

Pump it? Market Power and the Energy Transition in the Global Oil Market

Charlotte De Cannière ^{1,2}

January 2025

JOB MARKET PAPER

[Click Here for the Latest Version](#)

Abstract

This paper explores the interaction between market power and the energy transition in the global upstream oil industry. To align with the Paris Agreement's global warming target, a significant portion of world oil reserves needs to remain untapped. At the same time, the OPEC cartel in the global crude oil market exercises market power by strategically slowing down production to inflate prices. Using detailed micro-level data on global oil production, costs and reserves, I build and estimate a dynamic structural model of global oil production in a cartel-fringe setting. The model captures the trade-off faced by the cartel between exercising market power by slowing production and avoiding future devaluation of oil by accelerating production. My findings reveal that (i) OPEC exerted market power during the 1990 - 2019 period, although the cartel only partly exploited the full potential gains; (ii) increasing carbon taxes can accelerate production and emissions—a green paradox—causing a reallocation of production under imperfect competition; (iii) well-designed carbon taxes can significantly erode market power incentives.

Keywords: *Energy Transition; Market Power; Oil; OPEC*

¹I am deeply indebted to Jan De Loecker, Mar Reguant, and Frank Verboven for providing valuable support and advice throughout this project and beyond. I thank Jacob Bradt, Allan Collard-Wexler, Sebastian Fleitas, Chris Knittel, Per Krusell, Chloé le Coq, Gerard Llobet, Carles Mañó-Cabello, Lucija Muehlenbachs, Marten Ovaere, Florian Schuett, Jo Van Biesebroeck, and participants at EARIE, EAERE, LSE environment camp, CRESSE, IFN sustainable energy conference, EWI seminar, MCEE, IO Leuven seminar and BSE energy seminars for valuable comments. Financial support from a PhD Fellowship from Research Foundation Flanders (FWO) and the SofinaBoël-FWO fellowship is gratefully acknowledged. Any errors are my own.

²KU Leuven, Naamsestraat 69, 3000 Leuven, Belgium. Email: charlotte.decanniere@kuleuven.be

1 Introduction

The 2023 United Nations Climate Change Conference (COP28) agreement aims to signify “the beginning of the end” of the fossil fuel era.¹ To comply with the 2015 Paris Agreement, a significant portion of oil reserves must remain unextracted. To keep global warming limited to 2°C above pre-industrial levels with a 50% probability, [McGlade and Ekins \(2015\)](#) find that over one-third of all 2015 oil reserves should remain untapped. In anticipation of more stringent environmental policies, resource owners can be encouraged to pump more oil in the short run, a phenomenon described as the *Green Paradox* by [Sinn \(2008\)](#).

At the same time, the oil market is characterized by an international cartel, the Organization of Petroleum Exporting Countries (OPEC), that has been documented to slow down oil extraction ([Asker et al., 2019](#)). The cartel exercises market power by creating an artificial scarcity in the market. By not producing at full capacity, the cartel manages to raise oil prices, leading to welfare loss from production misallocation and deadweight loss.² The prospect of stricter climate regulation prompts questions about how resource owners with market power will adjust their production decisions.

This paper aims to assess the welfare implications of climate policy within the context of the oil market, where an oil cartel is present. In this context, OPEC balances maintaining high prices by withholding production and exercising its market power, against the risk of future oil devaluation, which may incentivize speeding up production. To quantify this trade-off, I develop and estimate a dynamic structural model that captures the primitives of the intertemporal production mechanisms of both the cartel and fringe producers. To this end, I use global field-level data on production, costs and reserves of oil fields from 1990 to 2019. Leveraging these empirical estimates, I conduct counterfactual analyses under various climate policy scenarios to evaluate their welfare effects.

My structural model incorporates several key features of the upstream oil industry. First, as oil is an exhaustible resource, the extraction problem is inherently dynamic: production decisions today directly affect future reserves. Following the seminal framework of [Hotelling \(1931\)](#), oil producers maximize intertemporal profits, taking into account the opportunity cost of selling oil today versus in the future. Second, the model explicitly accounts for the presence of an international oil cartel, OPEC, which aims to “coordinate and unify the petroleum policies of its member countries, and ensure the

¹Source: unfccc.int/news/cop28-agreement-signals-beginning-of-the-end-of-the-fossil-fuel-era, accessed on 31 July 2024.

²This second source of welfare loss will be mitigated by the fact that environmental damages are also lower at lower levels of consumption, see [Asker, Collard-Wexler, De Cannière, De Loecker, and Knittel \(2024\)](#)

stabilization of oil markets.”³ To model cartel behavior, I extend the dynamic cartel-fringe framework of [Salant \(1976\)](#), where a forward-looking OPEC cartel and a price-taking fringe simultaneously set quantities, with OPEC recognizing its influence on equilibrium prices. However, the literature debates the true extent of OPEC’s market power ([Baumeister and Kilian, 2016](#); [Lin Lawell, 2020](#)). To address this, I introduce an internalization parameter in the model, allowing for potential frictions, such as geopolitical tensions, that may prevent OPEC from fully optimizing its production decisions.

The model relies on cost, demand, and internalization parameters. First, I leverage detailed micro-level cost data from over 20,000 oil fields worldwide, spanning 1990 to 2019, provided by Rystad Energy. This dataset enables me to estimate yearly, field-level marginal costs with minimal structure imposed. Second, I estimate global upstream oil demand using an instrumental variable (IV) approach, where the detailed marginal cost data provide cost shifters to address endogeneity in demand estimation, following the approach of [Asker et al. \(2023\)](#). Third, I estimate the internalization parameter dynamically using the simulated method of moments (SMM), conditional on the estimated demand and cost parameters in the spirit of [Gowrisankaran et al. \(2024\)](#). This internalization parameter captures the degree of OPEC’s effectiveness, driven by various market frictions, including geopolitical constraints, or other failures in perfect coordination. A lower internalization parameter signals less effective cartel behavior. Furthermore, my model allows for the analysis whether cartel collusion is pro- or counter-cyclical, building on the work of [Rotemberg and Saloner \(1986\)](#) and [Green and Porter \(1984\)](#), by permitting the internalization parameter to vary over time with shifts in oil demand.

I rely on my estimated model to analyse the welfare impact of the energy transition in the upstream oil market. The global upstream oil market is characterized by the interaction of two key market failures. The first is the exercise of market power by the OPEC cartel, which distorts production decisions. Distorting equilibrium quantities leads to welfare losses through deadweight loss. On top of that, the presence of the cartel shapes *who* is producing, and oil will not be extracted in a cost-minimizing order ([Hassler et al., 2010](#)). This leads to welfare losses from productive inefficiencies. The second market failure arises from the negative environmental externalities associated with oil production and consumption, which can lead to over-extraction depending on the balance between market power and environmental damage. In this setting, increasingly stringent climate policy can render certain oil fields economically unviable, while also disincentivizing the cartel from withholding production, leading to a reallocation across cartel and fringe producers ([Benchekroun et al., 2019](#)). As a result, rising carbon taxes can reduce both environmental damages and welfare losses from production

³See [opec.org/opec_web/en/about_us/23](https://www.opec.org/opec_web/en/about_us/23), accessed on 24 January 2024.

inefficiencies, highlighting a novel channel through which climate policy can enhance welfare. A well-designed carbon tax, therefore, has the potential to mitigate welfare losses from both market failures.

I obtain three main findings. The first finding shows that the cartel does not operate fully effectively. The dynamic structural model reveals that the cartel could be roughly five times as effective in coordinating output from 1990- 2019. This finding is in line with intuition, given the significant amount of geopolitical or other frictions that are typically present in the oil market. However, given the highly dominant position of OPEC, even an ineffective cartel results in important departures from perfect competition. Given that demand is highly inelastic – about -0.16 mean demand elasticity –, small changes in quantities can have important price impacts. Further, my results show that cartel collusion is pro-cyclical. Hence, the cartel manages to collude more effectively during a demand boom. I find that cartel collusion has experienced an overall positive trend from the 1990s onwards.

The second main finding shows that anticipation of strong environmental policy can temporarily increase emissions in both a perfectly competitive setting, as well as in the cartel-fringe setting, consistent with a green paradox. However, under imperfect competition, welfare gains emerge through a reallocation of production from high-cost fringe fields to lower-cost cartel fields. I evaluate the welfare effects of the energy transition in the oil market by introducing carbon taxes in the global upstream oil market. In my counterfactual analysis, I levy a 50 USD carbon tax per tCO_2 in 2020 – following the Biden administration’s cost of carbon in 2021⁴ – with various levels of increasing stringency. My model presents evidence of a green paradox, which is stronger the steeper the carbon tax trajectory. Short-run emissions are exacerbated because of carbon taxes, while long-run emissions are reduced. However, in the imperfectly competitive context, an increasing carbon tax mainly results in higher production levels from the cartel. This is because its incentive to withhold capacity are reduced. I find that productive inefficiencies are reduced by 30.5% when the 50 USD/ tCO_2 carbon tax increases at 5%, as compared to an equivalent flat carbon tax – an additional welfare gain on top of the reduction of environmental damages.

The third main result reports that an optimal carbon tax can reduce welfare losses from the two market failures simultaneously, i.e. in the presence of a cartel and negative environmental externalities. Assuming a social cost of carbon of 190 USD per tCO_2 , in accordance with the most recent US Environmental Protection Agency (EPA) estimates (EPA, 2023), increasing at 3% annually, I find that the initial carbon tax in a cartel-fringe setting should be slightly lower than in the first-best setting – as equilibrium quantities are already reduced by the cartel. However, the tax should rise

⁴Source: [whitehouse.gov](https://www.whitehouse.gov), accessed on 1 September 2024.

more steeply than the social cost of carbon to counteract the cartel’s market power, fostering a welfare-enhancing reallocation of production from high-cost fringe to low-cost cartel fields. For a fully effective cartel, the tax should increase at 5.5% per year, starting at over 70% below the first-best initial level. The distributional consequences, however, compared to the first-best setting are very different. If market power is involved, tax revenues are considerably lower, to the benefit of the producer surplus of both the fringe and the cartel.

The implications of my findings are twofold. First, they suggest that tightening carbon taxes can help reduce the welfare losses from market power of the oil cartel, providing a pathway toward increased welfare benefits during the transition away from fossil fuels. Second, the results underscore the importance of considering the industrial organization of the oil market when designing environmental policies. The anticipated competition for remaining oil resources, or the “race to the last barrel”, indicates that sustainable energy alternatives may face stronger-than-expected resistance. Thus, policymakers should account for these market dynamics to effectively promote the transition to cleaner energy sources.

Related Literature This paper relates to three main strands of literature. First, this paper contributes to the literature on the interaction between environmental regulation and other market distortions. While corrective environmental policies are often analyzed for their welfare effects in markets with a second market failure, such as market power (e.g. [Buchanan, 1969](#); [Liski and Montero, 2005](#); [Fowle et al., 2016](#); [Reynaert, 2021](#); [Borenstein and Kellogg, 2022](#)), most studies on the welfare impact of environmental regulation in resource markets – particularly those examining the green paradox – assume perfect competition ([Sinn, 2008](#); [Gerlagh, 2011](#); [van der Ploeg and Withagen, 2012b](#); [Heal and Schlenker, 2019](#)). Also empirical assessments of the green paradox, such as [Lemoine \(2017\)](#) on coal markets, have largely focused on competitive settings.⁵ [Coulomb et al. \(2021\)](#) empirically evaluates carbon misallocation in the oil sector due to emission heterogeneity under inelastic demand. A notable exception on the perfectly competitive modelling are [Hassler et al. \(2010\)](#) and [Benchekroun et al. \(2019\)](#), who theoretically analyse the effect of backstop entry in the oil market. The closest related work is [Kellogg \(2024\)](#), who calibrates the effects of demand reductions on production levels in a cartel-fringe setting, incorporating investment decisions. In contrast, I quantify the welfare effects of

⁵Other literature quantifies the historical environmental damage of market power by comparing actual extraction to a perfectly competitive benchmark. Market power could frontload the extraction of more emission intensive fields [Asker et al. \(2024\)](#) highlights the opposing welfare effect of slower overall extraction under market power. In contrast to my work, they do not model oil extraction in the presence of market power.

environmental policy in the presence of market power using an empirical dynamic structural model of the global oil market, drawing on micro-level production and cost data. My work identifies and measures a novel channel through which environmental policy influences strategic players: increasingly stringent climate policies can erode the strategic incentives of oil cartels, resulting in welfare gains through production reallocation.

Second, there is a rich game-theoretic literature modeling resource extraction in an intertemporal framework. Earlier work primarily focused on perfect competition or monopolistic settings where a producer optimizes extraction over time (e.g. [Hotelling, 1931](#); [Stiglitz, 1976](#); [Dasgupta and Heal, 1974](#)). [Salant \(1976\)](#) was the first to introduce a more complex market structure, considering a cartel with a competitive fringe.⁶ [Anderson et al. \(2018\)](#) noted that empirical applications of these theoretical models have been limited in the context of micro-data on production costs and reserves.⁷ Their work was the first to bridge theory and data by reformulating Hotelling’s intertemporal trade-offs while focusing on Texan oil wells, avoiding the need to account for strategic production behavior. [Ahlvik et al. \(2024\)](#) study the environmental impact of supply-side taxation through changes in oil extraction, discoveries and exploration. In contrast, [Lin Lawell \(2020\)](#) built a structural model incorporating firm-level market power with aggregate data, providing evidence of OPEC’s strategic behavior. [Bornstein et al. \(2023\)](#) developed a general equilibrium model of the oil market, accounting for the impact of fracking and OPEC’s imperfect control over extraction rates. [Asker et al. \(2023\)](#) applied Hotelling’s rule to global oil well data to quantify welfare losses from market imperfections but relied on comparing actual behavior to a perfectly competitive benchmark without explicitly modeling the cartel. This paper differs by structurally modeling the cartel’s dynamic production decisions. Unlike previous work, I account for the possibility that production decisions may deviate from fully profit-maximizing behavior. Hence, to my knowledge, this paper is the first to empirically quantify a internalization parameter for the OPEC cartel. Further, I allow that degree of collusion within the cartel to vary over time, depending on the market conditions ([Green and Porter, 1984](#); [Rotemberg and Saloner, 1986](#); [Slade, 1992](#); [Igami and Sugaya, 2022](#)).

Third, this paper contributes to the literature on quantifying market inefficiencies stemming from

⁶This literature was later expanded by works like [Ulph and Folie \(1980\)](#) and [Lewis and Schmalensee \(1980\)](#) that incorporated cost heterogeneity, or [Pindyck \(1978\)](#) that quantified the gains from cartelization across commodities. Further, it is worth noting that a parallel theoretical perspective emerged, positing a first-mover advantage for the cartel, such as [Gilbert \(1978\)](#), [Newbery \(1981\)](#), or [Groot et al. \(1992\)](#). However, a limitation of this body of literature lies in the more limited tractability of equilibria, and dynamic inconsistency issues ([Groot et al., 1992](#)).

⁷Empirical work testing predictions of the Hotelling model do exist, such as [Hamilton \(2009\)](#).

market power. In energy markets, market power often manifests through capacity withholding. Building on the work of [Borenstein et al. \(2002\)](#), which introduced the idea of measuring market frictions by comparing observed production and costs to a benchmark allocation, this methodology has been widely applied in electricity markets to assess inefficiencies ([Puller \(2007\)](#); [Mansur \(2008\)](#); [Cicala \(2022\)](#); [Reguant \(2014\)](#)). In the oil market, studies like [Asker et al. \(2019, 2023\)](#) have similarly found significant inefficiencies due to market power. Unlike these papers, I empirically estimate a structural dynamic model of cartel and fringe production decisions, explicitly modeling the drivers of strategic behavior in this market. This approach enables me to perform counterfactual analyses of environmental regulations that incorporate the cartel’s market power, providing a deeper understanding of the implications on market efficiency of potential (environmental) policy interventions, such as an optimal carbon tax.

The remainder of the paper is organized as follows. Section 2 introduces the analytical framework of market power in the oil market and formalizes the cartel-fringe extraction model. Section 3 provides an overview of the data and industry background. In Section 4, I outline the empirical implementation of demand and marginal costs and provide discussion of the results. Section 5 presents the dynamic estimation of the internalization parameter. Section 6 goes into the counterfactual analyses, and Section 7 concludes.

2 Analytical framework

In this section, I begin by introducing a static framework to illustrate how market power is exercised in oil markets and explore its welfare implications when combined with another market failure—environmental externalities. Second, I formalize a dynamic structural model of resource extraction for a cartel, with a price-taking competitive fringe. Third, building on this dynamic model, I go into the impact of the energy transition on the extraction problem, with a particular focus on the role of carbon taxes in shaping overall emissions and the pace of emissions from oil extraction. I then discuss the welfare implications of environmental policies within the context of oil extraction.

2.1 Market power in oil markets: static analysis

The presence of the Organization of the Petroleum Exporting Countries (OPEC)⁸ cartel is a notable feature of the oil industry. Although the cartel’s effectiveness is often debated ([Baumeister and Kilian,](#)

⁸For a more in-depth description of the OPEC cartel, I refer to Section 3.

2016; Lin Lawell, 2020), it has been found to withhold part of the production capacity, in order to raise oil prices (Asker et al., 2019).

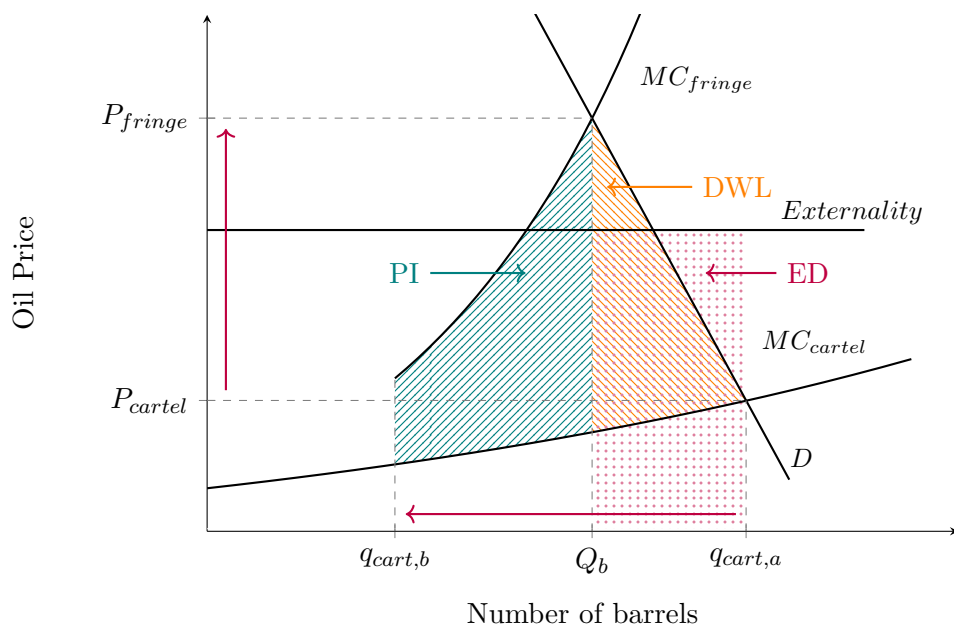
Figure 1 illustrates how a cartel can exercise market power for a homogeneous product in a static setting. Consider a scenario where the cartel has low marginal costs, MC_{cartel} , while the fringe firms in the market face much higher and more convex marginal costs, MC_{fringe} . Facing demand for oil D , in a perfectly competitive market, the cartel could satisfy the full demand by producing $q_{cart,a}$, leading to a market price P_a . The cartel, however, can increase its profits by withholding part of its production capacity, reducing output to $q_{cart,b}$. As a result, higher-cost fringe firms enter the market to meet the remaining demand, bringing the total quantity supplied to Q_b . This benefits the infra-marginal fields of the cartel, and raises its overall profits. In the new equilibrium, clearing prices are considerably higher, increasing from P_a to P_b .

