dynamic economic models". In: *Journal of Monetary Economics* 122, pp. 76–101.

Mirman, Leonard J., Kevin Reffett, and Marc Santugini (2016). "On learning and growth". In: *Economic Theory* 61.4, pp. 641–684.

Nachtigall, Daniel and Dirk Rübbelke (2016). “The green paradox and learning-by-doing in the renewable energy sector”. In: *Resource and Energy Economics* 43, pp. 74–92.

Nesbakken, Runa (1999). “Price sensitivity of residential energy consumption in Norway”. In: *Energy Economics* 21.6, pp. 493–515.

Nutz, Marcel and Florian Stebegg (2022). “Climate Change Adaptation under Heterogeneous Beliefs”. In: *Mathematics and Financial Economics* 16, pp. 481–508.

Rudik, Ivan, Derek Lemoine, and Maxwell Rosenthal (2018). *General Bayesian Learning in Dynamic Stochastic Models: Estimating the Value of Science Policy*. 2018 Meeting Papers 369, Society for Economic Dynamics.

Sinn, Hans-Werner (2008). “Public policies against global warming: a supply side approach”. In: *International Tax and Public Finance* 15.4, pp. 360–394.

— (2012). *The Green Paradox: A supply-side approach to global warming*. The MIT Press.

Smulders, Sjak, Yacov Tsur, and Amos Zemel (2012). “Announcing climate policy: Can a green paradox arise without scarcity?” In: *Journal of Environmental Economics and Management* 64.3, pp. 364–376.

Woike, Jan K., Ralph Hertwig, and Gerd Gigerenzer (2023). “Heterogeneity of rules in Bayesian reasoning: A toolbox analysis”. In: *Cognitive Psychology* 143, p. 101564.

A Appendix

A.1 Analytical characterization of market equilibrium

The producer of energy maximizes its intratemporal profit wrt K_t^c, K_t^d:

$$p_t A_y \left[\nu(A_t K_t^c)^\xi + (1-\nu)(K_t^d)^\xi\right]^{1/\xi} - p_t^c K_t^c - p_t^d K_t^d \to \text{ max.} \tag{A.1}$$

FOC:

$$p_t^c = p_t \nu (A_y A_t)^{\xi} \left(\frac{Y_t}{K_t^c} \right)^{1-\xi}, \tag{A.2}$$

$$p_t^d = p_t (1-\nu)(A_y)^{\xi} \left(\frac{Y_t}{K_t^d} \right)^{1-\xi}. \tag{A.3}$$

Each capiatal-producing firm ($i \in \{c, d\}$) maximizes its expected profit wrt I_t^i, K_t^i subject to the law of motion of corresponding capital stock, initial and transversility conditions (Kamihigashi 2005):

$$\mathbb{E}_0^i \sum_{t=0}^{\infty} \beta^t \left((p_t^i - \tau_t^i) K_t^i - (I_t^i)^{\alpha_i} \right) \rightarrow \ \max \tag{A.4}$$

$$\text{s.t. } K_t^i = (1-\delta_i) K_{t-1}^i + I_t^i, \tag{A.5}$$

$$\lim_{t \to \infty} \mathbb{E}_0^i \beta^t (p_t^i - \tau_t^i) K_t^i = 0, \quad K_{-1}^i > 0 \text{ is given.} \tag{A.6}$$

The corresponding Lagrangian is given by:

$$\mathcal{L} = \mathbb{E}_0^i \sum_{t=0}^{\infty} \beta^t \left((p_t^i - \tau_t^i) K_t^i - (I_t^i)^{\alpha_i} - \mu_t^i (K_t^i - (1-\delta_i) K_{t-1}^i - I_t^i) \right). \tag{A.7}$$

FOC:

$$-\alpha_i (I_t^i)^{\alpha_i - 1} + \mu_t^i = 0, \tag{A.8}$$

$$p_t^i - \tau_t^i - \mu_t^i + \beta(1-\delta_i) \mathbb{E}_t^i \mu_{t+1}^i = 0. \tag{A.9}$$

