

NBER WORKING PAPER SERIES

COASE, HOTELLING AND PIGOU:
THE INCIDENCE OF A CARBON TAX AND CO₂ EMISSIONS

Geoffrey Heal
Wolfram Schlenker

Working Paper 26086
<http://www.nber.org/papers/w26086>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2019

We are grateful for financial support under a SIPA faculty grant and to Soren Anderson and Joséphine Gantois for helpful comments on this paper as well as seminar participants at Berkeley, Harvard, INRA, the NBER spring meeting, and PSE. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2019 by Geoffrey Heal and Wolfram Schlenker. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Coase, Hotelling and Pigou: The Incidence of a Carbon Tax and CO₂ Emissions
Geoffrey Heal and Wolfram Schlenker
NBER Working Paper No. 26086
July 2019, Revised June 2020
JEL No. Q41,Q54

ABSTRACT

We use field-level cost estimates of all oil and natural gas fields to highlight dynamic aspects of a global carbon tax. Some of the initial reduction in consumption will be offset through higher consumption later on. Only high-cost reserves will be priced out of the market, e.g., at 200 dollars per ton of CO₂ cumulative emissions decrease by 4%. The tax incidence initially falls on consumers under a constant tax but eventually becomes negative as the lifetime of the resources is extended. An increasing tax over time reduces the initial incidence on consumers.

Geoffrey Heal
Graduate School of Business
516 Uris Hall
Columbia University
New York, NY 10027-6902
and NBER
gmh1@columbia.edu

Wolfram Schlenker
School of International and Public Affairs (SIPA)
Columbia University
420 West 118th St
New York, NY 10027
and NBER
wolfram.schlenker@columbia.edu

There is strong agreement amongst economists that a carbon tax is an effective method to reduce carbon emissions. For example, The Initiative on Global Markets (IGM) at the University of Chicago Booth School of Business maintains a representative panel of economists. A carbon tax was favored by almost all economists, and there was a greater divergence with views by the general public than for any other question (Paola Sapienza & Luigi Zingales 2013).¹ Every environmental economics text sees the internalization of external costs as a necessary step on the road to efficiency. Carbon emissions create externalities, and a tax will internalize them (Arthur Cecil Pigou 1920). The Pigouvian framework is the default setup when it comes to thinking about environmental policy, as a Pigouvian tax drives a wedge between producer and consumer prices and in a static one-period model generally reduces the equilibrium quantity.

However, fossil fuels are an exhaustible resource with a limited supply. Scarcity rents can be a significant portion of the price to ensure that the limited supply is optimally allocated between periods (Harold Hotelling 1931). For example, Saudi Arabia's production cost are in the range of \$5-8 per barrel, yet the oil price in 2019 was around \$60 per barrel. In the standard Hotelling model, all resources are used up, and a tax is paid out of the scarcity rents of producers. The tax might slightly shift consumption patterns over time, but does not change the cumulative use of the resource.

The point we are making in this paper is that the Pigouvian and Hotelling frameworks lead to rather different conclusions when it comes to thinking about the effectiveness of a carbon tax. Pigou emphasizes the impact of a tax on substitution between commodities, in this case between energy sources. Hotelling on the other hand emphasizes the impact of a tax on an exhaustible resource on the time-path of consumption of that resource. It can lead to the substitution from present to future consumption, so that less of the resource is consumed by any date but the same amount is consumed overall. One of the clear conclusions of the Hotelling model of equilibrium in a resource market is that if there is a substitute for the resource - think of renewable energy - available at a price in excess of the marginal extraction cost of the resource, then all of the resource will be consumed eventually, and

¹The statement "A tax on the carbon content of fuels would be a less expensive way to reduce carbon-dioxide emissions than would a collection of policies such as corporate average fuel economy requirements for automobiles" was agreed to by 92.5% of economists, while only 22.5% of the general public agreed, as measured by the Chicago Booth Kellogg School Financial Trust Index survey. Support for a carbon tax is growing among various policy circles. The New York Times reported that "Republican Group Calls for Carbon Tax" (2/7/17), and the Financial Times noted that "Leading Corporations Support US Carbon Tax" (6/20/17). The Carbon Pricing Leadership Coalition (www.carbonpricingleadership.org) is a coalition of international and national organizations and corporations dedicated to promoting a carbon tax.

a carbon tax can only change this under rather stringent conditions. Carbon taxes reduce carbon emissions less once these dynamic considerations are incorporated.

Our results build on an earlier literature that studies the optimal taxation of exhaustible resources. The origins of this literature predate the discussion of climate change regulation. The aim was to understand how to tax resource rents without generating a deadweight loss. The same logic that tells us that a tax on exhaustible resources does not generate a deadweight loss – a perfectly inelastic supply – also implies that a carbon tax might in effect not reduce cumulative carbon emissions. The contributions of our paper are:

First, we extend this theory to the setting of carbon tax in Section 1, starting with the most basic model of constant marginal extraction cost and a backstop technology, before relaxing several assumption, e.g., inducing heterogeneity in marginal extraction cost, fixed cost of field development, limited substitutability with the backstop technology, and a carbon tax that increases over time. The intuition remains the same in all cases. A constant carbon tax (in real terms) has two effects. It implies a shift in the time profile of consumption: a tax will delay consumption to future periods, but initial reductions in carbon emissions are offset through an extension of the period over which the resource is used. On the other hand, a tax might price high-cost reserves out of the market, which would lead to a permanent reduction in cumulative use. Which effect dominates is an empirical question. An increasing carbon tax limits the incentive to extend the time period over which the resource is consumed, as any extension would increase the tax further.

Second, we take the theoretical prediction to data. Section 2 utilizes a micro-level data set that gives the marginal extraction cost of all oil and natural gas fields around the world. As we briefly argue below, a carbon tax will make construction of new coal plants unprofitable and eliminate part of CO₂ emissions from coal. However, the effect on oil and natural gas is less clear. We start with the oil market, as oil is a standardized commodity traded globally. We study the effect of a carbon tax using proprietary data on the cost structure of oil fields from Rystad Energy’s UCube product and publicly available data on oil consumption from the Energy Information Agency. The oil market is interesting, as recent estimates have argued that consuming all oil would use up almost the entire carbon budget consistent with the world staying within 2°C warming (Richard J. Millar, Jan S. Fuglestedt, Pierre Friedlingstein, Joeri Rogelj, Michael J. Grubb, H. Damon Matthews, Ragnhild B. Skeie, Piers M. Forster, David J. Frame & Myles R. Allen 2017). Scarcity rents for oil are so high that only few oil fields will drop out of the market with moderate carbon taxes. For example, a carbon tax as high as \$200 will eliminate only 4% of oil production. An oil field is no longer

profitable if the extraction costs exceed the backstop (or choke) price minus the carbon tax. Lowering the backstop price (e.g., cheaper renewables) is equivalent to a carbon tax and might be used in combination with a carbon tax. About three quarters of a constant carbon tax will initially be passed on to consumers, but this incidence declines over time and even becomes negative as oil consumption is shifted from the present to the future by the carbon tax, decreasing the price of oil by the end of the century compared to a case without a tax. This makes the political economy of a global carbon tax difficult, as the costs are highest on immediate users. Producers and consumers roughly split the cost of a carbon tax in present-value terms, i.e., they face similar declines in surplus. The limited response in cumulative oil consumption implies that almost all losses in consumer and producer surplus are offset by higher tax revenue. Carbon taxes would be a way to raise revenue without deadweight loss. They might not be effective way to reduce emissions, but certainly could raise government revenue (Adele Morris 2013). If the tax is levied on consumption, net exporters of oil are predicted to see welfare declines, while net importers see welfare increases.

Third, we present the case of an increasing carbon tax equal to the social cost of carbon from the 2016 Interagency Working Group of the US government (United States Government 2016), which is increasing over time in real terms. An increasing tax gives yet another offsetting incentive to shift more of the consumption to the present to avoid a higher tax in the future. The initial incidence on consumers is reduced from 70% to 50% for the average social cost of carbon compared to a constant carbon tax that starts at the same level in 2019. Cumulative carbon emissions from oil fall by 2% if the carbon tax equals the average social cost of carbon over time. If the increase in the carbon tax over time becomes steeper than under SCC, an inverted U-shape is possible. For example, if the carbon tax increases three times faster than SCC, the incidence on consumers starts at 15% before peaking around 35% mid-century and then continually declining. In case of a carbon tax that rises exponentially at the rate of interest, the incidence falls almost entirely on producers, even in earlier periods.

Fourth, we conduct a sensitivity analysis that treats oil and natural gas as perfect substitutes. This might currently be a stretch: natural gas prices differ substantially around the globe as the commodity is difficult to trade, requiring the construction either of expensive pipelines or liquefaction plants and refrigerated ships. However, our analysis focuses on the long-term price dynamics, and one might argue that in future decades these two commodities could indeed become closer substitute. The joint availability of oil and natural gas greatly surpasses the carbon budget that would keep the Earth within +2C even assuming that there are no emissions from coal: only 68% of the resources (based on CO₂ content) could be

used. The dynamics of the oil market carry forward with very minor numerical differences: only a carbon tax of several hundred dollars can make a dent in carbon emissions.

Before we dive into the theory and empirical implementation, it might be worthwhile to give some basic intuition. We start with the simplest setup: constant marginal production cost coupled with a downward sloping linear demand curve. The top panel of Figure 1 shows the standard case for the static problem of a good that can be produced every period. The constant marginal production cost implies that there is no producer surplus, only consumer surplus. A tax will drive a wedge between consumer and producer price, lowering the equilibrium quantity from q_0 to $q_{\text{after tax}}$. The incidence of the tax is entirely on consumers, whose consumer surplus decreases from the grey triangle CS_0 to the blue triangle $CS_{\text{after tax}}$.

The bottom panel Figure 1 displays an equivalent setup for an exhaustible resource whose supply is fixed and which cannot be produced. The simplest setup is a two period model, and the available amount can be split between the two periods. The length of the horizontal axis gives the total supply of the exhaustible resource. Consumption in the second period is hence measured from the right. We utilize a comparable demand curve coupled with constant marginal extraction cost. For graphical simplicity, we plot the demand curve minus this constant marginal cost. The same demand functions are plotted from the right axis for consumption in period 2. The only difference is that they are discounted as consumption happens one period later – both the intercept and slope are lower by the discount factor. The price exceeds marginal cost as not enough of the resource is available to satisfy demand in both periods at a price equal to the marginal cost of extraction. This gives rise to scarcity rents for producers. Unlike the standard good case, there is substantial producer surplus in the market due to these scarcity rents, which ensures that some of the good is saved for the second period. Note how the results of a carbon tax flip compared to the top panel: while there might be a slight adjustment in how much is consumed in each period, the combined consumption is unaltered. It is given by the length of the horizontal axis. Consumer surplus doesn't change much, and almost all of the incidence is now on producers.

Our modeling section will introduce various modifications, e.g., heterogenous costs, but the main insight is given by the bottom panel of Figure 1: cumulative consumption of a resource that has significant scarcity rents is foremost given by its availability, and a tax might alter the consumption path between periods, but not the total overall consumption. Only high-cost reserves, where the sum of the marginal extraction cost and the tax rate exceeds the choke price of the demand function, will drop out of the market. As the empirical section of the paper shows, there are few high cost reserves, and very substantial carbon taxes

are required for reductions in carbon emissions.

Our findings are extension of an earlier literature. The impact of taxation on the pattern of resource use was discussed in the 1970s by Partha Dasgupta & Geoffrey Heal (1979) and Parth Dasgupta, Geoffrey Heal & Joseph Stiglitz (1980) using the Hotelling framework. These papers pre-date concerns about climate change and greenhouse gases, and focused on the impact of taxation on the time pattern of resource use in a continuous-time infinite-horizon competitive equilibrium. These papers showed that, to quote, “there exists a pattern of taxation which can generate essentially any desired pattern of resource usage” (Dasgupta, Heal & Stiglitz 1980). In other words, an appropriate system of taxation can produce any time pattern of use of a fossil fuel. But in all of these patterns, all of the fuel will be used up: cumulative use, and so emissions, will thus be the same in all. Only their distribution over time will differ from one case to the other. This is consistent with our finding that in the basic Hotelling model a carbon tax can change the time pattern of fuel use but not alter the total use and therefore not alter cumulative greenhouse gas emissions.

A later literature on the “green paradox” (Hans-Werner Sinn 2015, Hans-Werner Sinn 2012, Michael Hoel 2012, Michael Hoel 2010, Sven Jensen, Kristina Mohlin, Karen Pittel & Thomas Sterner 2015, Robert Cairns 2012) asks whether policies that are intended to reduce greenhouse gas emissions could in fact have the opposite effect: could they actually promote emissions? The literature arrives at a positive conclusion, noting that an expectation of rising taxes on fossil fuels will lead to an increase in the rate at which they are used in the present (Sinn 2012). Intuitively, any regulation that will make fuel use more costly or outright prohibit it in the future will shift some of the future emissions to the present and accelerate the depletion of fossil fuels and their emissions. This is consistent with earlier findings: Dasgupta Heal and Stiglitz find that “...the effects of tax structure on patterns of extraction are critically dependent on expectations concerning future taxation.” They show that a sales tax that rises over time will lead to more rapid use of an exhaustible resource, and vice versa, which is essentially the green paradox.

Reyer Gerlagh (2010) distinguishes between weak and strong green paradoxes: the weak paradox occurs when policies increase near-term carbon emissions, but not total emissions. The strong paradox is used for cases when total emissions are increased. In the models considered in this paper there are no strong green paradoxes, and weak ones occur only if there is an increase in the tax rate over time. Carbon taxes either have no impact on total emissions or reduce them. Rick van der Ploeg & Cees Withagen (2010) and Rick van der Ploeg & Cees Withagen (2015) show that the anticipation of a drop in the price of renewable

energy may also generate a green paradox, encouraging the more rapid use of fossil fuels. Hoel (2012) considers a model in which the cost of extraction of a fossil fuel depends on the cumulative extraction to date using the formulation of Geoffrey Heal (1976), and shows that in this case a carbon tax can reduce total greenhouse gas emissions. This is analogous to our results in sections 1.4, where we consider multiple grades of a fossil fuel differing in their extraction costs.

1 Model

We start in Section 1.1 with a basic model in which we explore the impact of a carbon tax on the time pattern of use of a fossil fuel facing competition from a renewable energy source which is a perfect substitute and is available at a price in excess of the marginal extraction cost of the fuel, and show that one of two outcomes must hold: either the tax has no impact on *cumulative* consumption of the fossil fuel, though it does delay consumption; or it prevents any consumption of the fuel at all. The two energy sources will never be used simultaneously. We then (Section 1.2) modify the model to reflect the fact that the renewable resource is only an imperfect substitute for the fuel. In this case we find that the fossil fuel and the renewable resource are used simultaneously, but the earlier basic conclusion still holds: a tax will either stop the consumption of the fuel altogether, or merely delay it. Section 1.3 looks at the consequences of introducing fixed costs in the extraction of fossil fuels on top of variable costs. In this case a carbon tax may lead to a reduction in the total consumption of fossil fuels because the net revenues from their sales no longer offer an adequate return on the investment in the fixed cost. In Section 1.4 we consider the more realistic, yet also more complex, case of multiple grades of the fossil fuel differing in their extraction costs. Here we find that a carbon tax may delay the consumption of the less expensive grades and eliminate from the market altogether the more expensive grades, thereby reducing greenhouse gas emissions. Section 1.5 further extends the model and includes increasing tax rates over time. The overall conclusion is that there are two dimensions to the impact of a carbon tax: delaying the consumption of fossil fuels, and eliminating expensive fuels (expensive in either fixed or variable costs) from the market. Only the latter reduces cumulative greenhouse gas emissions, and in some cases only the former mechanism will be effective. In Section 1.6 we extend our model to consider the impact of a cap and trade system on emissions from fossil fuels (an approach based on the ideas of Coase (Ronald Coase 1960) about the role of property rights in controlling externalities), and show that by fixing the allowable quantity

it attains the objective of reducing emissions, but even modest quantity reductions imply a steep permit price. If permits are auctioned off and not grandfathered, it has the effect of expropriating the scarcity rents associated with exhaustible fossil fuels.

1.1 Homogenous Resource with Backstop as Perfect Substitute

There is an initial stock $S_0 > 0$ of a fossil fuel, selling at a market price p_t at date t in a competitive market. Its marginal extraction cost is constant at $c_m > 0$ and its price p_t is given by

$$p_t = h_t + c_m + \tau \tag{1.1}$$

where τ is a per unit tax rate that must be paid on sales of the fuel. This is a carbon tax, meaning that it is calculated from the carbon released when the fuel is burned: it does not depend on the value of the product.² The scarcity or Hotelling rent on the fuel, h_t , is its net price after extraction cost and the tax are paid. We know that in a competitive market equilibrium this will rise exponentially at the prevailing interest rate r (Dasgupta & Heal 1979, chapter 6) or there would be intertemporal arbitrage opportunities.

$$p_t = h_0 e^{rt} + c_m + \tau \tag{1.2}$$

In addition to the fossil fuel there is a renewable resource available in unlimited amounts at constant marginal and average cost of $\bar{p}_b > c_m$. This is a perfect substitute for the fossil fuel and nobody would pay more for the fossil fuel than \bar{p}_b . It is a “backstop technology” in the terminology of Dasgupta & Heal (1979), so that if the fuel is consumed we must have

$$p_t \leq \bar{p}_b \tag{1.3}$$

Demand for the fuel is given by the demand function $q(p_t)$. We are interested in the competitive equilibrium dynamics of prices and demand for the fuel, and how these are affected by the carbon tax. We know that the market price of the fuel will rise exponentially away from $c_m + \tau$ at rate r , as given in (1.2), and that $p_t = h_0 e^{rt} + c_m + \tau \leq \bar{p}_b$ if the fuel is sold.

A competitive equilibrium is characterized by two conditions. First, the price of the fossil

²In the empirical section below we combine oil and natural gas fields, which have different carbon intensities, and the tax τ differs by fuel type.

fuel at the end of extraction p_T has to equal the price of the substitute $p_T = \bar{p}_b$

$$p_T = h_0 e^{rT} + c_m + \tau = \bar{p}_b \quad (1.4)$$

If it were lower ($p_T < \bar{p}_b$), a fuel producer could wait an infinitesimal time past T to sell at a discretely higher price $\bar{p}_b > p_T$ for an infinite return. A higher price ($p_T > \bar{p}_b$) is ruled out as the backstop is assumed to be a perfect substitute, and no consumer would pay more for the fossil fuel than the substitute. A sample price path is shown in the top of Figure 2.

Second, cumulative use over the period T will equal initial availability S_0

$$Q_T = \int_0^T q(p_t) dt = \int_0^T q(h_0 e^{rt} + c_m + \tau) dt = S_0 \quad (1.5)$$

Total consumption cannot be larger than availability by definition. It can also not be strictly less, as it would imply that some fossil fuel is left in the ground, yet the producer would prefer to sell it for a profit as $p_t > c_m$. Provided that the marginal extraction cost plus tax is less than the price of the renewable energy source, all of the fossil fuel will be consumed, as it will always be profitable to extract and sell it.

The two equations (1.4) and (1.5) characterize the competitive equilibrium and can be solved for T and h_0 . As long as $c_m + \tau < \bar{p}_b$ the competitive equilibrium will involve a period $[0, T]$ during which only the fossil fuel is consumed and then a period from T onwards during which only renewable energy is used, and that over the interval $[0, T]$ all of the fossil fuel will be consumed.

