The Grey Paradox: How fossil-fuel owners can benefit from carbon taxation.*

Renaud Coulomb † Fanny Henriet ‡

April 6, 2017

Abstract

This paper considers the distributional impact of optimal carbon taxation on fossil-fuel owners. A carbon-emitting exhaustible resource competes with a dirtier abundant resource (the dirty backstop) and a clean backstop. A time-dependent carbon tax is set to optimally use these resources under a cap constraint over CO_2 atmospheric concentration. As the cap is tightened, the dirtier resource becomes less competitive compared to the exhaustible resource (the "competition effect"), but the timing and duration of extraction of the exhaustible resource is modified (the "timing effect"). We provide analytical expressions of these effects, and determine conditions over size of reserves, pollution contents, extraction costs and demand elasticity such that the exhaustible-resource owners' profits increase as the ceiling is tightened. Calibrations for the transport and power sectors suggest that the profits of conventional-oil and natural-gas owners increase compared to a baseline without regulation for plausible carbon-ceiling values.

Keywords: Carbon Taxation, Externality, Global Warming, Non-renewable Resources, OPEC.

JEL Classification: H21, H23, Q31, Q41, Q54, Q58.

*We are grateful to Roger Guesnerie, Antony Millner, Michel Moreaux and seminar participants at Copenhagen University, ETH-Zurich, HKUST, HSE, LSE, Melbourne University, NES, Oxford University, PSE, Surrey University. We thank audiences at AFSE 2013, APET 2013, EEA 2013, IAEE 2014, SURED 2014, WCERE 2014 and at the 6th AWEEE in Spain. We acknowledge financial support from the ANR (ANR-16-CE03-0011).

†The University of Melbourne, Department of Economics, Level 4, Faculty of Business and Economics Building 105, 111 Barry Street Carlton VIC 3053. Corresponding author; email: renaud.coulomb@unimelb.edu.au.

‡Paris School of Economics-CNRS, PSE 48 bd Jourdan, 75014 Paris, France. Email: fanny.henriet@psemail.eu.

1 Introduction

Motivation and objective

Ongoing climate change is driven by the accumulation of greenhouse gases in the atmosphere, mostly CO_2 from fossil-fuel combustion (IEA 2014). This paper considers the impact of carbon taxation on countries endowed with fossil fuels, depending on the characteristics of the resources they own. Fossil fuels are spread unequally across the world and, in general, countries with large endowments of conventional oil and gas have small endowments of more polluting resources such as coal and unconventional oil, and vice versa. For example, OPEC-Gulf countries owned 49.6% and 39.8% of the world reserves of oil and natural gas respectively in 2011, but only 0.1% of world coal reserves, and their reserves of unconventional oil are negligible (Table 1).

Analysing the impact of carbon taxation on fossil-fuel owners' profits is important for at least three reasons. It first matters due to the size of the fossil-fuel production sector: fossil fuels accounted for 81% of the total energy supply in 2013 (IEA 2014). Second, some countries are very exposed to changes in the value of their fossil-fuels reserves. Oil rents represented 23.5% of Middle Eastern and North African oil countries' GDP in 2014; natural gas rents account for 8.7% of Qatar's GDP according to the World Bank. Finally, compensation and liability have been major issues in climate negotiations. It is thus crucial to identify the countries that could win or lose from international carbon regulation due to their resource endowment, and to understand countries' incentives when negotiating international agreements.

Sketch of the model and results

We build on Chakravorty et al. (2006, 2008) and construct a dynamic model of resource extraction with pollution control. Three energy resources are available: an exhaustible polluting resource (conventional oil or gas), a dirtier abundant resource (coal or unconventional oil) and a clean backstop. Environmental policy takes the form of a carbon cap over the atmospheric stock of CO_2, and optimal extraction is decentralised by a worldwide carbon tax.

We use a dynamic framework to reflect that the profitability of extracting the exhaustible resource depends on the price of other resources over periods of time where the exhaustible resource is not even used. Endogenous changes in the timing of use of the different resources due to carbon regulations are a key aspect of the problem.

Table 1: Shares of fossil-fuel reserves by world region and economic policy organisation in 2011, and fossil-fuel pollution content (direct combustion)

	Oil			**Natural gas**	**Coal**
	Total	*Conv.*	*Unconv.**		
Europe	1.0	*1.3*	*0.0*	2.2	8.6
CIS	8.1	*10.4*	*0.0*	31.9	20.7
Africa	8.3	*10.7*	*0.0*	7.5	3.5
Middle East	50.2	*64.7*	*0.0*	40.8	0.1
Australasia	2.6	*3.4*	*0.0*	8.6	40.3
North America	15.5	*3.8*	*56.1*	5.0	25.4
Latin America	14.3	*5.8*	*43.9*	3.9	1.4
World	100	*100*	*100*	100	100
OPEC-Gulf	49.6	*63.8*	*0.0*	39.8	0.1
OECD 2000	16.7	*5.3*	*56.1*	9.1	43.2
World Reserves in EJ	9032	*7014*	*2018+*	7415	21 952
Pollution content in gCO_2e/MJ		*92*	*106*	76	110

The data for reserves come from BGR (2012). "Reserves" are proven volumes economically exploitable at today's prices and using today's technology. "Conv." and "Unconv." stand for "Conventional" and "Unconventional". *Unconventional oil includes extra heavy oil, bitumen, oil shale, shale oil and light tight oil. "+" Reserves of unconventional oil are potentially very large as coal-based liquid supplies can be considered as a type of unconventional oil. Pollution contents are life-cycle greenhouse gas contents from Burnham et al. (2012). EJ = exajoules; MJ = megajoules; gCO_2e = grams of CO_2 equivalent.

Our main objective is to analyse the impact of optimal carbon taxation on the profits of the owners of a carbon-emitting exhaustible resource, and to determine under which conditions their profits may rise. We show that the profits of owners of some polluting exhaustible resources (like conventional oil or gas) will likely rise under optimal carbon taxation, even if the tax revenues are not redistributed to them. We label this phenomenon the "Grey Paradox" as it contradicts the claims from fossil-fuel owners—especially OPEC member countries[1]—that carbon taxation will reduce their profits. On the contrary, owners

[1]See for instance positions defended by the OPEC Secretary Generals, Rilwanu Lukman and Abdalla Salem El-Badri, at the 6^{th} and 19^{th} Conference of the Parties to the UN Framework Convention on Climate

of the dirtiest resources, such as coal or tar sands, will suffer under more stringent carbon regulation as it puts downward pressure on the value of their deposits. The value of the companies holding those assets may even fall to zero if their resource is not exhausted and is supplied on a competitive market.[2]

Tightening the carbon cap, or equivalently increasing the corresponding carbon tax in the decentralised economy, has two potential effects on the profits of the owners of an exhaustible polluting resource. When this resource competes with a more polluting resource at some future date, a higher tax tends to favour the owners of the less polluting resource at this date: the increase in the tax per unit of resource is higher for the more polluting resource, thus the after-tax price of the exhaustible resource must increase, i.e. current profits must increase. We call this the "competition effect". However, there is another effect which may work in the other direction and which reflects the dynamic nature of resource extraction. Carbon regulation modifies the timing of the extraction of the exhaustible resource. The extraction of the exhaustible resource may thus be slowed (reducing profits) or sped up (increasing profits) due to the regulation. We label this the "timing effect". Analytical expressions of these effects are derived.

We pay particular attention to the role played by resource endowments, pollution content, extraction costs and demand elasticity, and determine parameter conditions under which fossil-fuel profits rise as the ceiling falls. If the exhaustible resource is cheaper to extract than the dirtier resource, tightening carbon regulation increases the profits of exhaustible-resource owners if any of the following hold: (i) its demand elasticity is low enough; (ii) its extraction cost is close enough to that of the dirty backstop; (iii) its pollution content is low enough (compared to that of the dirty backstop); or (iv) its initial stock is small enough. As these conditions are independent, even if the polluting exhaustible resource has a high pollution content, its owners can benefit from carbon regulation. If the exhaustible resource

Change in 2000 and 2013.

[2] In April 2016, the largest private coal company, Peabody Energy, filed for U.S. bankruptcy protection. Its current situation is partly due to increased competition with renewable and natural gas in a world that has become more and more concerned with low-carbon energy transition.

is more expensive to extract than the dirty backstop, the only reason why we would want to use the exhaustible resource is its lower pollution content: more stringent regulation thus always increases its profitability as long as this resource becomes exhausted.

These analytical results hold with constant or increasing extraction costs for the exhaustible resource, with fixed reserves or endogenous reserves determined by early exploration, with or without technological progress in the clean-backstop technology, and with both a worldwide or a geographically-restricted tax.

Our second objective is to look for empirical evidence of the Grey Paradox. We calibrate our model for the transport sector, where conventional oil competes with unconventional oil. Introducing plausible carbon ceilings and limiting cumulative new emissions in this sector to a range of $322.7 - 637.5$ $GtCO_2$ increases the profits of conventional-oil owners. In the case of a tight carbon budget of 322.7 $GtCO_2$, profits are about 7% higher than in the reference scenario without carbon regulation. This figure can be as high as 30% in alternative scenarios. In the power sector, where natural gas competes with coal, introducing plausible carbon ceilings and limiting cumulative new emissions in this sector to a range of $390 - 1450$ $GtCO_2$ increases the profits of natural-gas owners. With a tight carbon budget of 585 $GtCO_2$, profits more than double compared to the no-regulation scenario. We carry out sensitivity analysis to ensure the validity of our findings, and take uncertainty concerning the value of some key parameters into account. All in all, our results suggest that OPEC-Gulf countries may become richer as a result of carbon taxation, contrary to countries with large endowments of more polluting resources.

Literature and contribution

The distributional impacts of carbon regulation on fossil-fuel owners have received some attention in the theoretical literature. This first focused on capturing rent from resource owners (Bergstrom 1982, Brander & Djajic 1983, Rubio & Escriche 2001, and Liski & Tahvonen 2004). In these papers, taxing resources—either directly or via the taxation of externalities—reduces the profits of resource owners. In our setting, the social planner has no bias towards a specific group of resource owners, and capturing resource rent is not a target *per se*. In addition, resource owners are in perfect competition, and are thus not involved in a strategic game with the regulator. We show that taxing a negative externality, here CO_2 emissions, may increase the profits of the owners of resources that generate it, without there being any possibility that the regulator capture this extra rent.

To analyse the Grey Paradox, we first characterise the optimal resource-extraction path. We determine how to optimally extract two carbon-emitting resources with different extraction costs and carbon contents under a carbon-cap policy. We consider the cases where the resource pollution contents are not ordered in the same way as their extraction costs, as well as the simpler case where they are. Both of these cases are applicable in the real world: in the transport sector, conventional oil is cheaper than its more-polluting competitor, unconventional oil; in the power-generation sector, natural gas is generally more expensive than its more polluting competitor, coal.

Our paper builds on the literature on exhaustible-resource extraction with pollution control.[3] Pollution was introduced in this literature by first assuming the existence of a unique polluting resource and an increasing damage function (Ulph & Ulph 1994, Tahvonen 1997) or a carbon ceiling on CO_2 (Chakravorty et al. 2006). More recent work has analysed resource extraction and/or carbon leakage with a number of polluting resources with distinct costs and/or pollution content (Chakravorty et al. 2008, Van der Ploeg & Withagen 2012, Michielsen 2014, Fischer & Salant 2014). Chakravorty et al. (2008) determine the optimal extraction of two exhaustible resources with different pollution contents and zero extraction costs under a carbon cap. They show that starting by extracting the most-polluting resource (to benefit from natural absorption) can be optimal, were the carbon ceiling not yet binding. Van der Ploeg & Withagen (2012) consider the optimal extraction of an exhaustible resource whose unit extraction costs increase with its cumulative extraction and which competes with a dirtier and abundant resource. The optimal policy is to extract more of the less-polluting

[3]For work on resource extraction without pollution control, see Hotelling (1931), Herfindahl (1967), Heal (1976), Kemp & Long (1980), Amigues et al. (1998), Gaudet et al. (2001), Gaudet & Lasserre (2011).

resource, and less of the dirtier one. Michielsen (2014) uses a two-period model with a polluting exhaustible resource and an abundant dirtier resource. He shows that anticipated future green policies can reduce short-term emissions if the substituability between resources rises over time. Fischer & Salant (2014) analyse the impact of different carbon policies on intertemporal[4] and spatial leakages in a calibration exercise with a number of resources. These papers do not focus on how resource owners' profits change with carbon policy. However Fischer & Salant (2014) note that if the pollution content of the exhaustible resource is close to zero, then its rent can rise with a carbon tax, thus preventing spatial leakage from this resource. Our analytical results are not restricted to the analysis of the rent of an almost-clean resource, and go beyond the effect of any difference in pollution content. We focus on the analytical derivation of the conditions under which the Grey Paradox occurs for very general functional forms. In particular, we show that this phenomenon can come about even when polluting resources have relatively similar polluting contents, if the substitutability between resources does not change through time, and even if the tax is optimally set. We highlight the roles of the elasticity of demand, extraction costs and the scarcity of the exhaustible resource, and carry out comparative-statics exercises on these parameters. Van der Ploeg (2016) focuses on leakages resulting from second-best policies in a general-equilibrium setting. He also checks in an extension that our main result–that oil owners may gain from carbon regulation at the expense of coal producers–holds in a general-equilibrium setting.

