



MIT CEEPR
Center for Energy and
Environmental Policy Research

**Working Paper
Series**

Two Wrongs Can Sometimes Make a Right: The Environmental Benefits of Market Power in Oil

J. Asker, A. Collard-Wexler, C. De Cannière,
J. De Loecker, and C. R. Knittel



MARCH 2025

CEEPR WP 2025-06

Working Paper Series.

Since 1977, the Center for Energy and Environmental Policy Research (CEEPR) has been a focal point for research on energy and environmental policy at MIT. CEEPR promotes rigorous, objective research for improved decision making in government and the private sector, and secures the relevance of its work through close cooperation with industry partners from around the globe. Drawing on the unparalleled resources available at MIT, affiliated faculty and research staff as well as international research associates contribute to the empirical study of a wide range of policy issues related to energy supply, energy demand, and the environment.

An important dissemination channel for these research efforts is the MIT CEEPR Working Paper series. CEEPR releases Working Papers written by researchers from MIT and other academic institutions in order to enable timely consideration and reaction to energy and environmental policy research, but does not conduct a selection process or peer review prior to posting. CEEPR's posting of a Working Paper, therefore, does not constitute an endorsement of the accuracy or merit of the Working Paper. If you have questions about a particular Working Paper, please contact the authors or their home institutions.

Two Wrongs Can Sometimes Make a Right: The Environmental Benefits of Market Power in Oil*

J. Asker, A. Collard-Wexler, C. De Cannière, J. De Loecker, C. R. Knittel

Abstract

Market power reduces equilibrium quantities and distorts production, typically causing welfare losses. However, as Buchanan (1969) noted, market power may mitigate overproduction from negative externalities. This paper examines this in the global oil market, where OPEC's market power affects oil production and carbon intensity. We estimate that from 1970 to 2021, OPEC's market power reduced emissions by over 67 $GtCO_2$, equating to \$4,073 billion in climate damages and 17.8% of the carbon budget needed for the 1.5° C Paris Agreement target. This environmental benefit outweighs the welfare loss from distorted production allocation.

Keywords: *Carbon emissions; market power; oil; OPEC*

*Authors are listed alphabetically. *Asker:* UCLA and NBER johnasker@econ.ucla.edu, *Collard-Wexler:* Duke University and NBER, collardwexler@gmail.com, *De Cannière:* KU Leuven charlotte.decanniere@kuleuven.be, *De Loecker:* KU Leuven and CEPR jan.deloecker@kuleuven.be, *Knittel:* MIT and NBER, knittel@mit.edu.

“The monopolist is the environmentalist’s best friend”

—often attributed to *Milton Friedman*

1 Introduction

The burning of fossil fuels leads to climate change. This represents a classic negative externality in economics where firms will have an incentive to overproduce, thereby reducing societal welfare. However, market structures across the different fossil fuels vary considerably. In particular, the oil market is characterized by the presence of a global cartel, which seeks to reduce the overall quantity of oil produced to increase profits. Buchanan (1969) was the first to note that in the presence of a negative externality, market power might have unexpected implications for overall welfare since it acts to undo the incentive to overproduce under the externality, echoing the more general discussion in Lipsey and Lancaster (1956). The presence of the Organization of the Petroleum Exporting Countries (OPEC) in the global oil market is an excellent example of this. The increase in market power due to the cartel acts to reduce the amount of greenhouse gas emissions from oil consumption. However, Asker et al. (2019) documents a second effect from OPEC that can impact greenhouse gas emissions; OPEC distorts where oil is produced inside and outside the cartel. This production distortion may also affect the average lifecycle emissions of oil consumption.

This paper sheds light on the complex dynamics between market structure, emissions, and overall welfare by investigating the relationship between OPEC’s presence and environmental externalities. We estimate how the presence of OPEC affects the greenhouse gas emissions associated with oil production and consumption. Using data from Rystad Energy (Rystad hereafter), a Norwegian-based energy consultant, and the lifecycle emissions from different types of oil worldwide, we disentangle how market power in the oil market affects total carbon emissions through (i) limiting the total consumption volume and (ii) changing the carbon intensity per barrel of production.

Production decisions in the upstream oil industry matter for global environmental outcomes. The 2015 Paris Agreement pledged to limit global warming well below 2°C and to make efforts to limit the temperature increase to 1.5°C compared to pre-industrial times. To achieve these objectives, about 25% of oil reserves would need to remain untapped (McGlade and Ekins, 2015). Hence, in order to align with the climate targets in the Paris Agreement, oil extraction needs to be slowed down. Without a globally coordinated agreement to curtail greenhouse gas emissions, other market frictions, such as market power, might limit oil consumption by increasing prices. But, as noted, there may be two countervailing effects of market power. This paper is the first to compare the welfare loss from market power to the potential environmental gains in the upstream oil sector. Our paper is not the first to empirically analyze the net

impact of externalities and existing price distortions, however. Fowlie et al. (2016) analyzes the efficiency impacts of several allowance-allocation methods of a cap-and-trade system applied to the cement industry. They find that, given the existence of unilateral market power, allocation methods that do not account for market power may actually reduce welfare relative to not regulating the market altogether. Davis and Muehlegger (2010) and Borenstein and Bushnell (2022) compare regulated retail prices for natural gas and electricity, respectively, to their social marginal costs and find that regulated prices can often exceed social marginal costs. In addition to the obvious difference in industry focus, our paper differs from these by studying the relative efficiency of an international cartel and perfectly competitive market and explicitly measuring the additional distortions resulting from shifting production from the cartel to a potentially dirtier competitive fringe.

To estimate the effects of OPEC’s market power, and more generally market power in the oil market itself, on global emissions, we compare the realized emissions to counterfactual emissions in a perfectly competitive setting. By comparing to a competitive market benchmark, we incorporate distortions other than those imposed by OPEC in the analysis. Thus, while OPEC is understood to be the primary distortion in the global oil market, residual distortions arising from ad-hoc taxes and subsidies, political economy considerations, and the like are also accounted for.¹ This is important, as the underlying objective of the paper is to understand how other classical market distortions mitigate the climate change externality arising from oil extraction and consumption. This requires several steps. First, we must fully characterize the global oil supply curve to construct a counterfactual production path. To this end, we hinge on yearly observed marginal cost and production capacity data on virtually all oil fields across the globe from 1970 onwards from Rystad, in the spirit of Asker et al. (2019). Second, since total oil production is an equilibrium outcome, we estimate consumers’ responsiveness to prices, leveraging cost shifters as in Asker et al. (2023). Third, we complement this analysis by including country-level data on the carbon intensity of oil extraction and refining from Masnadi et al. (2018). As noted by Coulomb et al. (2021), the lifecycle environmental outcome of oil depends on who produces it, given the geological and technological heterogeneity of oil fields. These three elements allow us to evaluate the value of the differences between the realized and counterfactual total environmental damages.

We can decompose the total environmental impact of market power in the upstream oil industry into two opposing effects, as noted by Bencheikroun et al. (2020). First, market power restricts total quantities to drive up prices. As a result, fewer barrels of oil are extracted and consumed, leading to environmental benefits; we refer to this as the *Volume-Effect*. Second, strategic production decisions affect aggregate oil extraction’s average lifecycle carbon intensity.

¹See, for example, Balke et al. (2015) and Aune et al. (2017) (taxes and subsidies), and Fattouh et al. (2013), Kilian and Murphy (2014), and Knittel and Pindyck (2016) (speculation).

This results from the fact that lifecycle carbon intensity positively correlates with extraction costs. Therefore, if the international OPEC cartel limits its total capacity of production, dirtier fields from fringe firms will be used to satisfy demand. Furthermore, the oil from these higher-cost fields also tends to be heavier, requiring additional energy to refine. This also increases lifecycle emissions. As a result, the market power of OPEC can lead to environmental losses through what we refer to as the *Composition-Effect*. On net, we find positive net environmental effects from OPEC.

We find that the market power has avoided releasing 67,738 $MtCO_2$ into the atmosphere from 1970-2021. This is a meaningful reduction of carbon emissions, the equivalent of three years of current oil consumption. The net present value of the total avoided damages amounts to \$4,073 billion, valued at a social cost of carbon of \$250/ tCO_2 , substantially larger than the total welfare loss from non-environmental related damages from market power, as quantified in Asker et al. (2023). Our results show that the environmental gain of market power through reduced overall oil extraction (i.e., the Volume-Effect) dwarfs the impact of increased carbon intensity of extraction (i.e., the Composition-Effect). The reduction in overall emissions can be mapped to changes in atmospheric temperature. Based on the DICE model (Nordhaus, 2014), we find that a 0.023 $^{\circ}C$ increase in atmospheric temperature anomaly has been avoided compared to the perfectly competitive scenario in the upstream oil market. This represents 17.8% of the remaining carbon budget for meeting the 1.5 $^{\circ}C$ target of the Paris Agreement of 2015 with a 50% probability using the estimates from Friedlingstein et al. (2022).

It is important to note that this paper does not advocate the abolition of environmental policy, nor does our paper advocate for the existence of OPEC. Environmental policy is undoubtedly required to address climate change. Furthermore, we make no statements about the impact of OPEC on societal welfare. For example, we do not estimate the effect of OPEC on consumer welfare or the transfers of wealth from oil-consuming countries to oil-producing countries; our focus, therefore, is on global welfare. These are important issues that would need to be addressed to make complete, policy-relevant welfare statements about OPEC and market power in the crude oil market more generally.