Clean and dirty capital markets clear iff:

$$p_t \nu (A_y A_t)^{\xi} \left(\frac{Y_t}{K_t^c} \right)^{1-\xi} - \tau_t^c = \mu_t^c - \beta(1-\delta_c) \mathbb{E}_t^c \mu_{t+1}^c, \tag{A.10}$$

$$p_t (1-\nu)(A_y)^{\xi} \left(\frac{Y_t}{K_t^d} \right)^{1-\xi} - \tau_t^d = \mu_t^d - \beta(1-\delta_d) \mathbb{E}_t^d \mu_{t+1}^d. \tag{A.11}$$

The law of motion for the stock of pollution and the energy demand schedule are given by:

$$D_t = (1-\delta)D_{t-1} + \gamma K_t^d, \tag{A.12}$$

$$p_t = Y_t^{-\eta}, \tag{A.13}$$

respectively.

A.2 Analytical characterization of the Social Planner's solution

The Social Planner maximizes the following Lagrangian wrt $Y_t, D_t, I_t^c, I_t^d, K_t^c, K_t^d$:

$$\begin{aligned}
\mathcal{L} = \mathbb{E}_0^p \sum_{t=0}^{\infty} \beta^t \Bigg[& \frac{Y_t^{1-\eta}}{1-\eta} - [(I_t^d)^{\alpha_d} + (I_t^c)^{\alpha_c}] - \kappa_d \frac{D_t^{1+\chi}}{1+\chi} \\
& - \Lambda_t \left(Y_t - A_y \left[\nu (A_t K_t^c)^\xi + (1-\nu)(K_t^d)^\xi \right]^{1/\xi} \right) \\
& - \Phi_t \left(D_t - (1-\delta) D_{t-1} - \gamma K_t^d \right) \\
& - \Phi_t^d \left(K_t^d - (1-\delta_d) K_{t-1}^d - (I_t^d)^{\alpha_d} \right) \\
& - \Phi_t^c \left(K_t^c - (1-\delta_c) K_{t-1}^c - (I_t^c)^{\alpha_c} \right) \Bigg] \to \ \max.
\end{aligned} \tag{A.14}$$

FOC:

$$Y_t^{-\eta} - \Lambda_t = 0, \tag{A.15}$$

$$-\kappa_d D_t^\chi - \Phi_t + \beta(1-\delta)\mathbb{E}_t^p \Phi_{t+1} = 0, \tag{A.16}$$

$$\Lambda_t \nu (A_y A_t)^\xi \left(\frac{Y_t}{K_t^c} \right)^{1-\xi} - \Phi_t^c + \beta(1-\delta_c)\mathbb{E}_t^p \Phi_{t+1}^c = 0, \tag{A.17}$$

$$\Lambda_t (1-\nu)(A_y)^\xi \left(\frac{Y_t}{K_t^d} \right)^{1-\xi} + \gamma \Phi_t - \Phi_t^d + \beta(1-\delta_d)\mathbb{E}_t^p \Phi_{t+1}^d = 0, \tag{A.18}$$

$$-\alpha_c (I_t^c)^{\alpha_c - 1} + \Phi_t^c = 0, \tag{A.19}$$

$$-\alpha_d (I_t^d)^{\alpha_d - 1} + \Phi_t^d = 0. \tag{A.20}$$

A.3 Proof of Proposition 1

From (A.20)

$$\Phi_t^d = \alpha_d (I_t^d)^{\alpha_d - 1}. \tag{A.21}$$

Plug this expression and (A.15) into (A.18) and get:

$$\begin{aligned} &Y_t^{-\eta}(1-\nu)(A_y)^{\xi}\left(\frac{Y_t}{K_t^d}\right)^{1-\xi} + \gamma\Phi_t - \alpha_d (I_t^d)^{\alpha_d-1} \\ &+ \beta(1-\delta_d)\mathbb{E}_t^p\left(\alpha_d (I_{t+1}^d)^{\alpha_d-1}\right) = 0. \end{aligned} \tag{A.22}$$