The empirical literature has made a number of attempts to evaluate the impact of long-term carbon regulation or the Kyoto Protocol on fossil-fuel prices and oil and gas revenues. No consensus has resulted. Most work based on simulations of energy-economic models has however concluded that OPEC will suffer losses from the implementation of a Kyoto-Protocol-like regulation (see Barnett et al. 2004 for a review of this literature), with the exceptions being Persson et al. (2007), Fischer & Salant (2014). None of these articles has provided analytical results allowing clear comparative statics. The mechanisms have thus

[4]See e.g. Sinn (2008) and Gerlagh (2011) on this question.

remained largely unexplained. Using a different approach, we present a simple model to disentangle the different mechanisms and establish analytical conditions under which some fossil-fuel owners benefit from carbon taxation.

Outline of the paper

The remainder of this paper is organised as follows. Section 2 describes the main model. Section 3 then characterises the optimal extraction path. Section 4 shows the effects of carbon taxation on the profits of exhaustible-resource owners, and Section 5 a calibration of our model for the transport and power-generation sectors. Last, Section 6 concludes.

2 The model

2.1 Assumptions and notation

We consider an economy with three energy resources, which are perfect substitutes in demand: a polluting exhaustible resource, of quantity X_e^0, an even more polluting abundant resource, and an abundant clean resource.[5] The labels e,d,b respectively stand for the "*e*xhaustible resource", "*d*irty abundant resource" (or dirty backstop) and "clean *b*ackstop". The extraction flow of resource i, $i \in \{e, d, b\}$, is $x_i(t)$. There is no stockpiling, so the extraction flow of a resource is also its current consumption. Preferences are quasilinear in money and the utility from energy consumption is $u(x_e(t) + x_d(t) + x_b(t))$, with u being twice continuously differentiable, strictly increasing and strictly concave. We also assume that $\lim_{g\to\infty} u'(g) = 0$ and $\lim_{g\to 0} u'(g) = +\infty$. We assume that the social discount rate, r, is constant, which is also the interest rate. We define θ_i as the pollution content of resource i: the use of one unit of resource i leads to θ_i units of CO_2. We assume that $0 < \theta_e < \theta_d$ and that $\theta_b = 0$. The extraction cost of resource i is c_i, and the clean backstop is the most expensive: $\max(c_e, c_d) < c_b$. In the baseline model, extraction costs are constant. In Sections 4.2 and 4.3, we allow the extraction cost of the exhaustible resource to rise with its

[5] Alternatively, we could think of this clean resource as a choke price, above which demand becomes zero.

cumulative extraction. In a separate extension, we let the price of the clean backstop fall over time.[6] The dirtiest resource is abundant enough not to be exhausted for the values of the carbon ceiling we consider. A large part of proven coal reserves must be left underground even with relatively high emission-stabilisation targets, as highlighted in Gerlagh (2011) and Van der Ploeg & Withagen (2012). The clean resource is available in infinite quantity. In the baseline, we also assume no exploration.[7] We relax this assumption in Sections 4.2 and 4.3. The initial amount of the exhaustible resource is X_e^0, and the change in its current stock is:

$$\dot{X}_e(t) = -x_e(t).$$

The combustion of the two polluting resources generates CO_2 emissions that accumulate in the atmosphere. We denote the atmospheric concentration of CO_2 by $Z(t)$. There is no natural decay of carbon, as in Van der Ploeg & Withagen (2012).[8] The change in the carbon stock over time is given by:

$$\dot{Z}(t) = \theta_e x_e(t) + \theta_d x_d(t).$$

The carbon stock, $Z(t)$, has to be kept under a carbon ceiling $\bar{Z}$, as in Chakravorty et al. (2006). A carbon-ceiling constraint is usually motivated in the literature as follows. It first can be considered as an exogenous constraint, for instance stemming from international negotiations. Second, recent climate-science literature has emphasised that crossing some CO_2 concentration thresholds could trigger abrupt climate change (e.g. Alley et al. 2003); attention has thus been put on reducing this catastrophic risk rather than dealing with smooth

[6] This assumption is similar to that used in Chakravorty et al. (1997) and Fischer & Salant (2014), amongst others.

[7] See e.g. Pindyck (1978) and Livernois & Uhler (1987) on this issue.

[8] No consensus exists over the form of natural dilution. It is, however, acknowledged that it is relatively small. Solomon et al. (2009) show that the increase in temperature due to high CO_2 concentration is largely irreversible for 1,000 years after the emissions stop. We thus assume that natural dilution is negligible.

manageable damage (e.g. Weitzman 2009). If damage is negligible below a certain carbon threshold but becomes extremely high or uncertain above it, the reduction of climate-change damage will boil down to keeping the carbon concentration below this critical level.[9] This second interpretation serves as a rationale for the stabilisation targets discussed in international negotiations. With negligible natural dilution the ceiling constraint is equivalent to a "carbon budget", a constraint on cumulative new emissions. We assume that $\bar{Z} > Z^0$. We restrict our analysis to the case where the carbon cap is tight enough to make the dirtiest resource not exhausted. This condition reads $\bar{Z} - Z^0 < \theta_d X_d^0$ where X_d^0 is the initial reserves of the dirtiest resource; for the sake of brevity we omit repeating it in the rest of the paper.

2.2 The welfare-maximisation program

The social planner wishes to identify the extraction $\{x_e(t), x_d(t), x_b(t)\}$ that maximises the net discounted social surplus under the environmental constraint. The social planner's problem is then:

$$Max \int_0^\infty e^{-rt} \Bigg(u(x_e(t) + x_d(t) + x_b(t)) - c_e x_e(t) - c_d x_d(t) - c_b x_b(t) \Bigg) dt$$

s.t.,

$$\begin{aligned}
&\dot{X}_e(t) = -x_e(t),\ X_e(t) \geq 0 \text{ and } X_e(0) = X_e^0 \\
&\dot{Z}(t) = \theta_e x_e(t) + \theta_d x_d(t),\ Z(t) \leq \bar{Z} \text{ and } Z(0) = Z^0 \\
&x_i(t) \geq 0,\ \text{for } i \in \{e, d, b\}
\end{aligned}$$

with $Z^0 < \bar{Z}$ and X_e^0 given.

Let $\lambda_e(t)$ be the shadow value of the remaining stock of the exhaustible resource $X_e(t)$

[9] Amigues et al. (2011) consider a carbon-cap constraint with constant marginal damage below the cap, and highlight that optimal-extraction patterns in this setting are similar to those in the baseline model with negligible damage below the ceiling.

and $-\mu(t)$ that of the pollution stock $Z(t)$.[10]

The transversality conditions are given by:

$$\lim_{t\to\infty} \lambda_e(t)e^{-rt}X_e(t) = 0 \tag{2.1}$$

$$\lim_{t\to\infty} \mu(t)e^{-rt}(\bar{Z} - Z(t)) = 0. \tag{2.2}$$

Equation 2.1 simply states that the exhaustible resource must be exhausted in the long run if the initial scarcity rent is strictly positive. Similarly, Equation 2.2 indicates that the carbon ceiling is binding in the long run if the initial shadow cost of pollution is strictly positive.

2.3 First-order conditions

We use control theory to solve this problem. Appendix A.1 provides more details on the solution, and discusses its existence and uniqueness.

The shadow value $\lambda_e(t)$ represents the current value of the scarcity rent of the exhaustible resource. As shown in Hotelling (1931), this increases at a rate of r: the discounted net marginal surplus of extraction must be constant. If the exhaustible resource is exhausted along the optimal extraction path, with $\lambda_e^0 \equiv \lambda_e(0)$, we have:

$$\lambda_e(t) \quad = \quad \lambda_e^0 e^{rt}.$$

The shadow value $\mu(t)$ represents the current value of the shadow cost of marginal pollution. This exhibits a familiar pattern reflecting the ceiling-type carbon regulation. If the ceiling does not bind (but will bind), the pollution cost rises at the rate of the discount rate. The intuition behind this result is similar to that of the Hotelling rule, as emitting CO_2 can be seen as extracting clean air from a reservoir with an initial stock of clean air defined by

[10] $\lambda_e(t)$ is the co-state variable associated with $X_e(t)$ and $\mu(t)$ is the co-state variable of $\bar{Z} - Z(t)$ used to solve the maximisation problem (see Appendix A.1).

$\bar{Z} - Z^0$. If the ceiling does not bind (but will bind), with $\mu^0 \equiv \mu(0)$, we have:

$$\mu(t) \quad = \quad \mu^0 e^{rt}.$$

Finally, along the optimal extraction path, the marginal utility from resource consumption must equal the true social cost of each resource:

$$u'(x_e(t) + x_d(t) + x_b(t)) = c_e + \lambda_e(t) + \theta_e\mu(t) - \gamma_e(t) \tag{2.3}$$

$$u'(x_e(t) + x_d(t) + x_b(t)) = c_d + \theta_d\mu(t) - \gamma_d(t) \tag{2.4}$$

$$u'(x_e(t) + x_d(t) + x_b(t)) = c_b - \gamma_b(t). \tag{2.5}$$

where $\gamma_i(t) \geq 0$ and $\gamma_i(t)x_i(t) = 0$, $i \in \{e, d, b\}$.

3 Optimal extraction paths

Since extraction costs partially determine the sequence and timing of resource use, we distinguish two cases: first, the exhaustible resource e is cheaper to extract than the dirty backstop d; second, the exhaustible resource e is more expensive to extract than d. Both of these cases are applicable in the real world: in the transport sector, conventional oil is cheaper than its more-polluting competitor, unconventional oil; in the power-generation sector, natural gas is generally more expensive than its more polluting competitor, coal. We restrict our analyses below to the cases where both resources are used and the exhaustible resource is exhausted.[11]

3.1 Exhaustible-resource extraction is cheaper

When the exhaustible resource is the cheapest to extract, it is never optimal to use resource d before exhausting resource e, as resource e is not only cheaper but also less polluting.

[11] If the exhaustible resource is not exhausted, the profits of the owners of this resource will be zero in the decentralised economy. Therefore, cases where this resource is not exhausted are of limited interest given our research question. We do however fully characterise the different cases in Appendices A.2 and A.3.

Both resources are used and resource e is exhausted along the optimal path if and only if the pollution from burning the entire stock of resource e is lower than the carbon budget: $X_e^0 < \frac{\bar{Z}-Z^0}{\theta_e}$. We assume hereafter that this condition holds.