2 Analytical framework

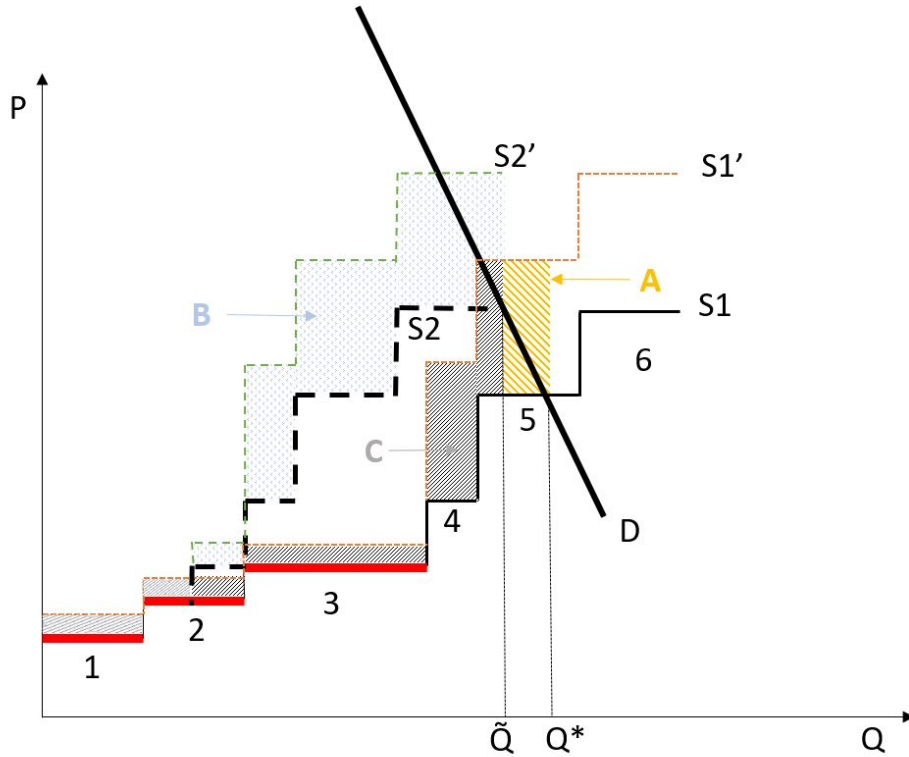
This section presents an analytical model of how an international cartel (OPEC) affects global emissions in the oil sector.

Market power shapes production decisions The stylized Figure 1 shows how market power in the upstream oil market can shape production decisions in the oil industry. Figure 1 compares the observed equilibrium in the upstream oil industry to the equilibrium that a

perfectly competitive market would reach. In this static example, the market consists of 6 oil fields with increasing marginal extraction costs. That is, field 1 has the lowest marginal cost of extraction, while field 6 has the highest cost of extraction. The fields owned by the cartel are indicated in red (i.e., fields 1, 2, and 3).

Supply schedule S1 shows the production decisions in a perfectly competitive upstream oil market. In that case, fields will operate according to marginal costs until the schedule intersects with the demand curve. In this setting, perfect competition implies that fields 1 - 4 extract at full production capacity while field 5 produces just under full capacity, until intersection with demand curve D. The total number of barrels extracted equals Q^* in the competitive equilibrium.

Figure 1: The emission impact of market power



Note: Observed (S2) and perfectly competitive (S1) supply schedule in the upstream oil market. Including emission cost into the respective supply schedules (S2' and S1') allows overall emissions to be decomposed into the (i) Volume-effect (A), and (ii) Composition-effect (B-C).

Now consider what happens in the presence of the cartel as indicated by supply schedule S2. The cartel can increase profits by not using field 3 and by only partially utilizing field 2. This shifts where field 4 and 5 are in the supply curve as they start producing at a lower quantity than before. More dramatically, it leads to production from field 6, which is now setting the

price.² The total extracted barrels decrease to \tilde{Q} , while the world oil price is inflated. Hence, in this stylized example, the cartel is raising overall profits by creating artificial scarcity.³

Two opposing effects of market power on emissions Production decisions in the upstream oil market matter for environmental outcomes. Market power can impact total emissions by two opposing effects. Firstly, market power can lower total emissions by decreasing the overall extracted quantity of oil—the *Volume-Effect*. By exercising market power, the cartel lowers overall equilibrium quantities from Q^* to \tilde{Q} . As a result, market power could decrease emissions associated with the production and consumption of oil by lowering the amount of barrels extracted. Secondly, market power can increase total emissions by changing *who* is producing—the *Composition-Effect*. Changing the production allocation will affect the environmental outcome if emission content is not homogeneous across oil fields. Overall, emission intensity from extraction at an oil field positively correlates with extraction costs as the higher marginal cost fields tend to be more energy intensive. In addition, oil’s viscosity tends to be positively correlated with extraction costs; high marginal cost oil tends to be heavier.⁴ Therefore, withholding cheaper, cleaner production capacity could increase total emissions.

A formal representation of the two opposing effects of market power on emissions

Let the set of extractable barrels J be ranked according to marginal costs. The vectors q_j^{obs} and q_j^{pc} indicate the production (0/1) of the j^{th} barrel according to the observed supply schedule S2, and the perfectly competitive supply schedule S1 respectively. This implies that $\sum_{j=1}^J q_j^{obs} = \tilde{Q}$, and $\sum_{j=1}^J q_j^{pc} = Q^*$. Let e_j be the emission intensity associated with the production and consumption of barrel j . Then, the two opposing effects of market power on total emissions can be denoted as follows.

$$\text{Total emission impact} = \underbrace{\sum_{j=1}^{\tilde{Q}} e_j \times (q_j^{obs} - q_j^{pc})}_{\text{Composition-Effect (-)}} + \underbrace{\sum_{j=\tilde{Q}}^{Q^*} e_j \times q_j^{pc}}_{\text{Volume-Effect (+)}} \quad (1)$$

We expect the *Volume-Effect*—the second term of Equation 1—to be positive. In a competitive setting, the equilibrium quantity is larger than in the observed equilibrium. As a result, this term captures the difference in emissions that would be released into the atmosphere if

²Note that there might also be frictions in the supply schedule of the fringe firms. For a detailed analysis of these frictions, see Asker et al. (2019).

³This anticompetitive behavior of the cartel results in two traditional sources of welfare loss, as described by Asker et al. (2023). First, the lower output and higher equilibrium prices result in a welfare loss due to deadweight loss. Second, there is a welfare loss due to production misallocation of production, as oil is not extracted according to marginal costs.

⁴We provide a detailed discussion of emission content of oil in Section 3.

the market would produce Q^* according to the perfectly competitive supply schedule S1, as compared to when the market would produce \tilde{Q} according to supply schedule S1. This implies an overall increase in emissions in the perfectly competitive setting, driven by the increased overall extracted quantity of oil as compared to the setting where market power is allowed.

The *Composition-Effect*—the first term of Equation 1—captures the increase in overall emissions due to the increased use of (dirtier) fringe fields, as compared to the perfectly competitive scenario. To this end, we compare the emissions associated with supply schedule S1 and S2, respectively, while keeping the total observed extraction levels at \tilde{Q} . Since production costs are largely correlated with emission intensity, we expect that the perfectly competitive supply schedule S1—that produces from lowest to highest marginal costs—will result in less emissions, given a certain production level.

Calculating the net effect of these two forces requires information on how competitive firms, in the absence of a cartel, make decisions, how the presence of the cartel changes firm decisions, and the life-cycle emissions associated with each oil field.

Figure 1 represents the welfare implications of these two opposing forces in a stylized setting. S1' and S2' include the carbon costs of the supply schedules S1 and S2 respectively. This results in a vertical shift of the supply schedules. Note, however, that the vertical shift due to the carbon cost is not uniform across fields. This captures the fact that OPEC fields (fields 1, 2 and 4) have a lower carbon intensity of production as compared to the non-OPEC fields. In Figure 1, the difference in emissions from S1' and S2' up to the point of actual production, \tilde{Q} , represents the *Composition-effect*. The yellow shaded area between \tilde{Q} and Q^* shows the *Volume-effect*. We next discuss our model of firm behavior, leaving the discussion of carbon intensities for Section 3.

Oil extraction in a dynamic setting Oil is a finite resource. Hence, the barrels of oil that are not extracted today may be extracted tomorrow. Calculating the emission impact of market power, therefore, requires a dynamic model. First, in this dynamic setting—as opposed to the static one—the main impact on emissions of the cartel is the *speed* at which they are released into the atmosphere, rather than *whether* they are released. A second difference with the static setting is that producers will consider the opportunity cost of the revenue forgone tomorrow when choosing to sell today. That means a Hotelling rent will be present in the dynamic setting, which means that even in the perfectly competitive setting, prices will not equal marginal extraction costs (Hotelling, 1931; Anderson et al., 2018).

We compare the observed path of emissions released in the oil sector to a perfectly competitive scenario. The perfectly competitive scenario maximizes the net present value of the gains

from trade. In a perfectly competitive market, the following function is maximized:

$$\max_{\{q\}_1, \dots, \{q\}_t, \dots, \{q\}_T} \sum_{t=1}^{\infty} \beta^{t-1} \int_0^{Q_t^*} D(x, \{q\}_t, t) - S(x, \{q\}_t, t) dx \quad (2)$$

$$\text{subject to: } \{q\}_t \in \mathcal{J}_t \quad \forall t \quad (3)$$

where \mathcal{J}_t are the barrels available to be produced in year t , $\{q\}_t$ denotes the set of barrels produced in year t , Q_t^* denotes the number of barrels produced in year t (a function of $\{q\}_t$), $S(x, \{q\}_t, t)$ returns the marginal cost of the x th barrel produced in year t , $D(x, \{q\}_t, t)$ returns the marginal value of the x th barrel produced in year t , and β is the discount factor.

In a perfectly competitive equilibrium, fields will be extracted according to marginal costs—we call this the *sorting property*. In the early years, cheaper fields will be extracted. In the later years, also the more expensive fields will also be used to satisfy demand. Further, in a perfectly competitive equilibrium, there should be no intertemporal arbitrage—the rents from the marginally extracted field should be equal across time periods. We denote the cost of the marginal barrel produced in a certain year $c_t^* = S(Q^*, \{q\}_t, t)$, associated to the clearing price $p_t^* = D(Q^*, \{q\}_t, t)$, such that the intertemporal no-arbitrage can be summarized as follows

$$(p_t^* - c_t^*) = \beta(p_{t+1}^* - c_{t+1}^*) \quad (4)$$

As in standard Hotelling models, this is a necessary condition for a dynamic competitive equilibrium. This coincides with the social planner’s solution to the problem in case the social planner does not take into account environmental outcomes.