Then, subtract (A.22) from (A.11) and taking (A.8, A.13) into account get:

$$-\tau_t^d - \gamma\Phi_t + \beta(1-\delta_d)\alpha_d\left[\mathbb{E}_t^d\left((I_{t+1}^d)^{\alpha_d-1}\right) - \mathbb{E}_t^p\left((I_{t+1}^d)^{\alpha_d-1}\right)\right] = 0. \tag{A.23}$$

Finally, from (A.16)

$$\Phi_t = -\kappa_d D_t^{\chi} + \beta(1-\delta)\mathbb{E}_t^p \Phi_{t+1} \tag{A.24}$$

so, (A.23) can be rewritten as:

$$\begin{aligned} \tau_t^d = {} & \kappa_d\gamma\left[D_t^{\chi} + \beta(1-\delta)\mathbb{E}_t^p\left(D_{t+1}^{\chi} + \beta(1-\delta)D_{t+2}^{\chi} + ...\right)\right] \\ & + \beta(1-\delta_d)\alpha_d\left[\mathbb{E}_t^d\left((I_{t+1}^d)^{\alpha_d-1}\right) - \mathbb{E}_t^p\left((I_{t+1}^d)^{\alpha_d-1}\right)\right]. \end{aligned} \tag{A.25}$$

The benevolent government mimics the Social Planner's solution by introducing taxes so optimal $\hat{\tau}_t^d$ is given by the previous formula. The result for $\hat{\tau}_t^c$ can be derived similarly by manipulating with (A.10) and (A.17). Since no other market is distorted by an externality, the first-best policy is fully characterized by $\hat{\tau}_t^d, \hat{\tau}_t^c$.

A.4 Proof of Proposition 2

By definition, for a generic variable X_t and generic indexes $i, j \in \{p, c, d\}$:

$$\mathbb{E}_t^i X_{t+1} = z_t^i X_{t+1}|_{A_{t+1}>A_t} + (1 - z_t^i) X_{t+1}|_{A_{t+1}=A_t} \tag{A.26}$$

where $z_t^i = \Pr(A_{t+1} > A_t | q_t^i) = \theta_H q_t^i + \theta_L(1 - q_t^i)$. So,

$$\begin{aligned} \mathbb{E}_t^i X_{t+1} - \mathbb{E}_t^j X_{t+1} &= (z_t^i - z_t^j)(X_{t+1}|_{A_{t+1}>A_t} - X_{t+1}|_{A_{t+1}=A_t}) \\ &= (\theta_H - \theta_L)(q_t^i - q_t^j)(X_{t+1}|_{A_{t+1}>A_t} - X_{t+1}|_{A_{t+1}=A_t}). \end{aligned} \tag{A.27}$$

From (2.8):

$$q_t^i - q_t^j = \begin{cases} \frac{\theta_H \theta_L (q_{t-1}^i - q_{t-1}^j)}{[\theta_H q_{t-1}^i + \theta_L(1-q_{t-1}^i)][\theta_H q_{t-1}^j + \theta_L(1-q_{t-1}^j)]}, & A_t > A_{t-1}, \\ \frac{(1-\theta_H)(1-\theta_L)(q_{t-1}^i - q_{t-1}^j)}{[(1-\theta_H) q_{t-1}^i + (1-\theta_L)(1-q_{t-1}^i)][(1-\theta_H) q_{t-1}^j + (1-\theta_L)(1-q_{t-1}^j)]}, & A_t = A_{t-1}. \end{cases} \tag{A.28}$$

Thus, $q_{t-1}^i - q_{t-1}^j \gtrless 0 \Rightarrow q_t^i - q_t^j \gtrless 0$. Hence, by iterating the previous inequality, $q_0^i - q_0^j \gtrless 0 \Rightarrow q_t^i - q_t^j \gtrless 0$. Overall,

$$\operatorname{sign}(\mathbb{E}_t^i X_{t+1} - \mathbb{E}_t^j X_{t+1}) = \operatorname{sign}([q_0^i - q_0^j][X_{t+1}|_{A_{t+1}>A_t} - X_{t+1}|_{A_{t+1}=A_t}]). \tag{A.29}$$

Part 1 of the Proposition 2 follows directly from the above formula.