Along the optimal path the exhaustible resource is used first: the corresponding social cost is $c_e + (\lambda_e^0 + \theta_e \mu^0)e^{rt}$. Once this is exhausted at date t^s, the dirty backstop is used to attain the ceiling, and the corresponding social cost is $c_d + \theta_d \mu^0 e^{rt}$. The ceiling is reached at date $\underline{t}$, after which the clean backstop is used. The extraction path is continuous. Finding the optimal extraction path requires that we determine the initial scarcity rent, λ_e^0, the initial shadow cost of pollution, μ^0, the date of the switch from the exhaustible resource e to the dirty backstop d, t^s, and the date the ceiling is reached, $\underline{t}$. The solution $(\lambda_e^0, \mu^0, t^s, \underline{t})$ must satisfy the following conditions:

$$c_e + \lambda_e^0 e^{rt^s} + \theta_e \mu^0 e^{rt^s} = c_d + \theta_d \mu^0 e^{rt^s} \tag{3.1}$$

$$c_d + \theta_d \mu^0 e^{r\underline{t}} = c_b \tag{3.2}$$

$$\int_0^{t^s} u'^{-1}(c_e + \lambda_e^0 e^{rt} + \theta_e \mu^0 e^{rt})dt = X_e^0 \tag{3.3}$$

$$\theta_e X_e^0 + \theta_d \int_{t^s}^{\underline{t}} u'^{-1}(c_d + \theta_d \mu^0 e^{rt})dt = \bar{Z} - Z^0. \tag{3.4}$$

Equations 3.1 and 3.2 reflect the continuity of the marginal surplus between phases. Equation 3.3 shows the exhaustion of resource e at date t^s, and Equation 3.4 that the environmental constraint binds at time $\underline{t}$.

3.2 Dirty-backstop extraction is cheaper

When the dirty backstop is cheaper to extract, the exhaustible resource will not be used without the carbon ceiling. The only reason to use the exhaustible resource is its lower carbon content in a carbon-regulated world. The carbon tax per unit of resource is higher for the dirty backstop. The exhaustible resource may be used if its after-tax price drops below that of the dirty backstop. Both polluting resources are used and the exhaustible resource is exhausted if it is neither too abundant nor too expensive. We show in Appendix A.3 that if $c_d < c_e$, there exists Z^* such that both polluting resources are used and the exhaustible resource is exhausted if and only if $c_b > \frac{\theta_d c_e - \theta_e c_d}{\theta_d - \theta_e}$ and $X_e^0 < \min(\frac{\bar{Z}-Z^0}{\theta_e}, \frac{Z^*-Z^0}{\theta_e})$. We assume hereafter that this condition holds.[12]

Along the optimal path, the dirty backstop is used until date t^s, and then the exhaustible resource that is exhausted once the ceiling is reached at date $\underline{t}$. The clean backstop is then used from this date. The extraction path is continuous. Finding the optimal extraction path requires that we determine the initial scarcity rent, λ_e^0, the initial shadow cost of pollution, μ^0, the date of the switch from the dirty backstop to the exhaustible resource, t^s, and the date the ceiling is reached, $\underline{t}$. The solution $(\lambda_e^0, \mu^0, t^s, \underline{t})$ satisfies:

$$\begin{aligned} c_e + \lambda_e^0 e^{rt^s} + \theta_e \mu^0 e^{rt^s} &= c_d + \theta_d \mu^0 e^{rt^s} && (3.5)\\ c_e + \lambda_e^0 e^{r\underline{t}} + \theta_e \mu^0 e^{r\underline{t}} &= c_b && (3.6)\\ \int_{t^s}^{\underline{t}} u'^{-1}(c_e + \lambda_e^0 e^{rt} + \theta_e \mu^0 e^{rt}) dt &= X_e^0 && (3.7)\\ \theta_d \int_0^{t^s} u'^{-1}(c_d + \theta_d \mu^0 e^{rt}) dt + \theta_e X_e^0 &= \bar{Z} - Z^0. && (3.8) \end{aligned}$$

4 Carbon taxes and profits in the decentralised economy

This section considers the decentralised economy. We analyse the impact of a tightening of the carbon ceiling, i.e. an increase in the optimal carbon tax, on the profits of the owners of the exhaustible carbon-emitting resource. We show under which conditions these profits can actually increase following tighter environmental constraints.

The planner can implement a worldwide tax on CO_2 emissions.[13] We assume that the social planner's tax scheme is credible for individuals with perfect foresight, and that fossil-

[12] The other extraction paths that may result are set out in Appendix A.3.

[13] Following the literature on spatial leakage (Hoel 2011, Fischer & Salant 2014), we can introduce some geographic restrictions on the implementation of the carbon tax. Our main results still hold under this assumption, as indicated in Remarks 1 and 2 and Web Appendix.

fuel owners and clean-energy providers are in perfect competition.[14] For a market price p_i of resource i, $i \in \{e, d, b\}$, the current marginal profits of an owner j of a deposit of resource i are $\pi_{i,j}(t) = p_i(t) - c_i - \theta_i \tau(t)$, where $\tau(t)$ is the current carbon tax. The profits are strictly positive only for owners of the exhaustible resource, since perfect competition wipes out the profits of the owners of resources that are not exhausted. The tax revenues are not redistributed to fossil-fuel owners but are rather lump-sum redistributed to consumers.[15]

4.1 Optimal carbon taxation

When only one polluting resource is used and exhausted, as in Chakravorty et al. (2006), there are an infinite number of ways of setting the carbon tax to bring about optimal extraction. For instance, the tax can be set equal to the optimal price net of extraction costs to fully capture profits.

This is no longer the case when two resources are used and only pollution can be taxed, i.e. when the tax base is CO_2 emissions. In our setting, the social planner has a limited capacity to capture rents from resource owners as only CO_2 emissions can be taxed. In other words, the unique carbon tax the social planner can implement to decentralise the optimum will leave some strictly positive profits for owners of the exhaustible resource. We show in Web Appendix[16] that the following lemma holds:

Lemma 1. *If it is optimal to use both resources and to exhaust resource e, the only carbon tax that can decentralise the optimum is that equal to the pollution cost, $\mu(t)$.*

[14] The empirical literature has produced contradictory evidence on the structure of oil markets. A number of pieces of empirical work (e.g. MacAvoy 1982, Verleger 1982) explain oil prices changes using a competitive model. While Alhajji & Huettner (2000) underline that 11 out of 13 statistical analyses of OPEC behaviour reject the cartel hypothesis, and find that OPEC behaviour is better explained by the assumption of a dominant firm (Saudi) with a competitive fringe during the 1970's, 80's and 90's, Huppmann & Holz (2012) suggest a recent shift in market structure and show that observed prices are closer to those in a competitive model since 2008.

[15] Were tax revenues to be entirely redistributed to resource owners, the exhaustible-resource owners would see their wealth increase were their resource to continue to be exhausted. This is due to the fact that their resource price must increase, so the sum of the profits and the tax must increase at each date.

[16] The Web Appendix is available on Renaud Coulomb's webpage.

We assume below that the tax is set to $\mu(t)$. The demand for energy in this decentralised world is denoted $D(p) \equiv u'^{-1}(p)$. The energy prices are:

$$p_e(t) = \quad c_e + \lambda_e(t) + \theta_e \mu(t) \tag{4.1}$$

$$p_d(t) = \quad c_d + \theta_d \mu(t) \tag{4.2}$$

$$p_b(t) = \quad c_b. \tag{4.3}$$

When the optimum is decentralised by a tax of $\mu^0 e^{rt}$, discounted total profits, Π_e, are:

$$\Pi_e = \int_0^\infty e^{-rt} D(p(t)) \lambda_e(t) dt = \lambda_e^0 X_e^0.$$

As discounted total profits are proportional to the initial scarcity rent, we analyse the effect of a lower ceiling on the initial scarcity rent.

Definition 1. *The Grey Paradox occurs if*

$$\frac{d\lambda_e^0}{d\bar{Z}} < 0.$$

4.2 Exhaustible-resource extraction is cheaper

In this part, we look at the impact of carbon regulation on the profits of exhaustible-resource owners when $c_e < c_d$ and both polluting resources are used to reach the ceiling. The optimum, given by Equations 3.1–3.4, is decentralised via a carbon tax of $\mu(t)$ at each date. In Figure 1 we plot resource prices before/after the lower ceiling. The black curves represent the prices for a given initial tax, i.e. a given ceiling value, and the grey curves those for a higher tax, i.e. for a lower ceiling. The symbol "*" indicates the new values after the lower ceiling. The price paths before the fall in the ceiling should be read using the (t, p)-diagram with 0 as the origin on the x-axis; those after the fall in the ceiling should be read using the (t^*, p^*)-diagram with 0^* as the origin on the x-axis. These two timelines allow us to keep t^s and t^{s*}, the dates of the switch from resource e to resource d, at the same point in the figure. The p^*-axis results from a shift of the p-axis to the left, as a lower ceiling postpones the date t^s.

The tightening of the ceiling has two effects on the profit of exhaustible-resource owners. First, there is a "competition effect" that comes from their resource competing with a dirtier resource at the switch date t^s: this can easily be seen in Equation 3.1.[17] A rise in the tax at the switch date $\mu^s \equiv \mu^0 e^{rt^s}$ entails an increase in the rent at the switch date $\lambda_e^s \equiv \lambda_e^0 e^{rt^s}$. Assume first that the switch date t^s is not affected by the lower ceiling. This is represented in Figure 1 by both the grey and black curves being read in the (t, p)-diagram with 0 as the origin on the x-axis. As the ceiling falls, μ^0 increases and so does the switch price $p^s \equiv c_d + \theta_d \mu^s$. Given the pollution differential, the rise in the switch price $\theta_d d\mu^s$ is larger than that of the tax on the exhaustible resource $\theta_e d\mu^s$. It follows that the current scarcity rent evaluated at the switch date rises as the carbon ceiling tightens, and $d\lambda_e^s = (\theta_d - \theta_e) d\mu^s$. If t^s were to be unchanged, the discounted profits after the fall in the ceiling would be given by the intersection between the y-axis and the $\lambda_e^*(t)$ curve (Point A), thus higher discounted profits would be represented by the segment AC. This is what would happen were energy demand to be totally inelastic, so that resource e would always be exhausted at the same date, whatever the price path.

If demand is instead elastic, the duration of the extraction of resource e is longer and t^s postponed as the increase in the switch price, p^s, entails higher prices over the whole price path. As shown in Appendix A.4, the change in the duration of the extraction of resource e is:

$$dt^s = \frac{D(p_0) - D(p^s)}{D(p_0)} \frac{dp^s}{r(p^s - c_e)}. \tag{4.4}$$

This longer duration depends on the size of the demand drop at dates 0 and t^s. After the lower ceiling, changes in λ_e^0, λ_e^s and t^s are linked by $d\lambda_e^0 = d\lambda_e^s e^{-rt^s} - r\lambda_e^0 dt^s$. The "timing effect", here a "slowdown" effect, is given by $-r\lambda_e^0 dt^s$. This effect is represented by segment AB in Figure 1, as the new price should be read using the (t^*, p^*)-diagram with 0^* as the

[17] At the switch date t^s, $\lambda_e^0 e^{rt^s} = c_d - c_e + (\theta_d - \theta_e)\mu^0 e^{rt^s}$.