The cartel's strategic behavior leads to three welfare implications when compared to a perfectly competitive market: changes in (i) productive inefficiencies (PI), (ii) deadweight loss (DWL), and (iii) environmental damages (ED). These three effects are shown in Figure 1. First, productive inefficiencies (PI) represent the additional overall production costs incurred due to the cartel's market power, which raises production costs above what they would be in a competitive market. As illustrated in Figure 1, without the cartel's strategic interventions, the demand could have been met by producing $q_{cart,a}$ at a lower production cost of MC_{cartel} . However, the cartel's actions result in the exploitation of more expensive fringe fields. The welfare loss due to productive inefficiency is the difference in production costs between these two scenarios, represented by the shaded area (PI) in Figure 1 between the fringe marginal cost curve MC_{fringe} and the cartel marginal cost curve MC_{cartel} . As noted by Hassler et al. (2010), productive inefficiencies are a result from the fact that market power in the oil sector violates the 'Herfindahl rule'. This rule dictates the principle that cheaper resources are extracted before the more expensive ones, in an exhaustible resource setting (Herfindahl, 1967).

Second, the cartel's strategic behavior results in welfare losses due to the creation of deadweight loss (DWL). As shown by the orange shaded area labeled DWL in Figure 1, the exercise of market power raises prices, which in turn reduces the quantity of oil demanded. This decline in demand generates a classic Harberger triangle, capturing the foregone gains from trade typically associated with market power.

Third, both the consumption and production of oil is associated with negative externalities, indicated by the line labeled 'Externality' in Figure 1. In my illustration, environmental externalities are assumed to be homogeneous across oil fields. A third welfare implication of the cartel's strategic behavior is a reduction in environmental damages (ED), which results in a welfare gain. As shown

Figure 1: Welfare Implications of Market Power in the Oil Sector Compared to Perfect Competition



Notes: A static illustration of how market power is exercised by the cartel in the oil market. The cartel, characterized by lower marginal costs MC_{cartel} , can exercise market power by reducing its production capacity from $q_{cart,a}$ to $q_{cart,b}$. This results in the entry of fringe oil fields, characterized by higher marginal costs MC_{fringe} , resulting in total production quantity Q_b . The cartel increases its profits by creating an artificial scarcity, inflating prices from P_a to P_b . Strategic withholding by the cartel results in three welfare changes. First, production allocation changes due to strategic production decisions, leading to a welfare loss from productive inefficiencies (PI). Second, the lower equilibrium quantity results in a deadweight loss (DWL). Third, at the same time, these lower quantities lead to lower emissions, resulting in decreased environmental damage (ED). The overall welfare effect of market power in the global oil market is ambiguous.

by the purple shaded area labeled ED in Figure 1 this gain arises from the decrease in externalities associated with lower oil consumption and production. Thus, by withholding production capacity, the cartel indirectly produces a welfare benefit through reduced emissions. This aligns with the argument made by [Dasgupta and Heal \(1974\)](#) that “the monopolist is the conservationist’s best friend”.⁹

When considering these three effects together, the overall welfare impact of the cartel’s market power is ambiguous. The net effect on welfare depends on the balance between productive inefficiencies, deadweight loss, and the value placed on environmental damages. If environmental damages are substantial, the welfare gains from reduced extraction may outweigh the losses from decreased productive efficiency and output. Conversely, if environmental damages are relatively minor, the implicit carbon tax imposed by higher prices under market power could exceed the socially optimal level. Additionally, the distributional effects of cartelization diverge significantly from a competitive market, as the economic rents accrue disproportionately to cartel members rather than being passed on to consumers as surplus in a more competitive environment.

2.2 Dynamic model of oil extraction: set-up

In this subsection, I describe a dynamic model of oil extraction in a cartel-fringe setting. Since tomorrow’s oil reserves depend on today’s production decisions, a dynamic setting is appropriate. I first formally derive the profit maximization problem of a cartel, complemented by a price-taking fringe firm in a dynamic context, following [Salant \(1976\)](#). Next, I provide the necessary equilibrium conditions for the cartel and fringe. A detailed derivation of the model can be found in Appendix B.

Model overview To model the dynamic extraction problem, I employ the framework of [Salant \(1976\)](#) with a coherent cartel and a large group of price-taking fringe firms. My modelling approach is motivated by the overarching goal of capturing extraction decisions in a parsimonious way, while being able to capture the main sources of welfare impact – deadweight loss, productive inefficiencies, and environmental outcomes – of market power. While internal stability within the cartel could influence production decisions, this parsimonious model does not explicitly address such dynamics.

The cartel owns a set of oil fields $i \in \Omega$. Each field $i \in \Omega$ is associated with constant marginal production costs c_{it} in year $t \in \mathbb{T}$, and an initial oil reserve $R_{i,0}$. The fringe fields $j \in \mathbb{J}$ – all individual price-takers – have constant marginal extraction costs c_{jt} and, are endowed with initial reserve level

⁹Note, however, that in a dynamic setting, market power merely slows down extraction. Ceteris paribus, all resources will still be exploited eventually, as noted by [Hotelling \(1931\)](#). Therefore, the cumulative emissions from oil will be the same as in the perfectly competitive scenario. Further details are provided in Section 2.3.

$R_{j,0}$. The cartel and the fringe simultaneously choose their production. Each fringe field takes the price path as given and chooses a path of extraction, whereas the cartel takes the aggregate production path of the fringe as given and determines a price path, and thereby its supply. The cartel realizes the impact of its production levels on equilibrium prices. I assume a linear demand curve $D_t(P_t)$ with an exogenous choke price α_t . The exogenous discount rate is denoted β , and the same for all players.

The cartel's problem is to maximize its net present value of profits over time by picking its optimal sales path. I assume the cartel knows the linear inverse consumer demand curve $P_t(Q_t)$ in each year $t \in \mathbb{T}$, and the aggregate sales path of the competitive fringe \bar{Q}_t^f . Since the cartel takes the aggregate production path of the fringe as given, it can construct a path of excess demand curves $(D_t - \bar{Q}_t^f)$. This parallel shift of the market demand curve indicates the amount of oil demanded, given the fringe supply path. The bar over the forecasted fringe quantities indicates that it is taken as given by the cartel. The cartel will then pick extraction levels q_{it} for each field $i \in \Omega$ in each year $t \in \mathbb{T}$, such that its net present value of profits are maximized while ensuring that it does not deplete its reserves in any given field.

More formally, the cartel's optimization problem is presented in Equations 1a - 1f. Equation 1a outlines the objective function of a coherent cartel dynamically optimizing yearly profits by choosing the production levels of each field q_{it} , subject to several constraints. The constraints shown in Equations 1b - 1c show the reserve constraint faced by the cartel for each of their fields $i \in \Omega$. Tomorrow's reserves depend on today's reserves, current production levels, and future discoveries.¹⁰ Total reserves of field i at time t , R_{it} , should remain non-negative. Next, Equation 1d captures the yearly flow constraint per field. This represents the idea that production levels are limited by well pressure, as documented by Anderson et al. (2018). Hence, resource owners can only extract fraction a κ_{it} of field reserves each year.

Constraint 1e - 1f show that the cartel optimizes with respect to the excess demand it is facing, defined as total demand $D_t(P_t)$ subtracted by the anticipated aggregate fringe production \bar{Q}_t^f . Total cartel quantity across all fields is denoted by Q_t^c . Cumulative extraction along the optimal path will always equal the total amount of economically recoverable reserves. Reserves are economically recoverable as long as their extraction costs are lower than the choke price, i.e. $c_{jt} \leq \alpha_t$ or $c_{it} \leq \alpha_t$.

¹⁰The discovery term is added for completeness. Following Salant (1976), exploration and discovery are not a separate decision variable in this model. For the remainder of this paper, discoveries will be treated as exogenous, and unanticipated to both the fringe and the cartel – therefore $\mathbb{E}(\text{disc}_{it+1}) = 0$ for all years, for all fields of the cartel, and fringe.

$$Max_{q_{it} \forall i \in \Omega, \forall t \in \mathbb{T}} \sum_{t=1}^T \beta^{t-1} \mathbb{E} \left(P_t(\bar{Q}_t^f, Q_t^c) Q_t^c - \sum_{i \in \Omega} c_{it} q_{it} \right) \quad (1a)$$

s.t.

$$\mathbb{E}(R_{it+1}) = R_{it} - q_{it} + \mathbb{E}(disc_{it+1}) \quad \forall i \in \Omega, \forall t \in \mathbb{T} \quad (1b)$$

$$R_{it} \geq 0 \quad \forall i \in \Omega, \forall t \in \mathbb{T} \quad (1c)$$

$$0 \leq q_{it} \leq \kappa_{it} \quad \forall i \in \Omega, \forall t \in \mathbb{T} \quad (1d)$$

$$Q_t^c = D_t(P_t) - \bar{Q}_t^f \quad \forall t \in \mathbb{T} \quad (1e)$$

$$\sum_{i \in \Omega} q_{it} = Q_t^c \quad \forall t \in \mathbb{T} \quad (1f)$$

At the same time, all fringe fields $j \in \mathbb{J}$, characterized by constant marginal extraction costs c_{jt} and reserves R_{jt} in year t , are optimizing their discounted profits as a price-taker. The cartel announces a price path \bar{P} , supported by cartel's sales, with respect to which the fringe field will optimize individually. Thus, each fringe firm $j \in \mathbb{J}$ will maximize the objective function outlined by Equation 2a, subject constraints 2b - 2e. The reserve constraints are denoted by Equation 2b - 2c. The fringe fields will always fully exhaust their economically recoverable stock. Equation 2d captures the yearly flow constraint dictated by well pressure. Aggregate fringe production Q_t^f in year t is the sum of production across all fields \mathbb{J} , as shown in Equation 2e.

$$Max_{q_{jt}, \forall t \in \mathbb{T}} \sum_{t=1}^T \beta^{t-1} \mathbb{E} (\bar{P}_t q_{jt} - c_{jt} q_{jt}) \quad \forall j \in \mathbb{J} \quad (2a)$$

s.t.

$$\mathbb{E}(R_{jt+1}) = R_{jt} - q_{jt} + \mathbb{E}(disc_{jt+1}) \quad \forall t \in \mathbb{T} \quad (2b)$$

$$R_{jt} \geq 0 \quad \forall t \in \mathbb{T} \quad (2c)$$

$$0 \leq q_{jt} \leq \kappa_{jt} \quad \forall t \in \mathbb{T} \quad (2d)$$

$$\sum_{j \in \mathbb{J}} q_{jt} = Q_t^f \quad \forall t \in \mathbb{T} \quad (2e)$$

In equilibrium, any competitive field $j \in \mathbb{J}$, acting as a price-taker, maximizes its discounted profits by adhering to the production path forecasted by the cartel. Simultaneously, the cartel maximizes

joint discounted profits by setting a production strategy that determines the market price path P_t . This equilibrium ensures that no player—whether within the cartel or the competitive fringe—has an incentive to deviate from the prescribed production levels. I assume that the cartel’s output levels are perfectly monitored by the fringe. This means that the fringe would always be able to react to (unannounced) output changes, and thus price path changes, of the cartel. In the absence of shocks, neither the cartel nor the fringe has an incentive to revise its optimal output path over time, as the model is dynamically consistent for a single strategic player (Benchekroun and Withagen, 2012).¹¹

Key necessary conditions In equilibrium, both the cartel and the fringe are indifferent between producing today and tomorrow, consistent with the classic model of Hotelling (1931). The cartel balances its perceived marginal revenue from selling an additional barrel today against the discounted marginal revenue from future sales, as shown in Equation 3. Similarly, the fringe evaluates its competitive rent today relative to its competitive rent tomorrow, represented by Equation 4.

Equation 3 shows that marginal revenues – which include scarcity rents as in the Hotelling (1931) model– minus production costs of the marginal barrel should raise at the rate of the exogenous discount rate β . If scarcity rents rise faster than the discount rate, the cartel will have an incentive to postpone extraction to a later period, as it is costless to leave oil unexploited for an additional year. Conversely, if scarcity rents grow more slowly than the discount rate, the cartel has an incentive to accelerate production today.

$$P_t(\bar{Q}_t^f + Q_t^c) + \theta_t \frac{\partial P_t}{\partial Q_t^c} * Q_t^c - c_{it} = \beta \mathbb{E} \left(P_{t+1}(\bar{Q}_{t+1}^f + Q_{t+1}^c) + \theta_{t+1} \frac{\partial P_{t+1}}{\partial Q_{t+1}^c} * Q_{t+1}^c - c_{it+1} \right) \quad (3)$$

The cartel’s first order condition, Equation 3, includes a internalization or conduct parameter denoted $\theta \in [0, 1]$, in contrast to the original Salant (1976) model. This conduct parameter allows for a flexible representation of collusion among OPEC producers. Specifically, OPEC producers are fully colluding if $\theta = 1$. If $\theta = 0$, they are not colluding. Any conduct parameter between 0 and 1 implies that the cartel is colluding, but in an ineffective way: they are leaving gains on the table. Thus, θ can in this context be interpreted as the actual market power of the cartel, as compared to the maximum attainable market power.¹² Thus, the marginal revenue of the cartel is composed of the price of oil, P_t , minus the *perceived* reduction in revenue on inframarginal units, $\theta_t \frac{\partial P_t}{\partial Q_t^c} Q_t^c$, due to

¹¹As discussed in more detail by Kellogg (2024), the "Open-Loop Nash-Cournot" equilibrium as described by Salant (1976) is a Markov Perfect Equilibrium in the case of one strategic player.

¹²This is a somewhat reduced-form way to capture the complex industrial organisation dynamics in the oil cartel.

the price decrease from additional production, and further reduced by the marginal production cost of the marginal barrel, c_{it} .

There are several reasons why the cartel may fail to fully exploit its market power. First, the oil market has historically been subject to significant geopolitical disruptions (Baumeister and Kilian, 2016), such as the 2003 Iraq War and more recently, the Russian oil price cap.¹³ Such events can introduce frictions that hinder the cartel from adhering to its optimal production rule. Second, the need to maintain cartel stability may also cause production levels to deviate from the first-order condition described in Equation 3. These stability concerns are not explicitly incorporated into the model, as doing so would compromise the tractability of the equilibrium. Third, actual realizations of demand or extraction costs in any given year may differ from expectations, leading to deviations from the optimal rule due to imperfect information.¹⁴ Finally, Andrade de Sá and Daubanes (2016) highlight that the presence of limit-pricing to deter potential entrants can result in behavior where the monopolist appears to forgo opportunities for higher pricing.

Within the cartel, it is always optimal to extract fields according to marginal costs. To this end, let us consider ordered available oil fields $i \in \Omega$ from the lowest to the highest cost. This implies that if field i produces in year t , it must be the case that field $i - 1$ has produced no later than year t . Each year is therefore characterized by a marginal field i_t , such that the cartel is indifferent between producing on additional barrel today or tomorrow – given the aggregate output of the fringe firms \bar{Q}_t^f .¹⁵ The intertemporal production path of the forward-looking cartel can therefore be summarized by a list of marginal fields in each year $\{i_1, i_2, \dots, i_T\}$, following the notation of Asker et al. (2023).¹⁶

Next, Equation 4 presents the first order condition for the competitive fringe. Similar to the cartel, Corts (1999) and Reiss and Wolak (2007) note that especially in a dynamic game context, conduct can be endogenous to e.g. investment or future competition, and can therefore vary depending on various factors. I try to address this concern by allowing the internalization parameter to change over time in the empirical implementation. This however does not mitigate the concern that conduct is not a truly structural parameter in the model. Further, as noted by Reiss and Wolak (2007), there is no clear theoretical model underpinning various cases of $\theta \in]0, 1[$ when assuming a coherent cartel. Miller and Weinberg (2017) have addressed these concerns by incorporating an ownership matrix in the supply-side specification, where a parameter captures the extent to which firms internalize their price externality to potential colluding parties, which can be interpreted as a conduct parameter under certain assumptions. They, however, do also not provide a formal theory on why the coordination effects are taking place in this specific way.

¹³Further details are provided in the Data section.

¹⁴The assumptions regarding expectations are discussed in the Empirical Section.

¹⁵Note that inframarginal producing fields are always facing a binding capacity constraint.

¹⁶Of course, there is also an intensive margin dimension to supply decisions: a marginal field does not necessarily exactly supply at full capacity. I abstract away from it in this theoretical framework for simplicity. Note that the empirical framework will take this into account.

the competitive fringe compares the scarcity rent on a marginal barrel produced today, $\bar{P}_t - c_{j_t t}$ with the discounted scarcity rent if that barrel were exploited tomorrow $\bar{P}_{t+1} - c_{j_t t+1}$. Thus, the fringe’s scarcity rents should rise according to the discount rate β . However, unlike the cartel, the fringe firms are price-takers and do not consider their influence on oil prices; therefore, no mark-up term is included in their first order condition. If scarcity rents grow faster than the discount rate, the fringe delays production. Conversely, if scarcity rents rise more slowly, the fringe has an incentive to increase current production.

The fringe firms will always produce from cheapest oil field to most expensive one. Therefore, the list of marginal fringe fields in each year $\{j_1, j_2, \dots, j_T\}$ implies a production path. In the optimum, for the marginal barrel of the fringe firms j_t in year t , the following condition should hold,

$$\bar{P}_t - c_{j_t t} = \beta \mathbb{E}(P_{t+1}^- - c_{j_t t+1}) \quad (4)$$

where $c_{f_t t}$ indicates the marginal extraction costs of the marginal field f in year t . The bar over the price variable indicates that the fringe firms take these as given.