To prove the Part 2 of the Proposition note that the sequence

$$\{I_{t+1}^i|_{A_{t+1}>A_t} - I_{t+1}^i|_{A_{t+1}=A_t})\}_{t=0}^{\infty} \tag{A.30}$$

is bounded since I_t^i is bounded[28]. Then, it is sufficient to show that the sequence $\{q_t^i - q_t^j\}_{t=0}^{\infty}$ converges to 0 a.s. To do that rewrite (2.8) as:

$$q_t^i = \begin{cases} \frac{1}{1 + \frac{\theta_L}{\theta_H}\left(\frac{1}{q_{t-1}^i} - 1\right)}, & A_t > A_{t-1}, \\ \frac{1}{1 + \frac{1-\theta_L}{1-\theta_H}\left(\frac{1}{q_{t-1}^i} - 1\right)}, & A_t = A_{t-1}. \end{cases} \tag{A.31}$$

[28]To see that note, firstly, that by imposing the transversality condition (A.6) we implicitly assume the finitness of the value of firm (the discounted sum of profits) at the optimum (Kamihigashi 2005). Then, by analogy with (Demers 1991, Lemma 1), it can be shown that there is an upper bound $\overline{I^i}$ on I_t^i and an upper bound $\overline{K^i}$ on K_t^i. Furthermore, the production function (2.4) does not allow capital stocks to be negative so $I_t^i \geq \min\{-(1-\delta_i)\overline{K^i}, -(1-\delta_i)K_{-1}^i\}$.

Iterating this formula backwards one gets:

$$q_t^i = \frac{1}{1 + \left(\frac{\theta_L}{\theta_H}\right)^m \left(\frac{1-\theta_L}{1-\theta_H}\right)^{t-m} \left(\frac{1}{q_0^i} - 1\right)}. \tag{A.32}$$

If the true state of the world is $\theta = \theta_j, j \in \{H, L\}$, then by the strong Law of large numbers for arbitrarily large t: $m \to \theta_j t$ a.s., so

$$\left(\frac{\theta_L}{\theta_H}\right)^m \left(\frac{1-\theta_L}{1-\theta_H}\right)^{t-m} \to \left[\left(\frac{\theta_L}{\theta_H}\right)^{\theta_j} \left(\frac{1-\theta_L}{1-\theta_H}\right)^{1-\theta_j}\right]^t. \tag{A.33}$$

It is easy to show that the expression in the square brackets (call it $g(j, \theta_L, \theta_H)$) is greater than 1 for $j = L$ and smaller than 1 for $j = H$, so for any q_0^i in the former case $q_t^i \to 0$ and in the latter case $q_t^i \to 1$ as $t \to \infty$.

If $j = L$, then for $\theta_H \in (\theta_L, 1)$:

$$\frac{\partial \log g(L, \theta_L, \theta_H)}{\partial \theta_H} = \frac{\theta_H - \theta_L}{\theta_H(1-\theta_H)} > 0. \tag{A.34}$$

If $\theta_H = \theta_L$, then $g(L, \theta_L, \theta_H) = 1$, so for any $\theta_H > \theta_L$: $g(L, \theta_L, \theta_H) > 1$.