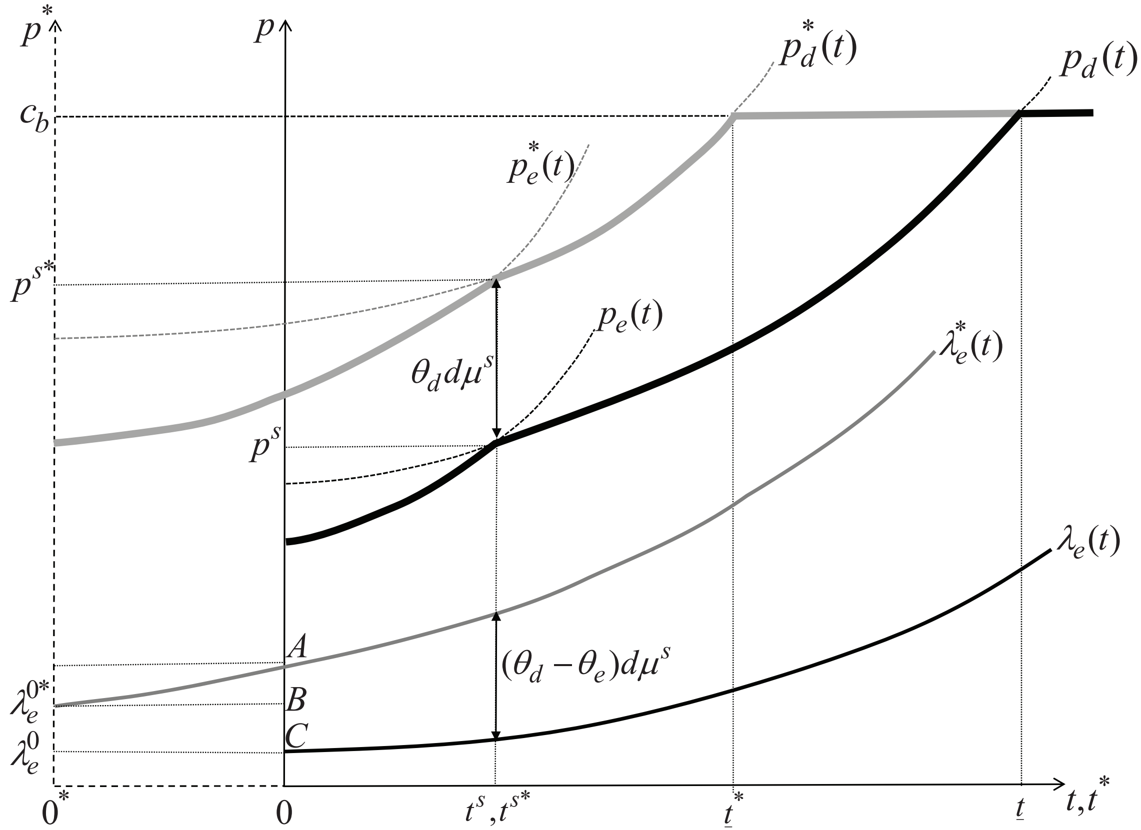

Figure 1: Price paths as the carbon ceiling is lowered, $c_e < c_d$

origin on the x-axis as the switch date changes. Were the value of the scarcity rent at the switch date to be unchanged, discounted profits would fall due to the longer extraction duration. This would be the case if the resource e was as polluting as the dirty backstop, $\theta_e = \theta_d$.

Figure 1 describes the case where the competition effect dominates the timing effect: the increase in discounted profits is represented by the segment BC. In general, the overall effect of a lower ceiling on discounted profits, i.e. on the initial scarcity rent, is ambiguous. Combining the two effects, we obtain the following proposition:

Proposition 1. $\forall \bar{Z} > \theta_e X_e^0 + Z^0$, *along the optimal path:*

$$d\lambda_e^0 = \underbrace{(1-\frac{\theta_e}{\theta_d})e^{-rt^s}dp^s}_{\textit{competition effect}} + \underbrace{(1-\frac{D(p^s)}{D(p_0)})\frac{-\lambda_e^0}{p^s - c_e}dp^s}_{\textit{timing effect}}$$

therefore:

$$\frac{d\lambda_e^0}{d\bar{Z}} < 0 \Leftrightarrow \frac{D(p(t^s))}{D(p(0))}\frac{\theta_d}{\theta_e} + (\theta_d - \theta_e)\frac{\mu^0}{\lambda_e^0} > 1.$$

Note that as $(\theta_d - \theta_e)\frac{\mu^0}{\lambda_e^0} > 0$, a sufficient condition for the Grey Paradox to hold is that $\frac{D(p(t^s))}{D(p(0))}\frac{\theta_d}{\theta_e} > 1$. When $D(p(t^s)) = D(p(0))$, it is always the case that $\frac{D(p(t^s))}{D(p(0))}\frac{\theta_d}{\theta_e} > 1$. This simple example underlines the role of the demand elasticity in the outcome. Without any further information on this elasticity, we can, however, show the following proposition:

Proposition 2. *If along the optimal extraction path the exhaustible resource is exhausted and both resources are used, tightening the carbon ceiling increases the profits of the exhaustible-resource owners if any of the following hold:*

(i) the demand elasticity is low enough;

(ii) the extraction cost of resource e is close enough to that of the dirty backstop;

(iii) the pollution content of resource e is low enough compared to that of the dirty backstop;

(iv) the exhaustible resource is scarce enough.

The proofs of Claims (i)-(iv) of Proposition 2 appear in Web Appendix. We provide some intuition for these results below. If the elasticity of demand is low, the effect of an increase in the switch price p^s is high, and this tends to exacerbate the competition effect. If the elasticity is low, the duration of extraction is not increased much after a fall in the ceiling, and so there is only little timing effect.

Consider the extreme case in which $c_e = c_d$. Under carbon regulation, the two resources are extracted simultaneously, so that $\lambda_e^0 = (\theta_d - \theta_e)\mu^0$. The initial scarcity rent is positive and increases with carbon taxation. This result continues to hold even with different extraction costs, if they are similar enough to each other.

If the pollution content of resource e is zero, the competition effect is $dp^s e^{-rt^s}$. This effect would clearly dominate the timing effect of $-(1 - \frac{D(p^s}{D(p_0)})e^{-rt^s}dp^s$.

Following a lower ceiling, the greater extraction duration resulting from higher prices falls as the stock of the exhaustible resource is reduced. The "timing effect" is only small in this case, as the change in the extraction duration of resource e, dt^s, is small.

Remark 1. Our main results are robust to the following changes in assumptions: the reserves of resource e are endogenous and depend on early exploration; unit extraction costs increase with cumulative extraction; the clean-backstop price falls over time to reflect technological progress; and part of demand is unregulated, that is the carbon tax is not implemented worldwide. Proofs and discussions are available in Web Appendix.

For general demand functions, it is not clear if the Grey Paradox is more likely to occur for low or high values of the carbon ceiling. We nevertheless do have the following result:

Proposition 3. *For concave or linear demand functions, the profits of owners of the exhaustible resource cannot exhibit a U-shape as the carbon ceiling is tightened, that is:*

$\forall Z^1, Z^2$ *such that* $Z^1 > \theta_e X_e^0 + Z^0$, $Z^2 > \theta_e X_e^0 + Z^0$ *and* $Z^1 > Z^2$, *if* $D'' \leq 0$,

$$\frac{d\lambda_e^0}{d\bar{Z}}_{|\bar{Z}=Z^1} > 0 \Longrightarrow \frac{d\lambda_e^0}{d\bar{Z}}_{|\bar{Z}=Z^2} > 0.$$

The formal proof is contained in Web Appendix.

4.3 Dirty-backstop extraction is cheaper

If the dirty backstop is cheaper to extract, without carbon regulation the exhaustible resource is never used. The only reason to use the exhaustible resource is because it is less polluting. Carbon regulation increases the price of the dirty backstop more than that of the exhaustible resource. The exhaustible resource may be used if its after-tax price drops below that of the dirty backstop. We assume that both polluting resources are used and the exhaustible resource is exhausted along the optimal extraction path (See Lemma 2 in Appendix A.3 for the conditions required for this case). The optimum, given in Equations 3.5–3.8, is decentralised via a carbon tax of $\mu(t)$ at each date. Since $d\lambda_e^0 = d\lambda_e^s e^{-rt^s} - r\lambda_e^0 dt^s$, we decompose the overall effect of the lower ceiling on the scarcity rent into two components as in the previous case.

The competition effect of $d\lambda_e^s e^{-rt^s}$ is actually zero when the exhaustible resource is more expensive to extract than the dirtier resource. The exhaustible-resource price is $p_e(t) = c_e + (p^s - c_e)e^{r(t-t^s)}$, with $p_e(\underline{t}) = c_b$, and its cumulative consumption is $\int_{t^s}^{\underline{t}} D(p_e(t))dt = X_e^0$. As long as this resource is exhausted, the price p^s does not depend on the value of the ceiling. It follows that the values of the tax and the scarcity rent at the switch date do not change as the ceiling is lowered, if the exhaustible resource becomes exhausted. The competition effect is thus zero.

The timing effect is $-r\lambda_e^0 dt^s$. If the switch date, t^s, is postponed, the cumulative consumption of the dirty backstop, $\int_0^{t^s} D(c_d + (p^s - c_d)e^{r(t-t^s)})dt$, would rise as p^s is unchanged. This would violate the ceiling constraint, as emissions from the exhaustible resource are unchanged. It follows that the switch date, t^s, is brought forward, i.e. extraction of resource e starts earlier. The timing effect thus increases discounted profits, contrary to the case where resource e is cheaper to extract.

Figure 2 describes how resource prices change after a fall in the ceiling. The black curves represent the prices for a given initial tax, i.e. a given ceiling value, and the grey curves those for a higher tax, i.e. for a lower ceiling value. The symbol "*" indicates the new values with a lower ceiling. The prices before the fall in the ceiling should be read using the (t, p)-diagram, with 0 as the initial value on the x-axis. The prices with the lower ceiling should be read using the (t^*, p^*)-diagram, using 0^* as the initial value on the x-axis. These two timelines allow us to keep t^s and t^{s*} at the same point in the figure. The p^*-axis is obtained by a shift in the p-axis to the right, as a lower ceiling brings the switch date t^s forward. The switch price does not depend on the value of the ceiling, so that the black and grey curves coincide from date 0^* onwards. The only effect of a lower ceiling is to reduce the extraction duration of the dirty backstop: in the figure, this extraction starts at date 0^* instead of 0. The prices of resources e and d cross at the unchanged value p^s, and the switch date is brought forward, it follows that the gap in initial prices must fall, so that the initial scarcity rent rises with the initial tax. The increase in profits is represented by segment AC in Figure 2.

The following proposition thus holds:

Proposition 4. *As long as both polluting resources are used and the exhaustible resource becomes exhausted along the optimal path, lowering the ceiling increases the profits of the exhaustible-resource owners, that is:*[18]

$$c_b > \frac{\theta_e c_d - \theta_d c_e}{\theta_e - \theta_d} \text{ and } X_e^0 < \min(X^*, \frac{\bar{Z} - Z^0}{\theta_e}) \Longrightarrow \frac{d\lambda_e^0}{d\bar{Z}} < 0.$$

[18] X^* is defined in Appendix A.3.

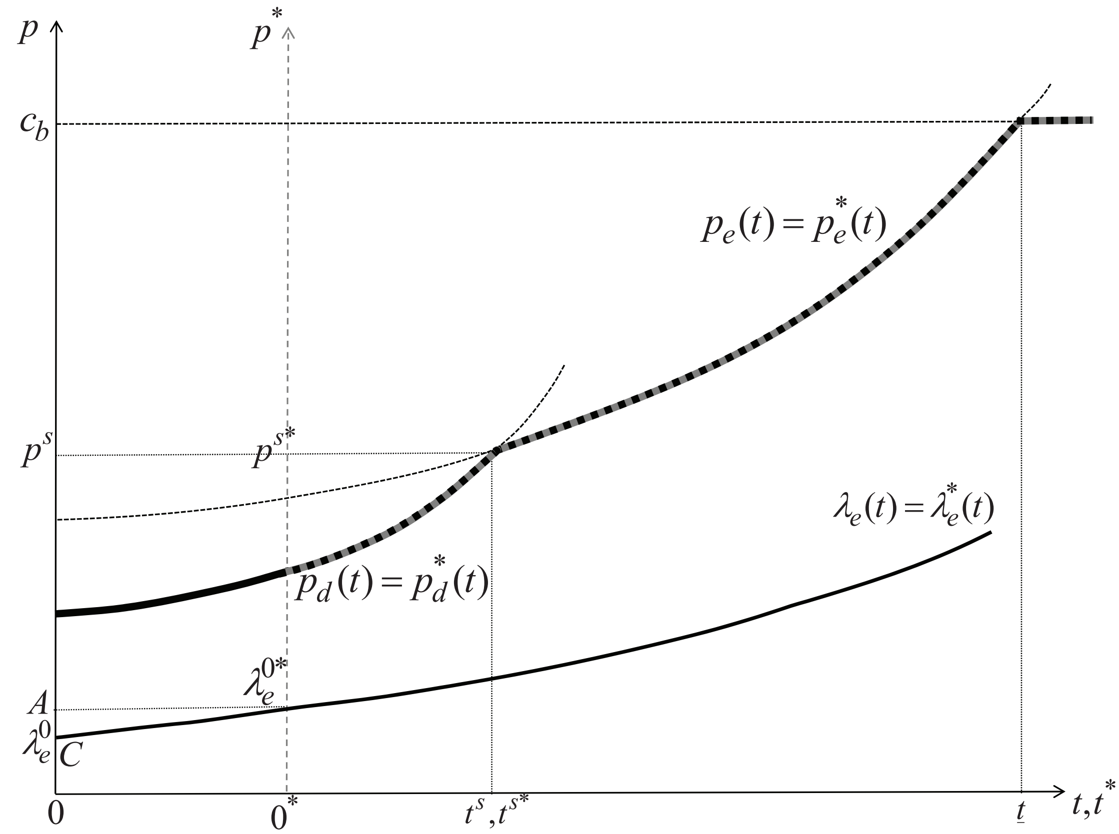

Figure 2: Price paths as the carbon ceiling is lowered, $c_e > c_d$

Remark 2. Our main results are robust to the following changes in assumptions: the reserves of resource e are endogenous and depend on early exploration; unit extraction costs increase with cumulative extraction; the clean-backstop price falls over time to reflect technological progress; part of demand is unregulated, so that the carbon tax is not implemented worldwide. The proofs and discussions appear in Web Appendix.