In this model, no economically viable barrel remains untapped, either for the fringe or the cartel. Thus, the key decision variable is the timing of extraction—determining when to produce. Additionally, based on the respective oil reserve endowments of the fringe and the cartel, fringe fields may be fully depleted before the cartel’s reserves are exhausted. In such a case, the cartel becomes the sole supplier to the market.¹⁷

2.3 Dynamic model of oil extraction: policy and welfare

Climate goals, carbon taxes and the green paradox In the absence of an energy transition, standard resource extraction models predict that all reserves will eventually be depleted, regardless of market conduct. Since cumulative fossil fuel extraction is the primary driver of global warming, the complete exploitation of currently known fossil fuel deposits would exceed any established climate target—whether it be 1.5°C, 2°C, or 2.5°C (Welsby et al., 2021). As a result, if countries aim to meet the goals set by the Paris Agreement, policymakers must accelerate the transition away from fossil fuels. A carbon tax represents the first-best policy tool to address the environmental externalities associated with fossil fuel consumption.

A carbon tax has two primary environmental effects in the context of exhaustible resource extraction (Heal and Schlenker, 2019). First, it influences the *cumulative* amount of oil extracted over time.

¹⁷For a detailed analytical exposition of the various scenarios, see Benchekroun and Withagen (2012).

Higher carbon taxes can render a greater share of oil reserves uneconomical to extract, thereby reducing cumulative emissions. Under a constant carbon tax, barrels remain untapped if their marginal extraction cost, including the tax, exceeds the choke price. This effect can be analytically illustrated with a simple demand formulation. Let the yearly linear demand curve be given by $P_t = \alpha_t - \beta(Q_t^f + Q_t^c)$, where α is the choke price. Both fringe firms and the cartel face a constant carbon tax τ , so their marginal extraction costs at time t are $c_{jt} + \tau$ for the fringe firms $\forall j \in \mathbb{J}$, and $c_{it} + \tau$ for the cartel $\forall i \in \Omega$. Consequently, fields for which $c_{it} + \tau \geq \alpha_t$ and $c_{jt} + \tau \geq \alpha_t$ can never be profitably exploited, leaving the oil in the ground. If all fields had identical marginal extraction costs, this would result in an all-or-nothing outcome. However, this paper allows for heterogeneity in extraction costs across fields.

Second, carbon taxes also affect the *speed* of extraction in the oil sector. In the standard Hotelling framework, [Dasgupta et al. \(1980\)](#) demonstrate that virtually any desired pattern of resource extraction can be achieved by adjusting the trajectory of (carbon) taxes. Intuitively, if carbon taxes are expected to rise over time, firms have an incentive to accelerate extraction, as future production becomes less profitable. This highlights the critical role of expectations regarding taxation in a dynamic context. Analytically, consider a modification to the equilibrium condition for the cartel to incorporate time-varying carbon taxes τ_t , as shown in Equation 5. When τ_t increases over time, future marginal revenues must rise even faster to justify delaying extraction. Given the same backstop price, this necessitates a complete adjustment to both the price and extraction path.

$$P_t + \theta_t \frac{\partial P_t}{\partial Q_{ct}} * Q_{ct} - c_{it} - \tau_t = \beta \mathbb{E} \left(P_{t+1} + \theta_{t+1} \frac{\partial P_{t+1}}{\partial Q_{ct+1}} * Q_{ct+1} - c_{it+1} - \tau_{t+1} \right) \quad (5)$$

This second effect suggests that future carbon taxes can unintentionally increase current emissions, a phenomenon termed the "Green Paradox" by [Sinn \(2008\)](#). Anticipating stricter future climate policies, resource owners may accelerate the extraction of resources, leading to a short-term *increase* in environmental damage from oil production. As a result, environmental regulation, intended to reduce emissions, may temporarily have the opposite effect. The extent and significance of the Green Paradox depend on several factors, including the type of environmental policy implemented, extraction costs, demand elasticity, market conduct, and the distribution of oil reserves, as discussed by [van der Ploeg and Withagen \(2012a\)](#).

Market power, and the run for the last barrel The cartel exercises market power by withholding production capacity, thereby creating artificial scarcity to elevate prices. However, in the context of increasingly stringent environmental policies, this strategy may become less attractive. Specifically, the profitability of delaying production diminishes as carbon taxes rise over time. In response, the cartel may shift its strategy by frontloading production to earlier periods, thereby reducing the extent to which it withholds capacity. Given the cartel’s relatively lower marginal costs, this frontloading strategy could displace production from higher-cost fringe producers (Benchekroun et al., 2019). Consequently, the trajectory of carbon taxes influences not only the quantity of oil produced but also *who* produces. This acceleration in production, aimed at avoiding stranded assets, may be likened to a ‘race for the last barrel of oil’ that remains economically viable to extract.

An increasing carbon tax can incentivize the cartel to frontload its production rather than withhold capacity. This shift impacts the three welfare effects from market power discussed in Section 2.1: (i) productive inefficiencies (PI), (ii) deadweight loss (DWL), and (iii) environmental damages (ED). First, a higher carbon tax can mitigate productive inefficiencies associated with market power. By pushing production forward from lower-cost fields and crowding out higher-cost fringe fields, the cartel’s production allocation moves closer to the Herfindahl principle, which prioritizes extraction from the lowest to the highest cost fields (Herfindahl, 1967). Consequently, a carbon tax—intended to address environmental externalities—can also reduce welfare losses due to imperfect competition.

Second, the effect of an increasing carbon tax on deadweight loss is ambiguous. In a static framework, implementing a carbon tax would reduce production and thus amplify deadweight loss due to market power. Conversely, in a dynamic context, a rising carbon tax might induce a temporary increase in production which could reduce deadweight loss in the short term. The overall impact of the carbon tax on the net present value of deadweight losses over time depends on the interplay between these effects and is influenced by the tax’s specific design and magnitude. Consequently, the equilibrium quantities and associated deadweight losses can vary in complex ways, making the welfare effect ambiguous.

Third, while an increasing carbon tax is expected to reduce cumulative emissions over time, it may also trigger a Green Paradox—a temporary surge in emissions from the oil sector that could temporarily exacerbate environmental damages. The net welfare effect of such a paradox depends on the magnitude of the short-term emission increase and the valuation of future environmental damages (Heal and Schlenker, 2019). Specifically, Gerlagh (2011) distinguishes between two types of Green Paradoxes: a *weak* Green Paradox and a *strong* Green Paradox. A weak Green Paradox refers to a situation where emissions increase temporarily but the overall welfare impact remains favorable over

the long term, as the cumulative environmental damage is still mitigated. Conversely, a strong Green Paradox occurs when the carbon tax leads to a net increase in the present value of environmental damages over time, indicating that the increase in emissions outweighs the benefits of the future reduction in emissions. Hence, the total welfare effect of a carbon tax on environmental damages is ambiguous in the context of exhaustible resources.

In sum, a carbon tax, designed to correct the market failure of environmental externalities, may also influence the welfare losses arising from another significant market failure—namely, the cartel’s anticompetitive behavior. The overall welfare impact remains ambiguous, contingent on factors such as resource endowments, production costs, demand conditions, market conduct, and the specific implementation of the carbon tax.

3 Data sources and institutions

Production, cost and reserve data The main dataset is obtained from the the Norwegian energy consultant Rystad, covering virtually the entire upstream oil industry. It provides comprehensive field-level data from 1990 to 2019, including annual production, reserves, and costs, along with geological and extraction technology details such as oil type, location, and whether the field is onshore or offshore. A field represents a distinct oil production area, which may comprise multiple wells and vary significantly in size. For instance, the ‘Ghawar Uthmaniyah’ field in Saudi Arabia had reserves of 72,036 million barrels (mB) and produced 611 mB in 1990, whereas the ‘Rostovitskoye’ field in Russia had a reserve capacity of just over 1 mB and reported no production in the same year. The dataset encompasses a total of 20,569 producing fields, which are employed in the primary analysis.

Ideally, the most accurate measure of reserves would be obtained by observing the entire production lifecycle of a field until depletion, with the total extracted oil representing the maximum reserve. However, since most fields in the dataset have not been fully exploited, I rely on reserve estimates provided by Rystad Energy. Reserves are classified by the oil industry into three probability levels: P90 (or P1), representing the quantity recoverable with a 90% probability given current technical and economic conditions; P50 (or P1 + P2), indicating reserves recoverable with a 50% probability; and P10 (or P1 + P2 + P3), corresponding to reserves with a 10% probability of recovery. These estimates are sensitive to fluctuations in oil prices. In this study, I use P50 reserves, assuming an oil price of 60 USD per barrel (in 2021 dollars), which approximates the historical average price of oil. Panels A and B in Table 1 display the distribution of field-level oil reserves in millions of barrels (mB) for the years 1990 and 2019, respectively. The final column in each panel shows the total reserves across all active

fields for the corresponding years.

Cost data is observed at the field-level. To construct the marginal cost of oil extraction c_{ft} of field f during year t , I leverage the detailed information on production costs in the dataset, following [Asker et al. \(2019\)](#). I divide the total cost of production across spending categories h by their annual reported production quantities: $c_{ft} = \frac{\sum_h Expenditure_{h,ft}}{q_{ft}}$. The spending categories include $h = \{ \text{Well Capital, Facility Capital, Abandonment cost, Production Operating, Transportation Operating, and SGA} \}$. Field-level marginal costs are displayed in [Table 1](#) across different percentiles. Overall, the 5th percentile extraction costs are about 7 times smaller than the 95th percentile extraction costs—an observation which is persistent throughout both Panels. The last column shows the production-weighted marginal costs for 1990 in Panel A, and 2019 in Panel B respectively.

Additional datasets I employ various supplementary datasets. For data on yearly oil prices, I use the EIA imported crude oil price.¹⁸, following [Asker et al. \(2023\)](#). The evolution of the oil price is shown in [Figure 2](#). Further, I download the GDP deflator from the US Bureau of Economics Analysis, from the Gross Domestic Product Implicit Price Deflator series. Third, I use global GDP data from the World Bank for OECD and non-OECD countries.

The upstream oil cartel The Organization of the Petroleum Exporting Countries (OPEC) was established in 1960 by five founding members: Iran, Iraq, Kuwait, Saudi Arabia, and Venezuela. Over time, 11 additional states joined the organization: Qatar (1961), Indonesia (1962), Libya (1962), United Arab Emirates (1967), Algeria (1969), Nigeria (1971), Ecuador (1973), Angola (2007), Gabon (1975), Equatorial Guinea (2017), and Congo (2018).¹⁹ OPEC’s stated objective is to “coordinate and unify petroleum policies among member countries, in order to secure fair and stable prices for petroleum producers”.²⁰ In [Table 1](#), OPEC members are defined as the countries listed above. Following the shale boom of the 2010s, OPEC entered into a Declaration of Cooperation in 2016 with 10 non-OPEC countries, collectively referred to as “OPEC plus,” aiming to “help rebalance the market, and support oil market stability.” The OPEC+ member states include Azerbaijan, Bahrain, Brunei, Kazakhstan, Malaysia, Mexico, Oman, Russia, South Sudan, and Sudan. The upstream oil industry consists of oil companies that are either (largely) state-run enterprises, or independent enterprises.

¹⁸Source: www.eia.gov/outlooks/steo/realprices/, downloaded in January 2023.

¹⁹Qatar terminated its membership in 2019. Indonesia suspended its membership in 2009, rejoined in 2016 and left later that year. Ecuador left in 2007 and rejoined in 2020. Gabon left in 1995, and rejoined in 2016.

²⁰Source: opec.org

The state-run (nationalized) companies can be split into those that are run by OPEC, and those that are from non-OPEC states. Every OPEC country has its own nationalized company, which controls production.²¹ For example, Saudi Aramco is the national oil company of Saudi Arabia, and the Nigerian National Petroleum Company operates in Nigeria. The largest independent companies are the '5 oil majors', i.e. ExxonMobil, Chevron, BP, Royal Dutch Shell and Total. In my analyses, these will be part of the fringe production.

As shown in Table 1, OPEC countries held just over half of global oil reserves in 1990. By 2019, this share had slightly declined to 45%. In that year, total global reserves amounted to 866,580 mB. OPEC countries generally possess larger oil fields, including some of the richest fields like Ghawar in Saudi Arabia and Kirkuk in Iraq. In terms of production, OPEC's market share has remained relatively stable at around 40% since the 1990s, as depicted in Figure 2. Saudi Arabia's market share has also remained constant over the 1990–2019 period, holding steady at 15%. In contrast, the U.S. has seen a significant increase in its market share, particularly in the last decade.

The field-level extraction costs are on average lower for OPEC countries. In the last Column, these are weighted by observed production levels of the respective years. There, it shows that OPEC extracts oil on average at roughly half the cost, compared to non-OPEC countries. This considerable efficiency difference between OPEC and non-OPEC members in oil extraction is persistent throughout the years. In 1990, the average marginal extraction cost of OPEC countries was 6.68 USD, while it was 12.23 USD for non-OPEC countries. The discrepancy is slightly larger in 2019, where the marginal OPEC extraction costs amounts to 5.61 USD, and the non-OPEC costs amount to 12.89 USD.


²¹Sometimes however, they contract with independents to operate specific facilities

Table 1: Summary Statistics in 1990 and 2019

	p5	p25	p50	p75	p95	Aggregate
<u>Panel A: Summary Statistics in 1990</u>						
Reserves (mB)	.07	1.62	9.14	44.95	352.28	1 248 796.80
OPEC	.18	3.12	14.35	77.06	977.31	669 159.71
non- OPEC	.07	1.49	8.52	41.49	304.79	579 636.80
Production (mB)	.00	.04	.25	1.20	11.36	22 563.49
OPEC	.00	.13	.71	2.77	32.68	8183.61
non- OPEC	.00	.04	.21	1.01	9.00	14 739.82
Marginal Cost (USD/bbl)	4.50	7.08	9.43	13.83	30.97	10.22
OPEC	4.36	7.24	9.41	13.61	26.81	6.68
non- OPEC	4.52	7.05	9.44	13.88	31.86	12.23
<u>Panel B: Summary Statistics in 2019</u>						
Reserves (mB)	0.00	0.03	1.01	11.54	170.98	866 580.28
OPEC	0.00	0.00	0.62	11.99	404.68	389 508.42
non- OPEC	0.00	0.04	1.07	11.53	153.42	477 071.86
Production (mB)	0.00	0.02	0.17	1.02	9.22	29 141.80
OPEC	0.00	0.03	0.23	1.67	23.12	10 570.29
non- OPEC	0.00	0.02	0.16	0.95	8.00	18 571.51
Marginal Costs (USD/bbl)	3.76	6.56	9.30	14.28	27.84	10.25
OPEC	3.76	6.24	8.78	13.21	26.92	5.61
non- OPEC	3.75	6.60	9.40	14.51	27.95	12.89

Notes: Unit of observation is oil field in a given year. Reserves and production are expressed in million barrels. Dollar values are deflated by US GDP deflator of 2021. The table summarizes the data for 1990 and 2019 in Panel A and B respectively. I use the 20,569 active fields from 1990 - 2019. In the Columns, the 5th, 25th, 50th, 75th, and 95th percentiles are shown. For reserves and production, the last Column 'Aggregate' shows the sum across all fields. For marginal costs, the Column 'Aggregate' displays the production-weighted average across fields. Countries are included OPEC in all years if they had ever had active membership between 1990 and 2019.

Figure 2: Market Shares and Oil Price



Notes: Market shares (left axis) and oil prices (right axis) across the sample period. Market shares (OPEC, Saudi Arabia and US) are expressed in percentages, the real oil price in 2021 USD per barrel. From the 1990s onwards, OPEC market shares remain stable at 40%, as does Saudi-Arabia’s market share at roughly 15%. The US market share has undergone a steep incline during the last decade.

4 Empirical Analysis: Cost and Demand

The objective of this section is to recover key primitives for marginal costs and demand, which will serve as inputs for the dynamic model presented in Section 2.2. First, I estimate marginal costs by leveraging annual field-level data on production and expenditures, allowing me to infer marginal costs with minimal structural assumptions. Second, I estimate global oil demand using an instrumental variable (IV) approach, where the instrument exploits variation in the aggregate oil supply curve to identify the slope of demand. My demand estimates reveal the highly inelastic nature of global oil consumption.

4.1 Marginal costs

4.1.1 Approach

The cost data in this analysis is structured at the field-year level, as outlined in Section 3. Specifically, production costs are reported for each year that a field is actively producing oil. However, for fields that are not producing in a given year, cost data is unavailable. Since environmental regulation can influence the extraction sequence of fields, it is necessary to estimate production costs for these non-producing (but discovered) fields. To address this gap, I assume that firms minimize production costs subject to the observed level of output. This cost-minimization framework allows for the estimation of costs even in years when a field is idle, ensuring a more comprehensive understanding of the cost structure across both producing and non-producing fields. Hence, the Leontief production function is given by

$$q_{ft} = \min\{\alpha_{ft}K_{ft}, \gamma_{ft}L_{ft}\} \quad (6)$$

s.t.

$$q_{ft} \leq R_{ft}, \quad R_{ft} = R_{ft-1} - q_{ft-1}, \quad R_{ft} \geq 0$$

where K represents the fixed inputs (e.g. drilling equipment, pipes), and L the variable inputs (e.g. production worker, energy) of production. R_f represents the reserves of field f , which could be both either owned by the cartel, or the fringe, i.e. $f = \{i, j\}$.²² The technology parameters α_{ft} and γ_{ft} are allowed to vary at the field-level. I assume that the price of capital units is given by r_{ft} , and the price of the variable input by w_{ft} . All input prices are exogenous in this framework. Hence, field-level costs can be summarized as

$$C_{ft}(q) = \left(\frac{w_{ft}}{\gamma_{ft}} + \frac{r_{ft}}{\alpha_{ft}} \right) q. \quad (7)$$

Following [Asker et al. \(2019\)](#), I put additional structure on the process that governs the evolution of the ratio of the input price and technology parameters. I assume both ratios evolve according to a common process μ_{ft} within a field, while allowing for Hicks-neutral productivity shocks across time.

²²To ease notation, I am abstracting away from the yearly capacity constraint that fields are facing due to well pressure constraints, i.e. $\kappa * R_{ft}$. This additional constraint, however, does not affect the conceptual framework of the cost specification.

Analytically, this process looks as follows

$$\frac{w_{ft}}{\gamma_{ft}} = \frac{w_f}{\gamma_f} \mu_{ft} \quad \text{and} \quad \frac{r_{ft}}{\alpha_{ft}} = \frac{r_f}{\alpha_f} \mu_{ft}. \quad (8)$$

Let μ_{ft} be governed by an exogenous martingale process, such that $E(\mu_{ft+k}|\mu_{ft}) = \mu_{ft}, \forall k \geq 1$. Intuitively, the μ captures long-run technological progress in oil extraction, changes in the costs of inputs, or technology-specific parameters.