Proposition 1. *If the tax rate τ in a competitive equilibrium with a perfect substitute at price \bar{p}_b is raised to $\tau' > \tau$, $c_m + \tau' < \bar{p}_b$ then all reserves are still consumed, albeit over an extended period $T_{\tau'} > T_\tau$. The initial scarcity rent $h'_0 < h_0$ is lowered, as are all scarcity rents for $t \leq T_\tau$, i.e., $h'_t < h_t$. If the tax is raised further such that $c_m + \tau' > \bar{p}_b$, then none of the fossil fuel is consumed.*

Proof. The equilibrium is given by the implicit function $f = \int_0^T q(h_0 e^{rt} + c_m + \tau) dt - S_0 = 0$. The implicit function theorem tells us that

$$\frac{dT}{d\tau} = -\frac{\frac{\partial f}{\partial \tau}}{\frac{\partial f}{\partial T}} = -\frac{\int_0^T q'(h_0 e^{rt} + c_m + \tau) dt}{q(h_0 e^{rT} + c_m + \tau)} = \frac{\int_0^T -q'(p_t) dt}{q(p_T)} > 0 \quad (1.6)$$

$$\frac{dh_0}{d\tau} = -\frac{\frac{\partial f}{\partial \tau}}{\frac{\partial f}{\partial h_0}} = -\frac{\int_0^T q'(h_0 e^{rt} + c_m + \tau) dt}{\int_0^T e^{rt} q'(h_0 e^{rt} + c_m + \tau) dt} = -\frac{\int_0^T -q'(p_t) dt}{\int_0^T -e^{rt} q'(p_t) dt} < 0 \quad (1.7)$$

Since demand is downward sloping $-q'(p) > 0$, and the integral is over only positive values. Furthermore, $h'_0 < h_0$ implies that $h'_t < h_t$. \square

Once the tax τ is so high that $c_m + \tau > \bar{p}_b$, the fossil fuel will never be consumed as it becomes uncompetitive relative to the perfect substitute at price \bar{p}_b .³ We conclude that there is a discontinuous effect of a carbon tax: low taxes shift consumption of the fossil fuel over time but do not change total cumulative consumption. No change in the tax rate - as long as it satisfies the condition $c_m + \tau < \bar{p}_b$ - will alter cumulative consumption. Another way of thinking about this is that with a normal produced good, a tax would drive a wedge between the consumer and the net price received, thereby reducing the output along the supply curve. With an exhaustible resource the supply curve is vertical: the resource is there whatever the price and is profitable as long as $c_m + \tau < \bar{p}_b$. Once the tax surpasses $\bar{p}_b - c_m$, the consumption of the fossil fuel abruptly drops to zero. There is no intermediate case in which the tax reduces the total consumption of the fossil fuel but not to zero.

1.2 Imperfect Substitutability

In the world around us, we see both renewable energy and fossil fuels in the market at the same time, rather than the abrupt switch from one to the other that the model predicts. There are several possible reasons for this discrepancy. Principal amongst them is that we have assumed that fossil fuels and renewable resources are perfect substitutes, so that demand switches completely from one to the other as the ordering of their prices changes.

In reality this is not the case: renewable energy is intermittent, which is a disadvantage relative to fossil energy, but is clean, producing no pollutants that damage the local environment and no greenhouse gases. Because of these factors we can imagine situations where renewable energy is used even if it is more expensive (situations where there is a need to reduce local pollution, or to reduce greenhouse gas emissions) and conversely situations where a fossil energy such as natural gas is used even though it is more costly (for example gas is used to back up intermittent renewable energy). To try to capture these possibilities, we now modify the demand for fossil fuels to show that it depends not only on its own price p_t but also on the price of renewable energy p_b : $q(p(t), p_b)$, $\partial q / \partial p_b > 0$. This admits the possible co-existence of both energy sources in the market simultaneously, with demand transferring from one to the other as the price difference changes. We assume the demand function to have a ‘‘choke price’’ $\bar{p}(p_b)$ such that demand for the fossil fuel falls to zero when

³See also Hoel (2012) for a discussion of this case: he refers to such a tax as a ‘‘high tax.’’

its price reaches $\bar{p}(p_b)$. So $q(\bar{p}(p_b), p_b) = 0$. Obviously, the choke price depends on the price of the substitute. In the previous analysis $\bar{p}(\bar{p}_b) = \bar{p}_b$.

It is still the case that in equilibrium the price of the fossil fuel will be given by equation (1.2), with the Hotelling rent rising exponentially at the interest rate. For all markets for the fuel to clear it is now necessary and sufficient that the time T at which the price of the fuel equals its choke price, $p_T = \bar{p}(\bar{p}_b)$. The two equations that define the equilibrium are equivalent to the previous section (see Appendix A1.1).

1.3 Fixed Costs of Extraction

So far, we have assumed that all the costs of extracting the fossil fuel are variable costs, with a marginal extraction cost of $c_m > 0$. Suppose in addition that there is a fixed cost $c_f > 0$ that must be incurred before the fuel can be extracted at a marginal cost of c_m . This could be the cost of finding and developing an oil or gas field, a cost that in practice can be substantial. Could this alter our conclusions?

In this case the fuel will only be produced if the price is high enough to cover the tax, extraction cost and fixed cost. The time path of the fuel price will still be given by equation (1.2), so now we require that the discounted rents cover the upfront fixed cost

$$\int_0^T e^{-rt} (p_t - c_m - \tau) dt = \int_0^T e^{-rt} h_0 e^{rt} = h_0 T \geq c_f \quad (1.8)$$

Market clearing conditions are still given by equations (1.4) and (1.5), with the additional constraint given in equation (1.8). This constraint make it possible that an increase in the tax rate reduces the Hotelling rent h_0 of undeveloped resources enough that $h_0 T < c_f$, which would imply that it would not be profitable to ever develop the resource. The introduction of fixed costs in the extraction technology therefore gives another mechanism via which a carbon tax might prevent the extraction of the fossil fuel. Once again we get a knife-edge case: either all resources are or none at all are used. If there were a range of resources with different fixed costs, then a tax might rule out only those with the highest fixed costs.

If there is an initial tax rate at which extraction is profitable - i.e., equation (1.8) is satisfied - but after extraction has begun the tax is increased to a point where this is no longer true (but the fixed development cost are sunk), extraction will continue provided that $c_m + \tau < \bar{p}_b$. In our empirical implementation below we therefore use only the marginal cost for developed fields (the fixed costs are sunk), while we include both fixed development and variable extraction cost for undeveloped fields.

1.4 Multiple Grades of Fossil Fuel

Another case of interest is that of multiple sources of the fossil fuel, with different extraction costs. Suppose we modify the model of Section 1.1 so that there are I different fuel sources each with marginal extraction cost $c_{m,i}$ and let them be numbered in increasing order of extraction costs, so that $c_{m,1} \leq c_{m,2} \leq c_{m,3} \leq \dots \leq c_{m,I}$ and further assume that $c_{m,I} < \bar{p}_b$ so all are less expensive than the renewable resource. If there are reserves where $c_{m,i} > \bar{p}_b$, they will never be used and can be neglected. The initial stock of the i th fuel is $S_{0,i}$. The competitive equilibrium outcome is that there exist dates T_i , $i = 1, 2, \dots, I$, $T_i < T_{i+1}$, and initial rents $h_{0,i}$, $i = 1, 2, \dots, I$ such that for all i ,

$$p_t = h_{0,i}e^{rt} + c_{m,i} + \tau, \quad T_{i-1} \leq t \leq T_i \quad (1.9)$$

and

$$\int_{T_{i-1}}^{T_i} q(p_t) dt = S_{0,i} \quad (1.10)$$

So each grade of fuel i is used over the interval $T_{i-1} \leq t \leq T_i$ and is used only during this interval and is used up by the end of this interval. The rent at the start of extraction of the i th grade is $h_i(T_{i-1}) = h_{0,i}e^{rT_{i-1}}$. The least expensive fuel is used up first and the most expensive last and the price moves continuously (Dasgupta & Heal 1979, page 172 section (iii)).

$$p_{T_i} = c_{m,i} + \tau + h_{0,i}e^{rT_i} = c_{m,i+1} + \tau + h_{0,i+1}e^{rT_i} \quad \forall i \quad (1.11)$$

This implies that the Hotelling rent decreases by $c_{m,i+1} - c_{m,i}$ in current value terms, or $h_{0,i} - h_{0,i+1} = e^{-rT_i} [c_{m,i+1} - c_{m,i}]$ in present value terms. By induction we know that

$$h_{0,i} > h_{0,j} \quad \forall i < j \quad (1.12)$$

We must have the last price of the fuel equal to that of renewable energy:

$$p_{T_I} = \bar{p}_b \quad (1.13)$$

A sample price path for two grades is shown in the bottom graph of Figure 2. Production switches from the low cost grade to the high-cost grade when the stock of the former is depleted. The price path is continuous at this point, otherwise there would be intertemporal arbitrage opportunities. The scarcity rent drops discontinuously by the difference in marginal extraction cost.

The figure gives an intuition why the resources are used in order of increasing marginal extraction cost. If grade i is currently being used, the price path $p_t = h_{0,i}e^{rt} + c_{m,i} + \tau$ implies that $\frac{dh_i(t)}{h_i(t)} = \frac{dp(t)}{p(t)} = r$, i.e., the Hotelling rent rises at the rate of interest. For another grade j , we have $\frac{dh_j(t)}{h_j(t)} = \frac{dp(t)}{h_j(t)} = \frac{rh_{0,i}e^{rt}}{h_{0,j}e^{rt}} = \frac{h_{0,i}}{h_{0,j}}r$. Recall from equation (1.12) that for a more expensive grade j , $h_{0,j} < h_{0,i}$ and its rent rises at *more* than the rate of interest while the cheaper grade i is being used and it is better to wait. Conversely, a less expensive grade j , $h_{0,j} > h_{0,i}$ and its rent rises at *less* than the rate of interest while the more expensive grade i is being used, i.e., it makes no sense to still hold that grade in the ground and it should have been depleted by the time grade i starting its exploitation.

Because of the existence of multiple grades of fuel we no longer have the earlier all-or-nothing impact of a tax rise: it can now lead to the elimination of some but not all of greenhouse gas emissions by pushing out of the market the more costly fossil fuels.⁴ This allows us to extend Proposition 1 to multiple grades

Proposition 2. *If the tax rate τ in a competitive equilibrium with heterogenous extraction cost and a perfect substitute at price \bar{p}_b is raised to $\tau' > \tau$, $c_m + \tau' < \bar{p}_b$ then not all reserves might still be consumed. The time period $T_{\tau'}$ can increase or decrease, but the initial scarcity rent $h'_0 < h_0$ is still unambiguously lowered, as are all scarcity rents for $t \leq T_{\tau}$, i.e., $h'_t < h_t$.*

Proof. We have shown in equations (1.6) and (1.7) that an increase in the tax rate will extend the lifetime T and reduce the initial Hotelling rate h_0 . Analogous comparative statistics with respect to the stock S_0 give $\frac{dT}{dS_0} < 0$ and $\frac{dh_0}{dS_0} < 0$.

In the case of heterogenous marginal extraction cost, the impacts of a carbon tax are essentially the same as before: provided that $c_{m,i} + \tau < \bar{p}_b$, a tax increase will merely delay the consumption of the fossil fuel by the logic outlined in equation (1.6), but will not alter cumulative consumption. However, it is possible that a tax increase could imply that the more expensive grades j are no longer profitable, either because $c_{m,j} + \tau > \bar{p}_b$, or in the case of fixed cost, that $c_{m,j} + \tau < \bar{p}_b$ but $h_{0,j}e^{rT_j-1} [T_j - T_{j-1}] < c_f$, i.e., the rents are not large enough to cover fixed cost of development. In either case, fossil fuel grade j will not be produced and cumulative emissions will fall. Whether the effect of a tax on the length of the extraction period in equation (1.6) outweighs the opposite effect of a reduction in cumulative availability S_0 of stocks that can profitable be explored is an empirical question, and we present empirical cases for both below. Intuitively, if there are few high-cost reserves

⁴Section A1.2 combines the results of Section 1.2 on imperfect substitutability with heterogeneous fuel grades.

that drop out, the former will dominate leading to a longer T . On the other hand, if there is a significant number of high-cost reserves that drop out, the latter will dominate leading to a shorter T .

There are similar opposite effects for the initial Hotelling rent $h_{i,0}$. Equation (1.7) established that a tax increase will reduce it as long as no reserves drop out of the market. If higher cost reserves drop out, we have now from $\frac{dh_0}{dS_0} < 0$ that scarcity rents should go up. However, it is possible to show the former effect always dominates. If the most expensive grade I is profitable under tax τ , we know that $p_{T_{I-1}} = c_{m,I-1} + \tau + h_{0,I-1}e^{rT_{I-1}} > c_{m,I} + \tau$, or $h_{0,I-1}e^{rT_{I-1}} > c_{m,I} - c_{m,I-1}$. If we continue to raise the tax until the I -th grade becomes unprofitable, so $c_{m,I} + \tau' > \bar{p}_b$, we know that grade $I - 1$ is now the final grade that will be exploited and hence $p_{T'_{I-1}} = c_{m,I-1} + \tau' + h'_{0,I-1}e^{rT'_{I-1}} = \bar{p}_b < c_{m,I} + \tau'$, or $h'_{0,I-1}e^{rT'_{I-1}} < c_{m,I} - c_{m,I-1}$ and $h'_{0,I-1}e^{rT'_{I-1}} < h_{0,I-1}e^{rT_{I-1}}$. \square

Finally, in Section 2.2 we combine oil and natural gas field. The carbon intensity of natural gas is lower than for crude oil. As a result the tax τ_{gas} is lower than that for oil τ_{oil} . There is a small modification to the above rule: fields with the lowest sum of marginal extraction cost and carbon tax will go first, i.e.,

$$c_{m,1} + \tau_1 < c_{m,2} + \tau_2 < \dots < c_{m,I} + \tau_I \quad (1.14)$$

1.5 Increasing Tax Rate

So far we have examined a constant carbon tax. Many carbon tax proposals call for a tax that is increasing in time $\tau(t)$ with $\frac{d\tau}{dt} > 0$. For example, the social cost of carbon is increasing in time in real terms, and a carbon tax that internalizes the externality associated with fossil fuel use should hence be increasing as well over time.

If the tax is rising over time, the resulting price path $p_t = h_0e^{rt} + c_m + \tau_t$ will rise even faster than before, as not only the Hotelling rent h_t but also the tax rate τ_t is increasing over time. The problem is again solved backwards as in previous sections. The final price $p(T)$ equals the backstop price \bar{p}_b . This implies that going backward the price path (red line in Figure 2) declines faster going back in time, and the quantity consumed will increase faster. A large fraction of consumption gets shifted towards the present again. Intuitively, consuming the exhaustible resource faster avoids a higher tax in the future. As we show in the empirical section below, this latter effect can even become dominating for a fast-increasing tax leading to a green paradox, i.e., a case where cumulative consumption will

initially increase compared to a case without a tax until the resource is exhausted at an earlier time and cumulative consumption drops (as high-cost reserves drop out of the market).

1.6 Cap and Trade

The widely-considered alternative to a carbon tax is a cap-and-trade system. Unlike a tax that levies a price on each unit, the cap-and-trade system requires entities to acquire a permit for each unit. Under certainty, the two are the same: one either sets the price or quantity, while the other adjusts accordingly.⁵

In our baseline model in Section 1.1 with constant marginal extraction cost, a tax τ with $c_m + \tau < \bar{p}_b$ will not reduce emissions but expropriate a fraction of the producer surplus. The equivalent is not possible with a permit regulation. Let the cap on cumulative emissions be K_0 . If $K_0 \geq S_0$, the cap is non-binding and the price would be zero, i.e., while there is also no reduction in use, the government cannot collect revenue and expropriate a fraction of the producer surplus. If the cap is set at $K_0 < S_0$, all producer surplus is expropriated as the scarcity rents now go to the permit owner (the limiting quantity) and not the owner of the resource.

In the case of heterogenous marginal extraction cost that gives rise to a continuous upward-sloping supply curve, there is a one-to-one relationship between a carbon tax and a cap-and-trade regulation. Any carbon tax will have an equivalent cap K_0 . As we mentioned in the introduction and will show in the empirical section below, the supply curve is very steep for higher prices, implying that carbon taxes that have been discussed would imply only minor reductions in carbon emissions. An implication is that a global cap-and-trade system that limits emissions would imply a very high permit price.

The intuition is similar to (Severin Borenstein, James Bushnell, Frank A. Wolak & Matthew Zaragoza-Watkins 2019), who examine the expected price of California's cap-and-trade system. With a steep marginal abatement cost curve for reductions in carbon emissions coupled with uncertainty about how much abatement is needed, there is a very high probability that the permit price will end up at the floor or ceiling price. A small increase in the required amount of abatement, e.g., because the economy is booming and carbon emissions are higher than expected yet the cap is fixed, leads to a sharp rise in the permit price.

⁵Martin L. Weitzman (1974) has shown that they might differ under uncertainty.

1.7 Other Modeling Extension

Our models so far have assumed a competitive market. The same intuition would apply even under a monopolist. While Hotelling (1931) has shown that in the case of a linear demand function, the monopolist will restrict output compared to the competitive equilibrium, the same cumulative emissions will occur, albeit over a longer time period. The opposite can also be true, e.g., for the case of quasi-fixed cost, the monopolist might extract at a faster rate (Tracy R. Lewis, Steven A. Matthews & Stuart Burness 1979). Finally, in our setup of an iso-elastic demand function, the monopolist and competitive equilibrium will be identical. In summary, while different demand function can lead to a fast, slower, or identical time period over which the monopolist extracts a resource, the cumulative use and the cumulative effect of a carbon tax are the same.

2 Numerical Analysis: Extraction Costs and Tax Rates

We now simulate the effect of a carbon tax on long-term oil consumption and prices. Some regions (e.g., British Columbia) or countries (Denmark, Finland, Sweden) have established carbon taxes. Several studies have argued that these taxes have led to significant reductions in CO₂ emissions. For example, Brian Murray & Nicholas Rivers (2015) find that a modest carbon tax of \$30 per ton of CO₂ has reduced emissions by 5-15%, while Boqiang Lin & Xuehui Li (2011) find mixed results for Scandinavian countries. Finland seems to have significantly reduced its emissions, while other countries do not see significant drop in emission, likely due to the fact that some emission intensive sectors are exempt. Gilbert E. Metcalf & James H. Stock (Forthcoming) conduct a systematic analysis of partial carbon regulation around the globe using time series methods. Similarly, recent micro-level studies find that the CO₂ emissions trading system of the European Union Earlier reduced carbon emissions, but not employment. For example, Sebastian Petrick & Ulrich J. Wagner (2014) study the effects on German manufacturing plants using micro-level plant data, while Jonathan Colmer, Ralf Martin, Mirabelle Muûls & Ulrich J. Wagner (2020) conduct a similar analysis for French manufacturing plants.

On the other hand, Joseph A. Cullen & Erin T. Mansur (2017) use short-term variation in the price ratios between natural gas and coal to estimate the reduction in CO₂ emissions from the electricity sector as prices rise and find a very low short-term elasticity, suggesting limited reductions from a carbon tax. Similarly, imperfect competition in the railway market might imply that not all the cost of a carbon tax will be passed on to coal purchasers (Louis

Preonas 2019).

What is common to all of these studies is that they look at partial regulation of small subset of the global economy. Their results are not at odds with ours. A partial regulation of a country might indeed reduce emissions of that country as firms in that countries shift away from energy as an input to other factors, or become more efficient in the use of energy. These partial regulations are not expected to have a sizable effect on global emissions and have the ramification we consider here: feedback on the optimal price and extraction path of an exhaustible resource.

These considerations arise when a global carbon tax is imposed. Our study focuses on such a global carbon tax. There is a catch 22: overall emissions are only meaningfully impacted if all major emission sources are regulated, but if we regulate them all, it would have ramifications on the extraction path that we emphasize. We show below that a carbon tax will reduce fossil fuel consumption in the near future, but increase it in the more distant future, so that a drop in current consumption in response to a tax is quite consistent with our results.

Our paper focuses on long-term implications of a global carbon tax, abstracting from market imperfections in the oil market. The dynamic aspects we emphasize, e.g., that oil will be used eventually if it can be extracted profitably and its effect on the expected price path, is especially important in the context of a *global* regulation, which is eventually required to solve the problem. We see our paper as a complement to the earlier literature on the effect of a carbon tax on oil use focusing on the long-term dynamics rather than the short-term impacts. We deliberately abstract from short-term influences, e.g., political unrest or demand shocks. For example, Soren T. Anderson, Ryan Kellogg & Stephen W. Salant (2018) have shown that once an oil field is set up for production, it is often costly to halt production, violating one of the assumptions of the classical Hotelling model that oil can be produced at any time. They show in their paper that development of new wells respond to prices, but production of existing wells does so to a lesser degree. This can lead to different short-term dynamics, e.g., the negative prices in spring 2020 when oil demand plummet as a result of the shutdown. Since we are interested in the optimal exploitation path over the next 100 years under various carbon taxes, we abstract from these short-term influences.