3 Data and Background

3.1 Data sources

This paper uses two main datasets.⁵ First, we use emission intensity data from engineering estimates from Jing et al. (2020), similar to Coulomb et al. (2021). Second, we complement the global data on carbon intensities with data on oil production, costs, and reserves obtained from Rystad.

Emission data In the context of the oil value chain, emissions can manifest at three main stages: upstream during extraction, midstream at the refinery, and downstream during combustion. Our first dataset contains carbon intensity information for upstream and midstream production processes based on public, engineering-based estimates from Masnadi et al. (2018)

⁵Additionally, (i) International Energy Agency (IEA) price data Western Texas Intermediate (WTI) and Brent if available, (ii) GDP deflator data, and (iii) global GDP time series are used.

and Jing et al. (2020).⁶ These data are all reported by the authors at the country-crude level.⁷ As described in Jing et al. (2020), the field-level carbon intensity estimates are weighted by volume to get to a country-level estimate of carbon intensity along the oil value chain.⁸ In contrast to the varying carbon intensities observed in the upstream and midstream stages, the utilization of refined oil products (such as gasoline, kerosene, or diesel fuel) yields a homogeneous carbon intensity per unit of energy. As a result, the carbon intensity of a barrel of oil (e_c) can be defined as follows, where α_{jc} signifies the weight of fields $j = 1, \dots, J$ in country c :

$$e_c = \sum_{j=1}^J \alpha_{jc} e_{jc} = \sum_{j=1}^J \alpha_{jc} (e_{jc}^{up} + e_{jc}^{mid}) + e^{down} \quad (5)$$

Emissions from upstream oil extraction (e_{jc}^{up}) are defined as emissions associated with oil production from 'well-to-refinery'. Upstream activities include developing, producing, and extracting crudes, transportation to the refineries, maintenance, waste disposal, and surface processing. The heterogeneity in carbon intensity can mainly be attributed to differences in operational, physical, chemical, and geological properties (Masnadi et al., 2018). For example, gas flaring or the direct venting of methane can considerably impact the carbon intensity. Algeria is one of the top 10 flaring countries, and as a result, Algerian oil is about six times more carbon-intensive in its upstream oil activities than Danish oil. The operational and physical characteristics of fields within a single country can vary, leading to differences in the carbon intensity of upstream production within a country. Overall, upstream emissions account for 11.1% of total emissions along the oil value chain. Masnadi et al. (2018) obtains carbon intensities by running the engineering-based model Oil Production Greenhouse Gas Emissions Estimator (OPGEE version 2.0). Emissions are measured in gCO_2e per MJ of crude petroleum.⁹ For both datasets, the reference year is 2015; therefore, variation in carbon intensity across time is excluded.

The midstream carbon intensity (e_{jc}^{mid}) captures the emissions associated with refining crude oil. Refineries transform crude oil into products such as gasoline or kerosene. Heavier crude oils require additional steps leading to greater carbon emissions to refine the oil into finished products. Emissions from this stage of production account for 6.5% of emissions along the oil supply chain. Jing et al. (2020) present an evaluation of the emission intensity of refining by running the engineering-based Petroleum Refinery Life Cycle Inventory Model (PRELIM) model. Depending on the crude quality or refinery configuration, the emission intensity can

⁶<https://www.nature.com/articles/s41558-020-0775-3> ; <https://www.science.org/doi/epdf/10.1126/science.aar6859>.

⁷Masnadi et al. (2018) and Jing et al. (2020) rely on proprietary data from WoodMackenzie for their field-level carbon intensity estimates, which cannot be disclosed due to confidentiality restrictions. As a result, field-level variation in carbon intensity within a country is not captured in our analysis.

⁸The reference year of these calculations is 2015.

⁹The energy content conversion factor used for one barrel of oil equivalent is 6.120 GJ.

differ across countries.¹⁰ For example, countries with a large share of light crudes, such as the United Arab Emirates ($35.7 \text{ kgCO}_2\text{e}/\text{bbl}$), tend to have considerably lower midstream emission intensity than countries with a large share of heavy oils, such as China ($50.9 \text{ kgCO}_2\text{e}/\text{bbl}$). In our main specification, we map the midstream carbon intensity of a field to the country where the barrel was extracted—not where the refinement took place.¹¹

While we rely on country-specific emission costs, we leverage variation across fields and crudes based on their relative weight in production. In Appendix A.2, we discuss the variation across these dimensions, and how it impacts our analysis. The take-away message is that our country-specific emission costs trace the underlying technology and type of crude, which ultimately are the major determinants of the emission costs.

Finally, for the downstream carbon intensity (e^{down}) accounts for the majority of carbon emissions along the value chain. About 82% of overall emissions are due to the combustion of refined fuel products. Similar to Coulomb et al. (2021), we assign a carbon intensity of $0.464 / \text{tCO}_2\text{e}/\text{bbl}$ to downstream combustion. At this stage of the value chain, there is little heterogeneity in the carbon intensity of fuel combustion per energy unit across the different refined products, e.g., kerosene or gasoline.

Production, cost, and reserve data We also use data on annual production, reserves, and costs, along with other geology or extraction technology information, across virtually all fields worldwide from 1970 to 2021. In total, there are 21,388 active fields in this dataset. A field is defined as a geographically homogeneous oil production area.

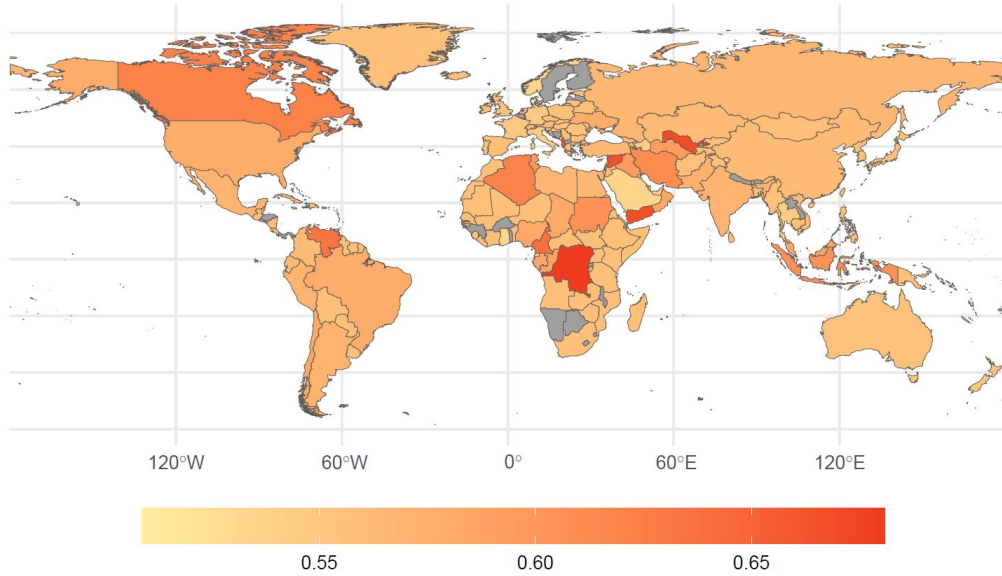
To construct the marginal cost of oil extraction c_{ft} of field f during year t , we follow the approach of Asker et al. (2019, 2023). That is, we divide the total cost of production across spending categories h by their annual reported production quantities, $c_{ft} = \frac{\sum_h \text{Expenditure}_{h,ft}}{q_{ft}}$. The spending categories include $h = \{ \text{Well Capital, Facility Capital, Abandonment cost, Production Operating, Transportation Operating, and SGA} \}$; all deflated by the US GDP deflator with 2009 as the base year.

Reserves of a field equal the (economically viable) barrels of oil that are still underground. The most accurate way to measure reserves at a point in time is to see the entire production life of a field. The total extracted oil is then the maximal reserve. However, most fields are not fully exploited in the data. Hence, we use industry reserve estimates. The oil industry reports reserves at three levels of probability. P90 (or P1) is the quantity that can be recovered with a 90% probability, given current technical and economic conditions. P50 (or P1 + P2)

¹⁰As a simplifying assumption in the main specification, we set the midstream carbon intensity of a country equals the same value from all crudes produced in that country, even if the refining occurs in a different country.

¹¹In an alternative specification, we attribute the midstream emissions to the country where the oil is refined, see Appendix C.2.

Figure 2: Country-level life-cycle emission intensity (tCO_2)



Notes: Life-cycle emission intensity (tCO_2) per barrel of oil equivalent at the country-level

denotes the total reserves recoverable with a probability of 50%. Finally, the P10 reserves (or $P1 + P2 + P3$) are total reserves economically viable with a 10% chance. This means that the magnitude of reserves will fluctuate with the oil price. In our paper, we use reserves measured as P50, assuming an oil price of \$60 (in 2021 dollars), the historical average price for oil.¹²

3.2 The upstream oil cartel

The global oil value chain is divided into three main activities. The upstream producers extract and explore crude oil. Refineries transform crude oil into petroleum products, such as gasoline. The main downstream use of oil in 2022 is in the transportation sector (67.2%), as reported by IEA.

The upstream oil industry consists of oil companies that are either (largely) state-run enterprises or independent enterprises. 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.¹³

OPEC is an intergovernmental organization founded in 1960. The founding members include Iran, Iraq, Kuwait, Saudi Arabia, and Venezuela. The membership of OPEC has expanded to

¹²As such, we do not endogenize reserves across the two different market regimes.