Hence, Equation 7 can be rewritten as Equation 9, where $c_f = \left(\frac{w_f}{\gamma_f} + \frac{r_f}{\alpha_f}\right)$. Unit costs of production are constant within a field f , within a given year t , up to their capacity constraint. This results in the following marginal cost function for field f in year t ,

$$c_{ft} = MC_{ft}(q) = AC_{ft}(q) = \begin{cases} c_f \mu_{ft} & \text{if } q \leq R_{ft} \\ +\infty, & \text{otherwise.} \end{cases} \quad (9)$$

Marginal (or average) costs of field f at time t , i.e. c_{ft} , can be decomposed in three elements. First, time-invariant marginal costs c_f , e.g. due to geology. Second, a technology-year-ownership specific cost shifter μ_{st} , where $\mu_{ft} = \mu_{st} \exp(\epsilon_{ft})$. Note that this differs from the specification of [Asker et al. \(2019\)](#), as their cost shifter is not ownership-specific. In total, I consider 5 technologies, that is (i) offshore deepwater oil, (ii) offshore shelf oil, (iii) oil sands, (iv) shale/ tight oil, (v) other onshore oil – as classified by Rystad. Ownership is defined as OPEC or non-OPEC.²³ Hence, a shock μ_{st} will hit e.g. all non-OPEC oil sands fields during a certain year. Third, the marginal costs contain measurement error $\exp(\epsilon_{ft})$. Thus, the unit cost can be rewritten in terms of these three elements, that is,

$$c_{ft} = c_f \mu_{ft} = c_f \mu_{st} \exp(\epsilon_{ft}) \quad (10)$$

In my model, oil extraction by field f at time t occurs at cost $c_f \mu_{ft}$ per barrel. The technology-year-ownership specific cost shifter, i.e. μ_{st} , is estimated as follows

$$\ln(\hat{\mu}_{st}) = \sum_{f \in s} \rho_{ft} \ln(c_{ft}) \quad (11)$$

where ρ_{ft} is the quantity-weight of a field in a given year's total output, equal to $\rho_{ft} = \frac{q_{ft}}{\sum_{f \in s} q_{ft}}$. The time-invariant marginal cost, c_f , is then estimated, using the following (within-field) regression using weighted least squares,

$$\ln(c_{ft}) - \ln(\hat{\mu}_{st}) = \ln(\hat{c}_f) + \epsilon_{ft}. \quad (12)$$

²³In total, there are 5x2 categories s . These consist of five technologies groups, each with 2 ownerships classes.

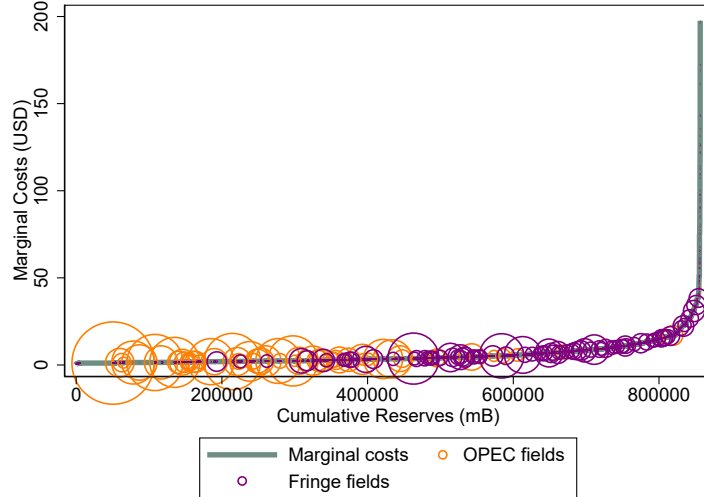
To account for potential estimation error in both the field fixed effects, c_f , and the technology-year-ownership shocks, μ_{st} , I employ a bootstrap procedure to estimate the standard errors of the cost estimates. The bootstrapped standard errors are derived through a two-step process. In the first step, I resample the full dataset with replacement k times to obtain μ_{st}^k . In the second step, within each field, I resample the field-year observations to estimate c_f^k , using the $\hat{\mu}_{st}^k$ estimates from the first step. In total, I use 50 bootstrap iterations.

4.1.2 Findings

Marginal costs and market power The cartel can exercise market power by withholding efficient production capacity. As described in Section 2.1, this artificial scarcity will spur the less efficient fringe firms to start producing. This mechanism increases oil prices, to the benefit of the infra-marginal capacity of the cartel. This behavior is enabled by the fact that the cartel can produce at lower marginal cost of extraction. Figure 3 shows the marginal extraction costs in 2019 for all remaining reserves. The shape of the line is highly convex, indicating a large heterogeneity in production extraction costs across fields. Larger OPEC fields, indicated with orange circles on the graph, are characterized by lower marginal costs than non-OPEC fields. The steeper part of the extraction curve consists out of mostly fringe fields.

Figure 4 presents the naive static price-cost margins for Saudi Arabia, Canada, and the United States, calculated as $\frac{P_t - c_{ft}}{P_t}$, where P_t is the oil price and c_{ft} represents the marginal cost of production. Note that these *naive* margins include inframarginal rents and hotelling rents. The dashed line represents the 75th percentile price-cost margin in each year, while the dotted line indicates the margin at the 90th percentile of production costs. Saudi Arabia, shown in purple, consistently exhibits margins exceeding 90% for both the 75th and 90th percentile, reflecting its position as a low-cost producer. In contrast, margins in the U.S. and Canada, both non-OPEC countries, are significantly lower. The U.S., in particular, displayed the lowest margins during the 1990s, though these margins have steadily increased over time. A notable feature for the U.S. and Canada is the wider gap between the 75th and 90th percentile margins, indicating greater heterogeneity in production costs within these countries. The spike in oil prices in the early 2010s led to relatively high margins across all three countries. The evolution of the 75th and 99th percentile of production costs throughout the years for the same countries is shown in Figure 10.

Figure 3: Marginal cost of OPEC and fringe fields (2019)




Notes: Cumulative reserves are shown according to their marginal extraction costs in 2019. Largest ($\geq 800\text{mB}$ reserves) OPEC and fringe fields are indicated by the circles along the marginal cost line. Large OPEC fields are characterized by lower marginal costs, while the larger fringe fields tend to position themselves at the more convex part of the cost curve. Marginal costs are expressed in 2021 USD.

The carbon budget Emissions of greenhouse gases, such as carbon dioxide (CO_2), are a primary driver of global warming. The Intergovernmental Panel on Climate Change (IPCC) published that the best estimate of the remaining carbon budget is 500 Gt CO_2 -eq for a 50% likelihood of limiting global warming below 1.5 °C, and 1150 CO_2 -eq for a 67% likelihood of limiting global warming to 2 °C as of 2020 (IPCC, 2023). In 2019, yearly anthropogenic emissions were estimated at 59 Gt of CO_2 -equivalent, 54% higher than in 1990. Given that *cumulative* emissions determine climate outcomes, this suggests that, under a business-as-usual scenario, the 1.5°C threshold would be surpassed within eight years of 2020. Fossil fuel combustion has been the largest contributor to greenhouse gas emissions in recent decades, with the fastest growth occurring in this sector (IPCC, 2023). The International Energy Agency (IEA) reports that global emissions from fossil fuel combustion reached 34 $GtCO_2$ -equivalent in 2019.²⁴ Approximately 30% of emissions from fossil fuel combustion, or about 11 $GtCO_2$ -equivalent

²⁴In this paper, I focus on the carbon dioxide *equivalent* emissions. This includes methane (CH_4), which is a greenhouse gas that is associated with oil production with a much higher ability to trap heat in the atmosphere, as compared to

Figure 4: Price-cost margin evolution (Saudi Arabia, US, Canada)



Notes: Naive price-cost margins ($\frac{P_t - c_{ft}}{P_t}$) throughout the years for Saudi Arabia, US, and Canada. The dashed lines represent the 75th percentile of the margin computed for each country. The dotted line represent the margin for the 90th percentile of marginal cost of production during each year. Saudi-Arabia has consistently large, and stable margins. Canada and US production has lower margins, and experiences more volatility throughout the years.

annually, are attributable to oil combustion.²⁵

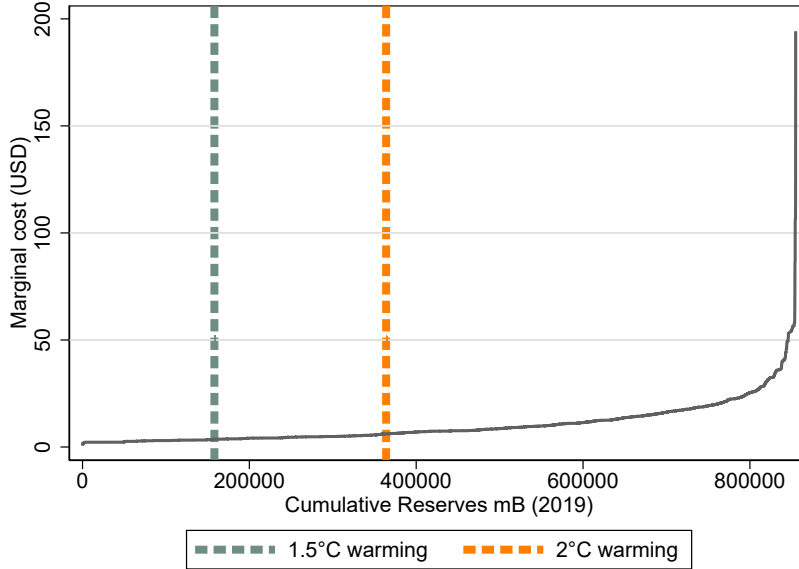
I compute that only around 86 $GtCO_2$ from the remaining global carbon budget can be "allocated" to oil combustion if the 1.5°C target is to be achieved, assuming the current fossil fuel usage ratios remain constant.²⁶ It is important to note, however, that this calculation is purely illustrative, as it

CO_2 . Methane, however, breaks down to CO_2 after 7-12 years in the atmosphere – which will remain in the atmosphere for hundreds years or more, according to NASA (climate.nasa.gov/vital-signs/methane/?intent=121). According to the IEA, the emissions associated with oil production amounted to at least 52 Mt CH_4 in 2019, the equivalent of 1,275 Mt CO_2 . While the climate consequences of methane and carbon dioxide vary widely, all emissions will be expressed in CO_2 -equivalents to reduce complexity.

²⁵Source: [iea.org/data-and-statistics/data-tools/greenhouse-gas-emissions-from-energy-data-explorer](https://www.iea.org/data-and-statistics/data-tools/greenhouse-gas-emissions-from-energy-data-explorer)

²⁶This estimate is based on approximately $500 * \frac{34}{59} * 0.3 GtCO_2$ of allowable oil combustion within the 500 $GtCO_2$ -

Figure 5: Marginal cost of global oil reserves and the carbon budget (2019)



Notes: Cumulative reserves are shown according to their marginal extraction costs in 2019. The vertical lines indicate the 1.5°C (50% probability) and 2°C (67% probability) global warming limit respectively, under the assumption that relative ratios within the carbon budget are constant. The 1.5°C line equals a budget of 86 $GtCO_2$ assigned to oil combustion, while the 2°C line indicates a 198 $GtCO_2$, assuming a uniform 0.544 tCO_2 emission per barrel. Marginal costs are expressed in 2021 USD.

is unlikely that the current proportions of fossil fuel use will remain constant throughout an energy transition. Figure 5 plots the total 2019 reserves of active oil fields, ranked by their marginal extraction costs (as described in Section 4.1). The green line represents the estimated 1.5°C oil carbon budget, while the orange line corresponds to the 2°C threshold, both following the IPCC’s 2020 calculations and assuming a constant share of oil in the carbon budget. These projections emphasize that a substantial portion of current oil reserves would need to remain untapped in order to comply with global climate goals.

There is strong marginal cost heterogeneity across oil fields, as shown by the highly convex marginal cost curve in Figure 5. The fields that can be exploited under the constraint of a 1.5°C carbon budget,

equivalent carbon budget for a 50% probability of staying below 1.5°C.

could all be extracted at the very low marginal cost of under 5 USD. Thus, different resource owners face various risks of ending up with stranded assets. The resource owners that have oil fields that are characterized by lower marginal cost, are less likely to become redundant in the face of the energy transition. Figure 6 shows asset at risk as a share of GDP per country, based on their marginal extraction costs, if the 2°C carbon budget were to be imposed, and fields would be tapped according to marginal costs.

OPEC countries, on average, have lower marginal production costs. At the same time, however, the fossil fuel industry is often an important part of the economy of OPEC countries. Hence, even though for Saudi Arabia only 4.5% of its total reserves are at risk, their stranded assets as share of the country’s GDP amounts to 65%. The US is facing one of the largest risks in absolute terms: 149 390 mB – equal to 98% of their total reserves. However, the value of these reserves is only 45% of their current GDP. For Russia, 86% of total barrels would be stranded, amounting to 97 507 mB of oil – equal to a value almost quadruple their GDP. The figure shows that most European countries are not importantly affected by stranded assets in the oil sector. The share of stranded assets, relative to total current reserves is shown in Figure 12 in the Appendix.

4.2 Demand

4.2.1 Approach

The global oil demand curve is an input into the dynamic oil extraction model. Following [Asker et al. \(2023\)](#), I estimate a linear demand curve for oil, specified in Equation 13

$$Q_t = \alpha + \beta P_t + \gamma GDP_t + g(t) + \epsilon_t. \quad (13)$$

In this specification, Q_t represents the annual quantity of extracted oil, and P_t denotes the annual crude oil price.²⁷ The regression specification includes world GDP (GDP_t) as a control variable to account for the influence of global economic activity on oil prices ([Kilian, 2009](#)). Additionally, a fourth-order polynomial time trend, $g(t)$, is included to control for macroeconomic trends. The error term is denoted by ϵ_t , and α represents the time-invariant constant.

A first main coefficient of interest is the price coefficient β , denoting the slope of the linear demand curve. Second, the constant term, GDP and time trend coefficients will jointly allow me to construct a

²⁷In practice, there is inherent heterogeneity among oil barrels due to differences in composition (e.g., heavy versus light oil), which can lead to variability in oil prices. However, the Rystad data is standardized to reflect the energy equivalent of a single Brent crude oil barrel, thereby largely mitigating the impact of compositional heterogeneity.

year-specific intercept of the demand curve. However, estimation of Equation 13 using ordinary least squares (OLS) may yield biased coefficients due to the endogeneity of oil prices. This endogeneity arises from simultaneity issues, as oil prices are determined in direct relation to total oil extracted. To obtain unbiased estimates, I employ an instrumental variable (IV) approach.

Specifically, I use a supply shifter as an instrument for oil prices. The identification of the slope of the demand curve relies on shifts in the supply curve, exogenous to oil demand. I track annual shifts in the global oil supply curve using detailed information on field-level marginal production costs. By ranking all oil barrels from lowest to highest marginal production cost, I can pin down the production cost of the Q th barrel of oil each year, where Q is set at 15 billion barrels—exceeding the annual minimum production. The marginal production cost of this Q th barrel serves as a valid instrument because it reflects shifts in the aggregate supply curve throughout the years and is correlated with oil prices. Furthermore, since the marginal cost of the Q th barrel is independent of global oil demand, the exclusion restriction is satisfied.

A threat to identification is the fact that energy prices are an input to marginal extraction costs, and therefore these costs would not be independent of global oil demand. To address this concern, the expenditure category associated with energy costs, i.e. 'Production Operating', can be dropped from the cost estimation. Hence supply shocks not driven by energy prices can be captured, in order to identify the slope of oil demand.

4.2.2 Findings

The estimation results of demand Equation 13 are presented in Table 2. Column (1) reports the results using OLS, while the remaining columns employ an Instrumental Variable (IV) approach, where oil prices are instrumented using the cost shifter, as outlined in Section 4.2.1. The F-statistic across all specifications indicates that the cost shifter is a strong instrument for oil prices. As expected, the OLS estimation in Column (1) shows an upward bias in the price coefficient due to the simultaneity between prices and quantities. Specifications (1), (3), and (4) include a polynomial time trend and world GDP as controls. In Column (2), I replace the polynomial time trend with decade-specific dummy variables to account for time-varying demand shifts. In the final specification, I extend the time series from 1990–2019 to 1970–2019. The price coefficient remains largely unchanged, but the precision of the estimate improves. First-Stage results are reported in Appendix C. The alternative instrument – a cost-shifter that does not include operating expenses – does not impact the results, as shown in Table 7.

Table 2: Demand Estimation Results

	Dependent Variable: Quantities (mB)			
	(1)	(2)	(3)	(4)
	OLS	IV	IV	IV
Price	-5.05 (9.21)	-48.84 (9.90)	-71.01 (54.66)	-69.66 (12.22)
Controls	✓	✓	✓	✓
2SLS		✓	✓	✓
N	30	30	30	50
F -statistic		34.26	69.61	41.53
Mean elasticity	-0.01	-0.11	-0.16	-0.16

Notes: The IV specification relies on the cost shifter IV computed using a production level of 15 billion barrels, as described in section 4. Controls in (1), (3), (4) include a polynomial time trend, and world gdp. In Column (2), the time trend is replaced by a dummy variable for each decade. In the last Column, the time series starts already in 1970. All reported standard errors are computed correcting for heteroskedasticity and autocorrelation using Newey-West, with a small sample correction, and rely on a 1 period lag autocorrelation structure. The 1st stage Stock and Yogo F -statistic is reported.

The estimated implied elasticity ranges from -0.11 to -0.16, indicating that oil demand is highly inelastic. This finding is consistent with the existing literature, as summarized by Kilian (2022) and Caldara et al. (2019), which often rely on structural VAR models. Figure 13 in the Appendix illustrates the evolution of oil demand elasticity over the 1990–2019 period. The elasticity remains relatively stable and highly inelastic throughout, with only slight variation. During the 2008–2012 period –the financial crisis–, demand becomes marginally more elastic, though the estimates remain comfortably above -0.3. The dashed lines in Figure 13 represent the 95% confidence intervals.

5 Empirical analysis: dynamic model of oil extraction

In the second step of the empirical implementation, I estimate the internalization parameters within the framework of the dynamic structural model described in Section 2.2. This process builds on the demand and cost estimates obtained in Section 4. First, I outline the inference approach used to

identify the cartel’s conduct. Next, I detail the assumptions regarding expectations for future demand, costs, and reserves. Third, I discuss the model’s parametrization. Finally, I provide a brief overview of the estimation algorithm employed to obtain the internalization parameters.

5.1 Approach

Simulated method of moments In the dynamic optimization model, I estimate the cartel internalization (or conduct) parameter θ using the simulated method of moments (SMM). The parameter θ captures the degree of effectiveness in the cartel’s production decisions, where $\theta = 1$ indicates full coordination among cartel members, and $\theta = 0$ implies no collusion. Intermediate values of θ reflect varying levels of effectiveness of cartel behavior.