To get a sense of the empirical significance of our analysis and understand the impact of a carbon tax on fossil fuels, we need to know how much CO₂ each type of fuel releases when burned. Table 1 gives this data for coal, gas and crude oil. For one metric ton of coal

(MT), one million BTU of gas (MMBTU), and one barrel of oil (BBL), it shows how much CO₂ is emitted when this is burned.⁶ There is a range of estimates for how much CO₂ will be released when one barrel of oil is burned. It depends on the exact composition of the fuel and the process by which it is burned. We give the baseline number underlying the Canadian carbon tax. The table also gives the 2019 US price in dollars, and the amount that a \$50 carbon tax would raise per unit of the fuel.

Looking at the numbers in Table 1, it is very clear that the effect of a \$50 carbon tax is potentially much greater in relative terms on coal than on oil: for coal the tax is \$143 per metric ton of coal, while the 2019 price is around \$50. The tax is almost three times the current price. The tax on natural gas is \$2.65 million BTU, while the wholesale price that is just under \$3, i.e., the tax almost equals the price. For oil, however, the tax is about \$17.6 per barrel and the market price around \$60, i.e., the tax equals around a third of the current price.

All three of these fuels are exhaustible resources, so that the earlier analysis is applicable to all of them. Whether reserves drop out of the market depends on the price of the backstop or choke price, whichever is lower. We therefore need to assess whether a carbon tax will increase the MEC - regarding the tax as a part of the MEC - to the point where it is unprofitable to extract the resource. For coal, adding a \$50 carbon tax would roughly quadruple the current price, very likely deeming it uncompetitive, especially relative to natural gas. For natural gas, the answer depends on the circumstances. It is widely assumed in the oil and gas industry that most gas producers are losing money at the current price of \$3 per MMBTU, implying that average costs exceed \$3, though the marginal costs of gas are generally low. One source gives operating costs as 34% of total costs for an average shale gas field, and if this field is breaking even at \$3 then we have an operating or marginal cost of \$1.⁷ In those few cases in which gas is an unintended byproduct of oil production (“associated gas”), one could make an argument that the gas effectively has a zero marginal cost. About 20% of US gas is associated gas⁸. Gas prices in Europe tend to be much higher, as they used to be in the US before the shale boom. Some natural gas might still be used even under carbon tax, but the transportation cost are higher than for oil and the market

⁶The exact carbon content varies by the exact compositions of the fuel and varies between varieties. We quote average estimates to highlight the order of magnitude. For gas see <https://www.eia.gov/tools/faqs/faq.php?id=73&t=11>, for coal see https://www.eia.gov/coal/production/quarterly/co2_article/co2.html, and for gasoline see <https://www.canada.ca/en/department-finance/news/2018/10/backgrounder-fuel-charge-rates-in-listed-provinces-and-territories.html>

⁷http://www.insightenergy.org/system/publication_files/files/000/000/067/original/RREB_Shale_Gas_final_20170315_published.pdf?1494419889

⁸See <https://www.forbes.com/sites/judeclemente/2018/06/03/the-rise-of-u-s-associated-natural-gas/#73e287c04bd7>

seems to be more regional.

Oil is an interesting case study. The world price (the price of Brent marker crude) is in the mid \$60s per barrel, and the US marker crude, West Texas Intermediate (WTI) is in the high \$50s (as of 6/28/2019). A \$50 carbon tax would imply a per-barrel charge that is roughly one third of the current price. The commodity is easily tradable, and the basic Hotelling framework of one global market applies. We therefore start with an analysis of the crude oil market. A later analysis in Section 2.2 will combine crude oil and natural gas into one market for energy.

2.1 Oil Market

Our empirical simulation of optimal oil extraction over time requires three important inputs: the marginal extraction costs of various oil fields (producer side), the price of the backstop technology or choke price (\bar{p}_b in the modeling sections above), and the demand function (demand elasticity, assuming an iso-elastic demand function).

For the production side, we use the proprietary data from Rystad Energy, a prominent source of micro-level data set of various oil fields around the globe. For example, it has recently been used by John Asker, Allan Collard-Wexler & Jan De Loecker (2019) to study the misallocation of oil production around the world. The data set gives estimates for roughly 15,000 discovered and 27,000 undiscovered oil “assets” around the world. An asset is the smallest geographic scale in the data. For example, portions of an oil field can be owned by different firms, and each one of the owners will be listed as a separate asset. Importantly for us, Rystad gives estimates of a breakeven price for each asset. For discovered oil fields, the cost c_m only includes the variable operating cost, as investments in exploration and development are sunk. For undiscovered assets, we include investment and exploration costs, as initial investments are required to access these assets.⁹

The left panel of Figure 3 shows the supply curves for these two categories of crude oil, those already discovered and in operation are shown in dark blue, and those not yet discovered but presumed to exist in light blue. We focus first on the assets already discovered and in operation. The supply curve becomes essentially vertical at a cost of around \$65 and a quantity of just under 1 trillion barrels. Annual global oil consumption is about 36 billion

⁹More precisely, Rystad models production by each asset in future years. It assumes that oil prices are rising 2.5% per year. Rystad estimates extraction cost for all future periods, and in case for undiscovered assets, the exploration and development cost, which are sunk and not included for producing assets. Future costs are discounted to the present using a 10% interest rate. The breakeven price is the current price that makes an asset profitable.

barrels, so the world has about 30 years of oil available at an MEC of \$65 or less. There are roughly an additional 0.8 trillion barrels available in undiscovered crude oil. These tend to be higher cost. As modeling Section 1.4 has shown, the optimal extraction path should first extract the cheaper oil fields, while more expensive ones are developed later. Estimates by Rystad list resources with cost up to \$250 per barrel as viable for future extraction, suggesting the backstop price \bar{p}_b is roughly four times the current price.¹⁰ In our simulation, we therefore assume in our baseline that $\bar{p}_b = 250$. To put a price of 250 into perspective, the price for a liter of gasoline in the United Kingdom at the end of 2019 (including all taxes and add-on cost) was roughly \$250 per barrel, yet there was positive demand. This choke price hence does not seem unrealistically high.

We follow James D. Hamilton (2009, Table 3) for the demand function and use a baseline long-term elasticity of -0.6, the average long-term elasticity given in the table, assuming iso-elastic demand curves. In sensitivity checks in Appendix Figure A1 we use the range of long-term elasticities that were listed in Hamilton (2009), ranging from -0.21 to -0.86. The elasticity has implications for the timeline of prices and quantity consumed, but not the total amount of oil that will be extracted, which only depends on the extraction cost, constant carbon tax, and the cost of the backstop technology.¹¹

Combining the three data sets allows us to construct the optimal extraction profile over time. We follow the theory of reserves with heterogeneous cost of Section 1.4. We know that the most costly reserves will be used last and that the price in the final period has to equal the cost of the backstop $\bar{p}_b = 250$. This allows us to solve the problem backwards, going from the most costly to the least costly reserves (which will produce first in time). We use a daily time step.¹² Rents have to rise at the rate of interest for reserves with the same marginal cost. Once reserves of a particular quality (marginal cost) are exhausted, the price stays continuous, but the rent $h()$ jumps discontinuously by the difference in marginal extraction cost. The exact steps of this backward analysis are given in Section A2. Our model is setup in *real* 2019 dollars, and we leave the demand function unchanged over time.¹³ The interest rate is in real (inflation-adjusted) terms. As the central estimate of the social cost of carbon we chose the case of the discount rate of $r = 3\%$ (Michael Greenstone, Elizabeth Kopits &

¹⁰There is a large mass point of reserves at 250 in the data.

¹¹The elasticity influence the time until depletion, and hence have an effect on cumulative consumption when the carbon tax is increasing over time.

¹²For simplicity every year is assumed to have 365 days.

¹³Alternatively, we could include inflation over time, but would have to shift the demand function over time to avoid money illusion, i.e., that demand is not altered by inflation. Both approaches lead to identical results.

Ann Wolverton 2013), and we use it as our baseline estimate as well.

There is one free parameter in our simulations, the parameter α of the iso-elastic demand function $q_t = \alpha p_t^\eta$. Once we fix this parameter, we simulate the problem backwards to obtain estimates for both the equilibrium price p_{2019} and quantity q_{2109} . We iterate over α to match current global consumption at 100 million barrels a day. Since we have two equilibrium outcomes but only one parameter, there is an implicit test of the other parameter assumptions of the model: they should give us a price p_{2019} that matches the current equilibrium price. For our baseline model, i.e. a demand elasticity of -0.6 , real interest rate of $r = 0.03$, and a backstop price of $\bar{p}_b = 250$, the equilibrium price of 64.42 closely aligns with the market price of oil in 2019. Using parameters from the literature gives results that are internally consistent. On the other hand, if we choose the lower bound of the elasticities $\eta = -0.21$, the simulated price of 82.35 seems too high, while the upper bound of the elasticities $\eta = -0.86$, the simulated price of 53.20 seems too low.¹⁴

2.1.1 Constant Carbon Tax

The baseline case showing the price and production path is shown as short black dashed line (carbon tax = 0) in the top row of Figure 4. Panel A shows the increasing price path over time, rising from 63.54 a barrel at the beginning of 2019 to 250 when the price equals the backstop price in the final period in the year 2095, at which point the production quantity, shown in blue in panel B, falls to zero. Demand after 2095 would only come from the backstop technology (i.e., renewables). Alternative scenarios for carbon taxes ranging from \$50 to \$400 per ton of CO₂ are added as well in different colors ranging from the lowest carbon tax (\$50 in blue) to the highest (\$400 in red). Not surprisingly, a carbon tax raises the price in 2019, as a portion of it is passed on to consumers. The higher price for consumers implies lower production, while the lower price for producers (consumer price minus the tax) implies lower resource rents $h()$ to producers. These lower resource rents now rise at the rate of interest, implying that oil prices grow more slowly than in the baseline case under no carbon tax. Interestingly, there is a point towards the end of the century when prices under the carbon tax become lower than in the baseline case without a carbon tax. The reason is that the carbon tax shifts some of the production from the present to later periods, implying a lower equilibrium price and higher production quantity. As shown in the modeling section,

¹⁴There might of course be other combinations of α , r , \bar{p}_b that give pairs for (p_{2019}, q_{2019}) that are consistent with the current market outcome, but we find it reassuring that the parameters from the literature seem to align with the current equilibrium.

the lifetime can be extended under a carbon tax, i.e., the final period will be 2098 under a \$50 carbon tax, and 2104 under a \$400 carbon tax, when production again falls to zero.

The relative change in prices is shown in panel C of Figure 4. It plots the share of the carbon tax that is passed on to consumers at each point in time by comparing consumer prices under a particular carbon tax to the price path without a carbon tax. This share is initially fairly high (between 70-80% in 2019), but declines continuously as oil prices under the new equilibrium path rise more slowly than under no carbon tax. The ratio eventually becomes negative towards the end of the 21st century when oil prices fall below the level they would have been without a carbon tax. In summary, the carbon tax will initially be passed through to consumers, leading to an immediate increase in oil prices, but the passthrough declines over time and even becomes negative in later years. The cost of a carbon tax would be felt most significantly right away, while future generations would even see a decline in prices.

The resulting reduction in quantity extracted is shown in panel D of Figure 4. The panel shows the cumulative reduction in oil use up to a given year. Since production initially declines, the curves show how cumulative extraction declines, i.e., the y-values are negative. However, prices under the carbon tax rise at a slower rate and the production decline becomes less over time. As a result, the cumulative savings start to level off. Towards the end of the century, when prices under a carbon tax are lower than under the counterfactual of no carbon tax, some of the cumulative reductions will be offset through higher production, i.e., the curve bends upward. Finally, the carbon tax extends the lifetime beyond 2095, the last year of extraction under no carbon tax. The curves show an almost linear upward trend for the additional years of production, which offset the majority of the initial cumulative savings. For example, a significant carbon tax of \$200 would decrease cumulative emissions by 13% in 2080, but these savings are offset through an extension of the production period. By the end, only 4% of the cumulative emissions are avoided. We find that the reallocation of current consumption into future periods is not only a theoretical concern, but empirically relevant.

Cumulative savings are small as the combined supply curve in left panel of Figure 3 is very steep. Any oil field will eventually be extracted under a carbon tax as long as the marginal extraction cost plus the carbon tax falls below the backstop price. In other words, only oil fields with a marginal cost higher than the backstop price minus the carbon tax will find it no longer profitable to extract oil. The convexity of the supply curve implies that as the carbon tax increases, the number of oil reserves that become no longer profitable increases

non-linearly. This is shown in Figure 5. Carbon taxes that have been proposed in the past – usually in the range of \$30-100 a ton of CO₂ – are projected to have very small reductions in cumulative oil use. For example, a \$100 tax reduces emissions by 1.8%. On the other hand, increasing the tax from \$500 to \$600 would reduce emissions by an additional 28%. Significant emission reductions are required if the world is to comply with the Paris Climate Agreement. The permissible cumulative emissions are 200GT of carbon (Millar et al. 2017), which is equivalent to roughly 2.1 trillion barrels of oil.¹⁵ So if the world were to use all of the 1.8 trillion barrels of oil, it would have almost entirely used up the carbon budget. This does not count emissions from coal and natural gas, methane, etc. We will present the joined emissions from oil and natural gas in Section 2.2 below. Meeting the Paris target can only be achieved if oil consumption is significantly reduced.

We present sensitivity checks under different demand elasticities in Appendix Figure A1, where the four panels of Figure 4 are shown as the four panels in the middle row. The top row uses a lower elasticity of -0.21, while the bottom row uses -0.86. Note how the overall emission changes do not depend on the demand elasticity in the case of a carbon tax that is constant in real terms, but the time path does. A larger demand elasticity leads to temporarily larger cumulative emissions reductions as the per-period consumption drops, but these are again offset through a further extension of the time period when the resource is used. For example, under a demand elasticity of -0.86 (bottom row), the cumulative emission reductions under a \$200 carbon tax reach almost 20% instead of the 13% in our baseline using an elasticity of -0.6, but in the end only 4% less of the oil is consumed in both cases. The demand elasticity is not an important driver of our overall results.

Similarly, we present sensitivity checks for different interest rates in Appendix Figure A2, where the four panels of Figure 4 are shown as the four panels in the middle row. The top row uses a real interest rate of 2%, while the bottom row uses 4%. Again, the cumulative savings are unaltered by the choice of the real interest rate, but the time path is not. Larger interest rates imply larger temporary reductions in carbon emissions, which are then offset through an extended extraction period.

There are several qualifications to our result that neither the interest rate nor the demand elasticity impact cumulative emissions. We implicitly assume that carbon emissions are the same at all times. First, technological progress, e.g., break-throughs in carbon capture and storage, might make it possible that burning a unit of fossil fuel in the future releases less

¹⁵200 GT of carbon are equivalent to 733 billion tons of CO₂ given the atomic mass of carbon (12) and oxygen (16). Using the estimate from Table 1 that each barrel of oil emits 0.35 tons of CO₂, the equivalent amount of oil is 2.1 trillion barrels.

greenhouse emissions than if it were burned today. In such a case, temporary reductions in fuel use would have a beneficial effect on cumulative releases. Second, delaying emissions temporarily does delay temperature changes accordingly, although this effect would not be expected to be very large.

Second, another important lever that we have held constant in our analysis so far is the price of the backstop \bar{p}_b . If this backstop price is reduced (e.g., as renewables become cheaper and storage becomes available), this would be equivalent to an increase in the carbon tax. Recall that fields will be extracted if marginal cost are less than $\bar{p}_b - \tau$. Increasing the tax τ or decreasing \bar{p}_b have equivalent effects. Each \$1 tax per ton of CO₂ implies a tax of roughly 35cents per barrel, so a reduction of $\bar{p}_b = 250$ to $\bar{p}_b = 145$ for $\Delta\bar{p}_b = 105$ is equivalent to an additional \$300 carbon tax. For example, a \$100 carbon tax as well as lowering the backstop from $\bar{p}_b = 250$ to $\bar{p}_b = 145$, would be equivalent to a \$400 carbon tax. There are alternative scenarios that would combine a carbon tax with investments in alternative energy to decrease $\bar{p}_b - \tau$. However, if this investment in alternative energy is anticipated by the oil market, it would shift the consumption of oil further to the present by the “green paradox.”

2.1.2 Increasing Carbon Taxes

Several of the proposed carbon taxes start at a lower level and are then increasing over time. First, there are theoretical reasons for doing so as the marginal damage of emissions is increasing over time (Greenstone, Kopits & Wolverton 2013). Second, policy makers might also prefer to phase the tax in to give industry time to adjust.

We start with a carbon tax that equals the social cost of carbon (SCC) in Figure 6. The top panel shows the social cost of carbon estimate. Estimates for each decade from 2010 through 2050 are marked as x. We convert the numbers into \$2019 dollars, the baseline of our analysis. We fit a linear line through the data points and extend them beyond 2050, the last year in the interagency report. The mean forecast is shown in solid green line, while the 10th and 90th percentile are added as dashed lines. The mean carbon tax starts at roughly \$50 per ton of CO₂ in 2019 and almost triples to \$150 in 2100.

The remaining four panels of Figure 6 replicate the same four panels of Figure 4, showing the the time path of consumer prices, quantity consumed, share of tax paid by consumers, and cumulative reduction in oil use. The cumulative reduction in oil use is given by the final price. While the extraction period was extended to 2100 under a \$100 constant carbon tax and 2104 under \$200, the extraction period ends in 2098 in case the tax is set to equal the average SCC. The increasing tax shifts some of the consumption back to earlier periods to

avoid the higher taxes in the end. Cumulative reductions are -2% in the end, which is the same as if there would have been a constant carbon tax that equals the final SCC price in 2098, or roughly \$150 per ton of CO₂. The large consumption quantity at the beginning implies a smaller price for consumers, which in turn implies that the share of the carbon tax that is paid by consumers initially in 2019 declines from 75% to 50%. This effect can become stronger the faster the carbon tax rises. Appendix Figures A3 increases the carbon tax at a faster rate. The incidence curve, i.e., the share of the tax that is paid by consumers, becomes hill-shaped if the tax increases 4 dollars per year. Initially, less than 10% is paid by consumers as most of the tax comes out of the scarcity rent of producers. Consumers pay about a third of the tax by 2045, before the incidence reduces again below 10% in 2090. The reason is that the price rises rapidly. Recall that it combines two increasing factors: the scarcity rent, which rises at the rate of interest for resources of the same marginal cost, and the tax, which rises linearly.

Figures A4-A6 present the case of a carbon tax that rises exponentially at a rate of 2%, 3% or 3.5%. Recall that our model uses a real interest rate of 3% for the scarcity rent. Similar to a linearly increasing carbon tax, an increase at 2% reduces the initial share of the carbon tax that is paid by consumers by shifting some of the consumption back to the present. Interestingly, if the carbon tax rises at the same rate as the scarcity rent in Figure A5, the share of the carbon tax that is paid by consumers is almost flat over time and close to zero, i.e., the tax is almost fully paid for out of the scarcity rents of producers. Finally, if the carbon tax rises at a rate that is faster than the scarcity rents in Figure A6, the share of the carbon tax that is paid by consumers is initially negative and becomes positive in later periods. The intuition is the same as under the “green paradox.” The carbon tax increases so fast that it is better to accelerate consumption in the present to avoid future taxes. As a result, the change in cumulative consumption is initially positive, i.e., the tax accelerates the use of the fossil fuel rather than delaying it.

An increasing carbon tax has two countervailing effects: on the one hand it reduces the burden on consumers, which policy makers might like. The cumulative reduction in fuel use is given by the final price, so a constant carbon tax in real terms reduces oil consumption just as much as an increasing carbon tax that culminates in the same tax level, yet the latter reduces that tax incidence on consumers. On the other hand, the increasing carbon tax will raise less revenues for the government than a constant tax that equals the final price of the increasing carbon tax. We examine this further in the Section 2.1.3 on welfare effects below.