¹³Sometimes, however, they contract with independents to operate specific facilities.

include several other major oil-producing nations, such as the United Arab Emirates, Libya, Algeria, Nigeria, and more.¹⁴ The cartel has been relatively stable over time, especially the membership of the core Middle-Eastern member countries. In 2016, OPEC reached formal agreements with Azerbaijan, Bahrain, Brunei, Equatorial Guinea, Kazakhstan, Malaysia, Mexico, Oman, Russia, Sudan, and South Sudan to further coordinate production, often referred to as ‘OPEC +’. In this paper, OPEC (which includes OPEC+ from 2017 onwards) members will include the countries described above. The largest independent companies are the ‘five oil majors’, i.e., ExxonMobil, Chevron, BP, Royal Dutch Shell and Total Energies.

At the start of our sample in 1970, OPEC held just over 64% of all oil reserves. By the end of our sample in 2021, this has decreased to 45.5% of total oil reserves. In terms of production market shares, there is quite some variation over time, as shown in Figure 3. Until the 1980s, OPEC served about half of the market. This dropped to just under 30% in the mid-1980s. Since then, it has been steadily on the rise until the mid-1990s to about 40%, and has remained there since. The largest producer of OPEC is Saudi Arabia, largely following the general pattern of OPEC production. The largest non-OPEC(+) producer is the United States, following a steady decrease from its 25% market share to about 8.6% in 2008. After 2008, the US market share steeply increased to just under 20% in 2019, explained by the US fracking boom.

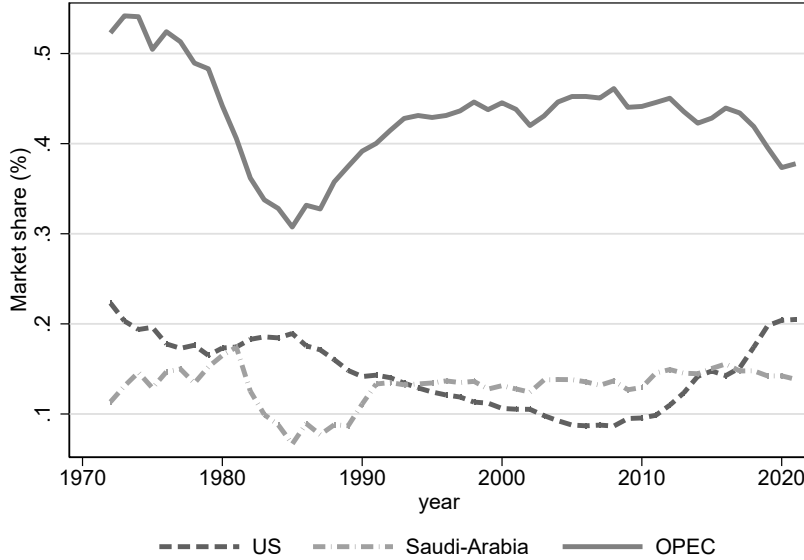
3.3 Preliminary Evidence

We document three key facts of the data in this section. First, we show that there is heterogeneity in carbon intensities and production costs in the oil deposits. Second, carbon intensities and production costs are positively correlated. However, most emissions are concentrated in the downstream sector (i.e., the consumption of the refined product), where there is no heterogeneity across fields. Third, we highlight the importance of the oil sector for overall climate goals.

Heterogeneity in carbon intensity and production costs Figure 4 illustrates the heterogeneity in carbon intensity within the upstream and midstream segments of the oil sector for both OPEC+ and non-OPEC+ nations. The data in the Figure indicate that major OPEC countries, including Saudi Arabia, Qatar, and the United Arab Emirates, exhibit comparatively lower carbon intensities in their up- and midstream oil activities, around 0.54 tCO_2 per barrel extracted. Conversely, nations outside the OPEC+ alliance, such as the United States and Canada, display higher carbon intensities, 0.57 tCO_2 and 0.63 tCO_2 respectively, reflecting relatively greater environmental impact associated with their oil production.

¹⁴OPEC member states include Algeria, Angola, Ecuador, Indonesia, Iran, Iraq, Kuwait, Libya, Nigeria, Qatar, Saudi-Arabia, UAE, and Venezuela.

Figure 3: Market shares in the upstream oil market



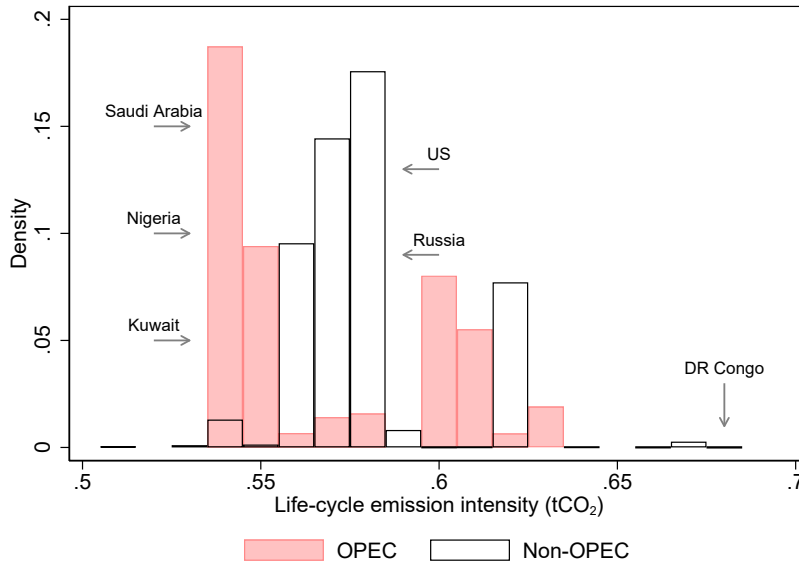
Note: Market shares indicate share of global upstream oil production. OPEC countries include Algeria, Angola, Ecuador, Indonesia, Iran, Iraq, Kuwait, Libya, Nigeria, Qatar, Saudi-Arabia, UAE, and Venezuela.

Further, there is a large dispersion of extraction costs across producers. Figure 5 compares the 99th percentile marginal cost across OPEC and non-OPEC countries. The black solid line indicates the yearly oil price. The figure shows that the range of costs for OPEC countries is consistently narrower and at a lower level than that of non-OPEC countries. For example, the 99th percentile of production costs amounted to over \$100 in 2012 for non-OPEC countries, while it about a third for OPEC countries.

All fields are dirty, but cheaper fields are, on average, cleaner It is an important feature of the data that the oil production within the OPEC cartel is generally characterized by lower production costs and carbon intensity, as illustrated in Figure 6.¹⁵ First, this correlation between cleanliness and production costs in the data implies that production distortions caused by the international cartel—withholding low-cost production units to create an artificial scarcity—will also have environmental implications. If lower-cost units are withheld, higher marginal cost fields will be reshuffled to earlier time periods. These higher marginal cost fields are likely to be dirtier, resulting in higher emission release in earlier time periods, i.e., the *Composition-Effect* of the oil cartel. This could result in a faster increase in atmo-

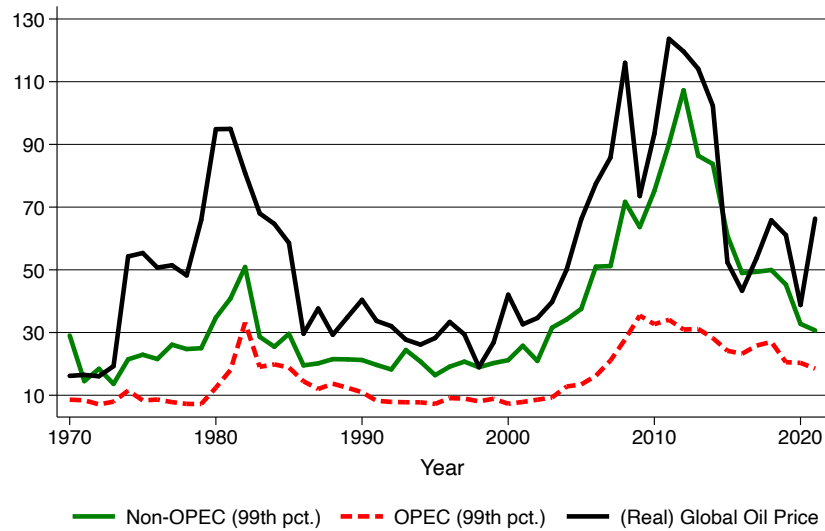
¹⁵Figure A2 demonstrates that this correlation persists when midstream emissions are assigned to the refinery country.

Figure 4: Life-cycle emissions intensity (tCO_2) across countries



Note: Distribution of life-cycle emissions intensity (tCO_2) per barrel for OPEC and non-OPEC, weighted by 2018 reserves.

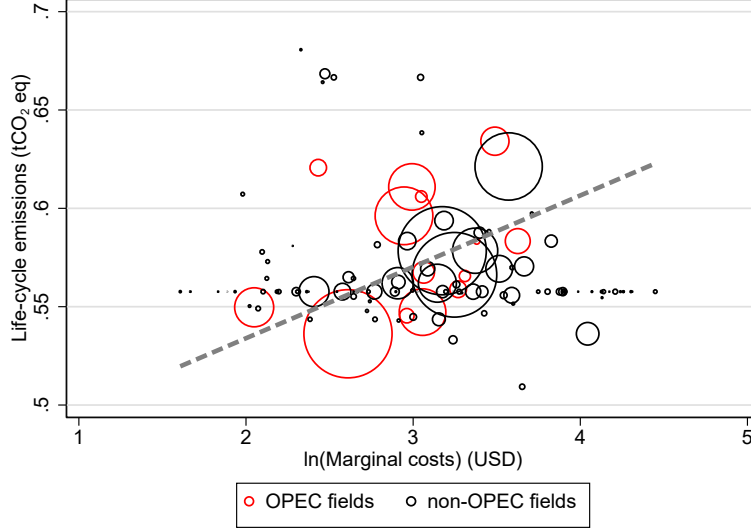
Figure 5: Marginal extraction costs and oil price (1970 - 2021)



Note: Comparison of 99th percentile marginal extraction costs in OPEC and non-OPEC countries. The series are deflated with the US GDP deflator (base year 2009).

spheric temperature overall, as compared to the perfectly competitive case. Because we do not account for variation in carbon intensity of production within countries, our estimate of the *Composition-Effect* should be interpreted as a lower bound.¹⁶

Figure 6: Marginal costs and life-cycle emission intensity



Note: Correlation of marginal extraction costs and carbon intensity of the midstream and upstream production process. Observations are weighted by reserves in 2018.