The estimation follows an indirect inference approach, as in [Gowrisankaran et al. \(2024\)](#), where I search over potential values of θ to identify the one that aligns the model’s predicted outcomes with key empirical moments. To implement this, I use the demand and cost estimates obtained in Section 4 as inputs for the dynamic model. I simulate annual production decisions for both the cartel and the competitive fringe, based on the cartel-fringe model from Section 2.2. The SMM procedure minimizes the distance between simulated and actual moments in the data, providing an estimate of θ . Hence, the objective function of the SMM estimation is the following:

$$\hat{\theta} = \underset{\theta}{\operatorname{argmin}} \quad g(\theta, \hat{\Gamma}, X_t)' \hat{W} g(\theta, \hat{\Gamma}, X_t) \quad (14)$$

where $g(\cdot)$ captures the mean differences in moments between the observed data (X_t) and the simulated data, as governed by the internalization parameter θ . In my main specification, the objective function is estimated using three key moments: (1) annual oil prices, (2) quantities produced by the fringe, and (3) total yearly extraction levels.²⁸ Each moment reflects the difference between the statistic predicted by the model and the corresponding value observed in the data (X_t). The simulations are based on the estimated demand and cost parameters, $\hat{\Gamma} = \{\hat{\alpha}_t, \hat{\beta}, \hat{c}_{ft}\}$, which were obtained in Section 4. The matrix \hat{W} represents the weights assigned to each moment, determining the relative importance of different moments in the estimation process. I employ a weighting matrix based on the inverse of the variance-covariance matrix of the moments, which is a standard choice to ensure efficient estimation. This approach gives more weight to moments with less variability, thus improving the precision of the conduct parameter θ estimation.

²⁸In alternative specifications, additional moments such as the square of quantities or market shares were included. Results are shown in Appendix C.

Expectations on Demand, Costs, and Reserves The cartel and the fringe form expectations about the future in order to decide on their optimal production paths. Specifically, expectations are formed on marginal extraction costs, demand, and reserves.

First, I follow [Asker et al. \(2019\)](#) by assuming that future cost shocks are unanticipated by market participants. Specifically, I model μ_{it} as following a martingale process, such that $\mathbb{E}(\mu_{it+k} | \mu_{it}) = \mu_{it}$ for $k \geq 1$. The term μ_{it} captures both long-term technological trends and variations in the absolute or relative costs of inputs or technology. This process is treated as exogenous, consistent with the partial equilibrium framework adopted in this paper. In a general equilibrium context, both input prices and technological progress could be endogenously determined by production decisions.

Next, market players form expectations on future demand. In the baseline model, I adopt the assumption that firms' perceptions of the demand levels D_t evolves according to a first-order Markov process, specified as $D_t = (1 + \rho)D_{t-1} + v_t$. Here, $\rho \geq 0$ reflects the anticipated average annual growth in demand and v_t is an iid shock with zero mean and constant variance. This specification implies that firms expect a consistent, exogenous upward shift in demand over time, reflecting factors such as population growth and economic expansion. Accordingly, the expected future demand is defined $\mathbb{E}(D_{t+k}|D_t) = (1 + \rho)^{k-1}D_t$ for $k \geq 1$. In my simulations, I set ρ equal to the historical growth rate in oil demand from 1990 to 2019, which is 1% per year.

Third, resource owners form expectations regarding future oil discoveries. In this paper, oil exploration is treated as exogenous to the model. In the main specification, I assume that firms do not anticipate any future discoveries when making production decisions; thus, any future discoveries are considered unanticipated.

In the dynamic estimation, after the realization of extraction costs, demand, and oil discovery shocks in year t , both the fringe and the cartel have the opportunity to re-optimize their production levels for each subsequent year until year T .

Conduct specification In practice, cartel conduct may not be constant over time, as demand fluctuations can influence the incentives to collude. [Green and Porter \(1984\)](#) suggest that negative demand shocks may lead to more competitive behavior, while [Rotemberg and Saloner \(1986\)](#) find that price wars are more common during economic booms. To account for these dynamics, I allow for time-varying cartel conduct by specifying the following model:

$$\theta_t = \theta^0 + \theta^1 \text{Demand Shift}_t \tag{15}$$

This specification permits the internalization parameter to vary over time with shifts in demand.

A positive estimate of θ^1 indicates that collusion intensifies during periods of high demand, whereas a negative θ^1 suggests that the market becomes more competitive following positive demand shocks. If θ^1 is not significantly different from zero, demand shifts do not have a meaningful impact on the effectiveness of the cartel.

The identification of the conduct parameter θ^0 is based on the extent to which the cartel strategically withholds capacity, given observed demand and cost conditions. Specifically, θ^0 captures the degree of market power exerted by the cartel, where higher observed price-cost margins correspond to a higher estimated level of θ^0 as described by the first-order conditions outlined in Section 2.2. Next, temporal exogenous variation in demand provides the key source of identification for time-varying internalization parameter, θ^1 . For instance, following a positive exogenous demand shock, a more restrictive output regime by OPEC would indicate procyclical cartel behavior, reflecting a higher degree of collusion during periods of high demand. The supply response to demand shifts is thus critical for understanding the strategic behavior of the cartel. However, while exogenous demand shocks serve as the primary source of variation, it is important to recognize that θ^1 may also be influenced by unobserved factors correlated with demand, such as technological advancements or geopolitical events. In such cases, θ^1 could capture not only the response to demand shifts but also the effect of these omitted factors.

Standard errors are obtained by bootstrapping the cost estimates, as described in Section 4.1. For each bootstrap procedure $s \in S$, a conduct parameter $\hat{\theta}^s$ is estimated by the inference method explained above. Repeating the estimation exercise S times, I obtain a vector of estimated structural parameters $\hat{\theta}^{s=1}, \hat{\theta}^{s=2}, \dots, \hat{\theta}^{s=S}$. Standard errors are calculated based on my $S=50$ bootstrapped internalization parameter estimates.

Parametrization Some auxiliary parametrizations are needed to fully specify the model. First, since simulating the production path requires intertemporal optimization, a joint annual discount factor is needed. This is set at $\beta = 0.95$ in the main specification. Further, the final time period is set at 2150, such that $|\mathbb{T}| = 160$.²⁹

Second, the members of the Organization of the Petroleum Exporting Countries (OPEC) must be exogenously assigned by the researcher, as they are not inferred directly from the model. In this analysis, I focus on the core members of OPEC in the main specification, which primarily includes countries from the Arabian Peninsula: Saudi Arabia, Kuwait, the United Arab Emirates, Algeria, Libya, Iran, and Qatar. This allows the analysis to concentrate on OPEC countries that tend to

²⁹The time period is set sufficiently large, such that extraction is close to zero during the final periods.

comply with agreed-upon production quotas.³⁰ These core members jointly control 80.2% of the total oil reserves of all OPEC members active at any point between 1990 and 2019.

Oil fields face an annual flow constraint on oil extraction, as emphasized by [Anderson et al. \(2018\)](#), due to the technical limitations of well pressure. Oil reserves cannot be extracted all at once. In the model, this constraint is represented by the parameter $\kappa_{it} = \kappa * R_{it}$. Additionally, other market frictions—beyond market power—may prevent wells from consistently operating at full technical capacity. I calibrate κ_{it} based on the observed extraction behavior of fringe firms. The data indicate that $\kappa = 0.053$, implying that resource owners are able to extract just over 5% of reserves annually, on average.^{31,32} The flow constraint κ is calibrated using only the behavior of fringe firms, as it could be confounded by market power incentives if cartel data were used. An implicit assumption in the model is that both the cartel and fringe face the same annual flow constraint.

To reduce the computational burden of optimizing across 20,569 assets over approximately 160 years (1990–2150), I group fields into bins based on their time-invariant marginal costs c_f , their extraction technology, and ownership. The time-invariant component of marginal costs is estimated using the methodology outlined in Section 4.1. It is then rounded to the nearest integer to construct cost-technology-ownership bins. The ownerships considered are OPEC and non-OPEC. The five technologies considered are deepwater offshore, shelf offshore, oil sands, shale field, or all other onshore. For example, the field "Hawtah" in Saudi Arabia is classified into the segment " $c_f = 3$ "—"OPEC"—"onshore". Each technology will experience a distinct year-specific shock to extraction costs, as described in Section 4.1. This binning process results in 140 classification bins, significantly reducing computational complexity while preserving most of the relevant information.

Estimation Algorithm The estimation procedure proceeds through two nested loops: an outer loop and an inner loop. The outer loop begins by selecting candidate values for the structural parameters, θ , from a predefined grid of possible parameter values. These candidate parameters serve as inputs into the cartel optimization problem in the inner loop.

The inner loop then iterates over the sample years (1990–2019). For each year, the model updates the information set and expectations, and the dynamic cartel-fringe game is solved for the remaining years, following the methodology outlined in Section 2.2. In this setting, the cartel and fringe resource

³⁰For a historical overview of OPEC compliance, see [Ghoddusi et al. \(2017\)](#).

³¹Engineers estimates find that in theory, $\kappa = 0.1$ should be feasible. In reality, however, this technical level of capacity is often not reached.

³²To make sure wells can get fully extracted eventually, I always allow a minimum of 10 extracted barrels per bin.

owners engage in a best-response game, iterating until no player can profitably deviate from their strategy. The optimum is solved with JuMP (Julia) using the Gurobi solver. Following [Asker et al. \(2023\)](#), the collection of fringe firms is modelled as a social planner maximizing the sum of producer and consumer surplus given cartel output, which is equivalent to modelling each fringe field as a price-taking competitive player. In the estimation, I abstract away from environmental damages in the optimization problem. A more detailed description of the inner loop is given in [Appendix D](#).

After completing the iterations across all sample years, the algorithm returns to the outer loop to select the next candidate parameter θ from the grid. The inner loop is then re-executed for this new candidate. As a result, for each candidate parameter, the algorithm generates a vector of simulated prices, cartel extractions, and fringe extractions.

In the final step, these simulated outcomes are compared to the actual observed data. The estimated parameter $\hat{\theta}$ is identified as the value that minimizes the distance between the simulated moments and the corresponding moments in the data.

5.2 Findings

The internalization or conduct parameter θ captures the effectiveness of the OPEC cartel in exerting market power. A value of $\theta = 1$, signifies that the cartel is fully profit-maximizing, acting cohesively to optimize joint profits. In contrast, $\theta = 0$ implies that the cartel members are not colluding. *Ceteris paribus*, a higher θ leads to higher expected initial oil prices and lower initial quantities produced by the cartel. As the cartel becomes more effective in coordinating output at higher levels of θ , it secures greater profits through enhanced market power.

Under the assumption of a constant level of conduct, the estimated internalization parameter is $\theta = 0.17$, as shown in the first column of [Table 3](#). This indicates that the cartel operates at roughly one-fifth of its potential effectiveness in coordinating joint output, in line with the findings of [Lin Lawell \(2020\)](#). Despite the relatively low conduct parameter, the cartel still manages to substantially impact market share and prices compared to a perfectly competitive market, as demonstrated in [Figure 16](#). My findings are robust to alternative discount factor and moment specifications, as reported in [Appendix C](#).

Several factors may contribute to limited cartel effectiveness. First, geopolitical frictions prevalent in the oil market can impede cohesive action among cartel members. Second, concerns about cartel stability—though not explicitly modeled in this framework—can significantly influence the output decisions of individual member states, preventing perfect profit optimization. Third, the cartel’s

Table 3: Internalization Parameter Results

	(1)	(2)	(3)	(4)	(5)	(6)
Conduct (θ^0)	0.17 (0.005)	0.16 (0.005)	0.20 (0.006)	0.16 (0.007)	0.15 (0.006)	0.19 (0.007)
Demand Shift $_t$ (θ^1) (mB)				0.005 (0.000)	0.003 (0.000)	0.004 (0.000)
N	30	30	30	30	30	30
Mean Conduct	0.17	0.16	0.20	0.19	0.17	0.21

Notes: SMM estimation results, $N=30$ for all estimations. Column (1) - (3) assumes a constant level of the cartel’s internalization parameter. In Column (4) - (6) the conduct is allowed to vary with the level of demand. In Column in (1) and (4), the baseline demand slope specification is used, i.e. $\hat{\beta} = -69.66$. In Column (2) and (5), the I use the lower bound estimate of the slope of demand in the estimation. In Column (3) and (6), the upper bound is used. Standard errors are computed using 50 bootstrap iterations.

pricing strategy may be influenced by limit-pricing considerations; by avoiding excessive price increases, the cartel deters fringe firms, which leads them to not fully exploit their potential influence on market prices (Andrade de Sá and Daubanes, 2016).

As shown in Column (4)-(6), the cartel colludes more effectively during periods of high demand, echoing the findings of Green and Porter (1984). These findings suggest that if overall potential gains are larger (i.e. demand is higher), players are more inclined to collude. This is indicated by the significantly positive coefficient θ^1 , suggesting that higher demand levels are associated with increased cartel conduct. Demand shifts are measured as deviations in million barrels of oil from average demand. Given that demand has generally risen over the sample period, the results imply that the cartel’s effectiveness has strengthened over time. This trend is illustrated in Figure 7, where the internalization parameter increases from approximately 0.16 in the early 1990s to 0.22 by the end of the sample period. The dashed lines represent the 95% confidence interval. The positive coefficient on the demand shift parameter may also capture other confounding factors correlated with demand over time, such as technological changes or geopolitical events that affect the degree of collusion.

The simulated moments from the model align well with the observed data moments. Figure 14 compares the actual oil prices and extraction quantities to those generated by the simulation of my parsimonious model. In Figure 14a the dashed line represents the simulated oil prices, while the solid

line depicts the realized price data. The model successfully captures the main price trends; however, it underestimates price volatility during the 2010s, particularly around the financial crisis. As a result, the simulated quantity response to the financial crisis is stronger than what is observed in the data, as shown in Figure 14b. This discrepancy is especially pronounced for fringe producers, indicated by the green line in Figure 14b. One possible explanation is the presence of startup frictions in reality. In the model, it is costless to turn oil wells on or off. In practice, firms may be reluctant to adjust production levels rapidly due to startup costs or the difficulty of halting production once a well is operational, as discussed by Anderson et al. (2018). Additionally, the financial crisis may have increased uncertainty, discouraging resource owners from accelerating extraction of their reserves. In Figure 15, I compare the yearly extraction of the cartel across different levels of conduct, confirming the good fit of the model to capture cartel behavior.

In terms of producer surplus, both the cartel and the fringe benefit from cartelization. This phenomenon is known as the *umbrella effect* in the literature on partial cartelization (see Bos and Harrington (2010)). I compare the profits in 2019 under two scenarios: a perfectly competitive market and a cartel-fringe market with an internalization parameter of $\theta = 0.17$. For the cartel, compared to the perfectly competitive scenario, static profits increased by 50%, as shown in Table 10. For the fringe, on the other hand, profits more than doubled due to cartelization, indicating that both parties gain from the artificially elevated prices.

6 Counterfactual analysis

In this section, I conduct counterfactual analyses to examine the impact of introducing carbon taxes in the oil sector starting from 2020, within the framework of an energy transition. I utilize the cartel-fringe model described in Section 2.2, employing as inputs (i) the estimated demand, cost, and internalization parameters derived from observed behavior in Section 4 and 5 and (ii) data on remaining oil field reserves, extraction costs, and ownership structures.

In the first counterfactual scenario, I assess the effects of implementing progressively increasing carbon taxes on the oil sector. The objective is to investigate the existence and magnitude of the Green Paradox, which posits that higher future carbon taxes might initially lead to lower oil prices and increased current emissions, thereby worsening environmental outcomes in the short term (Sinn, 2008; Heal and Schlenker, 2019). For this analysis, I impose a starting carbon tax equal to USD 50, following the 2021 Biden Administration, and allow it to escalate along various trajectories. The model enables me to evaluate the impact of increasing carbon taxes under different market structures,

namely perfect competition and the cartel-fringe setting. This approach allows for a comparison of productive inefficiencies across different carbon tax paths, highlighting how increasing carbon taxes may mitigate the cartel’s incentive to withhold capacity. As discussed in Section 2.3, more stringent future environmental policies could align the cartel’s extraction trajectory more closely with the competitive benchmark.

In the second counterfactual scenario, I explore the design of a welfare-maximizing carbon tax in an imperfectly competitive market. Specifically, I compare the overall welfare under a first-best setting—characterized by perfect competition and a carbon tax equal to the social cost of carbon—with various carbon tax policies in the second-best setting of a cartel-fringe market. I examine the optimal initial level and growth rate of the carbon tax in this context. While a rapidly escalating carbon tax might reduce the cartel’s market power incentives, it could also lead to a short-term increase in emissions due to the Green Paradox, generating negative welfare effects in the near term. This analysis aims to identify a carbon tax trajectory that balances these considerations to maximize overall welfare.

6.1 Carbon Taxes and the Green Paradox

Implementation To assess the impact of carbon taxes within the context of an energy transition, I simulate a global carbon tax by imposing a per-unit tax rate τ_t on oil production. This tax increases the marginal production costs for all producers, such that the marginal cost becomes $c_{it} + \tau_t$ for cartel fields $i \in \Omega$ at time t , and $c_{jt} + \tau_t$ for fringe fields $i \in \mathbb{J}$ at time t . The carbon tax is homogeneous across all fields,³³ but I consider different temporal trajectories for the tax, including constant and increasing rates over time.

In the counterfactuals, resource owners fully anticipate the entire future path of the carbon tax, allowing them to adjust their extraction strategies accordingly. I conduct these simulations under two market structures: (1) a perfectly competitive market where all producers are price-takers, and (2) a cartel-fringe market as described in Section 2.2, where non-OPEC countries act as price-taking fringe producers.

In the baseline scenario, following [Heal and Schlenker \(2019\)](#), I consider three carbon tax trajectories. The first is a constant tax over time ($\tau^1 = 0.00$), The second involves a tax increasing annually at a rate of 5 percent ($\tau^1 = 0.05$), representing a moderately escalating policy. The third trajectory features a tax increasing annually at a rate of 10 percent ($\tau^1 = 0.10$), reflecting a more aggressive policy stance. In all scenarios, the tax rate evolves according to $\tau_t = \tau^0 * (1 + \tau^1)^{t-1}$, where the initial

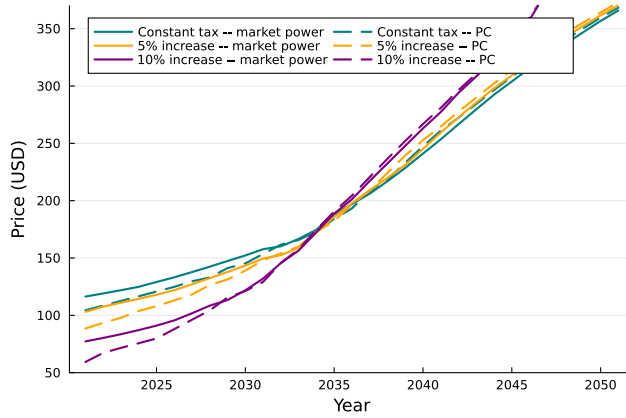
³³In reality, carbon emissions vary across fields. For an analysis of emission heterogeneity, see [Asker et al. \(2024\)](#).

tax level is set at $\tau^0 = 50$ USD per tCO_2eq , consistent with the Biden Administration’s social cost of carbon at the time. Using an average carbon content of 0.544 tCO_2eq per barrel of oil (Masnadi et al., 2018), this translates to an initial tax of 27.2USD per barrel extracted. The simulations begin in 2019—the final year of observed data. The simulation runs until $T = 2150$.