In summary, internalizing the externality by setting the carbon tax to equal the social

cost of carbon only reduces oil consumption by 2%. If the social cost of carbon is correctly calculated, our finding implies that the benefit of using fossil fuels outweighs its cost. There are however newer models that suggest that the social cost of carbon might be a lot higher (Marshall Burke, Solomon M. Hsiang & Edward Miguel 2015), which would indeed reduce carbon emissions, as might a policy that lowers the price of the backstop technology.

2.1.3 Welfare Effects

We have argued that only a sizable carbon tax, or a carbon tax together with advances in alternative energy that lower the cost of the backstop \bar{p}_b , have the potential to lower oil consumption significantly. What are the welfare consequences of various taxes? Below, we only count the direct welfare impacts in the oil market: we are not counting the externality reduction through limiting greenhouse gas emissions. We are interested in the ramifications for consumers and producers on top of that. We emphasize that aggregate welfare impacts are limited without the benefit of CO₂ reductions. Panel A of Figure 4 has shown that while consumers initially see a significant price increase, over time much of the tax is paid by producers.

Table 2 presents the net present value of various scenarios. The first row states again the cumulative amount of oil that will be extracted under various carbon taxes. It is simply the sum of all future extraction shown in panel B of Figure 4. The next three rows present producer surplus, consumer surplus, and tax revenue, all in net present value terms again assuming a discount rate of 3%. Producer surplus is the difference between the price in each period and the extraction cost as given by Rystad (recall that for undiscovered assets these include cost for exploration and development). Our backward solution gives us how much will be produced by each asset on each day over the next 100 years as well as the price. This allows us to take the simple difference and discount it. Consumer surplus is the area under the iso-elastic demand curve between the current price and the backstop of $\bar{p}_b = 250$, i.e., the surplus to consumers from having lower energy prices than under the backstop.¹⁶ We use quantity and price information from Figure 4, calculate the surplus under the iso-elastic demand curve, and discount it to 2019 with a interest rate of 3%. Finally, tax revenue is the quantity consumed times the carbon tax rate, again discounted to the present.¹⁷

First, note how for moderate carbon tax rates, e.g., up to \$100, the overall welfare impacts

¹⁶The formula for consumer surplus for the iso-elastic demand function is $\frac{\alpha}{1+\eta}[250^{1+\eta} - p^{1+\eta}]$

¹⁷Figure 3 shows the supply curve. There is a large mass at the upper bound of $c_m = \bar{p}_b = 250$. Our simulation uses oil fields where $c_m < \bar{p}_b$, and the amount extracted in Table 2 has a value of 1.79 trillion barrels under no carbon tax instead of 1.83 trillion when the reserves with zero profit are included.

are limited to at most 1.5%. This is the flip side of the fact that a carbon tax up to \$100 does not significantly reduce overall emissions, i.e., there is limited deadweight loss from taxation (again, not counting externality reductions). The roughly equal losses to producer and consumer surplus are offset by increased tax revenue. For example, a \$100 carbon tax reduces producer surplus by \$15 trillion, consumer surplus by \$14 trillion, but increases tax revenues by \$26 trillion, for a net surplus loss of less than 3 trillion.

Second, a carbon tax of \$500 would reduce carbon emissions by 27%, but expropriate most of the producers and consumer surplus. The reason is that the supply curve for oil is fairly flat for the first two thirds of oil reserves and producers find it still profitable to extract oil at much lower oil prices. At the same time, consumer prices (producer prices plus the tax) increase enough to also eliminate most of the consumer surplus. Combined producer and consumer surplus collapses from \$144 trillion to \$25 trillion, i.e., by more than 80%. This is again offset by \$91 trillion in tax revenue. The flat initial supply curve implies that significant reductions in oil use are only possible when most of the consumer and producer surplus is wiped out.

Third, increasing carbon taxes, as shown in the last three columns of the bottom panel of Table 2, do not decrease consumer surplus by nearly as much as a constant carbon tax. The extent of emissions reduction is determined by the level of the carbon tax τ_T when extraction ceases at time T . A constant carbon tax with the same reduction in fossil fuel use would have to be as high as τ_T for the entire extraction time. The carbon tax is initially lower for the case of an increasing tax, placing less of a burden on consumers, but instead taxing away scarcity rents. This is most striking for the case of a carbon tax that starts at US\$30 in 2019 and rises at 3.5% per year to US\$ 309 in 2085 when extraction ends, as shown in the last column. Figure A6 shows that the incidence on consumers is almost zero on average, and the consumer surplus in Table 2 declines only slightly from \$85.64 trillion under no tax to \$85.36 trillion under this increasing tax. Producer surplus decreases, but not as much as under a constant carbon tax that would equal the final tax rate of \$309 under the increasing case. On the other hand, tax revenue is much lower than under a constant carbon tax. The former two effects are stronger, implying that societal surplus is higher under an increasing carbon tax than a constant one that equals τ_T .

Table 3 reports producer surplus changes disaggregated by country. The reduction in producer surplus is not proportional but depends on the cost structure of each country. For example, Saudi Arabia is not only one of the biggest producers, but also has really low production cost, resulting in high producer surplus. A carbon tax of \$200 would eliminate

38% of that surplus. On the other hand, the same carbon tax would eliminate more than 50% of Canada's surplus, as the country extracts oil from high-cost tar sands, and a comparable reduction in price implies a large relative reduction in rents.

Since oil demand will likely shift significantly between countries in future years, e.g., a higher share will be consumed by developing countries, an analysis of consumer surplus by country for all future years is beyond the scope of this paper as we would have to simulate the shift in consumption. Instead we present an analysis for 2017, the last year for which the Energy Information Administration is providing data for most countries at the time of writing. Table 4 list the 25 countries with the highest decrease in overall surplus under a \$50 carbon tax, while Table 5 gives the 25 countries with the highest gains. All numbers are in billion dollars. Effects on producer surplus are split into two components. Column (1) gives the revenue effect, by multiplying the current production of each country by the decline in producer price that would result from the \$50 carbon tax. The carbon tax will drive a wedge between producer and consumer prices. While producer prices fall, consumer prices increase and demand will decrease. The drop in demand has to be matched by a drop in production. We present two counterfactuals: the first shown in column (2a) scales down the production of each country by the same relative aggregate drop in production, eliminating the reserves with the highest marginal cost in *each country*. On the other hand, column (2b) eliminates the production of the most expensive reserves *around the world*. For example, Saudi Arabia is a low-cost producer and would keep its production unchanged, while high-cost producers like Canada would reduce output by a higher ratio than the global reduction in output.

Consumer surplus changes are given in column (3), assuming the same iso-elastic demand function with an elasticity of -0.6 in each country and using 2017 consumption quantities as given by EIA. Column (4) is the tax revenue of each country, assuming that it is proportional to domestic consumption after the carbon tax is imposed, i.e., it assumes that each country imposes the same carbon tax on consumption and it is not imposed by producing countries. Columns (5a) and (5b) give the combined impact of producer surplus, consumer surplus, and the tax revenue. The difference between (5a) and (5b) is whether the producer surplus component (2a) or (2b) are used, respectively.

Intuitively, the biggest losers in Table 4 are countries that are net exporters of oil, e.g., Saudi Arabia. The drop in producer surplus is no longer offset by an increase in tax revenue, which occurs where oil is consumed. On the flip side, winners in Table 5 are generally net importers of oil, e.g., Japan, China and Germany. The increase in tax revenue more than

offsets the decrease in consumer and producer surplus.¹⁸ The tables also clearly show the high cost producers, e.g., Canada and Brazil. The producer surplus loss in column (2b) is much higher as most of a country's reserves should be shut down when the globally most expensive reserves are used to balance the implied demand reduction, while column (2a) reduces each country's output proportionally. Tables 4 and 5 is to stress that the aggregate impacts mask spatial heterogeneity.

2.2 Gas Market

We have focused so far on the crude oil market, as oil is an easily transportable commodity with a global market. Natural gas prices differ much more between regions. However, liquified gas is increasingly being shipped globally. Since we are focusing on a long-term analysis, this section assumes that oil and natural gas are substitutes, which might be more appropriate with further technological breakthroughs in the coming decades. Daily consumption of natural gas is equivalent to around 70 million barrels of oil per day.¹⁹ We count natural gas quantity in oil equivalents (barrels) when we talk about energy content. In terms of CO₂, a therm produced by natural gas releases only 75% of the CO₂ compared to crude oil on average. As a result, the carbon tax on natural gas assets is set to equal 75% of that on oil.

The supply curve including all fossil fuels in the Rystad data is shown in the right panel of Figure 3 above. The carbon equivalent of the natural gas reserves is the same as for 1.28 trillion barrels of oil. The combined CO₂ in oil and natural gas reserves equal the CO₂ of 3.1 trillion barrels of oil. Recall from above, that compliance with the Paris Climate Agreement is estimated to leave a CO₂ budget that equals 2.1 trillion barrels of oil, requiring at least a 32% reduction to stay within the budget, not even counting any emissions from coal.

The required tax per ton of CO₂ is shown as a blue line in Figure 5. It is almost identical to the case of crude oil alone, as the shape of the supply curves is comparable. A reduction of 32% requires a carbon tax of \$500 per ton of CO₂ in 2019 dollars. This level would have to be achieved at the end period of extraction to make the high-cost reserves unprofitable.

Since the world currently uses a smaller fraction of overall reserves for gas than for oil, combining the two into a single market implies a longer time until exhaustion. This is shown

¹⁸As previously mentioned, we used consumption quantity for 2017, the latest year for which EIA published demand estimates around the world at the time of writing. The United States have since become a net exporter.

¹⁹IEA lists global natural gas consumption at just under 4000 billion cubic meters. Using EIA's thermal content of 1036 BTU for one cubic feet of natural gas and 5.705 million BTU per barrel of oil, combined with the fact that there are 35.3 cubic feet per cubic meter and 365 days per year, the 4000 billion cubic meters are equivalent to 70 million barrels of oil per day based on energy content.

in Figure 7 extending into the 2120s without a carbon tax, and extended beyond 2040 with a carbon tax. Moreover, since natural gas is cheaper per unit of energy than oil, the weighted price in panel A is lower than for oil alone. The reason why we currently do not use more natural gas is that they are currently not perfect substitutes, e.g., there are a limited number of vehicles that run using gas (buses). The incidence on consumers in panel B becomes less smooth. The reason is that the carbon tax on natural gas is lower than on oil, and the denominator changes.²⁰

Sensitivity checks to the chosen demand elasticity and interest rate are given in Figures A7 and Figure A8. similar to before, the time path changes, but the cumulative emissions are not impacted.

3 Conclusions

In a static one-period framework a carbon tax is an obvious Pigouvian policy response to the global warming problem. However, the replacement of fossil fuels by alternatives will play out over several decades, which is long enough for intertemporal substitution to come into play. This is what is emphasized by the Hotelling model of extractive resource markets: equilibrium is a dynamic process not a static state. As a result, the effects of taxes are not immediately obvious. Taking the dynamics of resource use into account shows that a carbon tax may act in two ways: it can delay the consumption of a fossil fuel, leading to lower emissions of greenhouse gases at any date but the same emissions cumulatively over time. Alternatively, it may force a fossil fuel out of the market and so reduce total emissions and lead to the replacement of fossil by renewable energy. There are cases in which both of these effects will be seen, in particular the case where there are multiple grades of fossil fuel with varying extraction costs. In practice we can expect to see both effects of a carbon tax, with the balance between the two depending on how much fossil fuel has an extraction cost close to its backstop price. The latter effect is where a carbon tax will reduce fuel consumption and greenhouse gas emissions, and seems to be especially relevant for coal, which would be phased out under a carbon tax. The effect on crude oil is less clear. The remaining oil reserves are large enough that their use would release almost as much as CO₂ as the remaining carbon budget that would keep the world within 2°C. If other greenhouse gas emissions (natural gas use, methane emissions, agricultural uses, etc) are added, it becomes

²⁰Recall from the modeling section that fields are now used in order of increasing sum of the marginal extraction cost and the tax ($c_{m,i} + \tau_i$).

clear that staying within 2°C requires a reduction in oil use as well.

Applying our framework to empirical micro-level data on the MECs of crude oil suggests that a carbon tax would need to be much larger than is commonly suggested to have a significant impact on oil consumption. A carbon tax of \$100 would only reduce cumulative oil emissions by 1.8%. Some of the initial reductions in oil use are offset through an extended time of consumption. Around 70-80% of the tax will initially be passed on to consumers, but the passthrough is declining in time and even becomes negative in later years as the tax shifts oil consumption from the present to the future. In net present value terms, consumer and producer surplus in the oil market decline by equal amounts, most of which is offset by carbon tax revenues. Global welfare impacts in the oil market are limited: a carbon tax of \$100 reduces surplus in the oil market by less than 1.5%, not counting the externality of oil use. Given the convexity of the oil supply curve, significant reductions in oil use can only be achieved if most producer and consumer surplus are taxed away. An increasing carbon tax over time lower the incidence on consumers, preserves more of the consumer and producer surplus, but also generates less revenue than a constant carbon tax equal to the final price of the increasing carbon tax rate.

Another important lever when regulating oil consumption is the price of the backstop \bar{p}_b . If this backstop price becomes lower (e.g., as renewables become cheaper and storage becomes available), it would be equivalent to a carbon tax. Recall that fields will be extracted if marginal cost are less than $\bar{p}_b - \tau$. Increasing the tax τ or decreasing \bar{p}_b have equivalent effects. The result that the marginal reduction in oil use is highly convex in the carbon tax, implies equivalently that a carbon tax together with a lower backstop price (e.g., cheaper renewables) will decrease carbon emissions much more than either of the two policy levers by itself.


References

- Anderson, Soren T., Ryan Kellogg, and Stephen W. Salant.** 2018. “Hotelling under Pressure.” *Journal of Political Economy*, 126(3): 984–1026.
- Asker, John, Allan Collard-Wexler, and Jan De Loecker.** 2019. “(Mis)Allocation, Market Power, and Global Oil Extraction.” *American Economic Review*, 109(4): 1568–1615.
- Borenstein, Severin, James Bushnell, Frank A. Wolak, and Matthew Zaragoza-Watkins.** 2019. “Expecting the Unexpected: Emissions Uncertainty and Environmental Market Design.” *American Economic Review*, 109(11): 3953–3977.
- Burke, Marshall, Solomon M. Hsiang, and Edward Miguel.** 2015. “Global non-linear effect of temperature on economic production.” *Nature*, 527: 235–239.
- Cairns, Robert.** 2012. “The Green Paradox of the Economics of Exhaustible Resources.” McGill University, Department of Economics, Working Paper.
- Coase, Ronald.** 1960. “The Problem of Social Cost.” *Journal of Law and Economics*, III: 1–44.
- Colmer, Jonathan, Ralf Martin, Mirabelle Muûls, and Ulrich J. Wagner.** 2020. “Do Market-Based Mechanisms Reduce Carbon Emissions? Firm-Level Evidence from the European Union Emissions Trading Scheme.” *Working Paper*.
- Cullen, Joseph A., and Erin T. Mansur.** 2017. “Inferring Carbon Abatement Costs in Electricity Markets: A Revealed Preference Approach Using the Shale Revolution.” *American Economic Journal: Economic Policy*, 9(3): 106–133.
- Dasgupta, Partha, and Geoffrey Heal.** 1979. *Economic Theory and Exhaustible Resources*. Cambridge Handbooks in Economics, Cambridge University Press.
- Dasgupta, Parth, Geoffrey Heal, and Joseph Stiglitz.** 1980. “The Taxation of Exhaustible Resources.” In *Public Policy and the Tax System: Essays in Honour of James Meade*. , ed. Gordon Hughes and Geoffrey Heal. Harper Collins.
- Gerlagh, Reyer.** 2010. “Too Much Oil.” Fondazione Eni Enrico Mattei: Nota di Lavoro, <https://www.econstor.eu/bitstream/10419/43557/1/640265332.pdf>.
- Greenstone, Michael, Elizabeth Kopits, and Ann Wolverton.** 2013. “Developing a Social Cost of Carbon for US Regulatory Analysis: A Methodology and Interpretation.” *Review of Environmental Economics and Policy*, 7(1): 23–46.
- Hamilton, James D.** 2009. “Understanding Crude Oil Prices.” *Energy Journal*, 30(2): 179–206.

- Heal, Geoffrey.** 1976. "The Relationship between Price and Extraction Cost for a Resource with a Backstop Technology." *Bell Journal of Economics*, 7(2): 371–378.
- Hoel, Michael.** 2010. "Is There a Green Paradox?" CESifo working Paper 3168, <https://ssrn.com/abstract=1679663>.
- Hoel, Michael.** 2012. "Carbon Taxes and the Green Paradox." In *Climate Change and Common Sense: Essays in Honour of Tom Schelling.*, ed. Robert Hahn and Alistair Ulph. Oxford University Press.
- Hotelling, Harold.** 1931. "The Economics of Exhaustible Resources." *Journal of Political Economy*, 39(2): 137–75.
- Jensen, Sven, Kristina Mohlin, Karen Pittel, and Thomas Sterner.** 2015. "An Introduction to the Green Paradox: The Unintended Consequences of Climate Policies." *Review of Environmental Economics and Policy*, 9(2): 246–265.
- Lewis, Tracy R., Steven A. Matthews, and Stuart Burness.** 1979. "Monopoly and the Rate of Extraction of Exhaustible Resources: Note." *American Economic Review*, 69(1): 227–230.
- Lin, Boqiang, and Xuehui Li.** 2011. "The effect of carbon tax on per capita CO₂ emissions." *Energy Policy*, 39(9): 5137–5146.
- Metcalf, Gilbert E., and James H. Stock.** Forthcoming. "Measuring the Macroeconomic Impact of Carbon Taxes." *American Economic Review*.
- Millar, Richard J., Jan S. Fuglestedt, Pierre Friedlingstein, Joeri Rogelj, Michael J. Grubb, H. Damon Matthews, Ragnhild B. Skeie, Piers M. Forster, David J. Frame, and Myles R. Allen.** 2017. "Emission budgets and pathways consistent with limiting warming to 1.5C." *Nature Geoscience*, 10: 741–747.
- Morris, Adele.** 2013. "The Many Benefits of a Carbon Tax." In *15 Ways to Rethink The Federal Budget*.
- Murray, Brian, and Nicholas Rivers.** 2015. "British Columbia's revenue-neutral carbon tax: A review of the latest "grand experiment" in environmental policy." *Energy Policy*, 86: 674–683.
- Petrick, Sebastian, and Ulrich J. Wagner.** 2014. "The Impact of Carbon Trading on Industry: Evidence from German Manufacturing Firms." *SSRN Working Paper 2389800*.
- Pigou, Arthur Cecil.** 1920. *The Economics of Welfare*. London:Macmillan.
- Preonas, Louis.** 2019. "Market Power in Coal Shipping and Implications for U.S. Climate Policy." *Working Paper*.