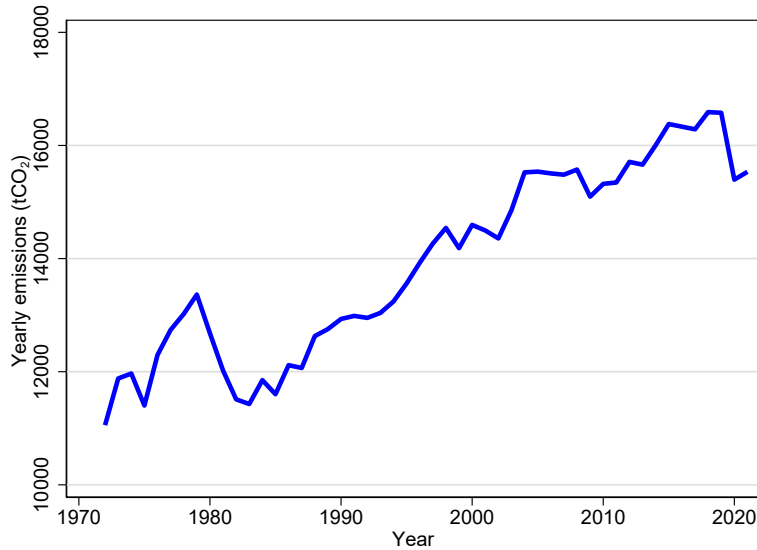
However, the data tell us that most emissions are concentrated in the downstream activities of the oil sector, i.e., combustion of the fuels. More than 80% of the carbon emissions are released at the time of fuel consumption. Therefore, when anticompetitive behavior brings down total extracted barrels, a large share of overall emissions will be saved—regardless of their origin. Since total combustion—which depends only on overall consumed barrels of oil—has such a large share in the total emissions, we expect the *Volume-Effect* to dominate the *Composition-Effect*. Even though market power will spur the earlier usage of dirtier fields, the emissions savings from bringing down overall demand by inflating prices are expected to be larger.

Emissions in the oil industry Currently, the global oil sector accounts for just over one-third of annual anthropogenic emissions (IEA, 2022). In 2021, the global annual carbon emis-

¹⁶To see this, note that while we rely on quantity-weighted average emission costs by country, the positive correlation between the cost of extraction and emission costs across crudes and technology, may in fact overestimate the emissions cost for (relatively) low-cost field, while underestimating it for (relatively) high-cost fields. We refer to Table B.1 in Coulomb et al. (2021).

sions from oil amounted to roughly 16 GtCO_2 , while total global carbon emissions equaled 37 GtCO_2 (IEA, 2022). Emissions from the oil industry are displayed in Figure 7. Further, it shows that yearly carbon emissions from the oil industry show an overall increasing trend throughout time—from 1970 to 2021, with an average annual growth rate of 1%. There is an exceptional dip in consumption caused by the COVID-19 pandemic at the end of the time series in 2020-2021. Another significant dip is shown in the early 1980s, following production reductions during the Iran - Iraq war.

Figure 7: Yearly emissions from the oil sector.



Note: Life-cycle emissions from the oil sector are expressed in tCO_2 . The assigned life-cycle emissions are heterogeneous across barrels, as described in Section 3. The horizontal axis is in years, from 1970 - 2021.

Friedlingstein et al. (2022) finds that to have a 50% chance of limiting average global temperature increase to 1.5°C , the remaining carbon budget from the beginning of 2023 is 380 GtCO_2 . To limit the average global temperature increase to 1.7° and 2°C carbon budgets are 730 and 1230 GtCO_2 , respectively. McGlade and Ekins (2015) finds that emissions content of recoverable reserves from oil alone, in 2015, would overshoot these targets. The findings of McGlade and Ekins (2015) are consistent with our estimates of the current carbon budget from all oil reserves of just under $1,000 \text{ GtCO}_2$ in 2021. This implies that the production behavior in the oil sector plays a vital role when considering climate policy.

4 Empirical Implementation

Evaluating the environmental impact of market frictions in the oil market, as described in Section 2, necessitates the modeling of a frictionless benchmark of oil extraction. To that end, we need to estimate marginal costs and demand and parametrize the model¹⁷. Next, the computational implementation is briefly discussed.

Extraction costs Marginal (or average) costs of field f at time t , i.e., c_{ft} , can be decomposed into three elements. First, time-invariant marginal costs, e.g. due to geology. Second, a technology-year-specific cost shifter μ_{st} . Technologies s considered are both onshore and offshore extraction. Third, the marginal costs contain measurement error $\exp(\epsilon_{ft})$. That is,

$$c_{ft} = c_f \mu_{ft} = c_f \mu_{st} \exp(\epsilon_{ft}) \quad (6)$$

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

$$\ln \hat{\mu}_{st} = \sum_{f \in s} \kappa_{ft} \ln c_{ft},$$

where κ_{ft} is the quantity weight of a field in a given year's total output, $\kappa_{ft} = \frac{q_{ft}}{\sum_{f \in s} q_{ft}}$.

The time-invariant marginal cost, c_f , is then estimated, allowing for measurement error, using the following (within-field) regression:

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

We estimate the time-invariant marginal costs using weighted least squares, with the weights being the proportion of total field output in that year. Further details on the estimation are provided in the appendix.

There are fields that have never produced in our dataset. Hence, they do not have observed production cost data. For these fields, we assign marginal costs equal to break even prices, as calculated by Rystad Enegy. They are defined as the flat oil price required for a positive net present value of continued operation. About 50% of all assets have never produced. For more details on the break even costs, we refer to the Data Appendix.

Estimating demand We estimate global annual (t) demand for oil as follows

$$Q_t = \alpha + \beta P_t + \gamma GDP_t + g(t) + \epsilon_t \quad (7)$$

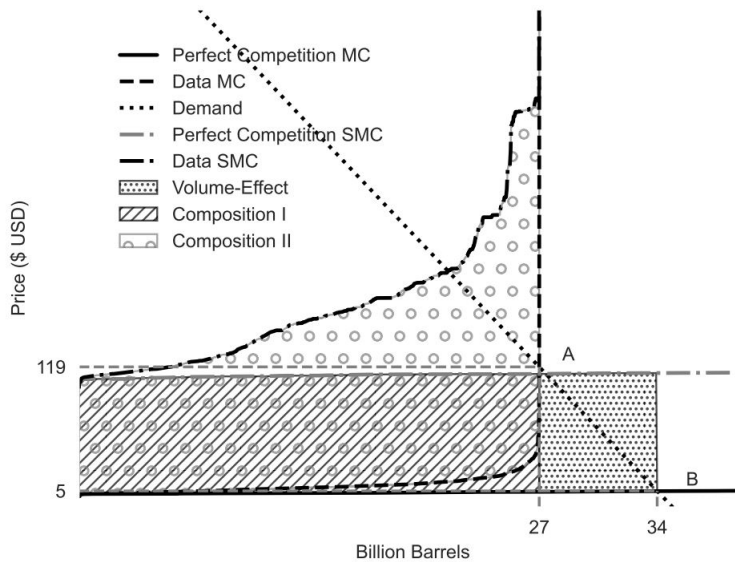
P_t and Q_t denote the observed annual global quantity and price, respectively. The $g(t)$ function denotes a fourth-order polynomial time trend in order to control for macroeconomic trends. To

¹⁷A similar approach to estimating cost and demand, and computation, is taken in Asker et al. (2023)

deal with price endogeneity, a supply shifter is used as a price instrument. The detailed cost data enable us to track shifts of the aggregate supply curve. That is, we can order all barrels from the lowest to the highest cost of production. For each year, we can select the unit cost of production of the Q^{th} barrel to keep track of the aggregate supply curve. We set $Q = 15$ billion bbl per year, which is greater than the lowest level of annual consumption. Details are provided in Appendix A.

Figure 8 shows the two opposing emission-related welfare effects due to market power in the upstream oil industry, as discussed in Section 2. Leveraging the cost and demand estimates as discussed in the previous paragraphs, the figure show the *Volume-* and *Composition-* effect in a static setting for 2012. The figure shows that the distance between the marginal cost with and without the cost of carbon is much wider for the observed supply curve, as compared to the perfectly competitive scenario. This is because the perfectly competitive supply curve uses the more efficient fields, which are generally cleaner and thus have a lower marginal social cost of carbon. The *Composition*-effect captures this characteristic. Further, since perfect competition leads to greater production (32 billion barrels) than what we actually observe (29 billion barrels), the *Volume*-effect captures the damages from the emissions associated with this production increase.

Figure 8: Static *Volume-* and *Composition-* effect in 2012



Notes: Comparison of the perfectly competitive equilibrium B and the observed equilibrium A in a static setting. The shaded areas capture the *Volume-Effect* and the *Composition-Effect* (Composition II - Composition I) for the year 2012. The social cost of carbon assigned here is \$190 per tCO_2 .

Computation This paper simulates oil extraction from a perfectly competitive market. Since oil is an exhaustible resource, the decision to model is *when* to produce, rather than *if*. Oil that is not extracted today can be extracted in a later period. To compute the competitive benchmark’s problem, we generate a production path that maximizes the net present value of all future gains from trade. The welfare effect from emissions is not considered when simulating the competitive benchmark’s production path.