5 Calibration of the sectoral model

5.1 Transport sector

We calibrate our theoretical model for the transport sector, where oil represents 93% of primary energy demand (IEA 2013*b*). In this sector, conventional oil is in competition with unconventional oil, that is more polluting and more expensive than conventional oil. This situation corresponds to the first theoretical case in Section 4.2, where the exhaustible resource is cheaper to extract. In this case, a lower carbon ceiling has an ambiguous effect on the profits of exhaustible-resource owners. Table 2 summarises the parameters in the baseline calibration, our sensitivity analyses and the sources we have used for the parametrisation. Web Appendix discusses the choice of parameters.

Table 2: Parameters of the reference scenario and sensitivity analysis

Parameter	**Unit**	**Reference value**	**Sensitivity analysis**	**Sources**
X_e^0	EJ	7014		BGR (2012), IEA (2012)
c_e	$ / barrel	30	2 deposits (20;40)	IEA (2013*a*)
c_d	$ / barrel	90	60; 120	IEA (2013*a*)
c_b	$ / MegaWattHour	170	100; 300	EIA (2014)
θ_e	gCO_2 / MJ	92		Burnham et al. (2012)
θ_d	gCO_2 / MJ	106		Burnham et al. (2012)
r	%	3		
Price elasticity		-0.3	-0.1; -0.6	Kilian & Murphy (2014), Brons et al. (2008)

EJ = exajoules; MJ = megajoules; gCO_2e = grams of CO_2 equivalent. The starting year is 2011.

5.1.1 Calibration results

Figure 3 plots the profits of conventional-oil owners (as a %-change from no-regulation profits) as a function of the carbon budget in the reference scenario, and decomposes the overall effect into the competition and timing effects.[19] The sectoral carbon budget is $\bar{Z} - Z^0$, using the notation from the theoretical model.

We find empirical evidence supporting the Grey Paradox in this sector.[20] For a rise in CO_2 concentration of between 41 and 81 ppm, profits are higher than those without carbon regulation. This range corresponds to a global ceiling range between 550 and 700 ppm of CO_2 if the relative contribution of each sector to global emissions is stable through time (See Web Appendix for more details). As a comparison, under the OECD baseline scenario, the CO_2 concentration is projected to be around 780 ppm in 2100 (OECD 2012).

In this reference simulation, conventional oil is exhausted in 42 years when the sectoral carbon budget is set to 41 ppm, with an initial carbon tax of 38.9$ per barrel, and 33 years when set to 230 ppm with a carbon tax of 0.5$ per barrel.

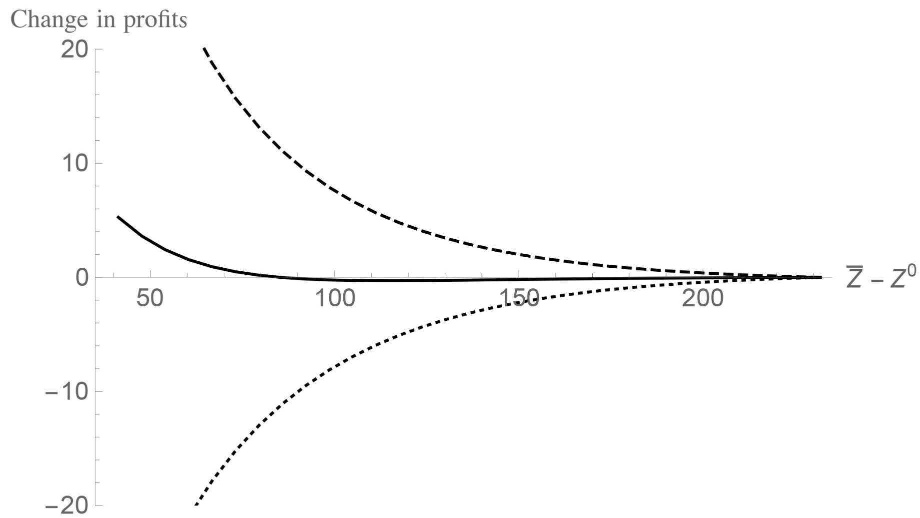

Figure 3: Change in profits (% of no-regulation profits) as a function of the carbon budget in the transport sector: baseline. Overall effect (solid line), competition effect (dashed line) and timing effect (dotted line)

[19]The decomposition is carried out using the fact that $\frac{\lambda_e^0(\bar{Z})}{\lambda_e^0(\infty)} = \frac{\lambda_e^s(\bar{Z})}{\lambda_e^s(\infty)} + [e^{r(t^s(\bar{Z}) - t^s(\infty))} - 1]\frac{\lambda_e^s(\bar{Z})}{\lambda_e^s(\infty)}$. The first term is the competition effect and the second term the timing effect.

[20] Fischer & Salant (2014) find that owners of a relatively expensive oil deposit can benefit from a carbon tax.

5.1.2 Sensitivity analysis

The upper-left panel of Figure 4 shows how the scarcity rent changes with the carbon ceiling for different values of the price elasticity of demand. For a price elasticity of -0.1 (dark grey), the Grey Paradox occurs for all values of the carbon ceiling, as long as conventional oil is exhausted i.e. for a carbon budget over 41 ppm. For a price elasticity of -0.6 (black), we see that any value of the carbon cap reduces the profits of conventional-oil owners compared to those without carbon regulation.

Technology improvements could reduce the extraction cost of unconventional oil. We run simulations with an extraction cost of unconventional oil of 60$ and 120$ per barrel. The upper-right panel of Figure 4 shows that introducing a carbon budget over 41 ppm increases profits if unconventional-oil costs 60$ per barrel (the dark grey line). For an unconventional-oil cost of 120$ per barrel (the black line), we find a small negative effect of carbon taxation on the scarcity rent.

Given the uncertainty over the cost of a clean-energy resource in the transport sector, we carry out sensitivity analysis regarding this parameter. The lower-left panel of Figure 4 plots the evolution of the scarcity rent as the carbon ceiling changes for different values of the clean-backstop price. The scarcity rent tends to the same value as the carbon ceiling tends to infinity for different clean-backstop costs. For a high ceiling, the clean backstop is used relatively late, and thus its value does not much affect the initial value of the scarcity rent. When the clean backstop costs 300$ per MegaWattHour (the black line), with a carbon budget under 100 ppm profits are larger with carbon regulation than without, with a maximum increase of about 20% for a carbon budget close to 41 ppm. Introducing a carbon budget over 100 ppm reduces profits, but this effect is only small (under 1%). When the clean backstop costs 100$ per MegaWattHour (the dark grey line), a carbon budget of over 80 ppm reduces profits compared to the case without regulation but this fall is only small (under 1%). For a carbon budget under 80 ppm, the rise in profits is small (slightly over 1%), and less than that for higher values of the clean-backstop cost.

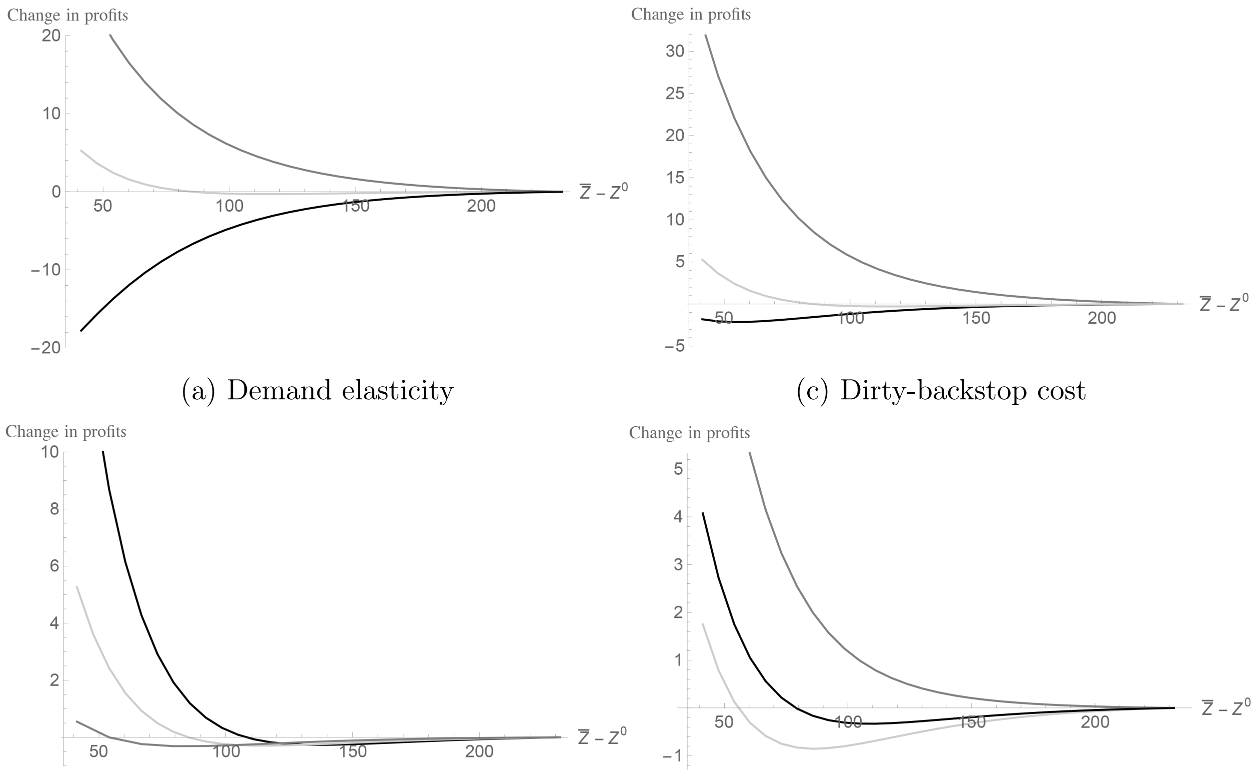

(a) Demand elasticity

(c) Dirty-backstop cost

(b) Clean-backstop cost

(d) Two deposits

Figure 4: Change in profits (% of no-regulation profits) as a function of the carbon budget: sensitivity analysis.
Upper-left panel: price elasticity of demand of -0.1 (dark grey), -0.3 (light grey) and -0.6 (black); Lower-left panel: clean-backstop cost of 300$ (black), 170$ (light grey) and 100$ (dark grey) per MegaWattHour; Upper-right panel: dirty-backstop cost of 60$ (dark grey), 90$ (light grey) and 120$ (black) per barrel; Lower-right panel: two-deposit case, profits of the 20$ per barrel deposit (light grey), profits of the 40$ per barrel deposit (dark grey), and sum of both profits (black).

Last, we decompose the exhaustible resource into two deposits with different extraction costs. The first includes around two-thirds of all conventional-oil deposits, and corresponds mainly to the Middle-Eastern and North African deposits. This deposit-extraction cost is set to 20$ per barrel (see IEA 2013*a*). The remaining third corresponds to other conventional-oil deposits, including deep-water fields. We set its extraction cost to 40$ per barrel. The lower-right panel of Figure 4 shows the rise in owners' profits for both deposits compared to no-regulation, as a function of the carbon budget. If the increase in CO_2 concentration is kept between 35 and 57 ppm, the owners of both deposits gain from environmental regulation. For a more lenient ceiling, the owners of the deposit with the higher extraction cost gain from environmental regulation, whereas the owners of the deposit with the lower extraction cost lose. The theory tells us that the owners of the most-expensive resource always gain more than the owners of the less-expensive resource for a marginal decrease in the ceiling (see Web Appendix). The simulation shows that that the profits of the owners of the more expensive deposit may rise while those regarding the less-expensive resource fall as the ceiling is lowered.