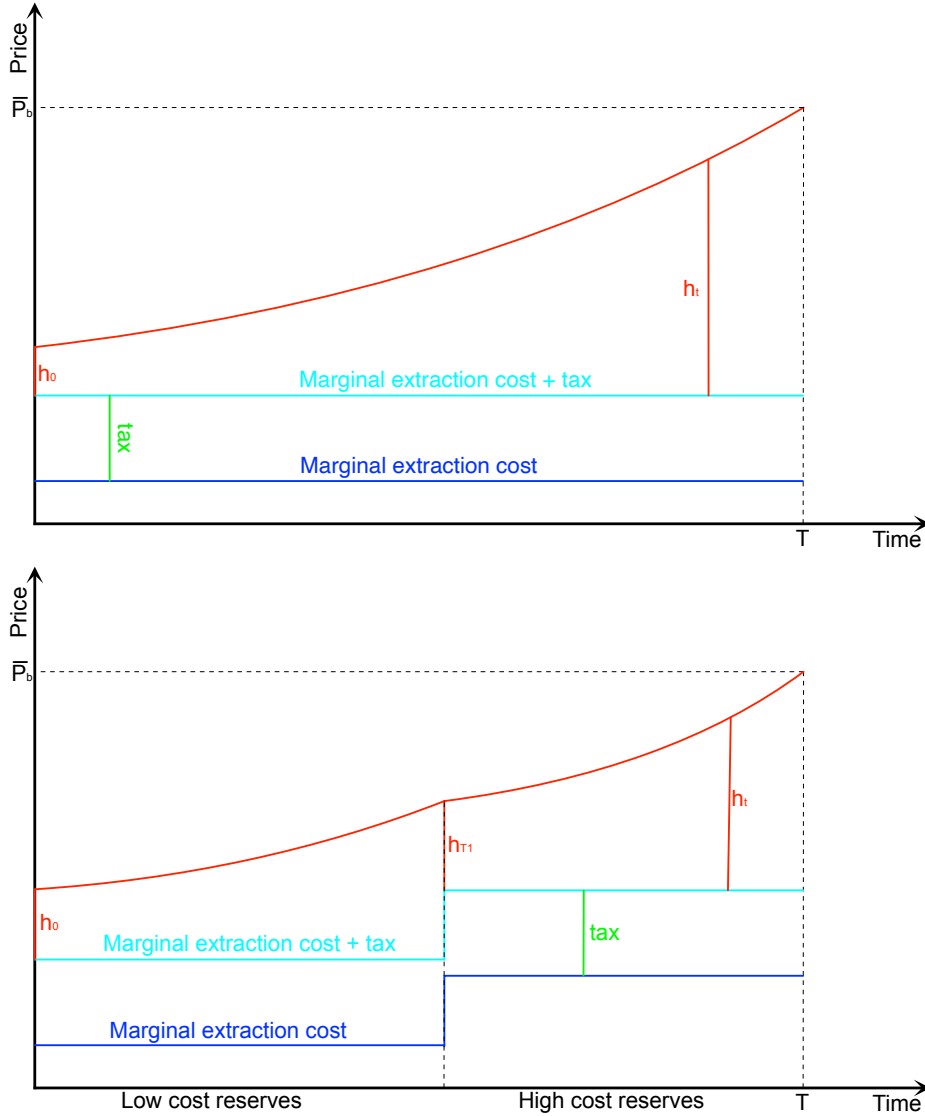
- Sapienza, Paola, and Luigi Zingales.** 2013. "Economic Experts versus Average Americans." *American Economic Review: Papers & Proceedings*, 103(3): 636–642.
- Sinn, Hans-Werner.** 2012. *The Green Paradox: A Supply-Side Approach to Global Warming*. Cambridge, Mass.:M.I.T. Press.
- Sinn, Hans-Werner.** 2015. "The Green Paradox: A Supply-Side View of the Climate Problem." *Review of Environmental Economics and Policy*, 9(2): 239–245.
- United States Government.** 2016. "Technical Support Document: Technical Update of the Social Cost of Carbon for Regulatory Impact Analysis Under Executive Order 12866." *Interagency Working Group on Social Cost of Greenhouse Gases*.
- van der Ploeg, Rick, and Cees Withagen.** 2010. "Is There Really a Green Paradox?" CESifo working Paper 2963, <https://ssrn.com/abstract=1562463>.
- van der Ploeg, Rick, and Cees Withagen.** 2015. "Global Warming and the Green Paradox: A Review of Adverse Effects of Climate Policies." *Review of Environmental Economics and Policy*, 9(2): 285–303.
- Weitzman, Martin L.** 1974. "Prices vs. Quantities." *Review of Economic Studies*, 41(4): 447–491.

Figure 1: Motivation: Taxation of Produced Good Versus Exhaustible Resource



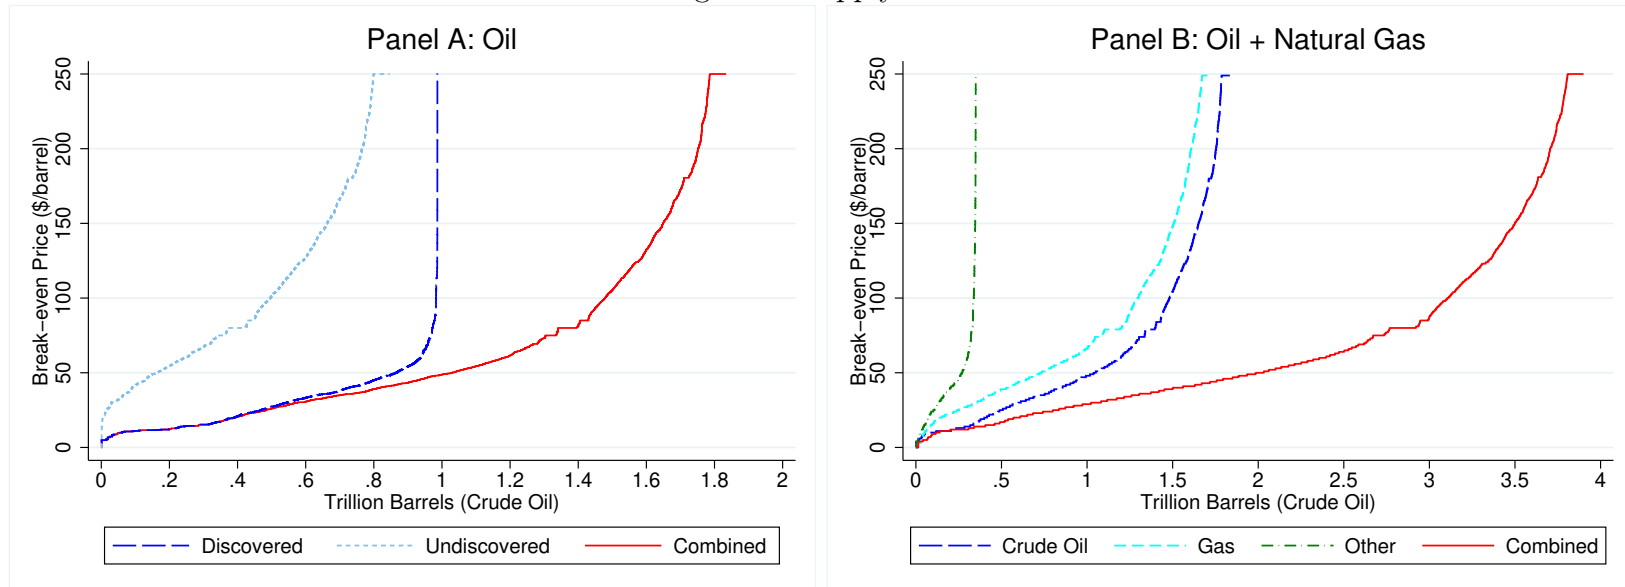
Notes: Figure displays the difference between taxing a good that is produced at constant marginal cost (top graph) and taxing an exhaustible resource with constant marginal extraction cost (bottom graph). In the top graph, the constant marginal production cost implies that the tax will reduce the equilibrium quantity from q_0 to $q_{\text{after tax}}$ and the incidence is entirely on consumers, who face a decrease in consumer surplus from CS_0 to $CS_{\text{after tax}}$. The bottom graph shows the case of a two-period model, where the length of the horizontal axis gives the overall availability of the exhaustible resource. Demand in the second period (discounted demand curve of period 1) is plotted from the right side. In this case with positive scarcity rents, the tax does not change the cumulative extraction quantity (still given by availability of the resource), but a reallocation between periods is possible. Producers see a significant decrease in scarcity rents, while consumers surplus remains largely unchanged.

Figure 2: Motivation: Homogenous vs Heterogenous Marginal Extraction Cost



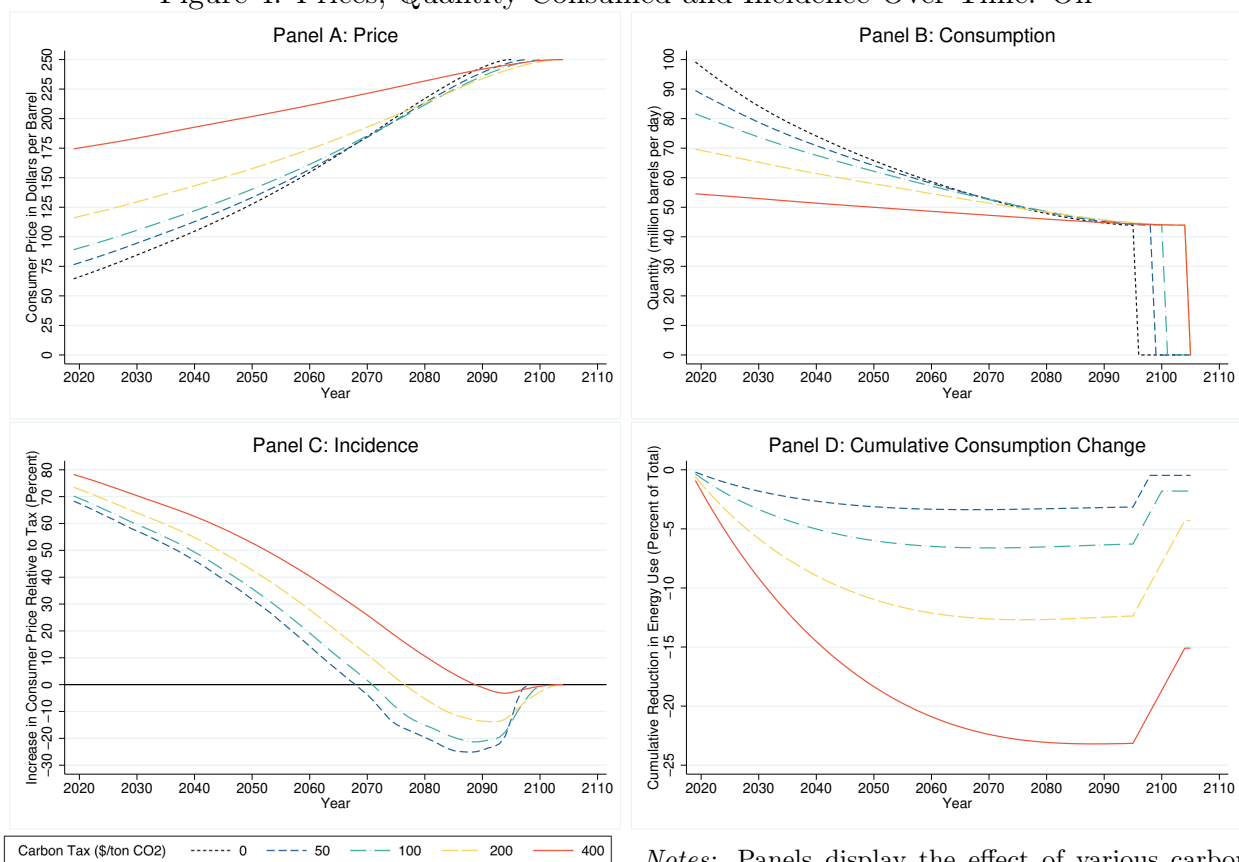
Notes: Figure displays price as well as scarcity rent $h(t)$ over time. The top graph shows the most simple case of a resource with constant marginal extraction cost, while the bottom graph shows the case when there are various fuel grades with different marginal extraction costs. Scarcity rents rise at the rate of interest when reserves from a particular grade are extracted, but jumps discontinuously when production is taken over from higher-cost reserves (which get extracted later). Prices are continuous, as there would be otherwise arbitrage opportunities by infinitesimally delaying production right before the jump.

Figure 3: Supply Curves



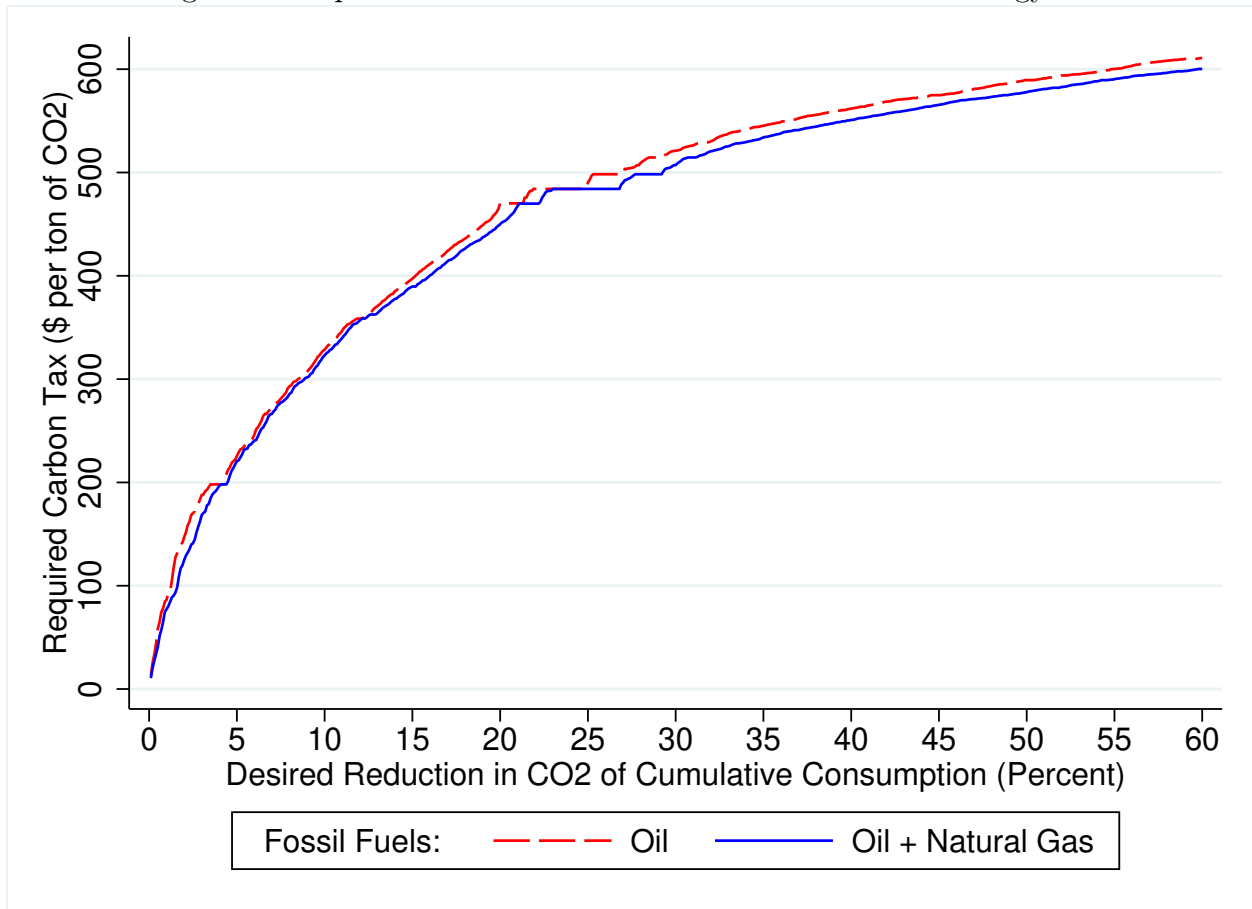
Notes: Figures display supply curves. Panel A uses only data for crude oil. It shows discovered (dark blue) and undiscovered (light blue) reserves. The red line combines the two. Break-even price for producing fields do not consider sunk exploration and set-up cost, while they are included for fields that need to be developed first. Panel B show the combined supply curves (discovered and undiscovered) for crude oil, natural gas as well as other liquids (NGL, Condensate, Unsold gas for flaring or injection). All sources of energy are counted in oil equivalents based on energy content (BTU) by Rystad. Note that natural gas has a carbon content per unit of energy of roughly 75%. Supply curves order fields from least to highest cost. The horizontal axis shows cumulative reserves.

Figure 4: Prices, Quantity Consumed and Incidence Over Time: Oil



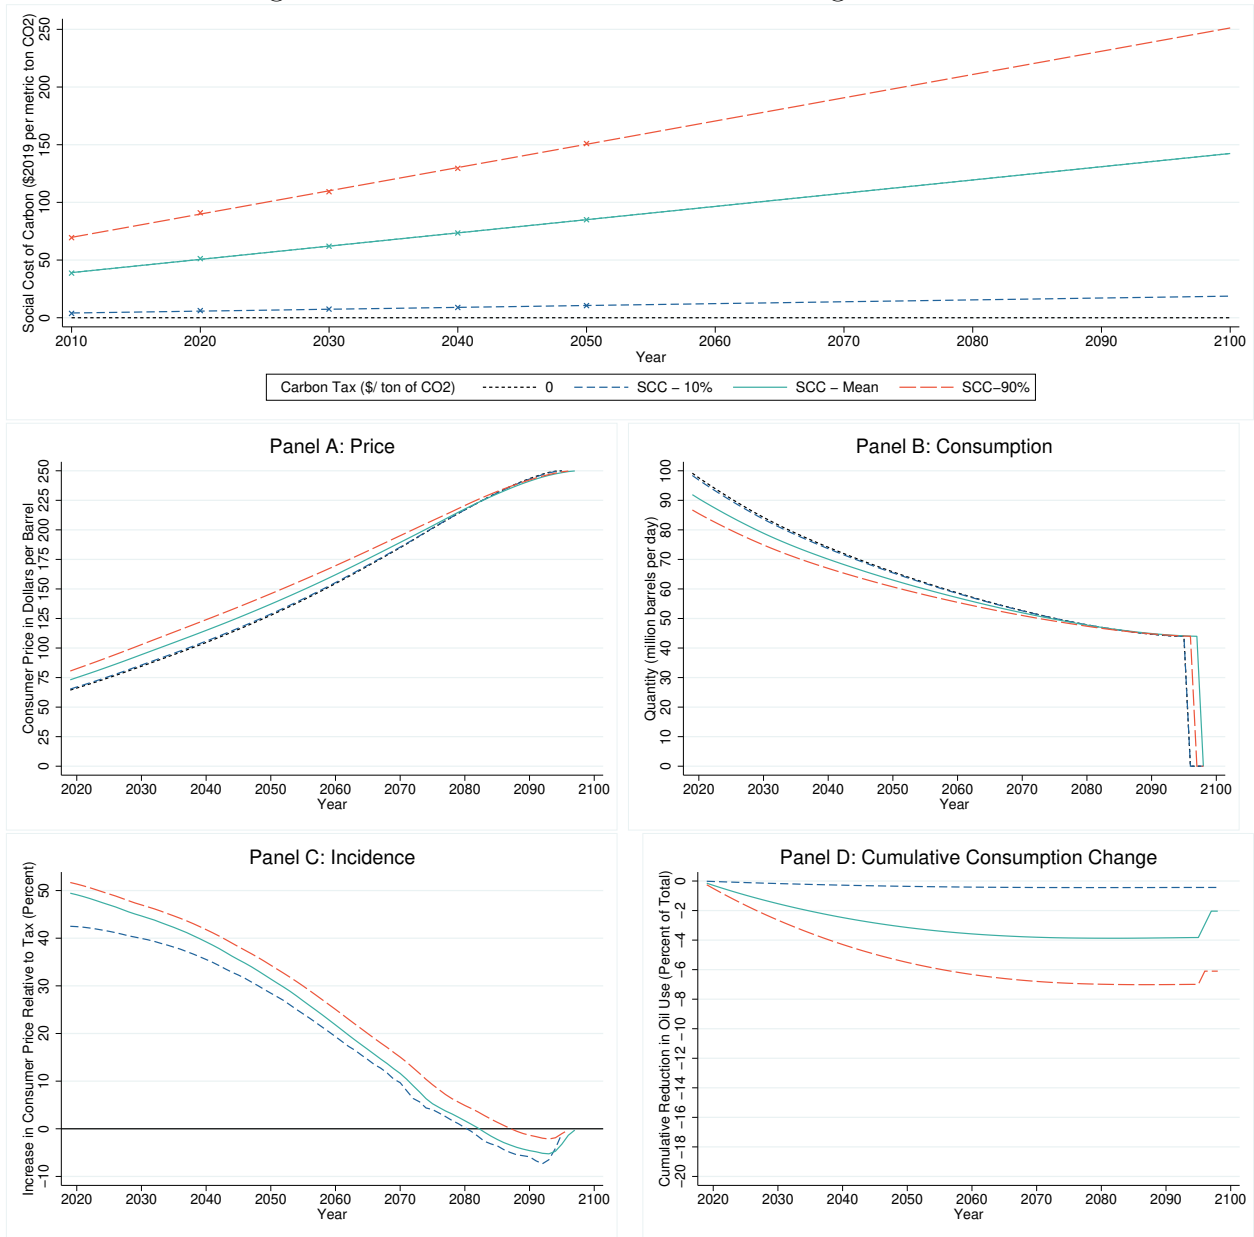
Notes: Panels display the effect of various carbon taxes. Panel A shows consumers prices and panel B the corresponding quantity consumed given an iso-elastic demand function. Panel C shows the share of the carbon tax paid by consumers, i.e., how much oil prices will be higher at each point in time compared to the case without a tax, relative to the carbon tax. Finally, panel D gives the cumulative reduction in energy consumption over time. Different colors indicated carbon taxes ranging from 50 to 400 dollars per ton of CO₂. A carbon tax of \$1 per ton of CO₂ implies a surcharge of 35 cents per barrel of oil.