For the perfectly competitive sorting algorithm, we order the fields i from the lowest to the highest cost, indicating their production order. This implies that if field i is produced in year t , it must be the case that field $i - 1$ is produced no later than year t . The purpose is to find the marginal field j_t for each year that is indifferent between producing and not producing. The intertemporal production path from perfect competition can therefore be characterized by a list of marginal fields in each year $\{j_1, j_2, \dots, j_T\}$, with an implied annual price path of $\{p_1, p_2, \dots, p_T\}$.

As a result, for each marginal field j_t , the following should hold:

$$P_t\left(\sum_i^{j_t} q_i\right) - c_{j_t} = \beta \left(P_{t+1}\left(\sum_i^{j_{t+1}} q_i\right) - c_{j_t} \right) \quad (8)$$

Equation 8 shows the well-known intertemporal no-arbitrage condition, as formulated by Hotelling (1931). This allows us to simulate a perfectly competitive benchmark production path.

Parametrization The algorithm described in the previous paragraph is used to simulate the production path of a perfectly competitive oil market. The inputs required for this counterfactual are (i) marginal costs $c_f \mu_{ft}$, (ii) field-level total reserves $R_{ft=1}$, (iii) demand parameters (β, α) . Additionally, a social annual discount factor is needed since simulating the competitive production path requires intertemporal optimization. This is set at 0.95.

Marginal costs and demand parameters are retrieved as described in the previous paragraphs. For this, however, some auxiliary modeling elements are also relevant.

The first auxiliary element is that the path of field discovery is assumed to be exogenous. Before field discovery, marginal costs are assumed to be infinite and thus excluded from the perfectly competitive production path. The second auxiliary element is the restriction on how much a field can extract yearly. If extraction is too large, well pressure might drop sharply. Therefore, these geological factors limit the proportions of reserves that can be extracted from a field in any given year. In the main specification, we assume that the upper limit on the rate at which a field can be extracted equals $\max\{x_f, 10\%\}$, where x_f is the maximal observed production proportion at any year for field f .

The third element concerns the demand estimates. Since all fields will produce until full extraction, assumptions need to be made for future demand. From 2022 onwards, global demand

is assumed to grow at a rate of 1.3% per year. This equals the average oil consumption growth rate observed from 1970 - 2021.

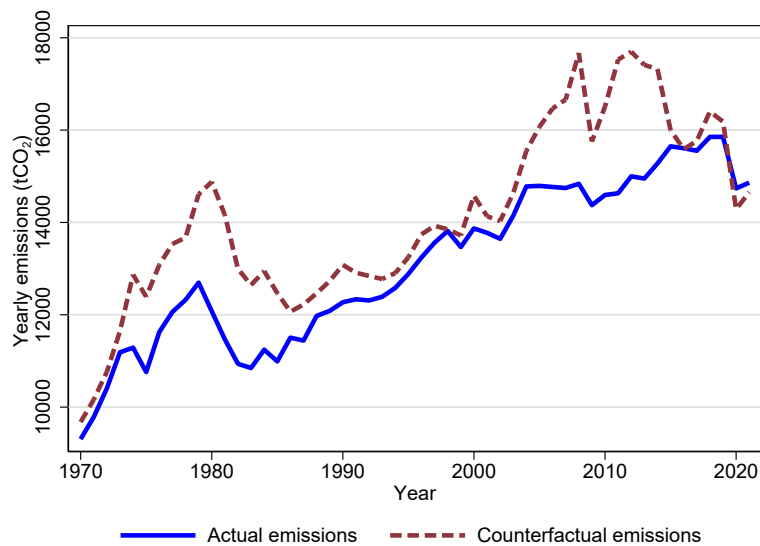
5 Quantifying the Environmental Effects

This section quantifies the environmental impact of market power in the oil market on global emissions. We do this by computing a counterfactual production path described in Section IV. We then compare the actual production decisions to the counterfactual production path. First, we describe the total impact on carbon emissions. Second, we discuss the implications of these findings on changes in atmospheric temperature.

5.1 Total Carbon Impact

To quantify the role of market power in the environmental effects of oil extraction, we need to compute the counterfactual path of extraction when firms are undistorted price takers. For this, we use the sorting algorithm as described in Section 4. For the counterfactual path, externalities to the environment are not taken into account when optimizing the intertemporal production path. First, we describe the environmental impact of OPEC without considering the heterogeneous carbon intensity across fields. Second, we decompose our analysis into the *Volume-* and *Composition-Effect* as described in Section 2, allowing the fields to differ in carbon intensity.

Figure 9: Total actual and counterfactual emissions under homogeneous carbon intensity



Homogeneous carbon intensity We first compute the perfectly competitive intertemporal production path from the full dynamic model. This allows us to compare the actual emissions to the emissions associated with the counterfactual production path. If we assume that each barrel is equally dirty¹⁸, the only factor that is driving a wedge between counterfactual and actual emissions is differences in total yearly extracted barrels. Figure 9 shows that counterfactual emissions are consistently larger than actual emissions, indicating that the competitive equilibrium is associated with higher extracted quantities. In a perfectly competitive equilibrium, only the lowest cost fields are produced, resulting in lower prices in equilibrium. For example, in 1980, actual prices amounted to \$94.9/bbl, while counterfactual prices equal only \$11.1/bbl. Since demand is not perfectly inelastic, this leads to higher total extracted volumes. The years 2016, 2020, and 2021, are exceptions during which counterfactual extractions are (slightly) lower than actual production. Overall, the counterfactual production path includes 69,775 $MtCO_2$ additional emissions.

Focusing on the 1970 to 2021 period, the net present value of the avoided emissions can be computed. This allows us to map the two emission paths to environmental damages in dollars. To this end, we need to quantify the dollar value of the negative externality associated with one additional ton of carbon emitted. That is, we need to establish the social cost of carbon. In our main specification, we assign the social cost of carbon to be \$250/ tCO_2 in 2021, following the estimations of Lemoine (2021).¹⁹ The results are presented in the first panel of Table A4. At a \$250/ tCO_2 social cost of carbon, the counterfactual production path would have caused \$5,406 billion in additional damages to the environment. To give a sense of scale, Asker et al. (2019) quantifies the productive inefficiencies caused by market power in the oil market to be \$916 billion.

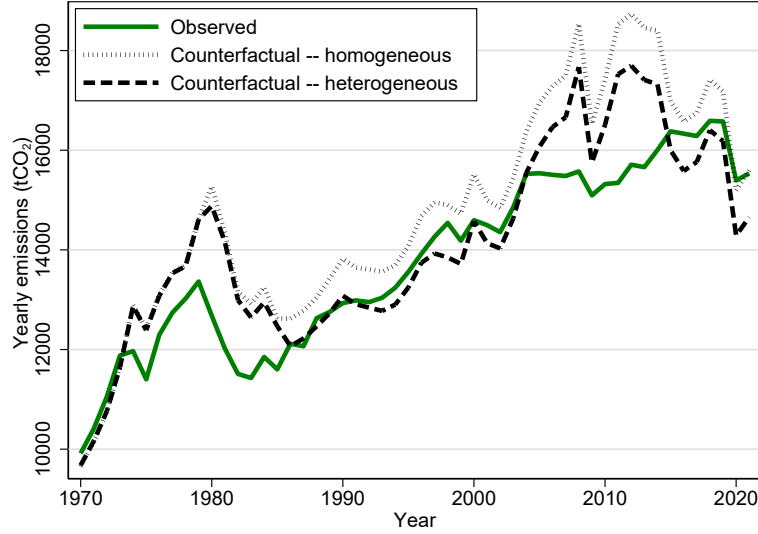
In our calculations, the net present values of environmental damages are linear with respect to the assumed social cost of carbon. Table A3 presents the results using the US Environmental Protection Agency’s latest estimate for the social cost of carbon, which is \$190/ tCO_2 (Environmental Protection Agency, 2023).

Heterogeneous carbon intensity Now, we allow barrels to have different carbon intensities. As described in Section 2, this allows us to decompose the total emission effect from market power in two components. First, the *Volume-Effect* captures the change in emissions due to distortions in the equilibrium quantities in the oil market due to market power. Second, the *Composition-Effect* quantifies the emission impact from distortions in the production path, *given* total yearly observed production. Market power alters which fields are dispatched to

¹⁸That is, 0.544 tCO_2 per million of extracted barrels

¹⁹Several estimates of the social cost of carbon exist. We do not take a stand on the “right” one but note that our estimates scale with the assumed social cost of carbon.

Figure 10: Total actual and counterfactual emissions (tCO_2) under heterogeneous and homogeneous carbon intensity



satisfy demand and will, therefore, affect total emissions associated with oil extraction due to heterogeneity in carbon intensity. Figure 10 shows the yearly emissions in the actual and counterfactual scenario, under homogeneous and heterogeneous carbon intensities. In both cases, the counterfactual emission path results in more cumulative emissions throughout the years. However, total emissions are smaller in the heterogeneous case, as compared to the homogeneous case. This is because perfect competition extracts the cheapest fields will first, that are generally cleaner. Thus, this leads to a lower per-barrel carbon intensive extraction path. This idea is captured by the *Composition-effect*. Figure 10 also shows that the counterfactual emissions are still higher than the observed emissions in the heterogeneous case, driven by the *Volume-effect*. This implies that even though the competitive supply path would employ, on average, cleaner fields, this positive environmental effect is offset by the increased equilibrium quantities in the competitive equilibrium. Overall, a 67,738 $MtCO_2$ emission difference is associated with the actual and perfectly competitive supply path. This equals four years of current oil consumption or 1.7 years of overall CO_2 emissions.