Green paradox Figure 8 illustrates the projected oil prices and total extracted quantities under different carbon tax trajectories imposed from 2019 onward. In both panels, the dashed lines represent the perfectly competitive scenarios, while the solid lines depict the outcomes under the cartel-fringe model. The different colors correspond to the various carbon tax evolution scenarios: constant, 5 percent annual increase, and 10 percent annual increase. All scenarios start with a base carbon tax rate of $\tau^0 = 50$ USD per tCO_2 .

Figure 8a shows that a more steeply increasing carbon tax initially *decreases* oil prices. Despite the carbon tax being systematically higher in the scenarios with a 5 percent or 10 percent annual increase compared to the constant tax, prices are lower during the first 15 years. This outcome arises because firms anticipate more stringent future climate policies and accelerate extraction to take advantage of lower current tax rates. This frontloading of extraction is depicted in Figure 8b. Notably, introducing a more rapidly increasing carbon tax leads to higher pollution in the short run. This finding confirms the possibility of the *Green Paradox*, a phenomenon coined by Sinn (2008), which posits that the anticipation of stronger future environmental policies can worsen environmental outcomes in the near term. However, in the long run, the carbon tax reduces overall extraction by making a portion of the oil reserves economically unviable, as shown in Figure 17. In Figure 8b, this effect is illustrated by extraction quantities declining to zero shortly after 2050.

Figure 8: Carbon tax impact under perfect competition and cartel-fringe



(a) Oil prices



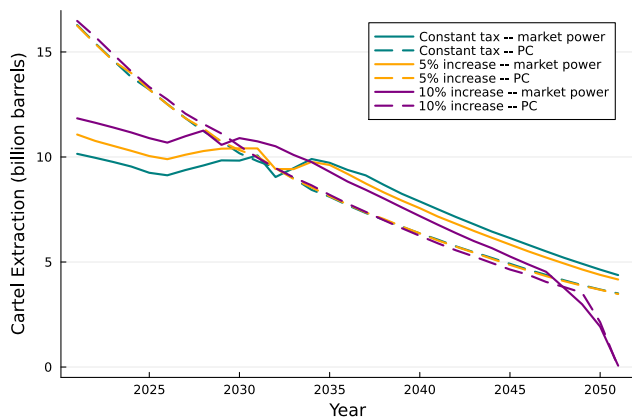
(b) Total extracted quantities

Notes: Oil prices and extracted quantities under perfect competition (dashed line) and cartel-fringe model (full line). In 2020, I impose a carbon tax of 50 USD which is constant (green), increasing at 5% (orange) or increasing at 10% (purple) yearly. The internalization parameter θ is set to 0.17 in the cartel-fringe scenario with market power. Note that the simulation runs until 2150 in all cases.

The prices under perfect competition are lower than those in the cartel-fringe scenario for the same tax rates, as shown in Figure 8a. However, the price gap narrows more quickly under the 10 percent tax increase rate (purple line) compared to the constant tax scenario (green line). The anticipation of tighter future environmental regulations prompts the cartel to expedite the extraction of its economically viable reserves. Consequently, the incentive to exercise market power by withholding

capacity to elevate future prices diminishes. As a result, the outcomes under the perfectly competitive scenario and the cartel-fringe model become more similar.

Figure 9: Cartel extraction quantities with an increasing carbon tax under perfect competition and cartel-fringe



Notes: Cartel extraction under perfect competition (dashed line) and cartel-fringe model (full line). In 2019, I impose a carbon tax of 50 USD which is constant (green), increasing at 5% (orange) or increasing at 10% (purple) yearly. This is a result of the cartel shifting production forwards to avoid stranded assets. The internalization parameter θ is set to 0.17 in the cartel-fringe scenario. Note that the simulation runs until 2150.

This "race to the last barrel" dynamic is particularly pronounced when examining the cartel's extraction behavior in Figure 9. In the perfectly competitive cases (dashed lines), the cartel's extraction is relatively consistent across the different carbon tax trajectories. The largest differences between the competitive scenarios occur by the end of the extraction period, where the more steeply increasing taxes result in an earlier final extraction period. In these scenarios, most of the additional production is undertaken by fringe firms. The opposite pattern emerges in the cartel-fringe cases. Under market power, the steeper the carbon tax increase, the greater the market share captured by the cartel. This is because delaying production becomes less profitable when the carbon tax escalates rapidly, incentivizing the cartel to extract more in the short term.

Productive inefficiencies Productive inefficiencies represent the welfare losses resulting from deviations from a cost-minimizing production allocation. When market power is exercised by withholding

efficient production capacity, higher-cost fields must be used to meet demand, generating productive inefficiencies. The introduction of carbon taxes can influence the allocation of production between the cartel and the fringe. As shown in Figure 9, a steeper increase in the carbon tax brings the cartel’s extraction path closer to that of a perfectly competitive market. As a result, we expect the welfare losses from productive inefficiencies to decrease. In a perfectly competitive market, resource owners always extract from the lowest-cost fields first, eliminating productive inefficiencies. Therefore, as the production profiles of the cartel-fringe model and the perfectly competitive case converge, productive inefficiencies decline.

Table 4 compares the productive inefficiencies resulting from different tax structures, specifically an increasing carbon tax versus a flat tax with the same quantity-weighted average. The inefficiencies are measured by comparing the net present value (NPV) of total production costs in a perfectly competitive market to those in a cartel-fringe market. In all cases, the NPV of production costs is higher under market power, indicating productive inefficiencies for both flat and increasing tax regimes. In Columns (1) and (2), I compare the effects of a carbon tax starting at 50 USD per tCO_2 and increasing at 5% annually (Column 2) with a flat tax of 156 USD per tCO_2 (Column 1). The results show that productive inefficiencies are 30.5% lower under the increasing tax. This occurs because the cartel accelerates production, aligning more closely with the perfectly competitive outcome. The incentive to withhold capacity and maintain high prices is diminished under the increasing tax structure. Columns (3) and (4) replicate the analysis for a steeper tax increase. In Column (4), a carbon tax starts at 50 USD per tCO_2 and rises by 10% annually, while Column (3) applies an equivalent flat tax of 206 USD per tCO_2 . Under this steeper tax regime, the cartel has even less incentive to delay production, resulting in a further reduction in productive inefficiencies by 73.9%.

The key takeaway from Table 4 is that carbon taxes can mitigate the welfare losses from productive inefficiencies, addressing two market failures simultaneously. In the following section, I will discuss the optimal levels of carbon taxes. It is important to note that in my model, productive inefficiencies arise solely from the misallocation of production between the cartel and the fringe. In practice, productive inefficiencies may also exist *within* the cartel and *within* the fringe firms, as documented by [Asker et al. \(2019\)](#).³⁴ Moreover, the anticipation of future carbon taxes could potentially reduce productive inefficiencies within these groups. Therefore, the estimates presented here should be interpreted as a *lower bound*.

³⁴Modeling internal frictions within the cartel and the fringe is beyond the scope of this paper.

Table 4: Impact of Carbon Taxes on Productive Inefficiencies

	Moderate Tax Increase		Steep Tax Increase	
Carbon Tax Design				
Tax base (τ^0)	156	50	206	50
Tax increase (τ^1)	0%	5%	0%	10%
Total Production Costs (b USD)				
Cartel-Fringe Model (A)	4,371	4,885	4,198	4,455
Perfect Competition Model (B)	4,116	4,784	4,006	4,405
Productive Inefficiency (b USD)				
(A) - (B)	144.2	100.3	192.5	50.1
% <i>Difference</i>		-30.5%		-73.9%

Notes: This table presents the net present value (NPV) of production costs under different carbon tax trajectories. Columns (1) and (2) compare production costs between a moderately increasing carbon tax (5%) and an equivalent flat tax. Columns (3) and (4) provide the same comparison for a steeply increasing carbon tax (10%) and its corresponding equivalent flat tax. The equivalent flat taxes are computed as the quantity-weighted averages of the increasing tax scenarios. Efficiency loss is defined as the difference in the NPV of production costs between the perfectly competitive model and the cartel-fringe model (with $\theta = 0.17$). A steeper tax increase results in a lower productive inefficiency.

6.2 The Optimal Carbon Tax

Implementation The optimal carbon tax maximizes social welfare, defined as the sum of producer surplus, consumer surplus, environmental damages, and government revenue. In the absence of market power, the optimal carbon tax equals the social cost of carbon, which represents the monetized value of future net damages from an additional ton of CO₂ emissions (Carleton and Greenstone, 2022). In 2020, the U.S. Environmental Protection Agency (EPA) estimates the social cost of carbon at 190 USD per tCO_2eq .³⁵ Under a first-best scenario with perfect competition, a carbon tax of 190 USD is levied, and since marginal environmental damages rise with emissions, the tax is expected to increase over time.

When market power is present, two key adjustments to the optimal carbon tax are necessary, as described in Section 2.3. First, in a static setting, market power reduces equilibrium quantities, so

³⁵Source: epa.gov/environmental-economics/scghg, accessed August 2024.

a lower carbon tax is required to reach the social optimum. Since firms with market power already restrict output, the tax can be lower than the social cost of carbon, and in some cases, a lower or no tax could improve welfare (Asker et al., 2024). Second, in a dynamic setting, market power affects the optimal growth rate of the tax. While rising carbon taxes in a competitive market accelerate production due to the Green Paradox, the slower extraction pace under market power means that a more rapidly increasing tax can encourage the cartel to accelerate production, thus reducing productive inefficiencies and increasing welfare.

In this counterfactual exercise, I aim to determine the optimal carbon tax trajectory in the presence of market power, representing a second-best scenario. I evaluate two cases: one with observed cartel behavior, where the internalization parameter $\theta = 0.17$, and another with a fully effective cartel, where $\theta = 1$. The tax is modelled using the function form $\tau_t = \tau_0(1 + \tau^1)^{t-1}$, where τ_0 is the base tax level, and τ^1 is the annual growth rate.

In the first-best scenario with perfect competition, the optimal tax equals the social cost of carbon, set at $\tau_0 = 190$, and $\tau_1 = 0.03$, following the EPA’s 2023 report. For robustness, I also consider an alternative specification with a lower social cost of carbon, $\tau_0 = 100$ while keeping $\tau_1 = 0.03$.

The welfare-maximizing carbon tax Table 11 presents the welfare-optimizing carbon tax trajectories across three scenarios: (1) the first-best case of perfect competition, (2) an imperfect cartel-fringe market with $\theta = 0.17$, and (3) a fully effective cartel-fringe market with $\theta = 1$. In Panel A, the social cost of carbon is set at USD 190, with an annual growth rate of 3%. In the perfectly competitive case (Column 1), the optimal tax trajectory follows $\tau_t = 190(1 + 0.03)^{t-1}$, and total environmental damages are estimated at 5.878 b USD for the period 2019–2150. By construction, tax revenues equal environmental damages in this case. In my main counterfactual, the demand curve slope is assumed constant. An alternative assumption is that the demand curve flattens over time as alternative energy sources become more substitutable. Results for this alternative approach are shown in Appendix E.

Column (2) of Panel A shows the optimal tax in the presence of an imperfect cartel ($\theta = 0.17$). The optimal initial carbon tax is reduced to 173 USD but increases at a faster annual rate of 3.3%. This steeper growth incentivizes the cartel to accelerate production, thereby eroding its market power. However, the NPV of tax revenues is lower than total environmental damages. While both the cartel and fringe benefit slightly from cartelization, consumers lose surplus, leading to a modest welfare loss in this second-best setting.

In Column (3), I explore the optimal tax in the case of a fully effective cartel ($\theta = 1$). Here, the optimal carbon tax is significantly lower, starting at 50 USD, but the growth rate is much steeper, at

5.5% annually. The rapid increase in the tax partially offsets the cartel's market power, but producer surplus is substantially higher than in the previous cases. The main welfare loss stems from reduced government revenues, as cartel profits capture a larger share of rents that would otherwise accrue to the government through taxes.

Panel B performs the same analysis under the assumption that the social cost of carbon is lower, at 100 USD per tCO_2 . In this case, welfare is generally higher due to lower environmental damages. Producers—both cartel and fringe—benefit from lower carbon taxes, while consumer surplus increases across all cases. The difference in optimal taxation between conduct levels is more pronounced when the social cost of carbon is lower. For the imperfect cartel scenario (Column 2), the initial tax decreases by 16%, from 100 USD to 84 USD. In the fully effective cartel case (Column 3), the optimal tax drops sharply by 89%, to 11 USD, with an annual growth rate of 6.0%, nearly double that of the social cost of carbon. As in Panel A, government revenues are largely transferred to producer surplus, and consumer welfare decreases by 5.2% in the fully effective cartel case compared to the first-best scenario.

Table 5: Optimal carbon tax under different market conducts

	First Best	Imperfect Cartel	Perfect Cartel
<u>Panel A: High Social Cost of Carbon</u>			
T_0	190	173	50
Tax increase (%)	3	3.3	5.5
CS	6.428	6.368	6.671
PS fringe	0.697	0.928	2.917
PS cartel	0.586	0.711	1.636
Tax Revenue	5.878	5.562	2.331
Env Damage	5.878	5.868	5.964
Total Welfare	7.713	7.703	7.592
<u>Panel B: Medium Social Cost of Carbon</u>			
T_0	100	84	11
Tax increase (%)	3	3.5	6.0
CS	7.093	7.014	6.720
PS fringe	2.047	2.072	4.047
PS cartel	1.436	1.441	2.346
Tax Revenue	3.245	3.483	0.547
Env Damage	3.245	3.278	3.207
Total Welfare	10.578	10.572	10.454

Notes: The first column displays the optimal tax rates by assumption in a first-best setting, and the NPV welfare distribution. The second column shows the case of an imperfect ($\theta = 0.17$) cartel-fringe case. The third column shows the optimal carbon tax when the cartel would operate fully effectively ($\theta = 1$), and its implied welfare consequences. In Panel A, the social cost of carbon is assumed to be 190 USD, while it is assumed to equal 100 USD in Panel B. I assume that the marginal NPV of damage of carbon increases 3% yearly in both cases.

7 Concluding remarks

To meet the climate goals set forth in the Paris Agreement, a substantial portion of existing oil reserves must remain untapped. This prospect may incentivize firms to accelerate production, anticipating more stringent future environmental regulations. Concurrently, the global crude oil market is heavily influenced by the Organization of Petroleum Exporting Countries (OPEC), an international cartel that exerts market power by restricting production to elevate prices. Consequently, climate policy, when applied in a market governed by a cartel, has the potential to mitigate environmental harm while also dampening the cartel’s incentives to restrict output.

This paper aims to assess the welfare implications of climate policy within the context of the oil market, where an oil cartel is present. To achieve this, I develop a dynamic structural model that captures the underlying mechanisms driving the production decisions of both the cartel and fringe producers. The model is estimated using comprehensive field-level data on production costs and reserves from the universe of oil fields spanning 1990 to 2020. Leveraging these estimates, I conduct counterfactual analyses under various climate policy scenarios to evaluate their welfare effects.

The global upstream oil market suffers from two major market failures: the exercise of market power and environmental externalities. Market power leads to welfare losses as it distorts production levels and causes productive inefficiencies, with oil often extracted out of a cost-minimizing sequence. Simultaneously, environmental externalities incentivize over-extraction, depending on the relative strength of market power and the extent of environmental damage. In this context, stringent climate policies—such as increasing carbon taxes—can address both issues. First, higher carbon taxes raise extraction costs, rendering some fields uneconomical to exploit and thereby reducing cumulative emissions. However, in the short term, these policies may increase emissions, a counterintuitive effect known as the green paradox. Second, because the cartel has less incentive to reduce production, some output shifts from higher-cost fringe producers to the cartel, improving productive efficiency. This reallocation partially mitigates the inefficiencies caused by market power, aligning extraction more closely with cost-minimizing production.

In this study, I develop a dynamic structural model that captures key features of the upstream oil industry. Recognizing oil’s exhaustible nature, the model incorporates dynamic extraction decisions, where producers weigh the opportunity cost of current versus future production in line with [Hotelling \(1931\)](#). I extend the dynamic cartel-fringe framework of [Salant \(1976\)](#) to account for OPEC’s influence on equilibrium prices, while allowing for potential frictions, such as geopolitical factors, that might limit OPEC’s ability to fully exercise market power.

Using granular data on field-level costs from Rystad Energy (1990-2019), I estimate marginal costs, upstream demand, and OPEC's internalization parameter. Marginal costs and demand are estimated following the approach of [Asker et al. \(2023\)](#), and serve as an input to the dynamic estimation. The internalization parameter, dynamically estimated via simulated method of moments (SMM), reflects OPEC's effectiveness and how it responds to demand cycles. I find that the cartel could be roughly five times as effective, and that cartel collusion is pro-cyclical. This approach provides a framework for analyzing market power dynamics and the potential efficiency gains from counterfactual climate policies in the global oil market.

The counterfactual analysis highlights two main findings. First, anticipating stringent climate policy can induce a green paradox, with short-run emissions rising as producers accelerate extraction in response to future carbon taxes. In the cartel-fringe setting, introducing a 50 USD carbon tax in 2020 that rises over time reduces long-run emissions but also reallocates production from fringe producers to the cartel, as the tax weakens the cartel's incentive to withhold output. This reallocation improves productive efficiency, reducing inefficiencies by 30.5% comparing a 5% increasing tax to a flat tax. Second, an optimal carbon tax can jointly address market power and environmental externalities. Assuming a social cost of carbon of 190 USD, the initial tax in a cartel-fringe setting should be slightly lower than in a first-best scenario, as the cartel already restricts output. However, to counteract the cartel's market power, the tax should rise more steeply, leading to further production reallocation from the fringe to the cartel. While this approach reduces both environmental damages and inefficiencies, it results in lower tax revenues, with surplus shifting to both the fringe and cartel producers.

This paper sheds light on the interplay between climate policy and market power in resource markets, emphasizing the implications of cartel behavior for policy effectiveness. The analysis suggests that optimal climate policy must account for the strategic incentives of an oil cartel to mitigate both environmental externalities and welfare losses from market distortions. By aligning producer incentives with environmental objectives, such policies can more effectively reduce emissions while enhancing allocative efficiency. These findings indicate that recognizing and addressing market power dynamics is crucial for designing policies that support a more efficient and sustainable energy transition.