Figure 5: Required Carbon Tax For Desired Reduction in Energy Use



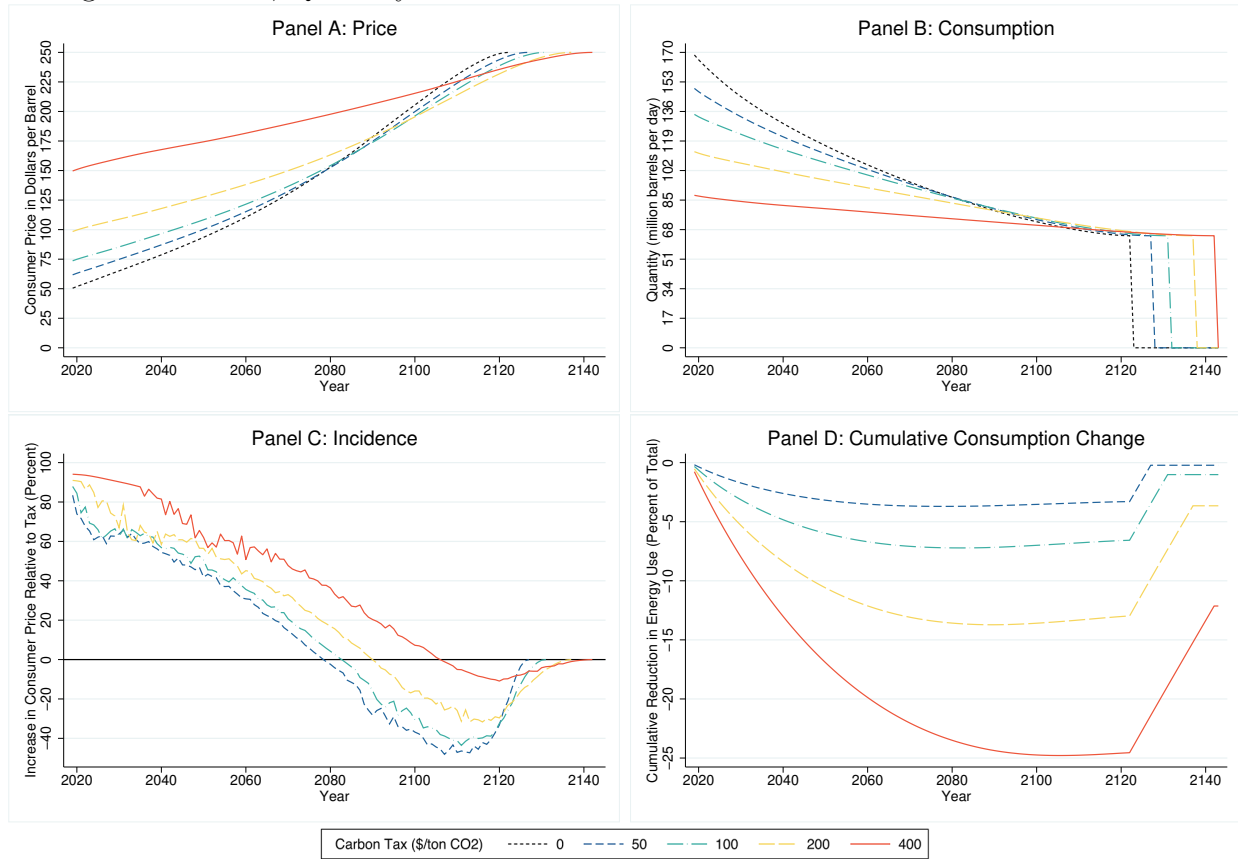
Notes: Figure displays the required carbon tax (\$ per ton of CO₂) for various desired reductions in cumulative oil use over all future years, as shown in the bottom row graph of Figure 4. Figure 4 aggregated fuel consumption based on energy content of natural gas (cubic feet of natural gas are converted to barrels based on similar amount of BTU). However, natural gas emits only around 75% of the CO₂ per BTU than oil. The current figure adjusts for this fact by multiplying carbon reductions from natural gas by 0.75.

Figure 6: Social Cost of Carbon - Increasing Carbon Tax



Notes: Graphs displays the results when the carbon tax is set to equal the social cost of carbon in real 2019 dollars. The top panel shows the social cost of carbon (mean, 10% and 90%) under a discount rate of 3%. The stars indicate the optimal level for 2010-2050. A linear time trend is fit to those values to extend them to 2100. The remaining four panels replicate the four panels of Figure 4 in case of this increasing carbon tax.

Figure 7: Prices, Quantity Consumed and Incidence Over Time: Oil + Natural Gas



Notes: Panels display the effect of various carbon taxes. Panel A shows consumers prices and panel B the corresponding quantity consumed given an iso-elastic demand function. Panel C shows the share of the carbon tax paid by consumers, i.e., how much oil prices will be higher at each point in time compared to the case without a tax, relative to the carbon tax. Finally, panel D gives the cumulative reduction in energy consumption over time. Different colors indicated carbon taxes ranging from 50 to 400 dollars per ton of CO₂. A carbon tax of \$1 per ton of CO₂ implies a surcharge of 35 cents per barrel of oil, and 26 cents per barrel of oil equivalent (based on BTU) for natural gas.

Table 1: Carbon Tax and Cost of Various Fuels

Fuel	Units	CO₂ Emissions (mt per fuel unit)	Current Price (\$ per fuel unit)	Carbon Tax (\$ per fuel unit)
Coal	mt	2.86	50	143
Gas	mmbtu	0.053	3	2.65
Oil	bbl	0.35	60	17.6

Notes: Table translates a uniform carbon tax of \$50 per ton into cost for various fuels. The first column lists the fuel type, the second column the common unit in which the fuel is measured: metric tons (mt), million BTU (mmbtu), or barrels (bbl). The third column shows the CO₂ emissions in metric tons for each unit of a fuel. The fourth column gives the current average price, while the last column shows the cost of a \$50 carbon tax on each unit of fuel.

Table 2: Simulated Cumulative Effects over all Future Years

Carbon Tax (Dollar per ton of CO ₂)	0	10	30	50	100	200
Oil Reserves Used (Billion Barrels)	1786	1785	1782	1778	1763	1710
Producer Surplus (Trillion Dollars)	58.45	56.79	53.58	50.55	43.68	32.51
Consumer Surplus (Trillion Dollars)	85.65	84.38	81.80	79.16	72.31	57.91
Tax Revenue (Trillion Dollars)	0.00	2.83	8.32	13.62	26.01	47.54
Total Surplus (Trillion Dollars)	144.10	143.99	143.70	143.32	142.00	137.96
Carbon Tax (Dollar per ton of CO ₂)	400	500	600	SCC	+2%	+3.5%
Oil Reserves Used (Billion Barrels)	1515	1303	804	1754	1755	1624
Producer Surplus (Trillion Dollars)	17.07	11.34	5.47	44.75	47.51	37.42
Consumer Surplus (Trillion Dollars)	28.68	14.86	3.57	78.37	82.60	85.36
Tax Revenue (Trillion Dollars)	80.26	91.30	88.91	20.06	13.62	19.85
Total Surplus (Trillion Dollars)	126.01	117.50	97.95	143.18	143.73	142.62

Notes: Table gives the value of all future global oil consumption, producer surplus, consumer surplus and tax revenues. First line of each panel lists the carbon tax, ranging from 10 to 600 dollars per ton of CO₂ in constant real 2019 dollars or increasing over time in the last three columns of the second panel. The tax equals the average social cost of carbon (SCC), or starts at 30 dollars and increases exponentially at 2% or 3.5% respectively). The second row of each panel gives total oil consumption over all future years. The remaining rows give the discounted net present value using a real discount rate of 3 percent. Producer surplus is the rent (price - marginal extraction cost), consumer surplus is the area under the demand curve from the current price to the backstop price of 250 dollars per barrel. Tax revenue is the quantity consumed times the carbon tax in each period.

Table 3: Discounted Net Producer Surplus over All Future Years

Carbon Tax	0	30	50	100	200	400	600	SCC	+2%	+3.5%
Saudi Arabia	10.89	10.11	9.63	8.53	6.72	4.14	1.87	8.82	9.29	7.93
United States	6.89	6.26	5.88	5.00	3.59	1.65	0.22	5.11	5.45	4.12
Russia	4.63	4.21	3.96	3.38	2.44	1.15	0.18	3.45	3.68	2.81
Iraq	3.86	3.59	3.42	3.03	2.39	1.49	0.75	3.12	3.28	2.76
Iran	3.65	3.37	3.19	2.80	2.14	1.20	0.44	2.88	3.05	2.51
Canada	2.66	2.41	2.25	1.89	1.31	0.55	0.06	1.91	2.05	1.48
UAE	2.54	2.36	2.25	2.00	1.59	1.01	0.50	2.06	2.17	1.86
Brazil	2.48	2.25	2.10	1.77	1.24	0.53	0.07	1.79	1.92	1.37
Kuwait	2.47	2.30	2.19	1.95	1.54	0.96	0.47	2.01	2.12	1.82
China	2.32	2.12	1.99	1.71	1.24	0.58	0.10	1.75	1.87	1.43
Venezuela	1.64	1.49	1.40	1.18	0.84	0.35	0.02	1.21	1.29	0.95
Mexico	1.61	1.47	1.39	1.19	0.87	0.42	0.09	1.22	1.30	1.00
Kazakhstan	1.47	1.35	1.28	1.12	0.85	0.46	0.14	1.15	1.22	1.00
Norway	0.90	0.83	0.78	0.67	0.50	0.25	0.07	0.69	0.74	0.58
Nigeria	0.77	0.70	0.65	0.55	0.39	0.16	0.03	0.56	0.60	0.44
Australia	0.73	0.67	0.63	0.54	0.40	0.20	0.04	0.56	0.60	0.47
Libya	0.71	0.65	0.62	0.54	0.40	0.21	0.07	0.55	0.58	0.46
Angola	0.45	0.41	0.38	0.32	0.23	0.10	0.02	0.33	0.35	0.24
United Kingdom	0.45	0.41	0.39	0.33	0.24	0.12	0.02	0.34	0.36	0.28
Azerbaijan	0.44	0.41	0.39	0.34	0.25	0.14	0.04	0.35	0.37	0.30
Algeria	0.43	0.39	0.36	0.30	0.21	0.08	0.01	0.31	0.33	0.22
India	0.40	0.37	0.35	0.30	0.23	0.12	0.03	0.31	0.33	0.27
Indonesia	0.39	0.34	0.32	0.26	0.18	0.08	0.02	0.26	0.28	0.20
Oman	0.30	0.28	0.26	0.23	0.17	0.09	0.02	0.23	0.25	0.19
Argentina	0.30	0.27	0.25	0.21	0.15	0.06	0.01	0.22	0.23	0.17

Notes: Table breaks the global producer surplus of all future oil production (third line of each panel in Table 2) by country and lists the 25 countries with the highest surplus under no carbon tax (column 1). Producer surplus is the rent (price - marginal extraction cost), discounted at 3 percent discount rate and given in trillion 2019 US dollars. Subsequent columns give the surplus under various carbon taxes ranging from 10 to 600 dollars per ton of CO₂. The last three columns have an increasing carbon tax, equaling either the average social cost of carbon (SCC), or starting at 30 dollars and increasing exponentially at 2% or 3.5%, respectively.

Table 4: Change in 2017 Surplus from 50 Dollar Carbon Tax - 25 Biggest Losers

	$\Delta_{ProducerSurplus}$			Δ_{CS}	Δ_{Tax}	Overall	
	(1)	(2a)	(2b)	(3)	(4)	(5a)	(5b)
Saudi Arabia	-20.03	-18.42	0.00	-13.92	19.22	-33.16	-14.73
Russia	-20.57	-16.68	-2.37	-15.41	21.29	-31.38	-17.07
Iraq	-9.11	-8.65	-0.01	-3.47	4.79	-16.43	-7.80
Iran	-7.62	-6.83	-0.07	-7.62	10.52	-11.54	-4.78
Kuwait	-5.42	-5.25	0.00	-1.87	2.58	-9.96	-4.70
UAE	-5.87	-5.21	-0.14	-3.79	5.23	-9.64	-4.56
Canada	-7.75	-3.87	-21.70	-10.13	13.99	-7.77	-25.59
Venezuela	-3.66	-2.65	-6.92	-2.06	2.85	-5.53	-9.80
Angola	-3.20	-2.48	-0.61	-0.53	0.73	-5.47	-3.60
Kazakhstan	-3.08	-2.18	-2.97	-1.33	1.84	-4.76	-5.55
Norway	-3.16	-1.80	-2.18	-0.87	1.21	-4.62	-5.01
Nigeria	-2.84	-1.58	-1.96	-1.79	2.48	-3.73	-4.11
Brazil	-5.18	-3.02	-4.58	-12.67	17.50	-3.37	-4.93
Mexico	-3.87	-2.67	-1.44	-8.45	11.67	-3.32	-2.10
Algeria	-2.09	-1.87	0.00	-1.77	2.44	-3.29	-1.42
Libya	-1.64	-1.44	-0.00	-0.84	1.16	-2.76	-1.32
Azerbaijan	-1.44	-1.14	-0.12	-0.41	0.56	-2.43	-1.41
Oman	-1.77	-0.69	-0.98	-0.71	0.98	-2.19	-2.48
Colombia	-1.62	-0.93	-2.67	-1.41	1.95	-2.01	-3.76
Qatar	-1.23	-1.01	-0.15	-1.23	1.70	-1.77	-0.91
Ecuador	-1.06	-0.97	0.00	-1.07	1.48	-1.63	-0.65
South Sudan	-0.30	-0.24	0.00	-0.03	0.04	-0.53	-0.29
Chad	-0.26	-0.23	0.00	-0.01	0.01	-0.49	-0.25
Malaysia	-1.10	-0.54	-0.98	-3.05	4.22	-0.48	-0.93
Equatorial Guinea	-0.30	-0.19	-1.60	-0.02	0.03	-0.48	-1.89

Notes: Table gives the effect of a US\$50 carbon tax for the most recent year in which EIA list consumption data: 2017. It separates overall surplus change into changes in producer surplus, consumer surplus, and tax revenues raised. Column (1) gives the change in revenue from a price decline holding output constant $q_{i0}(p_p - p_0)$. Column (2a) gives the change in producer surplus from a constant proportional change in quantity produced by all countries. Columns (2b) replicate (2a) but no longer require a proportional reduction in every country but instead retires the fields with the highest cost in the entire world. Column (3) gives the change in consumer surplus assuming a common demand elasticity of -0.6 using a countries consumption from EIA. Tax revenues are given in column (4), which are simply the after-tax consumption times the tax rate. Overall effects of proportional production adjustments are given in columns (5a), which is the sum of (1), (2a), (3), and (4). Overall effects when the globally most costly fields are retired are given in columns (5b), which is the sum of (1), (2b), (3), and (4). All numbers are in billion US\$.

Table 5: Change in 2017 Surplus from 50 Dollar Carbon Tax - 25 Biggest Winners

	$\Delta_{ProducerSurplus}$			Δ_{CS}	Δ_{Tax}	Overall	
	(1)	(2a)	(2b)	(3)	(4)	(5a)	(5b)
Czech Republic	-0.00	-0.00	0.00	-0.90	1.24	0.33	0.34
Austria	-0.03	-0.02	0.00	-1.10	1.52	0.37	0.39
Israel	-0.01	-0.01	0.00	-1.03	1.42	0.38	0.39
Morocco	-0.00	-0.00	0.00	-1.23	1.69	0.47	0.47
Greece	-0.01	-0.00	-0.00	-1.27	1.75	0.47	0.47
Indonesia	-1.39	-0.88	-0.78	-7.33	10.13	0.52	0.62
Chile	-0.00	-0.00	0.00	-1.49	2.06	0.56	0.56
Pakistan	-0.12	-0.12	0.00	-2.39	3.30	0.67	0.79
Philippines	-0.01	-0.00	-0.03	-1.81	2.50	0.68	0.65
Poland	-0.04	-0.03	-0.03	-2.75	3.80	0.98	0.98
South Africa	0.00	0.00	0.00	-2.59	3.58	0.99	0.99
Turkey	-0.10	-0.08	-0.06	-4.14	5.72	1.40	1.42
Netherlands	-0.04	-0.01	-0.08	-3.96	5.47	1.46	1.39
United Kingdom	-1.56	0.51	-3.30	-6.68	9.22	1.50	-2.31
Australia	-0.28	-0.05	-0.39	-4.90	6.77	1.54	1.20
Taiwan	-0.00	-0.00	0.00	-4.25	5.87	1.62	1.62
Thailand	-0.28	-0.12	-0.36	-5.56	7.68	1.72	1.48
Italy	-0.17	-0.06	-0.05	-5.20	7.19	1.76	1.76
Spain	-0.00	-0.00	-0.01	-5.43	7.49	2.06	2.05
France	-0.03	-0.02	-0.09	-7.26	10.03	2.72	2.64
Germany	-0.09	-0.06	-0.51	-10.25	14.15	3.75	3.31
United States	-17.21	-10.52	-14.35	-83.49	115.28	4.07	0.24
India	-1.36	-1.05	-0.28	-18.14	25.05	4.49	5.26
Japan	-0.01	-0.01	0.00	-16.42	22.67	6.23	6.24
China	-7.66	-5.05	-19.70	-56.75	78.37	8.90	-5.74

Notes: Table gives the effect of a US\$50 carbon tax for the most recent year for which EIA list consumption data: 2017. It separates overall surplus change into changes in producer surplus, consumer surplus, and tax revenues raised. Column (1) gives the change in revenue from a price decline holding output constant $q_{i0}(p_p - p_0)$. Column (2a) gives the change in producer surplus from a constant proportional change in quantity produced by all countries. Columns (2b) replicate (2a) but no longer require a proportional reduction in every country but instead retire the fields with the highest cost in the entire world. Column (3) gives the change in consumer surplus assuming a common demand elasticity of -0.6 using a countries consumption from EIA. Tax revenues are given in column (4), which are simply the after-tax consumption times the tax rate Overall effects of proportional production adjustments are given in columns (5a), which is the sum of (1), (2a), (3), and (4). Overall effects when the globally most costly fields are retired are given in columns (5b), which is the sum of (1), (2b), (3), and (4). All numbers are in billion US\$.

A1 Model

A1.1 Imperfect Substitutability

The two equations that define the equilibrium (time T and initial Hotelling rent h_0) are

$$p_T = h_0 e^{rT} + c_m + \tau = \bar{p}(\bar{p}_b) \quad (\text{A1})$$

$$Q_T = \int_0^T q(p_t, p_b) dt = \int_0^T q(h_0 e^{rt} + c_m + \tau, p_b) dt = S_0 \quad (\text{A2})$$

These equations are the same as equation (1.4) and (1.5) except that the price of the renewable resource has been replaced by the choke price, a function of the price of the renewable resource.²¹ As in the earlier case, these two equations have two unknowns, h_0 and T , and can be solved for these.

This framework leads to similar conclusions to the previous one, except that the transition from the fossil fuel to the renewable resource is now smooth rather than abrupt.