5.2 Power sector

We calibrate our theoretical model for the power sector, where gas, coal and renewables are in competition. In this sector the levelised cost of generating electricity (LCOE: this is the net present value of the cost of generating electricity over the lifetime of a plant) is higher with gas power plants than with coal power plants (see IEA/NEA 2015). These two types of classical thermal plants are good substitutes to renewable.[21]

However, the LCOE for each technology varies across regions and countries due to local regulations (such as those regarding air pollution) and other cost factors (see IEA 2013*a*). As a result, gas is already used in some regions of the world. Based on the current use of coal and gas in power generation, we assume that the LCOE of coal is above that of renewables (or equivalently that the use of coal is prohibited) for one third of total sectoral demand. This calibration exercise is thus at the frontier between the two theoretical cases in Sections 4.2 and 4.3. The calibration of the power sector produces more complicated effects than the two main theoretical cases. However, the intuition in the simple theoretical model helps us to understand the mechanisms at play.

Table 3 summarises the calibration parameters and their sources.[22] These are discussed in Web Appendix.

[21] All gas power plants and most coal power plants were originally designed or can be modified for flexible output: the ability to "ramp" on an hourly basis to much less than full output, and "cycle" on and off on a daily basis (Martinot 2016).

[22] The pollution contents of coal and gas in Table 3 are estimates for electricity generation. They differ from those in Table 1, which refer to direct combustion.

Table 3: Parameters of the reference scenario and sensitivity analysis

Parameter	**Unit value**	**Reference**	**Sensitivity analysis**	**Sources**
X_e^0	TWh	2.04 10^6	1.53 10^6; 2.6 10^6	BGR (2012)
c_e	\$ / MWh	80	60; 89	IEA/NEA (2015)
c_d	\$ / MWh	90		IEA/NEA (2015)
c_b	\$ / MWh	150	120; 300	IEA/NEA (2015)
θ_e	tCO_2eq / MWh	0.47		Heath et al. (2014)
θ_d	tCO_2eq / MWh	0.98		Heath et al. (2014)
r	%	3		
Price elasticity		-0.25	-0.1; -0.6	Alberini et al. (2011)

5.2.1 Calibration results

Figure 5 plots the profits of gas owners (as a %-change from no-regulation profits) as a function of the carbon budget in the reference scenario. The sectoral carbon budget is $\bar{Z} - Z^0$, using the notation of the theoretical model. We find empirical evidence of the Grey

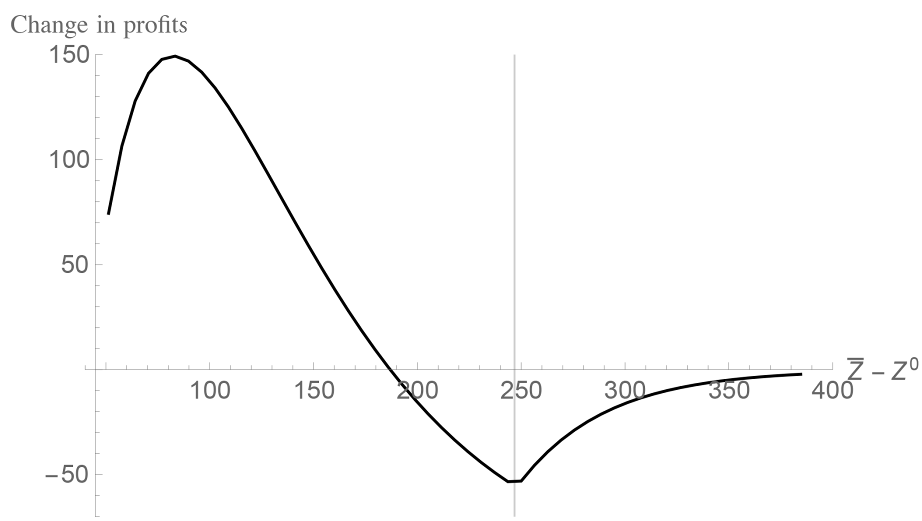

Figure 5: Change in profits (% of no-regulation profits) as a function of the carbon budget of the power sector.

Paradox in this sector as well. For a rise in CO_2 concentration of between 50 and 190 ppm, profits are higher than those without carbon regulation. This range corresponds to a global ceiling of between 500 and 780 ppm of CO_2 if the relative contribution of each sector to global emissions is stable over time (See Web Appendix for more details).

In this reference simulation, coal is totally replaced by gas, and gas is not exhausted when the sectoral carbon budget is set below 50 ppm.[23] In this case, the profits of gas owners are zero.

For a ceiling value of $50 - 246$ ppm, coal then gas is used in regions where coal is relatively cheaper, and only gas is used in the rest of the world. For these values of the carbon ceiling, the impact of a marginal fall in the ceiling on profits is not monotonic as the lower ceiling reduces cumulative gas consumption in areas where only gas is used. This makes gas relatively more abundant in the other region, tending to reduce the pure Grey Paradox effect described in Section 4.3. For a ceiling value of $50 - 85$ ppm, the fall in gas demand in areas where coal is never used is strong enough to make the impact of a tightening of the carbon constraint on gas profits negative.

For a more lenient ceiling (over 246 ppm), only gas is used in the areas where coal is "prohibited", whereas only coal is used in areas where the LCOE of gas is above that of coal. With this lenient ceiling, coal and gas are thus not in competition, and the result is straightforward: the more lenient the ceiling, the higher the profits of gas owners (right of the vertical line in Figure 5).

5.2.2 Sensitivity analysis

The results of the sensitivity analysis appear in Figure 6. As expected, a lower elasticity of demand and higher costs for the clean and dirty backstops are more favourable for gas owners. When gas reserves are higher than in the reference scenario (for instance reflecting larger unconventional-gas reserves), the Grey Paradox occurs for higher ceiling values than in the reference scenario.

[23]This corresponds to theoretical case 4 in Lemma 2 in Appendix A.3.

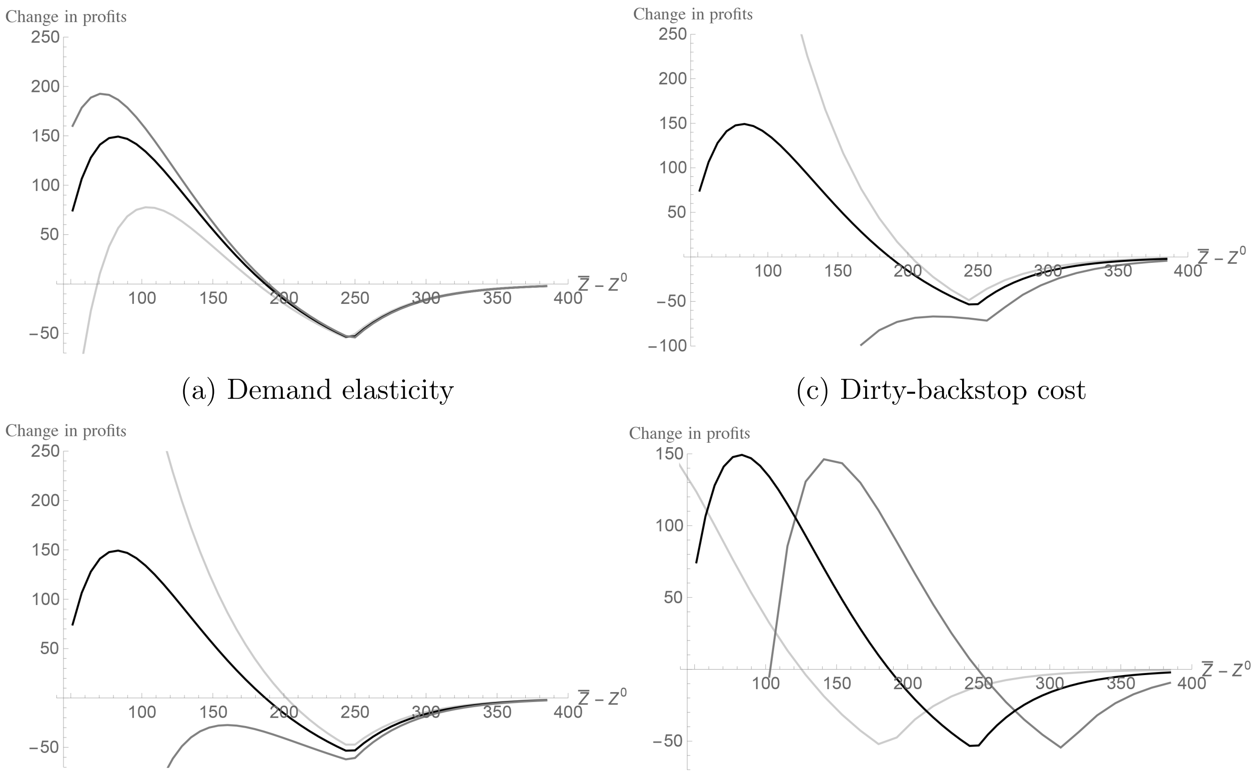

(a) Demand elasticity

(c) Dirty-backstop cost

(b) Clean-backstop cost

(d) Gas reserves

Figure 6: Change in profits (% of no-regulation profits) as a function of the carbon budget: sensitivity analysis.

Upper-left panel: price elasticity of demand of -0.1 (dark grey), -0.25 (black) and -0.6 (light grey); Lower-left panel: clean-backstop cost of 300$ (light grey), 150$ (black) and 120$ (dark grey) per MegaWattHour; Upper-right panel: dirty-backstop cost of 60$ (dark grey), 89$ (light grey) and 80$ (black) per MegaWattHour; Lower-right panel: Gas reserves of 2.04 10^6 TWh (black), 2.6 10^6 TWh (dark grey) and 1.53 10^6 TWh (light grey).

6 Conclusion

This paper has shed some light on the distributional effects of optimal carbon taxation, showing that the owners of a carbon-emitting resource may benefit from regulation if a dirtier abundant resource and a clean backstop are (will be) also used, even when the tax revenues are not redistributed to resource owners. In a dynamic model of resource extraction, we have separated the overall effect of a higher carbon tax into two effects: the "competition effect", representing the rise in profits at the date at which both resources are used simultaneously, and the "timing effect" that represents the shift over time of the extraction of the exhaustible resource after a tax rise.

We have calibrated our model for the transport and power-generation sectors. Limiting new emissions in these sectors to values that are consistent with plausible global ceiling targets increases the discounted profits of the owners of conventional oil and gas, compared to the no-regulation baseline.

Overall, our results lead us to reconsider the debate over the compensation for losses in oil-export revenues due to carbon taxation claimed by OPEC countries. Conventional oil and gas exporters, mostly OPEC-Gulf countries, may be more favourable to a worldwide carbon tax as they may directly benefit from it. On the contrary, coal or unconventional-oil exporters are likely to remain insensitive to pro-mitigation arguments as long as their losses are not at least partially compensated. Canada, for instance, withdrew from the Kyoto Protocol in 2011 after experiencing a period of growth, partly fuelled by oil sands.

A Appendix

A.1 Solving the maximisation problem

Following Forster (1980), who applied optimal control to the problem of extracting a single exhaustible resource that generates pollution, we use optimal control techniques to solve the social-planner's problem. The inequality constraints on the state variables are not trivial in optimal-control theory. We can use Theorem 1 p.317 in Seierstad & Sydsaeter (1987) here, which is for fixed-final-time problems with mixed and pure state constraints. It is straightforward to rewrite our problem as a two-step problem with a fixed-final-time problem as the first step. Take T as given:

$$Max \int_0^T e^{-rt}\Bigg(u(x_e(t)+x_d(t)) - c_e x_e(t) - c_d x_d(t)\Bigg)dt$$

s.t.

$$\begin{aligned}
&\dot{X}_e(t) = -x_e(t),\ X_e(t) \geq 0 \text{ and } X_e(0) = X_e^0\\
&\dot{Z}(t) = \theta_e x_e(t) + \theta_d x_d(t),\ Z(t) \leq \bar{Z} \text{ and } Z(0) = Z^0\\
&x_i(t) \geq 0,\ \text{for } i \in \{e, d\}
\end{aligned}$$

with $Z^0 < \bar{Z}$ and X_e^0 given.