References

- Ahlvik, L., Andersen, J., Hamang, J., and Harding, T. (2024). Quantifying supply-side climate policies. *Working Paper*.
- Anderson, S. T., Kellogg, R., and Salant, S. W. (2018). Hotelling under pressure. *Journal of Political Economy*, 126(3):984–1026.
- Andrade de Sá, S. and Daubanes, J. (2016). Limit pricing and the (in)effectiveness of the carbon tax. *Journal of Public Economics*, 139:28–39.
- Asker, J., Collard-Wexler, A., De Canniere, C., De Loecker, J., and Knittel, C. R. (2024). Do two wrongs make a right for the environment? the environmental benefits of market power in the oil sector. *NBER Working Paper*.
- Asker, J., Collard-Wexler, A., and De Loecker, J. (2019). (mis)allocation, market power, and global oil extraction. *American Economic Review*, 109(4):1568–1615.
- Asker, J., Collard-Wexler, A., and De Loecker, J. (2023). The welfare impact of market power: The opec cartel,. *Working Paper*.
- Baumeister, C. and Kilian, L. (2016). Forty years of oil price fluctuations: Why the price of oil may still surprise us. *Journal of Economic Perspectives*, 30(1):139–60.
- Benckroun, H., van der Meijden, G., and Withagen, C. (2019). An oligopoly-fringe non-renewable resource game in the presence of a renewable substitute. *Journal of Economic Dynamics and Control*, 105:1–20.
- Benckroun, H. and Withagen, C. (2012). On price taking behavior in a nonrenewable resource cartel–fringe game. *Games and Economic Behavior*, 76(2):355–374.
- Borenstein, S., Bushnell, J. B., and Wolak, F. A. (2002). Measuring market inefficiencies in california’s restructured wholesale electricity market. *American Economic Review*, 92(5):1376–1405.
- Borenstein, S. and Kellogg, R. (2022). Carbon pricing, clean electricity standards, and clean electricity subsidies on the path to zero emissions. Working Paper 30263, National Bureau of Economic Research.

- Bornstein, G., Krusell, P., and Rebelo, S. (2023). A world equilibrium model of the oil market. *The Review of Economic Studies*, 90(1):132–164.
- Bos, I. and Harrington, J. E. (2010). Endogenous cartel formation with heterogeneous firms. *The RAND Journal of Economics*, 41(1):92–117.
- Buchanan, J. (1969). External diseconomies, corrective taxes, and market structure. *American Economic Review*, 59(1):174–77.
- Caldara, D., Cavallo, M., and Iacoviello, M. (2019). Oil price elasticities and oil price fluctuations. *Journal of Monetary Economics*, 103:1–20.
- Carleton, T. and Greenstone, M. (2022). A guide to updating the us government’s social cost of carbon. *Review of Environmental Economics and Policy*, 16(2):196–218.
- Cicala, S. (2022). Imperfect markets versus imperfect regulation in us electricity generation. *American Economic Review*, 112(2):409–41.
- Corts, K. S. (1999). Conduct parameters and the measurement of market power. *Journal of Econometrics*, 88(2):227–250.
- Coulomb, R., Henriot, F., and Reitzmann, L. (2021). Bad oil, worse oil, and carbon misallocation. *Working Paper*.
- Dasgupta, P. and Heal, G. (1974). The optimal depletion of exhaustible resources. *The Review of Economic Studies*, 41:3–28.
- Dasgupta, P., Heal, G., and Stiglitz, J. E. (1980). The taxation of exhaustible resources. Working Paper 436, National Bureau of Economic Research.
- EPA (2023). Report on the social cost of greenhouse gases: Estimates incorporating recent scientific advances.
- Fowle, M., Reguant, M., and Ryan, S. P. (2016). Market-based emissions regulation and industry dynamics. *Journal of Political Economy*, 124(1):249–302.
- Gerlagh, R. (2011). Too much oil. *CEPR Economic Studies*, 57(1):79–102.
- Ghoddusi, H., Nili, M., and Rastad, M. (2017). On quota violations of opec members. *Energy Economics*, 68:410–422.

- Gilbert, R. J. (1978). Dominant firm pricing policy in a market for an exhaustible resource. *The Bell Journal of Economics*, pages 385–395.
- Gowrisankaran, G., Langer, A., and Reguant, M. (2024). Energy transitions in regulated markets. *Working Paper*.
- Green, E. J. and Porter, R. H. (1984). Noncooperative collusion under imperfect price information. *Econometrica*, 52(1):87–100.
- Groot, F., Withagen, C., and Zeeuw, A. d. (1992). Note on the open-loop von stackelberg equilibrium in the cartel versus fringe model. *The Economic Journal*, 102(415):1478–1484.
- Hamilton, J. D. (2009). Understanding crude oil prices. *The Energy Journal*, 30(2):179–206.
- Hassler, J., Krusell, P., and Olovsson, C. (2010). Oil monopoly and the climate. *American Economic Review*, 100(2):460–64.
- Heal, G. and Schlenker, W. (2019). Coase, hotelling, and pigou: the incidence of a carbon tax and co2 emissions. *NBER working paper*.
- Herfindahl, O. C. (1967). Depletion and economic theory. *Extractive Resources and Taxation: Proceedings of a Symposium at the University of Wisconsin-Milwaukee*, Madison: University of Wisconsin Press:63–90.
- Hotelling, H. (1931). The economics of exhaustible resources. *Journal of Political Economy*, 39(2):137–175.
- Igami, M. and Sugaya, T. (2022). Measuring the incentive to collude: the vitamin cartels, 1990–99. *The Review of Economic Studies*, 89(3):1460–1494.
- IPCC (2023). Climate change 2023:synthesis report. contribution of working groups i, ii and iii to the sixth assessment report of the intergovernmental panel on climate change. Technical report.
- Kellogg, R. (2024). The end of oil. *Working Paper*.
- Kilian, L. (2009). Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *American Economic Review*, 99(3):1053–69.
- Kilian, L. (2022). Understanding the estimation of oil demand and oil supply elasticities. *Energy Economics*, 107:105844.

- Lemoine, D. (2017). Green expectations: Current effects of anticipated carbon pricing. *The Review of Economics and Statistics*, 99(3):pp. 499–513.
- Lewis, T. R. and Schmalensee, R. (1980). On oligopolistic markets for nonrenewable natural resources. *The Quarterly Journal of Economics*, 95(3):475–491.
- Lin Lawell, C. (2020). Market power in the world oil market: Evidence for an opec cartel and an oligopolistic non-opec fringe. *Working Paper*.
- Liski, M. and Montero, J.-P. (2005). A note on market power in an emission permits market with banking. *Environmental and Resource Economics*, 31:159–173.
- Mansur, E. T. (2008). Measuring welfare in restructured electricity markets. *The Review of Economics and Statistics*, 90(2):369–386.
- Masnadi, M. S., El-Houjeiri, H. M., Schunack, D., Li, Y., Englander, J. G., Badahdah, A., Monfort, J.-C., Anderson, J. E., Wallington, T. J., Bergerson, J. A., Gordon, D., Koomey, J., Przesmitzki, S., Azevedo, I. L., Bi, X. T., Duffy, J. E., Heath, G. A., Keoleian, G. A., McGlade, C., Meehan, D. N., Yeh, S., You, F., Wang, M., and Brandt, A. R. (2018). Global carbon intensity of crude oil production. *Science*, 361(6405):851–853.
- McGlade, C. and Ekins, P. (2015). The geographical distribution of fossil fuels unused when limiting global warming to 2 c. *Nature*, 517(7533):187–190.
- Miller, N. H. and Weinberg, M. C. (2017). Understanding the price effects of the millercoors joint venture. *Econometrica*, 85(6):1763–1791.
- Newbery, D. M. (1981). Oil prices, cartels, and the problem of dynamic inconsistency. *The Economic Journal*, 91(363):617–646.
- Pindyck, R. S. (1978). Gains to producers from the cartelization of exhaustible resources. *The Review of Economics and Statistics*, 60(2):238–251.
- Puller, S. L. (2007). Pricing and firm conduct in california’s deregulated electricity market. *The Review of Economics and Statistics*, 89(1):75–87.
- Reguant, M. (2014). Complementary bidding mechanisms and startup costs in electricity markets. *The Review of Economic Studies*, 81(4):1708–1742.

- Reiss, P. C. and Wolak, F. A. (2007). Chapter 64 structural econometric modeling: Rationales and examples from industrial organization. volume 6 of *Handbook of Econometrics*, pages 4277–4415. Elsevier.
- Reynaert, M. (2021). Abatement strategies and the cost of environmental regulation: Emission standards on the european car market. *The Review of Economic Studies*, 88(1):454–488.
- Rotemberg, J. J. and Saloner, G. (1986). A supergame-theoretic model of price wars during booms. *The American Economic Review*, 76(3):390–407.
- Salant, S. (1976). Exhaustible resources and industrial structure: A nash-cournot approach to the world oil market. *Journal of Political Economy*, 84(5):1079–93.
- Salant, S. W. (1982). Imperfect competition in the international energy market: a computerized nash-cournot model. *Operations research*, 30(2):252–280.
- Sinn, H.-W. (2008). Public policies against global warming: a supply side approach. *International tax and public finance*, 15:360–394.
- Slade, M. E. (1992). Vancouver’s gasoline-price wars: An empirical exercise in uncovering supergame strategies. *The Review of Economic Studies*, 59(2):257–276.
- Stiglitz, J. E. (1976). Monopoly and the rate of extraction of exhaustible resources. *The American Economic Review*, 66(4):655–661.
- Ulph, A. M. and Folie, G. M. (1980). Exhaustible resources and cartels: An intertemporal nash-cournot model. *Canadian Journal of Economics*, 13(4):645–58.
- van der Ploeg, F. and Withagen, C. (2012a). Is there really a green paradox? *Journal of Environmental Economics and Management*, 64(3):342–363. 2010 Monte Verita Conference on Sustainable Resource Use and Economic Dynamics(SURED).
- van der Ploeg, F. and Withagen, C. (2012b). Too much coal, too little oil. *Journal of Public Economics*, 96(1):62–77.
- Welsby, D., Price, J., Pye, S., and Ekins, P. (2021). Unextractable fossil fuels in a 1.5° c world. *Nature*, 597(7875):230–234.

A Additional Cost Descriptives

Figure 10 compares the marginal production cost of Saudi-Arabia with the median production cost of the US and Canada throughout the years. The dashed lines represent the 75th percentile of extraction costs of producing fields. The dotted line shows the 99th percentile. The thick line indicates the yearly oil price. The Figure shows that that the 99th percentile production costs of Saudi Arabia, the leading member state of OPEC, is still considerably lower than the 75th percentile of production cost of important fringe countries, such as the US and Canada. The Norwegian production costs are the highest of the three countries, as their fields are mostly offshore. Offshore fields typically have higher production costs than conventional onshore shelf oil sources – the main production source of Saudi Arabia.

Another notable fact on Figure 10 is that production cost for Saudi have been remarkably stable throughout the time period of 1990 - 2019. This is in contrast to the production cost of the fringe countries US or Canada. Their production costs, especially in the higher percentiles, experience more fluctuations throughout the years. Figure 11 shows the aggregate production for OPEC and fringe countries during the 1990- 2019 period, together with the total incurred production costs. The Figure shows that production levels of both OPEC countries and fringe countries have been steadily growing since 1990. Production costs, however, show a much more volatile profile. This is, however, only true for the fringe countries. Total OPEC oil production costs have been much more stable throughout the years.

B Cartel and fringe maximization problem

A more detailed derivation of the optimality conditions of the fringe and cartel are provided in this Section. For simplicity, I assume that the cartel operates effectively for this derivation exercise, i.e. $\theta = 1$. This assumption can be adjusted in a straightforward way by changing the marginal revenue first order condition. I follow the model of [Salant \(1976\)](#), as described in [Section 2.2](#).

B.1 Fringe Optimization problem

In my setting, fringe fields are perfectly competitive. Hence, they take the price – as announced by the cartel as given when optimizing their surplus. Hence, for the competitive fringe fields, indexed by j , the objective is to maximize the present discounted value of their profits over time, treating the price as exogenous.

The Lagrangian formulation for the fringe fields j is:

$$\mathbb{L} = \sum_{t=1}^T \beta^t \mathbb{E} [P_t q_{jt} - c_{jt} q_{jt} - \psi_{jt} (q_{jt} - \kappa_{jt}) - \gamma_{jt} (R_{j,t-1} - q_{jt} - R_{jt}) + \mu_j R_{j,T}] \quad (16)$$

Where:

- P_t : The market price at time t , treated as given by the fringe.
- q_{jt} : Production of fringe field j at time t .
- c_{jt} : Marginal production cost for fringe field j .
- κ_{jt} : Capacity limit of fringe field j , and ψ_{jt} is the Lagrange multiplier for the capacity constraint.
- R_{jt} : Resource stock remaining in field j at time t , with γ_{jt} as the multiplier for the stock depletion constraint.
- μ_j : Multiplier for the terminal condition at T , ensuring the final resource level $R_{j,T}$ is economically viable.

As with the cartel fields, the fringe fields are subject to:

1. **Capacity constraint:** Fringe production cannot exceed κ_{jt} , the capacity of field j .
2. **Resource stock constraint:** Resource stock depletes over time, $R_{jt} = R_{j,t-1} - q_{jt}$.

3. **Terminal condition:** Ensuring that resources at the end of the time horizon T are either fully extracted or economically unrecoverable.

The first-order condition for production q_{jt} for optimality for fringe fields can be derived as follows:

$$\frac{\partial \mathbb{L}}{\partial q_{jt}} = P_t - c_{jt} - \psi_{jt} - \gamma_{jt} = 0 \quad (17)$$

For producing fringe fields, the marginal revenue (which is simply the price P_t in this case, since fringe firms take prices as given) equals the marginal cost of production, adjusted for capacity and resource constraints (ψ_{jt} and γ_{jt}).

If (non-empty) fringe field j is not producing, then:

$$\frac{\partial \mathbb{L}}{\partial q_{jt}} \leq 0 \quad \text{for } t = 1, \dots, T \quad (18)$$

The Lagrange multiplier ψ_{jt} ensures that the production from any fringe field j does not exceed its capacity κ_{jt} :

$$\frac{\partial \mathbb{L}}{\partial \psi_{jt}} = \kappa_{jt} - q_{jt} = 0 \quad (19)$$

As with the cartel fields, this constraint binds for all non-marginal fringe fields producing at capacity.

The shadow price γ_{jt} reflects the value of the remaining resource in each fringe field. For fringe fields that are marginally producing (i.e., not constrained by capacity), Equation 20 can be rewritten as 21 – scarcity value should rise at the rate of interest.

$$\frac{\partial \mathbb{L}}{\partial R_{jt}} = -\gamma_{jt} + \beta \gamma_{j,t+1} = 0 \quad \text{for } t = 1, \dots, T - 1 \quad (20)$$

$$P_t - c_{jt} = \beta(P_{t+1} - c_{jt+1}) \quad (21)$$

This ensures that the marginal fringe fields are indifferent between extracting now or later, as the shadow price of the resource increases at the rate of interest

At the end of the time horizon, fringe fields must account for the remaining resources, ensuring that they are either fully extracted or economically recoverable. This gives the following condition for the final period T :

$$\frac{\partial \mathbb{L}}{\partial R_{jT}} = -\gamma_{jT} + \mu_j = 0 \quad (22)$$

The fringe firms optimize their production based on the exogenously given market price. The first-order conditions highlight that, for the fringe, the trade-off is simply between the current price and the opportunity cost of production, subject to capacity and resource constraints. The scarcity rent increases at the rate of interest, and the final resource extraction decision at T depends on the terminal value of the remaining stock.

B.2 Cartel Optimization Problem

The cartel optimizes its production decisions over time, taking into account the market demand, aggregate fringe production, production costs, capacity constraints, and the resource reserve dynamics. The cartel's objective is to maximize the present discounted value of its profit over a finite time horizon T , while adhering to both capacity and resource constraints.

The Lagrangian formulation of the cartel's problem is given by:

$$\mathbb{L} = \sum_{t=1}^T \beta^t \mathbb{E} \left[P_t(\bar{Q}_t^f, Q_t^c) Q_t^c - \sum_{i \in \Omega} (c_{it} q_{it} + \psi_{it}(q_{it} - \kappa_{it}) + \gamma_{it}(R_{i,t-1} - q_{it} - R_{it}) + \mu_i R_{i,T}) \right] \quad (23)$$

Here, the key terms represent:

- β : Discount factor.
- $P_t(\bar{Q}_t^f, Q_t^c)$: Market price as a function of total output, where \bar{Q}_t^f is the output from the competitive fringe, and Q_t^c is the cartel's output.
- q_{it} : Production from cartel field i in period t .
- c_{it} : Marginal cost of production for field i at time t .
- κ_{it} : Capacity limit of field i , and ψ_{it} is the Lagrange multiplier for the capacity constraint.

- R_{it} : Resource stock remaining in field i at time t , with γ_{it} being the multiplier for the stock depletion constraint.
- μ_i : Multiplier for the terminal condition, where $R_{i,T}$ is the stock remaining at the end of the time horizon.

The Lagrangian incorporates the following key constraints:

1. **Capacity constraint:** Production from any field i cannot exceed its capacity κ_{it} .
2. **Resource stock constraint:** The stock of resources depletes over time as production occurs, so $R_{it} = R_{i,t-1} - q_{it}$.
3. **Terminal condition:** The stock of resources in the final period T must either be zero (if fully extracted) or economically unrecoverable.

The Kuhn-Tucker theorem provides the necessary conditions for optimality. These include the following first-order condition for production q_{it} for producing fields:

$$\frac{\partial \mathbb{L}}{\partial q_{it}} = P_t + \frac{\partial P_t}{\partial q_{it}} q_{it} - c_{it} - \psi_{it} - \gamma_{it} = 0 \quad (24)$$

This condition states that, for a producing field, the marginal revenue (price plus the price effect of additional output) must equal the marginal cost of production, adjusted for the capacity and resource constraints (represented by ψ_{it} and γ_{it}).

If non-empty field i is not producing, we have:

$$\frac{\partial \mathbb{L}}{\partial q_{it}} \leq 0 \quad \text{for } t = 1, \dots, T \quad (25)$$

The Lagrange multiplier ψ_{it} is associated with the capacity constraint, which ensures that production does not exceed the field's capacity κ_{it} :

$$\frac{\partial \mathbb{L}}{\partial \psi_{it}} = \kappa_{it} - q_{it} = 0 \quad (26)$$

This constraint is binding for all non-marginal fields that are producing at full capacity. At most one producing field can have a slack capacity constraint (i.e., not producing at full capacity) at any point in time.