In Table A4, the NPV of the (avoided) environmental damages under heterogeneous carbon intensity is displayed in Panel B. Overall, market frictions in the oil market have avoided \$4,073 billion carbon emission externalities. The increased equilibrium quantities in the competitive scenario, as compared to the actual quantities, resulted in a \$5,586 billion welfare loss all else equal, i.e., the Volume-Effect. However, part of this welfare loss is offset by the environmental gain of the perfectly competitive supply path due to the increased production of cleaner (low-

cost) fields, i.e., the Composition-Effect. The composition-effect equals \$1,512 billion.²⁰

	NPV emissions (\$B)	Difference (\$B)
<u>Panel A: Homogenous emissions</u>		
Observed production	55 669	
SP production	61 075	5406
<u>Panel B: Heterogenous emissions</u>		
Observed production	58 726	
SP production	62 800	4073
<i>SP(a): Composition</i>	<i>57 214</i>	<i>- 1512</i>
<i>SP(b): Volume</i>	<i>5586</i>	<i>5586</i>

Table 1: Net Present Values of environmental damages, valued at a social cost of carbon of \$250/ tCO_2

²⁰In an alternative specification, displayed in Table A4, midstream emissions are allocated to the country where the refinement takes place, and the results are invariant to this change.

5.2 Impact on Atmospheric Temperature

This section maps the avoided carbon emissions to their impact on atmospheric temperature. Fewer emitted tons of CO_2 will result in a lower atmospheric temperature. Following Covert et al. (2016), we use two climate models to simulate the climate system impact, i.e., the DICE model Nordhaus (2014) and a Climate-Carbon Response model (Matthews et al., 2009).

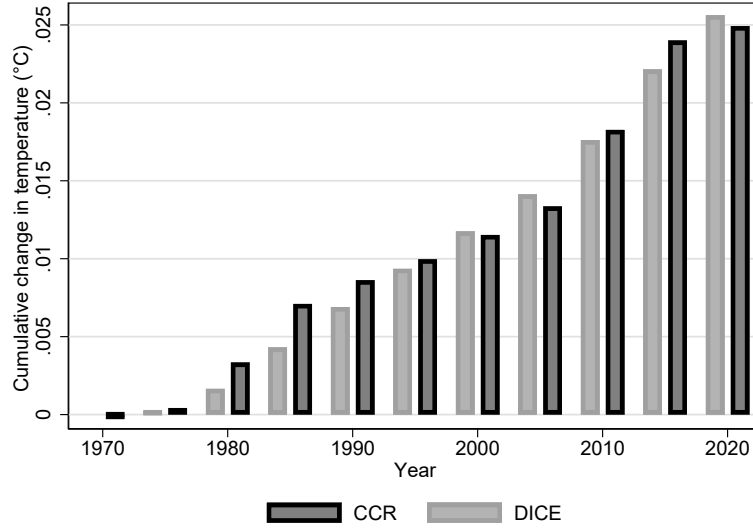
As documented in the previous section, market frictions in the oil market have saved over 67 $GtCO_2$. First, the differences in yearly emissions can be used as an input in the DICE model, an integrated assessment model (IAM). We use Nordhaus' DICE2013R IAM model, complemented by actual atmospheric carbon concentrations from 1970-2021 obtained from the NOAA GML (National Oceanic and Atmospheric Administration Global Monitoring Laboratory) satellite data.²¹ The DICE model allows us to simulate counterfactual carbon concentrations. This affects radiative forcings logarithmically, leading to atmospheric and lower ocean temperature changes. Overall, the DICE model predicts an avoided temperature increase of about $0.028^\circ C$ in 2021. The DICE model is able to capture the gradual, lagged effect of additional emissions on global warming over time, as shown in Figure 11. We discuss the details of the implementation in Appendix D.

Alternatively, we can follow the more straightforward estimates from the climate-carbon response (CCR) model of Matthews et al. (2009). They rely on the idea that global warming is approximately proportional to total carbon emissions over time. This assumption implies that a specific change in emissions will result in a specific change in temperature, regardless of the 'stock' of carbon emissions. We follow Council (2011) calibration of the linear approximation. That is, a 1,000 $GtCO_2$ increase is associated with a $1.75^\circ C$ increase in global temperature. As a result, a 67,738 $GtCO_2$ increase in emissions results in a $0.032^\circ C$ temperature increase, as plotted in Figure 11. This is similar to the estimates from the DICE Model. In contrast to the DICE model, changes in atmospheric temperatures materialize instantaneously in the CCR model.

An alternative way to interpret our results is as a percentage of the remaining carbon budget to hit certain targets. Returning to the estimates of Friedlingstein et al. (2022), our results suggest that the reduction in emissions from market power in the upstream oil industry from 1970 to 2021 is 17.8% of the 2022 carbon budget remaining to have a 50% chance of meeting the $1.5^\circ C$ increase in global mean temperatures.

²¹(https://gml.noaa.gov/webdata/ccgg/trends/co2/co2_annmean_mlo.txt)

Figure 11: Counterfactual atmospheric temperature change



Note: Cumulative increase in atmospheric temperature change ($^{\circ}\text{C}$) in the perfectly competitive counterfactual scenario (as compared to the actual scenario), as calculated by the DICE model, and the CCR model. Results are shown per 5-year period, from 1970 - 2020.

6 Concluding remarks

To evaluate the environmental impact of market power frictions in the global oil sector, we compare the environmental outcome of a perfectly competitive benchmark to observed production behavior in the upstream industry. We assess the environmental implications of market power in the sector by incorporating both the effect on equilibrium quantities and changes in the carbon intensity of oil extraction due to differential production allocation.

We find that market power decreases overall emissions in the oil sector in a sizable way. In a perfectly competitive scenario, overall emissions would increase by 67,738 $MtCO_2$. This result is driven by the increased equilibrium quantities in a perfectly competitive market (‘Volume-Effect’). Since lower-cost reserves also tend to be cleaner, the increased emissions are partly mitigated by a drop in carbon intensity of oil extraction (‘Composition-Effect’). Our result implies an NPV environmental welfare gain of \$4,073 billion due to market frictions in the oil market, evaluated at a social cost of carbon of \$250/ tCO_2 . Our results echo Buchanan (1969) and the more general theory of the second best from Lipsey and Lancaster (1956), pointing at the relevance of considering the market structure for (environmental) policy implementation.

Notably, we can map the differential emissions paths of the actual and counterfactual exercise to different climate change scenarios. Using the DICE model (Nordhaus, 2016), we find that market frictions resulted in a 0.032 $^{\circ}\text{C}$ avoided global temperature anomaly from 1970 to

2021. Consequently, this counterfactual exercise highlights the climate gains from elevated oil prices. The purpose of this exercise is not to argue in favor of OPEC or market power in the oil industry more generally, but instead to highlight the impact of market structure on global emissions and climate change.

References

- ANDERSON, S., R. KELLOGG, AND S. SALANT (2018): “Hotelling under Pressure,” *Journal of Political Economy*, 126, 984 – 1026.
- ASKER, J., A. COLLARD-WEXLER, AND J. DE LOECKER (2019): “(Mis)Allocation, Market Power, and Global Oil Extraction,” *American Economic Review*, 109, 1568–1615.
- (2023): “The Welfare Impact of Market Power: The OPEC Cartel,” NBER Working Paper Series XXX.
- AUNE, F. R., K. GRIMSRUD, L. LINDHOLT, K. E. ROSENDAHL, AND H. B. STORRØSTEN (2017): “Oil consumption subsidy removal in OPEC and other Non-OECD countries: Oil market impacts and welfare effects,” *Energy Economics*, 68, 395–409.
- BALKE, N. S., M. PLANTE, AND M. YUCEL (2015): “Fuel Subsidies, the Oil Market and the World Economy,” *The Energy Journal*, 36, 99–128.
- BENCHEKROUN, H., G. VAN DER MEIJDEN, AND C. WITHAGEN (2020): “OPEC, unconventional oil and climate change - On the importance of the order of extraction,” *Journal of Environmental Economics and Management*, 104, 102384.
- BORENSTEIN, S. AND J. B. BUSHNELL (2022): “Do Two Electricity Pricing Wrongs Make a Right? Cost Recovery, Externalities, and Efficiency,” *American Economic Journal: Economic Policy*, 14, 80–110.
- BUCHANAN, J. M. (1969): “External Diseconomies, Corrective Taxes, and Market Structure,” *The American Economic Review*, 59, 174–177.
- COULOMB, R., F. HENRIET, AND L. REITZMANN (2021): “‘Bad’ Oil, ‘Worse’ Oil and Carbon Misallocation,” Pse working papers, HAL.
- COUNCIL, N. R. (2011): *Climate Stabilization Targets: Emissions, Concentrations, and Impacts over Decades to Millennia*, Washington, DC: The National Academies Press.
- COVERT, T., M. GREENSTONE, AND C. R. KNITTEL (2016): “Will We Ever Stop Using Fossil Fuels?” *Journal of Economic Perspectives*, 30, 117–38.
- DAVIS, L. W. AND E. MUEHLEGGGER (2010): “Do Americans consume too little natural gas? An empirical test of marginal cost pricing,” *The RAND Journal of Economics*, 41, 791–810.
- ENVIRONMENTAL PROTECTION AGENCY (2023): “Report on the social cost of greenhouse gases: Estimates incorporating recent scientific advances,” .