Call $V(T)$ the value of the objective at the optimum; $V(.)$ is continuous and differentiable.

The solution to the overall problem can be found by maximising, with respect to T, $V(T) + e^{-rT}(u(u'^{-1}(c_b)) - c_b u'^{-1}(c_b))$. This expression has a unique maximum. This produces the continuity in the extraction paths at T. Theorem 1 p.317 in Seierstad & Sydsaeter (1987) is thus sufficient to ensure that the solution we derive is the optimum.

We define the current-value Hamiltonian as:

$$\begin{aligned} H(t) = u(x_e(t) + x_d(t) + x_b(t)) - c_e x_e(t) - c_d x_d(t) - c_b x_b(t) + \beta(t) X_e(t) + \nu(t)(\bar{Z} - Z(t)) \\ + \gamma_e(t) x_e(t) + \gamma_d(t) x_d(t) + \gamma_b(t) x_b(t) - \lambda_e(t) x_e(t) - \mu(t)(\theta_e x_e(t) + \theta_d x_d(t)). \end{aligned}$$

This has the following slackness conditions:

$$\nu(t) \geq 0 \quad \text{and} \quad \nu(t)(\bar{Z} - Z(t)) = 0 \tag{A.1}$$
$$\beta(t) \geq 0 \quad \text{and} \quad \beta(t) X_e(t) = 0 \tag{A.2}$$
$$\gamma_e(t) \geq 0 \quad \text{and} \quad \gamma_e(t) x_e(t) = 0 \tag{A.3}$$
$$\gamma_d(t) \geq 0 \quad \text{and} \quad \gamma_d(t) x_d(t) = 0 \tag{A.4}$$
$$\gamma_b(t) \geq 0 \quad \text{and} \quad \gamma_b(t) x_b(t) = 0. \tag{A.5}$$

For any control $\{x_e(t), x_d(t), x_b(t)\}$, there exist co-state variables, $\lambda_e(t)$ and $\mu(t)$, that must satisfy the following conditions, as well as the transversality and slackness conditions:

$$\dot{\lambda}_e(t) = r\lambda_e(t) - \frac{\partial H(t)}{\partial X_e(t)} \iff \dot{\lambda}_e(t) = r\lambda_e(t) - \beta(t) \tag{A.6}$$
$$\dot{\mu}(t) = r\mu(t) - \frac{\partial H(t)}{\partial Z(t)} \iff \dot{\mu}(t) = r\mu(t) + \nu(t) \tag{A.7}$$
$$\frac{\partial H(t)}{\partial x_e(t)} = 0 \iff u'(x_e(t) + x_d(t) + x_b(t)) = c_e + \lambda_e(t) + \theta_e \mu(t) - \gamma_e(t) \tag{A.8}$$
$$\frac{\partial H(t)}{\partial x_d(t)} = 0 \iff u'(x_e(t) + x_d(t) + x_b(t)) = c_d + \theta_d \mu(t) - \gamma_d(t) \tag{A.9}$$
$$\frac{\partial H(t)}{\partial x_b(t)} = 0 \iff u'(x_e(t) + x_d(t) + x_b(t)) = c_b - \gamma_b(t). \tag{A.10}$$

If $c_e < c_d$ and $X_e^0 < \frac{\bar{Z}-Z^0}{\theta_e}$, we can check that the unique solution to the set of Equations 3.1-3.4 satisfies the necessary conditions A.1-A.10 and the transversality conditions. If $c_e > c_d$, and $c_b > \frac{\theta_d c_e - \theta_e c_d}{\theta_d - \theta_e}$ and $X_e^0 \leq min(\frac{\bar{Z}-Z^0}{\theta_e}, \frac{Z^*-Z^0}{\theta_e})$, we can check that the unique solution of the set of Equations 3.5-3.8 satisfies the necessary conditions A.1-A.10 and the transversality conditions.

The social cost of resource i ($i \in \{e, d, b\}$) is of the form: $c_i + d_i e^{rt}$, $d_i \in \mathbb{R}^+$, so that these social costs can at most cross once, given the different c_i's. The different resources must be used sequentially along the optimal path. If $c_i > c_j$ ($i, j \in \{e, d, b\}$, $j \neq i$) and if there exists t^* such that $c_i + d_i e^{rt^*} < c_j + d_j e^{rt^*}$, then $d_i e^{rt^*} < d_j e^{rt^*}$ and $c_i + d_i e^{rt} < c_j + d_j e^{rt}$ for all $t > t^*$. In other words, a resource cannot be used before another resource if the former resource is more expensive to extract than the latter. This result conforms to the Herfindahl rule (Herfindahl 1967).

A.2 Ordering resource extraction when $c_e < c_d$

If $c_e < c_d$, then resource e is necessarily used first. Assume, by way of contradiction, that resource d is used first. Recall that the scarcity rent and the pollution cost increase at the rate of r. If $c_e < c_d$ and $c_e + \lambda_e^0 + \theta_e\mu^0 > c_d + \theta_d\mu^0$, which is the case if resource d is used first, then it is easy to see that, for all t, $c_e + \lambda_e(t) + \theta_e\mu(t) > c_d + \theta_d\mu(t)$: only resource d is used until the ceiling is reached. But then, lowering λ_e^0 and increasing μ^0, so that resource e is used first, would be preferred by the social planner, as resource e is cheaper to extract and less polluting. As resource e is strictly preferred to resource d, resource d cannot be used before exhausting resource e. A necessary and sufficient condition to exhaust resource e and to use resource d is that the initial carbon stock augmented by the pollution coming from the burning of the entire stock of resource e ($\theta_e X_e^0$) is below the carbon ceiling $\bar{Z}$, i.e. $\theta_e X_e^0 < \bar{Z} - Z^0$. See the main text for the description of the solution when resource e and resource d are used, and resource e is exhausted. When it is optimal to use only resource e, the solution is as described by the solution to the one-resource case.

A.3 Ordering resource extraction when $c_d < c_e$

As shown above, resource d cannot be used after resource e. Four cases pertain. First, both resources d and e are used to reach the ceiling: resource d is used first, and resource e becomes exhausted (Case 1). Second, both resources d and e are used to reach the ceiling: resource d first and resource e is not exhausted (Case 2). Third, only resource d is used to reach the ceiling (Case 3). Last, only resource e is used to reach the ceiling (Case 4). The relevant case for analysis is Case 1, where both resources are used and resource e is exhausted. We show the parameter conditions for this case to apply in the following lemma.

Lemma 2. *If resource d is cheaper to extract than resource e, $c_d < c_e$, the following cases arise:*

- *If $c_b \leq \frac{\theta_d c_e - \theta_e c_d}{\theta_d - \theta_e}$, only the dirty backstop is used to reach the ceiling (Case 3).*
- *If $c_b > \frac{\theta_d c_e - \theta_e c_d}{\theta_d - \theta_e}$, resource e is used when the ceiling is about to bind and there exists Z^* such that:*
  - *If $\bar{Z} \leq Z^*$ and $X_e^0 \geq \frac{\bar{Z} - Z^0}{\theta_e}$, only resource e is used to reach the ceiling (Case 4);*
  - *If $\bar{Z} > Z^*$ and $X_e^0 > \frac{Z^* - Z^0}{\theta_e}$, the dirty backstop is used, and resource e is then used to reach the ceiling but not exhausted (Case 2);*
  - *Otherwise, if $X_e^0 < min(\frac{\bar{Z} - Z^0}{\theta_e}, \frac{Z^* - Z^0}{\theta_e})$ the dirty backstop is used at the beginning, and then resource e is used to reach the ceiling and is exhausted (Case 1).*

Proof. We first consider the conditions under which only resource d is used to reach the ceiling (Case 3). First note that if resource d is being used when the ceiling is reached, then resource e is never used, as if $c_e + (\theta_e\mu^0 + \lambda_e^0)e^{r\underline{t}} > c_d + \theta_d\mu^0 e^{r\underline{t}}$ then we have that, for all $t \leq \underline{t}$, $c_e + (\theta_e\mu^0 + \lambda_e^0)e^{rt} > c_d + \theta_d\mu^0 e^{rt}$ (given that $c_d < c_e$). Resource d is being used when the ceiling is reached if and only if at date $\underline{t}$, defined by $c_d + \theta_d\mu^0 e^{r\underline{t}} = c_b$, the price of resource e satisfies $p_e(\underline{t}) = c_e + \theta_e\mu^0 e^{r\underline{t}} > c_b$. We can equivalently write that only resource d is used to reach the ceiling if and only if $c_b \leq \frac{\theta_d c_e - \theta_e c_d}{\theta_d - \theta_e}$. Assume now that $c_b > \frac{\theta_d c_e - \theta_e c_d}{\theta_d - \theta_e}$, then resource e is used at the binding-ceiling date. The carbon tax τ_t while resource e is being used must satisfy $c_e + \theta_e \tau_t \leq c_d + \theta_d \tau_t$, which can be rewritten as $\tau_t \geq \frac{c_e - c_d}{\theta_d - \theta_e}$. The lowest possible price path of resource e is thus $p(t) = c_e + \theta_e \frac{c_e - c_d}{\theta_d - \theta_e} e^{rt}$. Along this path, the price reaches the backstop price c_b at date T^* defined by $c_e + \theta_e \frac{c_e - c_d}{\theta_d - \theta_e} e^{rT^*} = c_b$. The maximum amount of resource e that can be consumed, if $c_b > \frac{\theta_d c_e - \theta_e c_d}{\theta_d - \theta_e}$, is then: $X^* = \int_0^{T^*} D(c_e + \theta_e \frac{c_e - c_d}{\theta_d - \theta_e} e^{rt}) dt$. Note that X^* does not depend on $\bar{Z}$. If $X_e^0 > X^*$, resource e is not exhausted. If $X_e^0 > X^*$ and $\bar{Z} > Z^0 + \theta_e X^* \equiv Z^*$, then resource d is used first, an amount X^* of resource e is then used to reach the ceiling, and this resource is not exhausted (Case 2). If $X_e^0 < X^*$ and $\bar{Z} < Z^*$, then only resource e is used to reach the ceiling and resource e is not exhausted (Case 4). If $X_e^0 < X^*$ and $\bar{Z} > Z^0 + \theta_e X_e^0$, then resource d is used first, and resource e is then used to reach the ceiling and is exhausted (Case 1). □

For Case 2, the equation set is similar to Equations 3.5–3.8, except that Equation 3.7 is dropped and the scarcity rent is set to 0. For Cases 3 and 4, the solution is as described by the solution to the one-resource case.

A.4 Change in the duration of the extraction of resource e

Proof. We can apply backward reasoning. Define $\tilde{t} = t - t^a$, where t^a is the date at which the extraction of the exhaustible resource ends. With this convention, extraction ends at date 0 and starts at date $-t^a$ and the price path is, for $\tilde{t} \in [-t^a, 0]$, $p_{\tilde{t}} = c_e + (p^s - c_e) e^{r\tilde{t}}$. This price is represented by the grey solid line (line L_A) in Figure 7. With an increase in p^s, the new price is $p^*_{\tilde{t}} = c_e + (p^s - c_e)(1 + \frac{dp^s}{p^s - c_e}) e^{r\tilde{t}}$, which can be rewritten, defining time u by $u = \frac{1}{r} \frac{dp^s}{p^s - c_e}$, as: $p^*_{\tilde{t}} = c_e + (p^s - c_e) e^{r(\tilde{t} + u)}$. The new price path $p^*_{\tilde{t}}$ is thus a left shift of the former price path $p_{\tilde{t}}$, with some starting date $-t^{a*}$ before $-t^a$. This new price path is represented by the black solid line (line L_B) in Figure 7.