If $j = H$, then for $\theta_H \in (\theta_L, 1)$:

$$\frac{\partial \log g(H, \theta_L, \theta_H)}{\partial \theta_H} = \log\left(\frac{\theta_L(1-\theta_H)}{\theta_H(1-\theta_L)}\right) < 0. \tag{A.35}$$

If $\theta_H = \theta_L$, then $g(H, \theta_L, \theta_H) = 1$, so for any $\theta_H > \theta_L$: $g(H, \theta_L, \theta_H) < 1$.

Overall, for any $i, j \in \{p, c, d\}$, for any q_0^i, q_0^j:

$$q_t^i \underset{t\to\infty}{\longrightarrow} q_t^j \underset{t\to\infty}{\longrightarrow} \begin{cases} 1, & \theta = \theta_H \\ 0, & \theta = \theta_L \end{cases} \quad \text{a.s.} \tag{A.36}$$

The desired conclusion then follows from (A.27).

A.5 Proof of Proposition 3

Consider p_t^d as a function of A_t only: $p_t^d = p_t^d(A_t)$. Firslty, let us prove that this function is strictly convex. Let $B = (1-\nu)(A_y)^\xi (K^d)^{\xi-1} > 0$. Then, dropping the time index t henceforth, the formulas (A.1)-(A.3), (A.13) imply

$$\begin{aligned} p^d &= BY^{1-\xi-\eta}, && \text{(A.37)} \\ \frac{dp^d}{dA} &= B(1-\xi-\eta)Y^{-\xi-\eta}\frac{\partial Y}{\partial A} < 0, && \text{(A.38)} \\ \frac{d^2p^d}{dA^2} &= B(1-\xi-\eta)(-\xi-\eta)Y^{-1-\xi-\eta}\left(\frac{\partial Y}{\partial A}\right)^2 && \text{(A.39)} \\ &\quad + B(1-\xi-\eta)Y^{-\xi-\eta}\frac{\partial^2 Y}{\partial A^2} > 0. && \text{(A.40)} \end{aligned}$$

Note that $\partial Y/\partial A$ is K^c times the marginal product of AK^c (call it $MPAK^c$) which is positive and decreasing (due to the properties of CES production function) and so

$$\frac{\partial^2 Y}{\partial A^2} = K^c\frac{\partial MPAK^c}{\partial A} = K^c\frac{\partial MPAK^c}{\partial (AK^c)}\frac{\partial (AK^c)}{\partial A} = (K^c)^2\frac{\partial MPAK^c}{\partial (AK^c)} < 0. \tag{A.41}$$

Given $\xi + \eta > 1$, the signs of (A.38), (A.40) follow.

Now, the Law of total expectation, the Jensen's inequality and the fact that $\mathbb{E}[A|A^T] = \mathbb{E}[A^T|A^T]\mathbb{E}[A^C|A^T] = A^T$ together imply (note that A^T_{t+1} is a deterministic function of m_{t+1}, $A^T_{t+1} = A^T_{t+1}(m_{t+1})$):

$$\begin{aligned} \mathbb{E}^i_t p^d_{t+1}(A_{t+1}) &= \sum_{k=0}^{t+1} \mathbb{E}\left[p^d_{t+1}(A_{t+1})|A^T_{t+1}(m_{t+1}=k)\right]\Pr(m_{t+1}=k|q^i_t) \\ &> \sum_{k=0}^{t+1} p^d_{t+1}\left(\mathbb{E}\left[A_{t+1}|A^T_{t+1}(m_{t+1}=k)\right]\right)\Pr(m_{t+1}=k|q^i_t) \\ &= \sum_{k=0}^{t+1} p^d_{t+1}(A^T_{t+1}(m_{t+1}=k))\Pr(m_{t+1}=k|q^i_t) \\ &= \mathbb{E}^i_t p^d_{t+1}(A^T_{t+1}). \end{aligned} \tag{A.42}$$

It remains to be noted that $\mathbb{E}^i_t p^d_{t+1}(A^T_{t+1})$ is the expected price of dirty capital in the basic model.