Proposition 3. *Assuming imperfect substitutability between the fossil fuel and renewable energy reflected in the demand function $q(p_t, p_b)$ with choke price $\bar{p}(\bar{p}_b)$, increasing the tax rate in a competitive equilibrium to $\tau' > \tau$, $c_m + \tau' < \bar{p}(\bar{p}_b)$ will alter the fuel consumption path, but not the overall amount of the fuel that is consumed. If the tax is so high that $c_m + \tau' > \bar{p}(\bar{p}_b)$ then the fossil fuel is never consumed.*

Proof. For all markets for the fuel to clear it is necessary and sufficient that the time T at which $p_t = \bar{p}(\bar{p}_b)$ and the initial Hotelling rent h_0 satisfy the two equations (A1) and (A2) that are analogous to (1.4) and (1.5) before.

The rest of the argument is as in Proposition 1, except that it is now possible that the fossil fuel and renewable energy are used simultaneously. \square

The important point here is that even with imperfect substitutability and the co-existence of both products in the market, a carbon tax will not affect the total cumulative consumption of the fossil fuel. The intuition is exactly as before. Renewable energy may be substituted for the fossil fuel, but this will merely spread out the consumption of the fuel over time and will not reduce total consumption. We can also show, as in Proposition 1, that an increase in the tax rate will increase T and lower the initial rent h_0 .

A1.2 Multiple Grades of Fossil Fuel - Imperfect Substitutability

We can combine the results of Proposition 3 of Section A1.1 on imperfect substitutability with those of Section 1.4 to consider the effect of taxation when there are multiple grades of fossil fuel, all of which are perfect substitutes for each other but imperfect substitutes for renewable energy, as in Section 1.2. Because the different grades are perfect substitutes for each other, they must sell at the same price, which means that only one can be on the

²¹Depletion of the fuel before its choke price is reached is inconsistent with profit-maximization.

market at any time. As in Section 1.2 there is a choke price $\bar{p}(p_b)$ for the fuel (the same for all grades as they are perfect substitutes). Now we have an equilibrium in which different grades of the fuel are exhausted sequentially from least to most expensive, with the use of some of them overlapping with that of the renewable energy source. So an equilibrium is characterized by dates T_i , $i = 1, 2, \dots, I$, $T_i < T_{i+1}$, and initial rents $h_{0,i}$, $i = 1, 2, \dots, I$ such that for all i ,

$$p_{i,t} = c_{m,i} + \tau + h_{0,i}e^{rt}, \quad T_{i-1} \leq t \leq T_i \quad (\text{A3})$$

$$\int_{T_{i-1}}^{T_i} q(p_t, p_b) dt = S_{0,i} \quad (\text{A4})$$

and continuity of prices with the last price of the fuel being its choke price:

$$m_i + \tau + h_{i,0}e^{rT_i} = p_{i,T_i} = p_{i+1,T_i} = m_{i+1} + \tau + h_{i+1,0}e^{rT_i} \quad \forall i, \quad p_{I,T_I} = \bar{p}(\bar{p}_b) \quad (\text{A5})$$

In this case the tax will lead to lower emissions at any date and to lower emissions in total over time if it displaces one or more of the expensive grades of the fuel. In Section A1.3 we look at the case of a fossil fuel whose extraction costs today are a function of cumulative extraction to date, a framework that leads to conclusions similar to those of Section 1.4: total extraction may be reduced.

A1.3 Extraction-Dependent Costs

An extension of heterogenous cost is a fuel whose extraction cost is a function of cumulative extraction to date. The motivation for such an assumption is clear: there are many grades of the resource that vary in extraction costs, and the lowest cost grades, those that are easiest to extract, are removed first, driving up costs as extraction increases. This is similar to the case considered in the last section, except that the problem is formulated in a continuously variable framework and there is an explicit dependence of current costs on past extraction, implying that current policies can alter future costs and this needs to be considered in deciding how much to extract now. We assume that the resource extraction at date t is given by $q_t \geq 0$, and that cumulative extraction is denoted $Q_t = \int_0^t q_\kappa d\kappa$. As before \bar{p}_b denotes the cost of a renewable substitute for the resource. Extraction costs at time t , $c_{m,t}$, are given as follows:

$$c_{m,t} = g(Q_t), \quad g'(Q_t) = \frac{dg}{dQ} > 0 \quad (\text{A6})$$

So the cost of extraction is given by the increasing function $g(Q_t)$ as long as it is less than the cost of the renewable resource $g(Q_t) \leq \bar{p}_b$. After this point, only the renewable resource is used. This is the formulation used in Heal (1976), and also in Hoel (2012), who also studies the effect of a carbon tax in this framework, focusing on the consequences of a tax that changes over time.

Similarly to previous setup, the resource will first be used exclusively as long as $g(0) < \bar{p}_b$. All of the resource will be used if $g(S_0) \leq \bar{p}_b$, otherwise, the resource will be used until

$Q_t < S_0$ where $g(Q_t) = \bar{p}_b$. After this point, only the backstop will be used.

As before let the carbon tax rate be τ , so that the total cost of bringing the resource to market is $g(Q_t) + \tau$. Let p be the market price of the resource and p_o the price of a generic output good produced from the resource. Then we can establish the following using the proof in (Heal 1976).

Proposition 4. *The market price of the resource in the first regime satisfies the following equation*

$$\frac{\dot{p}}{p} = \delta \left(\frac{p - g(Q_t) - \tau}{p} \right) + \frac{\dot{p}_o}{p_o} \frac{g(Q_t) + \tau}{p} \quad (\text{A7})$$

This proposition has a simple interpretation. The resource price rises at a rate which is a weighted average of the discount rate and the rate at which the output price is increasing, where the weight on the discount rate is the fraction of the price made up of rent and the weight on the rate of change of the output price is the fraction of price made up of costs. So if extraction costs are zero we have the pure Hotelling case, and if extraction costs are non-zero but constant, as in Section 1.1, the output price is constant and we have the rent rising at the discount rate. The resource price will rise according to this rule until either the resource is exhausted or the price reaches that of the renewable resource and society switches to that: if this happens before resource exhaustion then unused stocks of the resource remain.

In this context the impact of a carbon tax is easily understood: it raises the combined extraction cost and tax $g(Q_t) + \tau$. The fossil fuel will cease to be used as soon as its marginal cost including tax exceeds that of the renewable resource, i.e. as soon as

$$g(Q_t) + \tau \geq \bar{p}_b \quad (\text{A8})$$

or

$$S_0 \geq Q^* = g^{-1}[\bar{p}_b - \tau] \quad (\text{A9})$$

As g is increasing, so is g^{-1} , so an increase in the tax rate τ may reduce Q^* the level of cumulative extraction at which the fossil resource ceases to be competitive. There are two cases: if $g(Q^*) + \tau \leq \bar{p}_b$ then the tax has no impact on the amount of the fossil fuel used, as it is not sufficient to raise the extraction cost above the cost of the renewable resource. If however $g(Q^*) + \tau > \bar{p}_b$ then the tax does reduce total consumption of the fossil resource, setting a bound on cumulative extraction at \tilde{Q} where $g(\tilde{Q}) = \bar{p}_b - \tau$, $\tilde{Q} < Q^*$.

A1.4 Cap-and-Trade

We now review a cap-and-trade model in the context of a Hotelling model in more detail. We first work with a simplified version of the basic model of Section 1.1, and then consider the impact of various refinements. There is a stock $S_0 > 0$ of a fossil fuel, selling at a market price p_t at date t in a competitive market. There is no carbon tax and we take marginal extraction costs to be zero for the moment. The price satisfies $p_t = p_0 e^{rt}$ where the initial

price p_0 satisfies

$$\int_0^\infty q(p_0 e^{rt}) = S_0 \quad (\text{A10})$$

Consumption of a unit of the fossil fuel emits one unit of greenhouse gas, and an environmental authority imposes a cap of K_0 units on the total cumulative emissions of greenhouse gases. This implies that

$$\int_0^\infty q(p_0 e^{rt}) \leq K_0 \quad (\text{A11})$$

This formulation means that permits can be banked, that is carried over freely from one period to the next, so that the constraint is on total cumulative emissions and not on period-by-period emissions. Clearly one of the equations (A10) and (A11) is redundant: if $S_0 < K_0$ then the emissions constraint is redundant, and in the more likely case that the reverse is true, namely $K_0 < S_0$, some of the fossil fuel will be left unused and the binding constraint will be that $\int_0^\infty q(p_0 e^{rt}) = K_0$. In this case the scarcity rent associated with the constraint (A10) will be zero, but a positive scarcity rent will be associated with the emissions constraint (A11). So in a market equilibrium, the price of the fossil fuel will be zero but there will be a price for emissions permits. As such permits are an exhaustible resource, their price will move exactly as the price of such a resource. Letting the permit price be π_t , this will satisfy $\pi_t = \pi_0 e^{rt}$ and $\int_0^\infty q(\pi_0 e^{rt}) = K_0$. The key point to understand here is that the presence of a binding cap on emissions from the fossil fuel reduces the rent on the resource to zero and all of the rent is now captured by the permit price. So the agency that auctions permits now captures all of the scarcity rent that previously accrued to the resource owners. Financially speaking, the resource has been fully expropriated.

Now suppose that as in Section 1.1 there is a positive cost $c_m > 0$ to extracting the fossil fuel. In the absence of a cap and trade system, Proposition 1 would hold, and the rent on the resource would rise at the interest rate, with the stock of the resource being exhausted at exactly when the price first equals that of the backstop technology if there is one. But if as in the previous paragraph there is a cap and trade system with the cap on emissions tight enough that not all of the fossil fuel can be consumed, matters are again more complex. Letting π_t be as before the price of a permit at time t , in selling a unit of fossil fuel at time t the owner incurs costs of c_m to extract it and π_t to buy a permit, so that her cost is $c_m + \pi_t$. Permits are as before an exhaustible resource, so that their price will rise at the interest rate, so that the resource seller's costs move over time as $c_m + \pi_0 e^{rt}$, where the initial permit price π_0 will as before be chosen so that $\int_0^\infty q(\pi_0 e^{rt}) = K_0$. Once again the scarcity rent on the fossil fuel is reduced to zero and is replaced by the scarcity value of the emission permits, so again the fuel is effectively expropriated.

If there is heterogeneity in extraction cost $c_{m,i}$ among reserves (Section 1.4), owners of the cheaper reserves will retain some of their rents, as the price of the permit is given by reserve owner who is on the margin between producing or not producing. As we will show in the empirical section, the convexity of the marginal cost curve implies that a modest reduction in cumulative oil consumption would expropriate a significant share of the scarcity rents.

Finally, we consider a more complex case: above the emissions permits were infinitely

bankable, that is could be used at any point in time. In reality permits generally have a finite life, so we analyze the outcome in this case. To be precise, we assume that the environmental authority issues two sets of permits: one set are valid from time zero to time T , and the others from T onwards forever. Permits issued at time zero lose all value at time T , and cover in total K_0 units of emissions. The permits issued at date T cover a total of K_T units of emissions. We will take the marginal extraction cost to be zero, so that $c_m = 0$. Let q_t^* be the competitive equilibrium consumption of the fuel at date t in the absence of any policy interventions, i.e. with no cap and trade system or tax, and let $Q_0^T = \int_0^T q_t^* dt$, $Q_T^\infty = \int_T^\infty q_t^* dt$. We will for the moment take it that $K_0 = \infty$, and $K_T < Q_T^\infty$, so that there is no constraint on emissions from zero to T and the cap on emissions after T is less than would be consumed on the competitive path from that date onwards. In this situation, what is the competitive path of consumption (and emissions) from zero to T , assuming that all players in the market at date zero are aware of the cap that comes into effect at T ? The total amount of fuel available for consumption over $[0, T]$ is $S_0 - K_T$ and the competitive path is one on which just this amount is consumed over that time period. So the price path \tilde{p}_t satisfies

$$\int_0^T q(\tilde{p}_0 e^{rt}) dt = S_0 - K_T \quad (\text{A12})$$

(\tilde{p}_0 is the only unknown in this equation, which we assume to have a solution.) In this case the amount left at time T is exactly equal to the cap under the C&T system and the price of the fuel post- T will rise at the interest rate as in a competitive equilibrium. There will be a drop in consumption and a jump in the price at T , which will be fully anticipated but will not give rise to arbitrage as no fuel can be transferred from before to after T because of the cap.

Now suppose that $K_0 < S_0 - K_T$ so that the solution we have just described is not permitted. In the earlier period $[0, T]$ the permit constraint is binding, not the resource constraint. In this case the resource price will be zero and the permit price will be positive. Permits for the period $[0, T]$ are an exhaustible resource over that period, and their competitive price will rise at the interest rate from 0 to T from an initial level such that the stock K_0 of $[0, T]$ permits is just exhausted at T . Once again, the C&T system transfers value from the resource market to the permit market. After T the emissions constraint is again binding, as $K_T < S_0 - K_0$, so that again the price of the resource is zero and all scarcity rent is captured in the permit market.

A2 Empirical Deviation of Equilibrium

The iso-elastic demand function is $q_t = \alpha p_t^\eta$ and the inverse demand function is $p_t(q_t) = \left[\frac{\alpha}{q_t} \right]^{\frac{-1}{\eta}}$

The final price will either be the backstop or the choke price, whichever one is lower. We call this \bar{p}_b . Since we are solving the problem backwards, we start with $p_T = \bar{p}_b$ and solve for

prices p_t backward for $t < T$ until all reserves are extracted. The final step of this backward simulation gives the most current price and quantity, i.e., p_{2019}, q_{2019} .

Our baseline model uses a demand elasticity of $\eta = -0.6$, the average estimate of long-term elasticities in the literature (Hamilton 2009, Table 3), and sets the interest rate $r = 0.03$. We adjust the constant α of the demand function so the demand at the start of extraction process (the end of the backward simulation, i.e., corresponding to 2019) matches the observed demand quantity of 100 million barrels per day. In a first step we solve the below algorithm repeatedly until the simulated quantity we obtain from the backward simulation \widehat{q}_{2019} deviates at most 0.001 from 100, i.e., falls within $[99.999, 10.001]$. We do this by adjusting α upward if the q_{2019} is too low and vice versa until convergence occurs. Specifically, we multiply the old α by $\frac{100}{q_{2019}}$.

Below are the steps how we solve the problem backwards: We use the results from the section on heterogenous extraction cost (Section 1.4), which showed that the cheapest reserves will be extracted first and the most expensive last. Our backward induction starts with $i = I$ (most expensive reserves) down to $i = 1$ (cheapest reserves). Recall that $t = T_i$ is the time when all reserves of quality i are extracted. Since cheapest reserves are extracted first, we get $T_i < T_{i+1} < T_I$. The carbon tax is τ .

Looping over reserves $i = I, I - 1, I - 2, \dots, 1$:

1) By the continuity of prices the final price for reserves I will be \overline{p}_b . Start at step (1a) below


1a) If $i = I$: For the final reserve when we get $p_{T_I} = c_{m,I} + \tau + h_I(T_I)$. This can be solved for $h_I(T_I) = \overline{p}_b - c_{m,I} - \tau$. Go to step 2.

1b) If $i < I$: For all but the final reserve we get by the continuity of prices that at the time when reserves i are exhausted, the final price equals the new starting price of the next reserves, or $p_{T_i} = c_{m,i} + \tau + h_i(T_i) = c_{m,i+1} + \tau + h_{i+1}(T_i)$. This can be solved for $h_i(T_i) = h_{i+1}(T_i) + c_{m,i+1} - c_{m,i}$.


2) The resources rents $h_i(t)$ have to rise at the rate of interest. Since we are solving backwards in time we get $h_i(t < T_i) = h_i(T_i)e^{-rt}$ and hence prices $p_t = c_{m,i} + \tau + h_i(t) = c_{m,i} + \tau + h_i(T_i)e^{-rt}$ and quantity consumed $q_t = \alpha p_t^\eta$. We solve this on a daily time step $t = \frac{1}{365}$ and add up the daily demands until all reserves with marginal cost $c_{m,i}$ are used up. Keeping note of the number of daily time steps Δt we know that $T_{i-1} = T_i - \Delta t$. The remaining demand that could not be satisfied on the last day when reserves i are exhausted is carried over to the next reserve quality $i - 1$. If $i > 1$ go back to step (1b) and decrease i by one, otherwise go to step (3)

3) This gives us the extraction time for reserves $i = 1 \dots I$ and T_I . We renormalize time so that the current price / consumption are labeled p_{2019}, q_{2019}


Figure A1: Sensitivity to Demand Elasticity: Oil
Demand Elasticity of -0.21



Demand Elasticity of -0.6




Demand Elasticity of -0.86




Carbon Tax (\$/ton CO₂) - - - - - 0 - - - - 50 - - - 100 - - 200 - 400

Notes: Figure shows a sensitivity analysis of the four panels of Figure 4 to the chosen demand elasticity. The baseline is shown in the middle row using a demand elasticity of -0.6, while the top row uses an elasticity of -0.21 and the bottom row an elasticity of -0.86. Different colors indicated carbon taxes ranging from 0 - 400 dollars per ton of CO₂. A carbon tax of \$1 per ton of CO₂ implies a surcharge of 35 cents per barrel of oil.


Figure A2: Sensitivity to Interest Rate: Oil
Interest Rate of 2%



Interest Rate of 3%



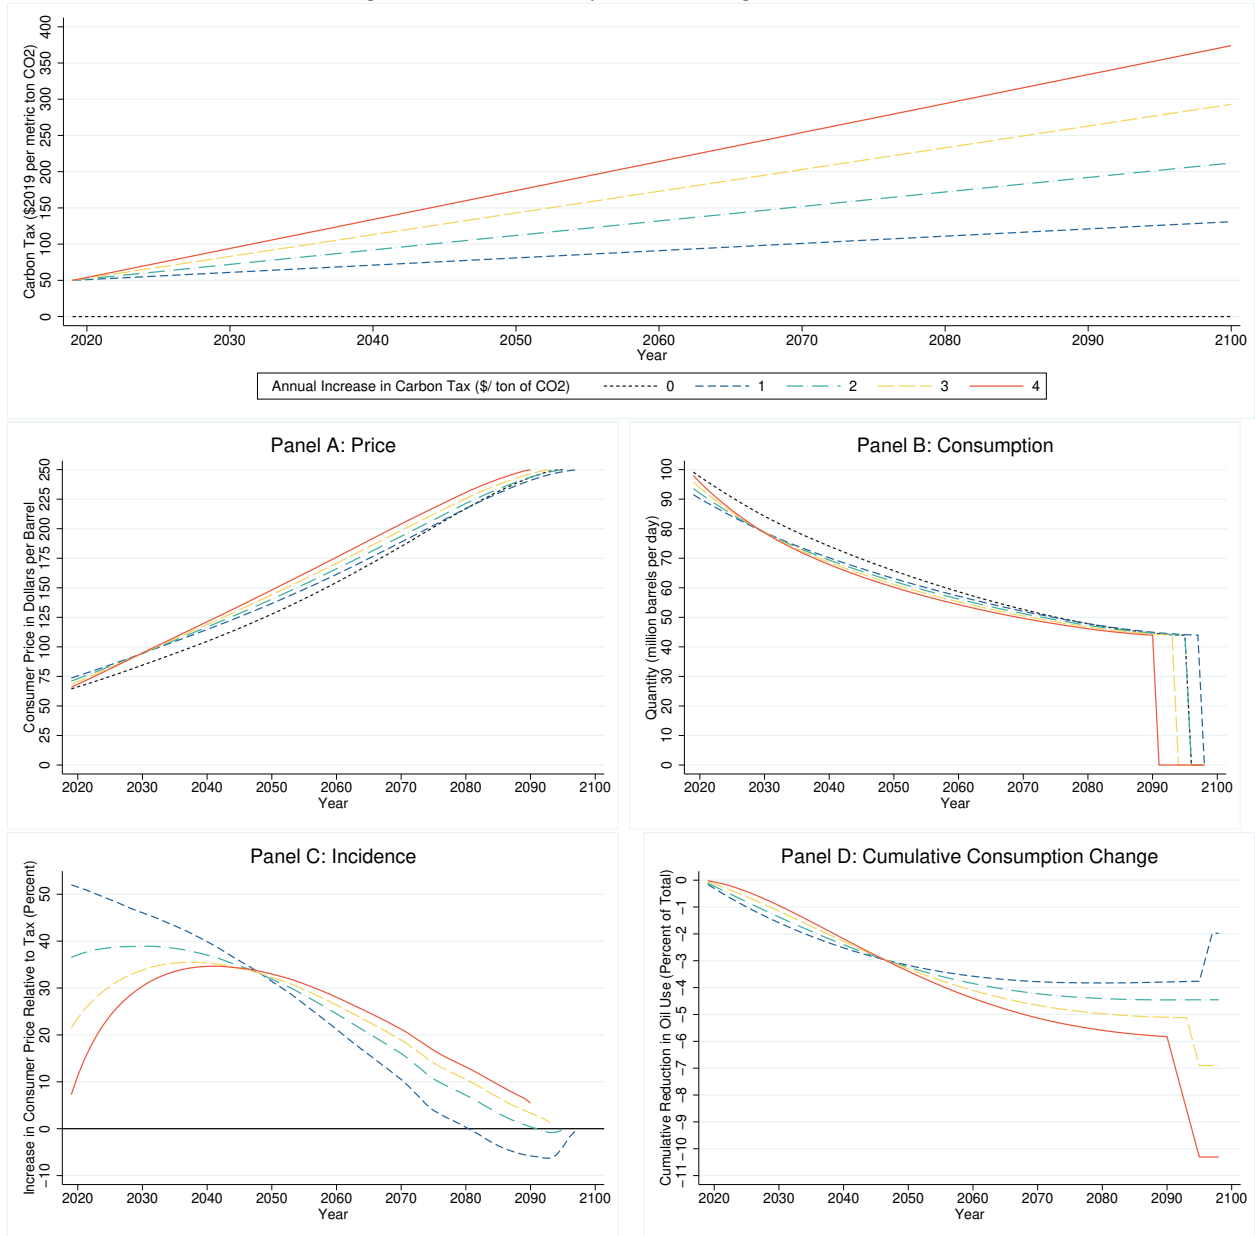
Interest Rate of 4%



Carbon Tax (\$/ton CO₂) - - - - - 0 - - - - 50 - - - - 100 - - - - 200 - - - - 400