As the new price is a translation of the former price, the cumulative consumption under price path $p_{\tilde{t}}$ between dates $-t^a + u$ and 0 (X_A) is the same as the cumulative consumption under price path $p^*_{\tilde{t}}$ between dates $-t^a$ and $-u$ ($X_B = X_A$) The quantity consumed over price path $p^*_{\tilde{t}}$ between dates $-u$ and 0 is $\delta_1 X_B = D(p^s) u$. The quantity consumed over price path $p_{\tilde{t}}$ between dates $-t^a$ and $-t^a + u$ is $\delta X_A = D(p_{-t^a}) u$. As the total quantity consumed over both price paths has to be the same, then it must be the case that the quantity consumed between date $-t^{a*}$ and $-t^a$ is $\delta X_A - \delta_1 X_B = (D(p_{-t^a}) - D(p^s)) u$, so that $t^{a*} - t^a = \frac{D(p_{-t^a}) - D(p^s)}{D(p_{-t^a})} u$. Using the normal timeline and the former notation, we obtain Equation 4.4.

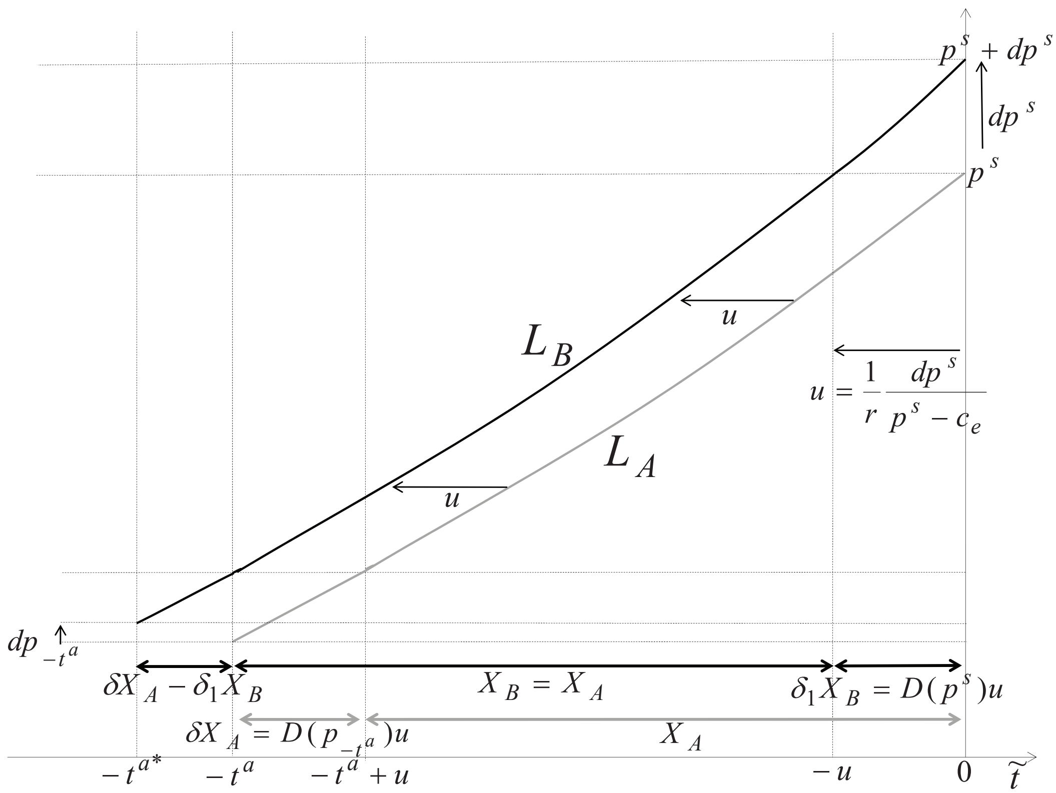

Figure 7: The change in resource e price path as the ceiling is lowered, $c_e < c_d$

□

References

Alberini, A., Gans, W. & Velez-Lopez, D. (2011), 'Residential consumption of gas and electricity in the U.S.: The role of prices and income', *Energy Economics* **33**(5), 870–881.

Alhajji, A. & Huettner, D. (2000), 'Opec and other commodity cartels: a comparison', *Energy Policy* **28**(15), 1151–1164.

Alley, R. B., Marotzke, J., Nordhaus, W. D., Overpeck, J. T., Peteet, D. M., Pielke, R. A., Pierrehumbert, R. T., Rhines, P. B., Stocker, T. F., Talley, L. D. & Wallace, J. M. (2003), 'Abrupt climate change', *Science* **299**(5615), 2005–2010.

Amigues, J.-P., Favard, P., Gaudet, G. & Moreaux, M. (1998), 'On the optimal order of natural resource use when the capacity of the inexhaustible substitute is limited', *Journal of Economic Theory* **80**(1), 153–170.

Amigues, J.-P., Moreaux, M. & Schubert, K. (2011), 'Optimal use of a polluting non renewable ressource generating both manageable and catastrophic damages', *Annales d'Economie et de Statistique* (103-104), 107–142.

Barnett, J., Dessai, S. & Webber, M. (2004), 'Will OPEC lose from the Kyoto protocol?', *Energy Policy* **32**(18), 2077–2088.

Bergstrom, T. C. (1982), 'On capturing oil rents with a national excise tax', *American Economic Review* **72**(1), 194–201.

BGR (2012), Energy resources 2012, Technical report, BGR, German Federal Institute for Geosciences and Natural Resources, Hannover, Germany.

Brander, J. & Djajic, S. (1983), 'Rent-extracting tariffs and the management of exhaustible resources', *Canadian Journal of Economics* **16**(2), 288–98.

Brons, M., Nijkamp, P., Pels, E. & Rietveld, P. (2008), 'A meta-analysis of the price elasticity of gasoline demand. a SUR approach', *Energy Economics* **30**(5), 2105–2122.

Burnham, A., Han, J., Clark, C. E., Wang, M., Dunn, J. B. & Palou-Rivera, I. (2012), 'Life-cycle greenhouse gas emissions of shale gas, natural gas, coal, and petroleum', *Environmental Science & Technology* **46**(2), 619–627.

Chakravorty, U., Magné, B. & Moreaux, M. (2006), 'A hotelling model with a ceiling on the stock of pollution', *Journal of Economic Dynamics and Control* **30**(12), 2875–2904.

Chakravorty, U., Moreaux, M. & Tidball, M. (2008), 'Ordering the extraction of polluting nonrenewable resources', *American Economic Review* **98**(3), 1128–1144.

Chakravorty, U., Roumasset, J. & Tse, K. (1997), 'Endogenous substitution among energy resources and global warming', *Journal of Political Economy* **105**(6), 1201–34.

EIA (2014), *Annual Energy Outlook*, US Energy Information Administration, Paris.

Fischer, C. & Salant, S. W. (2014), 'Limits to limiting greenhouse gases: Intertemporal leakage, spatial leakage, and negative leakage'.

Forster, B. A. (1980), 'Optimal energy use in a polluted environment', *Journal of Environmental Economics and Management* **7**(4), 321–333.

Gaudet, G. & Lasserre, P. (2011), 'The Efficient Use Of Multiple Sources Of A Nonrenewable Resource Under Supply Cost Uncertainty', *International Economic Review* **52**(1), 245–258.

Gaudet, G., Moreaux, M. & Salant, S. W. (2001), 'Intertemporal Depletion of Resource Sites by Spatially Distributed Users', *American Economic Review* **91**(4), 1149–1159.

Gerlagh, R. (2011), 'Too much oil', *CESifo Economic Studies* **57**(1), 79–102.

Heal, G. (1976), 'The relationship between price and extraction cost for a resource with a backstop technology', *Bell Journal of Economics* **7**(2), 371–378.

Heath, G., O'Donoughue, P., Arent, D. & Bazilian, M. (2014), 'Harmonization of initial estimates of shale gas life cycle greenhouse gas emissions for electricity generation', *Proceedings of the National Academy of Sciences of the United States of America* **111**(31), 3167–3176.

Herfindahl, O. C. (1967), *"Depletion and Economic Theory" In Extractive Resources and Taxation ed. Mason Gaffney, 63-90*, University of Wisconsin Press.

Hoel, M. (2011), 'The supply side of CO_2 with country heterogeneity', *Scandinavian Journal of Economics* **113**(4), 846–865.

Hotelling, H. (1931), 'The economics of exhaustible resources', *Journal of Political Economy* **39**(2), 137–175.

Huppmann, D. & Holz, F. (2012), 'Crude Oil Market Power – A Shift in Recent Years?', *The Energy Journal* **33**(4), 1–22.

IEA (2012), *World Energy Outlook 2012*, IEA, Paris.

IEA (2013*a*), *Resources To Reserves 2013: Oil, Gas and Coal Technologies for the Energy Markets of the Future*, IEA, Paris.

IEA (2013*b*), *World Energy Outlook 2013*, IEA, Paris.

IEA (2014), *The IEA Oil Market Report*, IEA, Paris.

IEA/NEA (2015), *Projected Costs of Generating Electricity 2015*, OECD/IEA/NEA, Paris.

Kemp, M. C. & Long, N. V. (1980), 'On two folk theorems concerning the extraction of exhaustible resources', *Econometrica* **48**(3), 663–673.

Kilian, L. & Murphy, D. P. (2014), 'The Role Of Inventories And Speculative Trading In The Global Market For Crude Oil', *Journal of Applied Econometrics* **29**(3), 454–478.

Liski, M. & Tahvonen, O. (2004), 'Can carbon tax eat OPEC's rents?', *Journal of Environmental Economics and Management* **47**(1), 1–12.

Livernois, J. R. & Uhler, R. S. (1987), 'Extraction costs and the economics of nonrenewable resources', *Journal of Political Economy* **95**(1), 195–203.

MacAvoy, P. W. (1982), *Crude Oil Prices as Determined by OPEC and Market Fundamentals*, Ballinger Publishing Company, Cambridge, MA.

Martinot, E. (2016), 'Grid integration of renewable energy: Flexibility, innovation, and experience', *Annual Review of Environment and Resources* **41**(1), 223–251.

Michielsen, T. O. (2014), 'Brown backstops versus the green paradox', *Journal of Environmental Economics and Management* **68**(1), 87–110.

OECD (2012), 'OECD environmental outlook to 2050: Key findings on climate change'.

Persson, T. A., Azar, C., Johansson, D. & Lindgren, K. (2007), 'Major oil exporters may profit rather than lose, in a carbon-constrained world', *Energy Policy* **35**(12), 6346–6353.

Pindyck, R. S. (1978), 'The optimal exploration and production of nonrenewable resources', *Journal of Political Economy* **86**(5), 841–861.

Rubio, S. J. & Escriche, L. (2001), 'Strategic Pigouvian taxation, stock externalities and polluting non-renewable resources', *Journal of Public Economics* **79**(2), 297–313.

Seierstad, A. & Sydsaeter, K. (1987), *Optimal Control Theory with Economic Applications*, North-Holland, Amsterdam.

Sinn, H.-W. (2008), 'Public policies against global warming: a supply side approach', *International Tax and Public Finance* **15**(4), 360–394.

Solomon, S., Plattner, G.-K., Knutti, R. & Friedlingstein, P. (2009), 'Irreversible climate change due to carbon dioxide emissions', *Proceedings of the National Academy of Sciences* **106**(6), 1704–1709.

Tahvonen, O. (1997), 'Fossil fuels, stock externalities, and backstop technology', *Canadian Journal of Economics* **30**(4), 855–874.

Ulph, A. & Ulph, D. (1994), 'The optimal time path of a carbon tax', *Oxford Economic Papers* **46**(Special Issue), 857–868.

Van der Ploeg, F. (2016), 'Second-best carbon taxation in the global economy: The green paradox and carbon leakage revisited', *Journal of Environmental Economics and Management* **78**, 85–105.

Van der Ploeg, F. & Withagen, C. (2012), 'Too much coal, too little oil', *Journal of Public Economics* **96**(1), 62–77.

Verleger, P. K. (1982), 'The determinants of official OPEC crude prices', *The Review of Economics and Statistics* **64**(2), 177–182.

Weitzman, M. L. (2009), 'On modeling and interpreting the economics of catastrophic climate change', *Review of Economics and Statistics* **91**(1), 1–19.