The shadow price γ_{it} reflects the value of the remaining resource in each field. Equation 27 implies that for marginally producing fields (fields not facing binding capacity constraints), the marginal revenue should increase at the rate of interest, as shown in Equation 28.

$$\frac{\partial \mathbb{L}}{\partial R_{it}} = -\gamma_{it} + \beta\gamma_{i,t+1} = 0 \quad \text{for } t = 1, \dots, T-1 \quad (27)$$

$$P_t + \frac{\partial P_t}{\partial q_{it}} q_{it} - c_{it} = \beta(P_{t+1} + \frac{\partial P_{t+1}}{\partial q_{it+1}} q_{it+1} - c_{it+1}) \quad (28)$$

This implies that the cartel is indifferent between extracting today or delaying extraction to future periods for marginally producing fields. The shadow price of the resource grows at the rate of interest, reflecting the increasing opportunity cost of delaying extraction.

At the end of the time horizon, the stock of resources in each field must either be depleted or economically recoverable, represented by the following condition for the final period T :

$$\frac{\partial \mathbb{L}}{\partial R_{iT}} = -\gamma_{iT} + \mu_i = 0 \quad (29)$$

This ensures that any remaining resources at the end of the horizon are valued appropriately and extracted optimally.

Table 6: First-stage Demand Estimation Results

	Dependent Variable: Prices (USD)			
	(1)	(2)	(3)	(4)
Cost Shifter _{<i>t</i>}	19.65 (4.23)	19.64 (4.09)	7.94 (5.06)	5.315 (1.19)
Controls	✓	✓	✓	✓
<i>N</i>	30	30	30	50
<i>F</i> -statistic	34.01	34.26	69.61	41.53

Notes: the First-Stage results of the estimation of Equation 13. In Column (1), GDP is the only control variable. In Column (2), a dummy variable is added for each decade. In Column (3) and (4), I a polynomial time trend, and world GDP. In the last Column, the time series starts already in 1970. All reported standard errors are computed correcting for heteroskedasticity and autocorrelation using Newey-West, with a small sample correction, and rely on a 1 period lag autocorrelation structure. The 1st stage Stock and Yogo *F*-statistic is reported.

C Alternative specifications empirics

C.1 Demand

Table 7: Demand Estimation Results – Alternative Instrument

	Dependent Variable: Quantities (mB)			
	(1)	(2)	(3)	(4)
	OLS	IV	IV	IV
Price	-5.05 (9.21)	-51.41 (12.45)	-65.68 (92.67)	-68.57 (12.10)
Controls	✓	✓	✓	✓
2SLS		✓	✓	✓
N	30	30	30	50
F -statistic		25.10	63.60	39.67
Mean elasticity	-0.01	-0.12	-0.13	-0.15

Notes: The IV specification relies on the alternative cost shifter IV, which does not include operating expenses. I compute the shifter for a production level of 15 billion barrels, as described in Section 4. Controls in (1), (3), (4) include a polynomial time trend, and world gdp. In Column (2), the time trend is replaced by a dummy variable for each decade. In the last Column, the time series starts already in 1970. All reported standard errors are computed correcting for heteroskedasticity and autocorrelation using Newey-West, with a small sample correction, and rely on a 1 period lag autocorrelation structure. The 1st stage Stock and Yogo F -statistic is reported.

C.2 Internalization Parameter Estimates

Table 8: Robustness check internalization parameters: discount rates

	(1)	(2)	(3)
Conduct (θ_0)	0.17 (0.005)	0.21 (0.006)	0.16 (0.005)
N	30	30	30
Discount rate	0.95	0.93	0.97

Notes SMM estimation results, $N=30$ for all estimations. Column (1) assumes a discount rate of 0.95, which is the baseline specification. Column (2) displays the result for a discount rate of 0.93, while Column (3) shows the results for a 0.97 discount rate.

Table 9: Robustness check internalization parameters: moment selection

	(1)	(2)	(3)	(4)
Conduct (θ_0)	0.17 (0.005)	0.21 (0.006)	0.017 (0.006)	0.17 (0.005)
N	30	30	30	30
Discount rate	0.95	0.95	0.95	0.95

Notes SMM estimation results, $N=30$ for all estimations. Column(1) shows the baseline specification, where prices, total output and fringe output are targeted moments. In Column (2) I target cartel market share as a moment, instead of fringe quantity. In Column (3), I add squared fringe and total quantities to the main targeted moment. In Column (4), I target the baseline moments, squared fringe and total quantities, and cartel market share.

Table 10: Comparison of Profits in 2019: the Umbrella-Effect

	Perfect Competition	Cartelization	% Change
Cartel Profits (in b\$)	798	1,244	55%
Fringe Profits (in b\$)	2,134	6,395	199%

Notes: The table compares cartel and fringe NPV profits before and after cartelization in 2019, expressed in billion USD. Both the fringe, as well as the cartel gain from cartelization, known as the umbrella-effect. The first column shows profits under perfect competition, the second column under an imperfect cartel, i.e. $\theta = 0.17$.

D Computational details

D.1 Algorithm Cartel-Fringe game

This section outlines the estimation algorithm. The outer loop searches over a predefined grid of θ values. Within each iteration, an inner loop calculates the best responses for the cartel and fringe, repeated annually to incorporate updated information.

1. Construct a predefined grid of potential values for θ , denoted as $\theta = \theta^1, \dots, \theta^A$. Begin by selecting an initial value θ^i from this grid.
2. Given θ^i , initialize the information and expectation set for the first year, $t = 1$. Specify an initial guess for the fringe production path $Q^{0,f}$ and the cartel production path $Q^{0,c}$ across all subsequent time periods.
3. Based on the initial fringe production guess $Q^{0,f}$, solve the cartel's optimization problem with respect to excess demand using mixed integer programming, ensuring that Equation 3 holds for the marginal field, following a similar approach as in [Asker et al. \(2023\)](#).³⁶ The output includes the cartel's extraction path $Q^{1,c}$ and the associated equilibrium price path $P^{1,c}$.
4. Using the cartel production path $Q^{1,c}$, compute the competitive fringe's response, as outlined in [Asker et al. \(2023\)](#). This step yields the fringe's output path $Q^{1,f}$ and price path $P^{1,f}$.
5. Check if the resulting price paths satisfy $|P^{1,f} - P^{1,c}| \leq tol$. If they do, proceed to the next step. If not, update the quantity guesses moderately and return to step 3.

³⁶Alternatively, an artificial maximization problem formulation may be employed, as in [Salant \(1982\)](#).

6. Update the information set for the following year, and repeat the steps from 2 onward for each subsequent period through 1991–2019. Keep track of reserves.
7. After solving for all time periods in the sample, return to step 1, select the next value θ^{i+1} from the grid, and repeat the process.

Algorithm 1 Pseudocode of the Cartel-Fringe Game

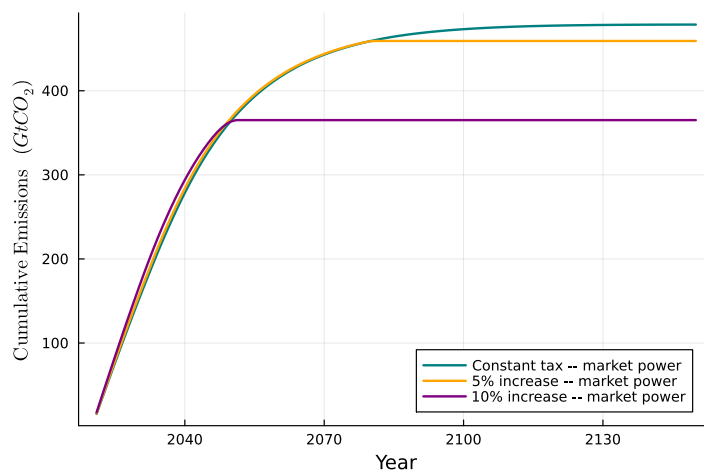
```

Initialize  $guess \leftarrow zeros(2, T)$ 
Set  $iter \leftarrow 1$ 
Set  $crit \leftarrow 10000$ 
while  $iter \leq maxiter$  and  $crit > tol$  do
   $oldguess \leftarrow guess$ 
  for  $n \in \{c, f\}$  do
     $newguess[n] \leftarrow \arg \max_{q_n} \sum_{t=1}^T \Pi_n(q_n, guess_{-n})$ 
  end for
   $crit \leftarrow ||newguess - guess||$ 
   $guess \leftarrow newguess$ 
   $iter \leftarrow iter + 1$ 
end while

```

E Additional Counterfactual

Figure 17: Cumulative emissions ($GtCO_2$) from the oil sector under different carbon tax regimes



Notes: The cumulative emissions with under different carbon tax regimes, where the baseline is $\tau^0 = 50$ USD. The market is modelled as a cartel-fringe game, where $\theta = 0.17$. The annual increase of the carbon tax matters for emissions in the long run.

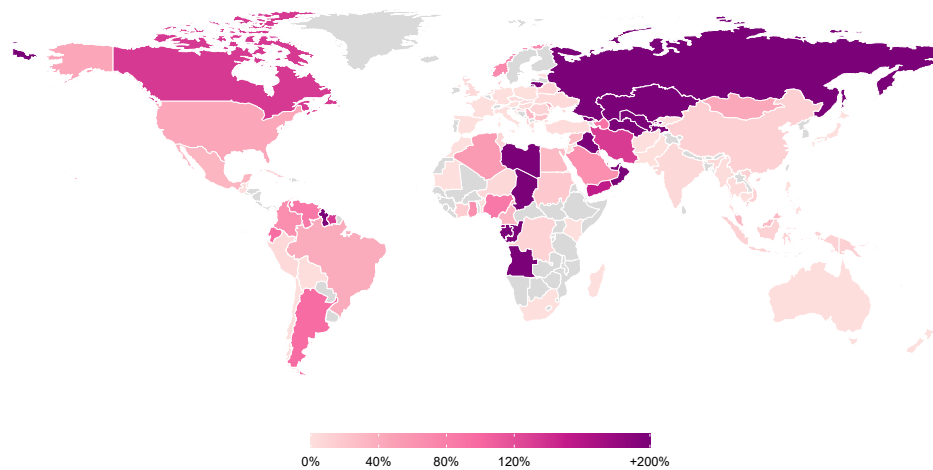
Table 11: Robustness checks: Optimal carbon tax

	First Best	Imperfect Cartel	Perfect Cartel
<u>Panel A: Increasingly elastic demand</u>			
T_0	190	175	100
Tax increase (%)	3	3.4	6.0
CS	3.523	3.470	3.790
PS fringe	0.058	0.203	0.725
PS cartel	0.123	0.191	0.209
Tax Revenue	3.670	3.455	2.538
Env Damage	3.670	3.637	3.725
Total Welfare	3.704	3.684	3.538

Notes: In this robustness check, the demand curve rotates, such that the slope of the demand becomes 5% more elastic each year. The first column displays the optimal tax rates by assumption in a first-best setting, and the NPV welfare distribution. The second column shows the case of an imperfect ($\theta = 0.17$) cartel-fringe case. The third column shows the optimal carbon tax when the cartel would operate fully effectively ($\theta = 1$), and its implied welfare consequences. The social cost of carbon is assumed to be 190 USD, and I assume that the marginal NPV of damage of carbon increases 3% yearly.


F Additional Figures

Figure 6: Oil Reserves at Risk as a Share of GDP Under a 2°C Carbon Budget




Notes: Stranded oil reserves by country when imposing a 2°C carbon budget as of 2019, as a percentage of GDP. To compute this measure, I assume that oil is extracted according to marginal costs of the fields. Oil is valued at 65 USD per barrel. Saudi-Arabia faces a risk of 65% of their GDP, while the Russia's and the US' risk is 370% and 45% of their 2019 GDP respectively. In absolute terms (million barrels of oil), Russia and the US are facing the largest risks. Only active fields are included. I assume that the share of oil within the carbon budget remains identical.

Figure 7: Yearly Cartel Internalization Parameter



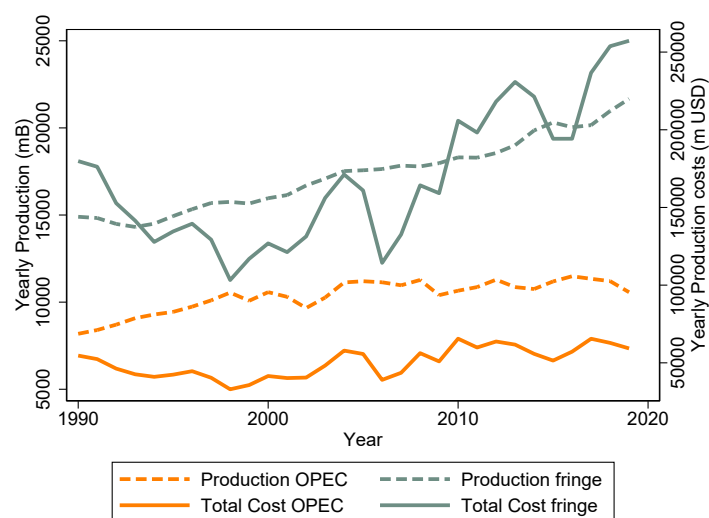
Notes: Yearly implied conduct level with the estimates of Table 3, i.e. $\theta = \theta_0 + \theta_1 \text{ Demand Shift}_t$. In this figure, the changes in conduct are driven by demand shifts. A perfect cartel would imply that $\theta = 1$. The dashed lines show the 95% confidence interval, computed from 50 bootstrap iterations.

Figure 10: Oil Price and Marginal Cost Evolution (Saudi Arabia, US, Canada)



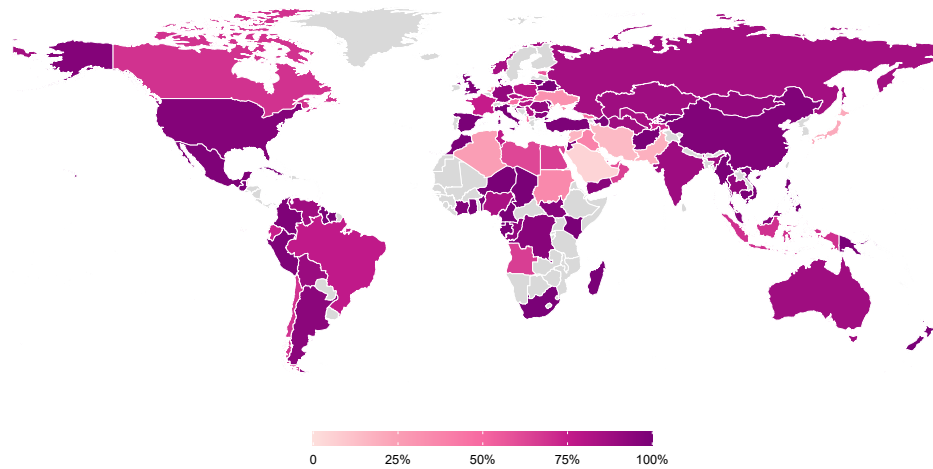
Note: The thick line shows the yearly real oil price. The dashed lines indicate the 75th percentile of marginal cost of production, the dotted line the 99th percentile of Saudi Arabia, US and Canada.

Figure 11: Total Production and Total Costs 1990- 2019




Notes: Yearly production (mB, left axis) and yearly overall production costs (USD, right axis) are shown for OPEC and fringe countries. The production of both OPEC and non-OPEC countries are fairly stable throughout the years. Production costs are much more volatile for non-OPEC countries than OPEC countries. Total production costs are calculated by multiplying the marginal extraction costs with produced quantities, at the field-level.

Figure 12: Percentage of reserves (%) at risk by country at a 2 degree carbon budget



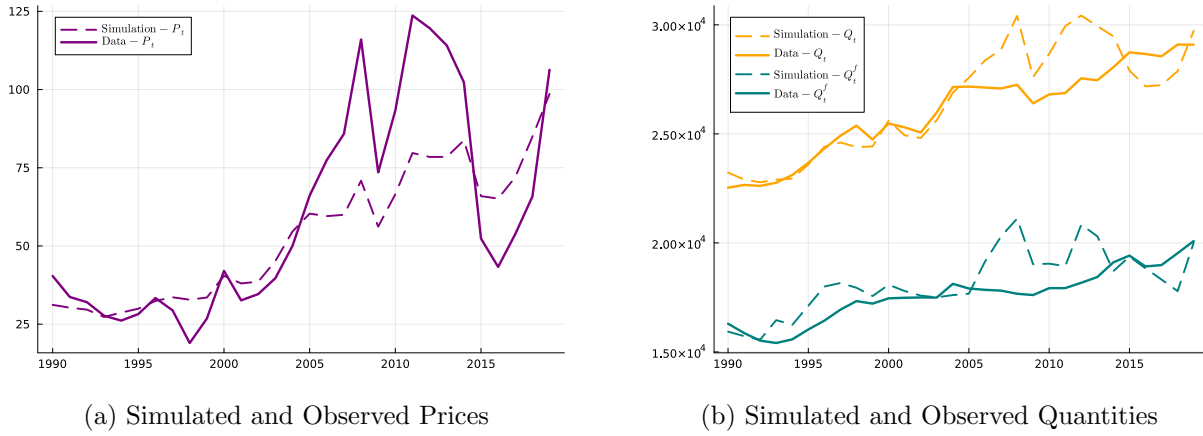
Notes: Stranded oil reserves (% of total reserves) by country when imposing a 2°C carbon budget as of 2019, and extracting according to marginal costs of the fields. Saudi-Arabia only has 4.5% of reserves at risk, while 98% of US oil reserves would be stranded. In absolute terms (million barrels of oil), Russia and the US are facing the largest risks. Only active fields are included. I assume that the share of oil within the carbon budget remains identical.

Figure 13: Yearly Demand Elasticity (1990-2019)




Notes: Yearly implied demand elasticity with an estimated oil price coefficient of $\hat{\beta} = -48.84$. Demand for oil remains highly inelastic throughout the considered time period from 1990 -2019. The 95% confidence interval is given by the dashed lines, as computed correcting for heteroskedasticity and autocorrelation using Newey-West, with a small sample correction.

Figure 14: Model fit key moments 1990 - 2019




Notes: Figure presents the key moments from the data and the model simulated at the estimated parameter values. In the "data" lines, I use the observed prices and quantities. In the "simulated" lines, I apply the time varying $\theta = \theta_0 + \theta_1 \text{ Demand Shift}_t$ following the estimates of Table 3. The left hand panel shows yearly oil prices. The right hand panel shows total extracted quantities (orange, Q_t), and total fringe quantities (teal, Q_t^f). In the simulation, the final period T is set to 2150.

Figure 15: Yearly extraction oil barrels by the cartel (1990-2019)



Notes: Yearly cartel quantity under different levels of conduct. The stronger the cartel effectiveness, the lower the extracted quantities by the cartel. In reality, however, $\theta = 0.17$ captures the reality best.

Figure 16: Yearly oil prices (1990-2019)



Notes: Yearly oil prices under different levels of conduct. The stronger the cartel effectiveness, the higher oil prices. Note that in perfect competition, shocks seem to be more smoothed out, as compared to the case where $\theta = 1$. Despite the relatively low internalization parameter, the influences on oil prices is important.