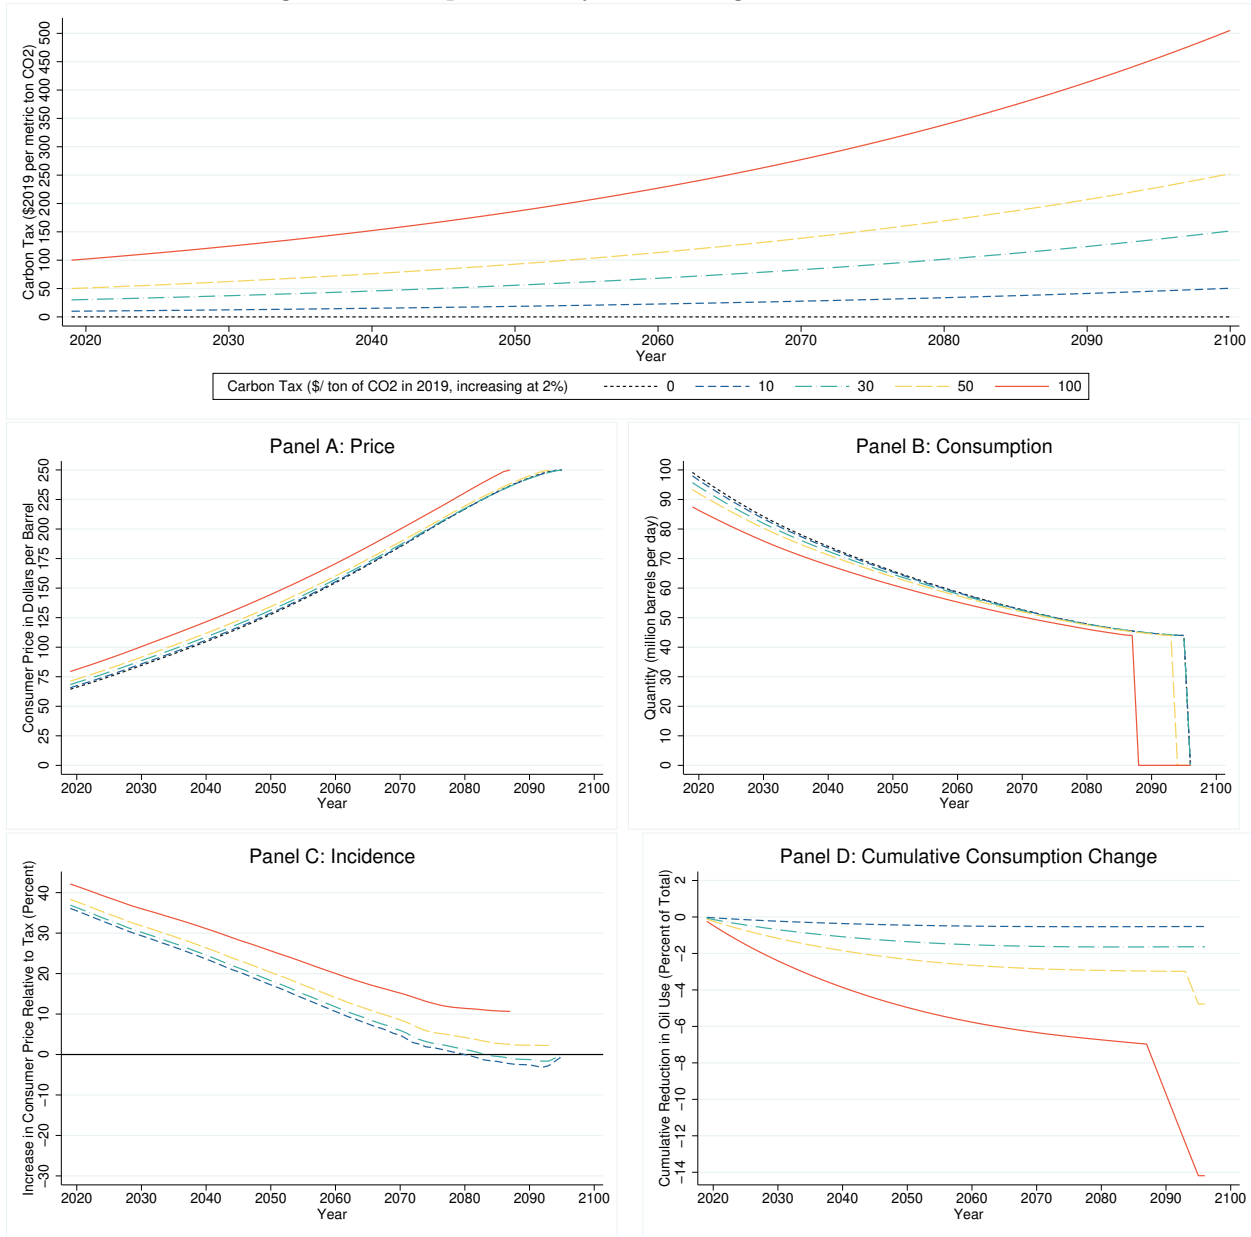
Notes: Figure shows a sensitivity analysis of the four panels of Figure 4 to the chosen interest rate. The baseline is shown in the middle row using an interest rate of 3%, while the top row uses an interest rate of 2% and the bottom row an interest rate of 4%. Different colors indicated carbon taxes ranging from 0 - 400 dollars per ton of CO₂. A carbon tax of \$1 per ton of CO₂ implies a surcharge of 35 cents per barrel of oil.

Figure A3: Linearly Increasing Carbon Tax



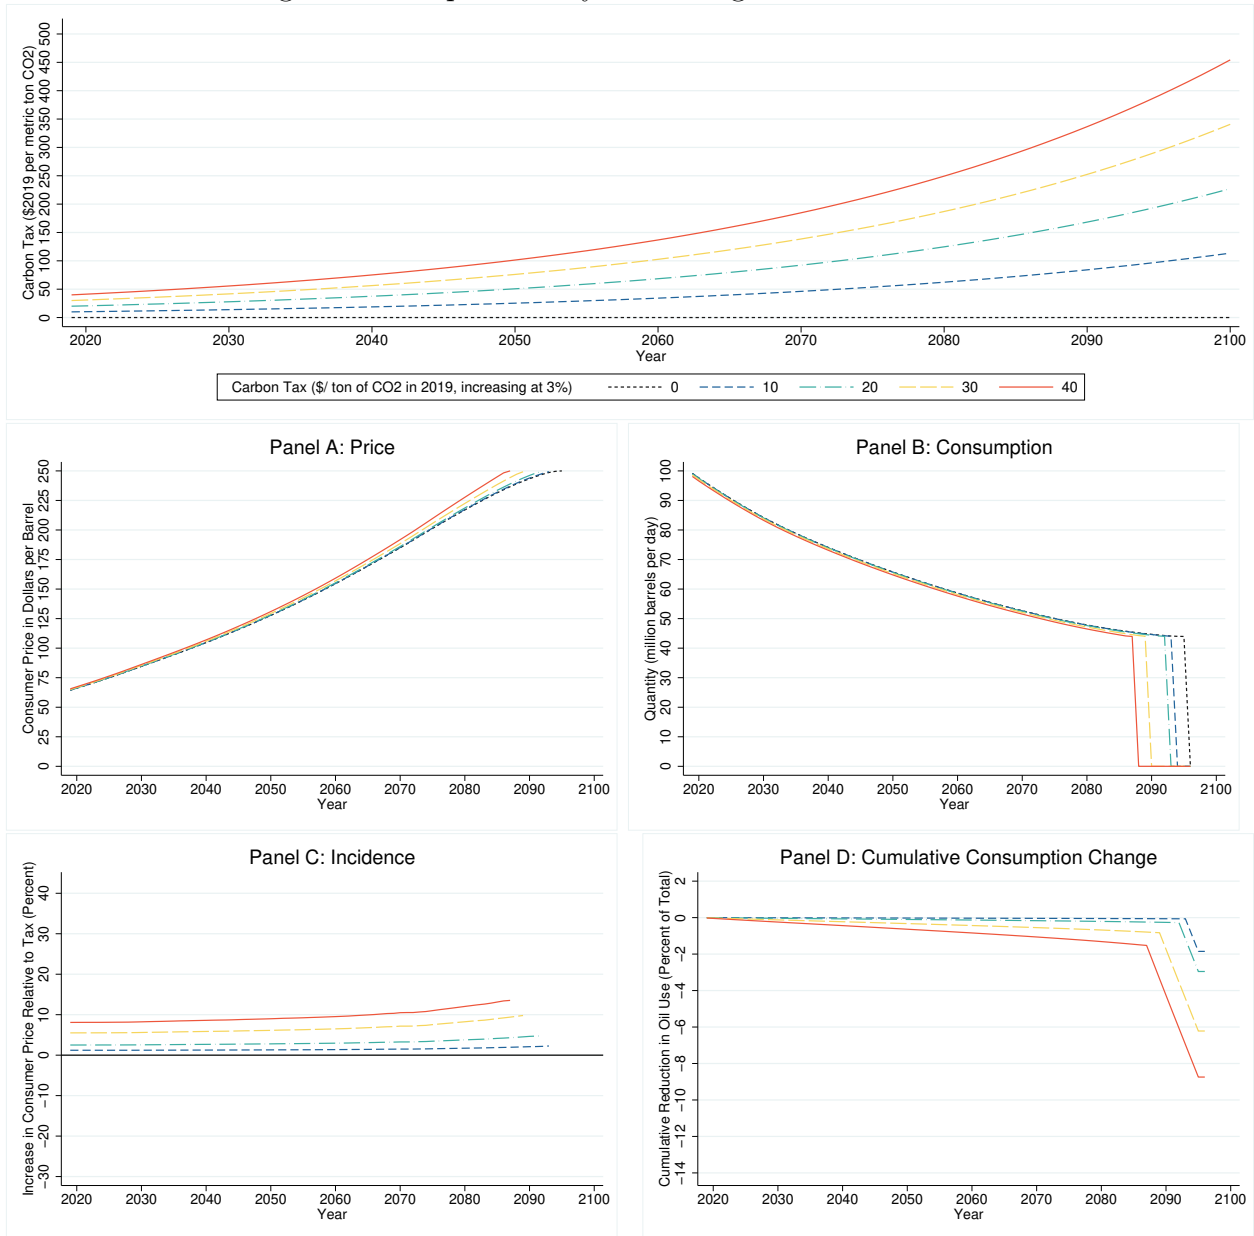
Notes: Graphs displays the results when a carbon tax of \$50 per ton of CO₂ is rising at between \$1 and \$4 per ton of CO₂ and year. The top panel shows the resulting carbon tax over time. The remaining four panels replicate the four panels of Figure 4 in case of these exponentially increasing carbon taxes.

Figure A4: Exponentially Increasing Carbon Tax at 2%



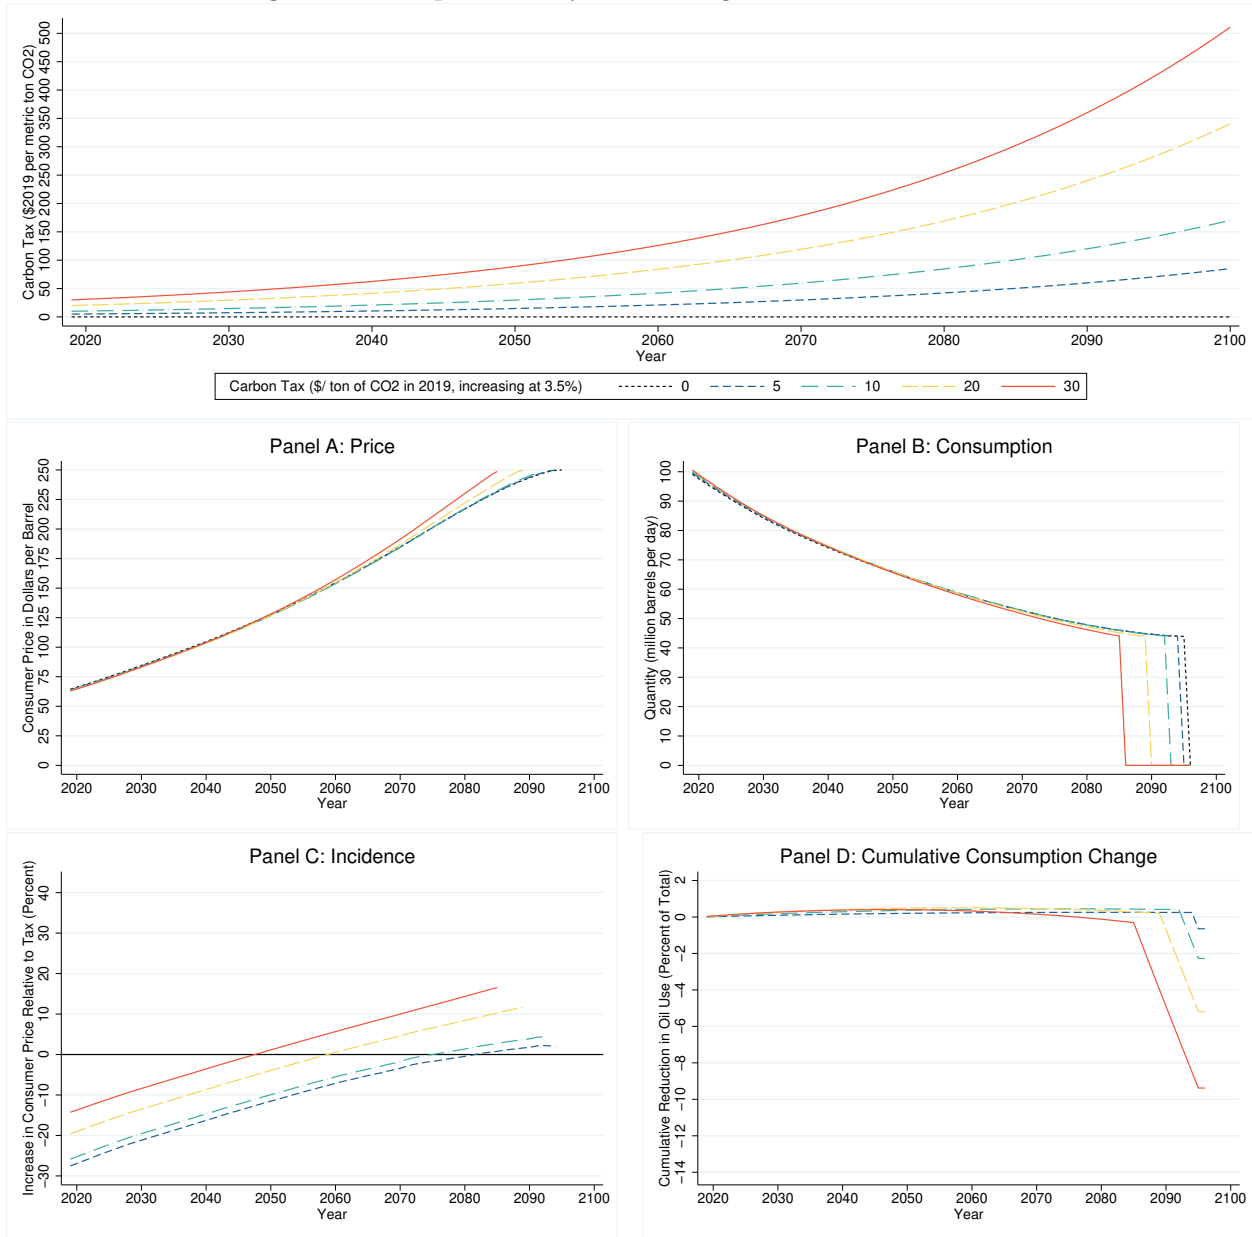
Notes: Graphs displays the results when the carbon tax ranging from \$10 to \$100 in 2019 are rising at 2%, slower than the scarcity rents at 3%. The top panel shows the resulting carbon tax over time. The remaining four panels replicate the four panels of Figure 4 in case of these exponentially increasing carbon taxes.

Figure A5: Exponentially Increasing Carbon Tax at 3%




Notes: Graphs displays the results when the carbon tax ranging from \$10 to \$40 in 2019 are rising at the same 3% as the scarcity rents. The top panel shows the resulting carbon tax over time. The remaining four panels replicate the four panels of Figure 4 in case of these exponentially increasing carbon taxes.

Figure A6: Exponentially Increasing Carbon Tax at 3.5%




Notes: Graphs displays the results when the carbon tax ranging from \$5 to \$30 in 2019 are rising at 3.5%, faster than the scarcity rents at 3%. The top panel shows the resulting carbon tax over time. The remaining four panels replicate the four panels of Figure 4 in case of these exponentially increasing carbon taxes.


Figure A7: Sensitivity to Demand Elasticity: Oil + Natural Gas
Demand Elasticity of -0.21



Demand Elasticity of -0.6



Demand Elasticity of -0.86




Carbon Tax (\$/ton CO₂) - - - - - 0 - - - - 50 - - - 100 - - - 200 - - - 400


A13

Notes: Figure shows a sensitivity analysis of the four panels of Figure 7 to the chosen demand elasticity. The baseline is shown in the middle row using a demand elasticity of -0.6, while the top row uses an elasticity of -0.21 and the bottom row an elasticity of -0.86. Different colors indicated carbon taxes ranging from 0 - 400 dollars per ton of CO₂. A carbon tax of \$1 per ton of CO₂ implies a surcharge of 35 cents per barrel of oil.


Figure A8: Sensitivity to Interest Rate: Oil
Interest Rate of 2%



Interest Rate of 3%



Interest Rate of 4%



Carbon Tax (\$/ton CO₂) - - - - - 0 - - - - 50 - - - 100 - - - 200 - - - 400

Notes: Figure shows a sensitivity analysis of the four panels of Figure 7 to the chosen interest rate. The baseline is shown in the middle row using an interest rate of 3%, while the top row uses an interest rate of 2% and the bottom row an interest rate of 4%. Different colors indicated carbon taxes ranging from 0 - 400 dollars per ton of CO₂. A carbon tax of \$1 per ton of CO₂ implies a surcharge of 35 cents per barrel of oil.