- FATTOUH, B., L. KILIAN, AND L. MAHADEVA (2013): “The Role of Speculation in Oil Markets: What Have We Learned So Far?” *The Energy Journal*, 34, 7–33.
- FOWLIE, M., M. REGUANT, AND S. P. RYAN (2016): “Market-Based Emissions Regulation and Industry Dynamics,” *Journal of Political Economy*, 124, 249–302.
- FRIEDLINGSTEIN, P., M. O’SULLIVAN, M. W. JONES, R. M. ANDREW, L. GREGOR, J. HAUCK, C. LE QUÉRÉ, I. T. LUIJKX, A. OLSEN, G. P. PETERS, W. PETERS, J. PONGRATZ, C. SCHWINGSHACKL, S. SITCH, J. G. CANADELL, P. CIAIS, R. B. JACKSON, S. R. ALIN, R. ALKAMA, A. ARNETH, V. K. ARORA, N. R. BATES, M. BECKER, N. BELLOUIN, H. C. BITTIG, L. BOPP, F. CHEVALLIER, L. P. CHINI, M. CRONIN, W. EVANS, S. FALK, R. A. FEELY, T. GASSER, M. GEHLEN, T. GKRTZALIS, L. GLOEGE, G. GRASSI, N. GRUBER, O. GÜRSSES, I. HARRIS, M. HEFNER, R. A. HOUGHTON, G. C. HURTT, Y. IIDA, T. ILYINA, A. K. JAIN, A. JERSILD, K. KADONO, E. KATO, D. KENNEDY, K. KLEIN GOLDEWIJK, J. KNAUER, J. I. KORSBAKKEN, P. LANDSCHÜTZER, N. LEFÈVRE, K. LINDSAY, J. LIU, Z. LIU, G. MARLAND, N. MAYOT, M. J. MCGRATH, N. METZL, N. M. MONACCI, D. R. MUNRO, S.-I. NAKAOKA, Y. NIWA, K. O’BRIEN, T. ONO, P. I. PALMER, N. PAN, D. PIERROT, K. POCKOCK, B. POULTER, L. RESPLANDY, E. ROBERTSON, C. RÖDENBECK, C. RODRIGUEZ, T. M. ROSAN, J. SCHWINGER, R. SÉFÉRIAN, J. D. SHUTLER, I. SKJELVAN, T. STEINHOFF, Q. SUN, A. J. SUTTON, C. SWEENEY, S. TAKAO, T. TANHUA, P. P. TANS, X. TIAN, H. TIAN, B. TILBROOK, H. TSUJINO, F. TUBIELLO, G. R. VAN DER WERF, A. P. WALKER, R. WANNINKHOF, C. WHITEHEAD, A. WILLSTRAND WRANNE, R. WRIGHT, W. YUAN, C. YUE, X. YUE, S. ZAEHLE, J. ZENG, AND B. ZHENG (2022): “Global Carbon Budget 2022,” *Earth System Science Data*, 14, 4811–4900.
- HOTELLING, H. (1931): “The Economics of Exhaustible Resources,” *Journal of Political Economy*, 39, 137–175.
- JING, L., H. EL-HOUJEIRI, J.-C. MONFORT, A. BRANDT, M. S. MASNADI, D. GORDON, AND J. BERGERSON (2020): “Carbon intensity of global crude oil refining and mitigation potential,” *Nature Climate Change*, 10, 1–7.
- KILIAN, L. AND D. P. MURPHY (2014): “THE ROLE OF INVENTORIES AND SPECULATIVE TRADING IN THE GLOBAL MARKET FOR CRUDE OIL,” *Journal of Applied Econometrics*, 29, 454–478.
- KNITTEL, C. R. AND R. S. PINDYCK (2016): “The Simple Economics of Commodity Price Speculation,” *American Economic Journal: Macroeconomics*, 8, 85–110.

- LEMOINE, D. (2021): “The Climate Risk Premium: How Uncertainty Affects the Social Cost of Carbon,” *Journal of the Association of Environmental and Resource Economists*, 1, 27–57.
- LIPSEY, R. G. AND K. LANCASTER (1956): “The General Theory of Second Best,” *The Review of Economic Studies*, 24, 11–32.
- MASNADI, M. S., H. EL-HOUJEIRI, D. SCHUNACK, Y. LI, J. ENGLANDER, A. BADAHDAH, J.-C. MONFORT, J. ANDERSON, T. WALLINGTON, J. BERGERSON, D. GORDON, J. KOOMEY, S. PRZESMITZKI, I. AZEVEDO, X. BI, J. DUFFY, G. HEATH, G. KEOLEIAN, C. MCGLADE, AND A. BRANDT (2018): “Global carbon intensity of crude oil production,” *Science*, 361, 851–853.
- MATTHEWS, H. D., N. P. GILLETT, P. A. STOTT, AND K. ZICKFELD (2009): “The proportionality of global warming to cumulative carbon emissions,” *Nature*, 459, 829–832.
- MCGLADE, C. AND P. EKINS (2015): “The geographical distribution of fossil fuels unused when limiting global warming to 2 °C,” *Nature*, 517, 187–190.
- NORDHAUS, W. (2014): “Estimates of the social cost of carbon: concepts and results from the DICE-2013R model and alternative approaches,” *Journal of the Association of Environmental and Resource Economists*, 1, 273–312.

A Data Appendix

Data were obtained from Rystad Energy (Rystad hereafter), an energy consultancy based in Norway that covers the global oil industry. Various ancillary data sets are also used. Price data are drawn from the EIA imported crude oil price (refiner average imported crude oil acquisition cost) <https://www.eia.gov/outlooks/steo/realprices/> accessed January 4, 2023, as data from other commonly use prices for oil, namely West Texas Intermediate and Brent, only exist from the early 1980s onwards. GDP deflators are drawn from U.S. Bureau of Economic Analysis, Gross Domestic Product: Implicit Price Deflator (series USAGDPDEFQISMEI). Global GDP is taken from the World Bank (series NYGDPMKTPCDWLD).

The data record all significant oil fields across the globe from 1970 through 2021. A field, in the data, is defined as a geologically homogeneous oil production area. This often coincides with common management and ownership. Fields vary considerably in the number of wells and the associated infrastructure. For instance, in the data, the Gullfaks offshore field in Norway is decomposed into two separate oil fields; Gullfaks, which has three oil platforms, and Gullfaks South, which has a single platform. On the other hand, the Ghawar Uthmamiyah onshore field, which is one of the largest fields in the world, is composed of many hundred wells. Different fields can, of course, be owned by a single owner.

There are 21,388 active fields in the data across the entire sample. However, we use only 15,017 active fields, since we drop fields with no reported discovery year or with missing information on reserves. We complement the fields for which we observe production over the 1970-2021 with 25,991 fields that did not produce over this period, and indeed some of these fields may not have been discovered before 2021, but have positive reserves. Thus, in total, the analysis utilizes 41,008 fields.

For each field, the data include annual production, reserves and a breakdown of operating and capital costs, as well as the characteristics of the field, such as the location, geology and climate zone.²² The distinction between a production unit (field) and its smaller components (wells) is important since, in our data, we observe cost and production information at the field level.

A.1 Measuring Costs and Reserves

The per barrel cost of production is recovered by dividing the total cost of production of each field by the reported production, q_{ft} , (in million bbl/day), computed as $c_{ft} = \frac{\sum_h \text{Expenditure}_{hft}}{q_{ft}}$,

²²There is some heterogeneity across oil crudes produced at various locations. The data measure output in energy equivalent barrels, where the benchmark is one barrel of Brent Crude. Hence, the measure of quantity accounts for the compositional heterogeneity of crudes. Another issue is that different crudes trade at different premia and discounts related to their composition.

where the various expenditure categories are $h=\{\text{Well Capital, Facility Capital, Abandonment cost, Production Operating, Transportation Operating, and SGA}\}$, and all expenditures are deflated by the US GDP deflator with 2009 as the base year. An important cost category that is omitted from our analysis is royalties and production taxes, which form about half of total expenditures. From the perspective of the total social cost of oil extraction these expenses are not relevant, so we exclude them. However, even in a perfectly competitive market, these taxes would make oil production differ from the perfectly competitive benchmark, as they can be distortionary.

These unit costs, denoted c_{ft} , are decomposed into three elements: 1) the time-invariant marginal cost, c_f ; 2) a technology-year specific cost shifter, μ_{st} , where s indexes the technology (onshore and offshore); and 3) measurement error, $\exp(\epsilon_{ft})$. That is,

$$c_{ft} = c_f \mu_{st} = c_f \mu_{st} \exp(\epsilon_{ft}). \quad (9)$$

In counterfactuals, production undertaken by field f in year t is taken to have occurred at cost $c_f \mu_{st}$ per barrel. The technology-year specific cost shifter, μ_{st} , is estimated as

$$\ln \hat{\mu}_{st} = \sum_{f \in s} \kappa_{ft} \ln c_{ft},$$

where κ_{ft} is the quantity weight of a field in a given year's total output, $\kappa_{ft} = \frac{q_{ft}}{\sum_{f \in s} q_{ft}}$. Observations are weighted by production, as opposed to giving all fields equal weighting, since a field is an already aggregated unit of production, with the extent of aggregation varying across fields.

The time-invariant marginal cost, c_f , is then estimated, allowing for measurement error, using the following (within-field) regression:

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

Estimation is conducted using weighted least squares, with the weights being the proportion of total field output done in that year.

Since we need to make predictions about the path of production and prices after our data ends, we will need to produce estimates of cost for fields that are potential reserves that are not utilized by 2021. We use Rystad's break even cost, that is the oil price that would be needed for a field to break even. These break even costs are substantially higher than the costs we estimate for operating fields, which is not surprising as the fields that are extracted earlier should have lower costs.

The data also report reserves. Reserves are the unextracted, but recoverable, quantity of oil remaining in the ground in a field. The most reliable way to measure the reserve at a point

in time is to see the entire production life of a field. The total extracted oil is the maximal reserve. Most fields are not fully exploited in the data. Hence, industry reserve estimates need to be used. The oil industry reports reserves at different levels of extraction probability. There are three levels. P90 is the quantity able to be recovered with a 90% probability, given current technical and economic conditions. The P90 reserve is the asset value that can be reported on company balance sheets under U.S. GAAP. P50 are the reserves recoverable with a 50% probability. Finally, Pmean is the expected total reserves recoverable. The level of P90, P50 and P10 can vary significantly within a field. For instance, in the North Ward Estes field in West Texas , P90, P50 and P10 in 1975 were estimated at 62.8, 89.3 and 92.7 million barrels, respectively given a \$60 price per barrel. These definitions mean that reserves will fluctuate with the oil price. In the data used here, reserves are measured and reported as P50 assuming an oil price of \$130 (in 2021 dollars).

Table A1: **Summary Statistics For Producing Fields**

variable	Mean	p50	p5	p95
Year	2001	2003	1975	2020
Revenue	76.81	4.63	0.03	240.63
Production	3.12	0.22	0.00	10.21
Reserve	108.62	3.50	0.01	257.25
Discovery year	1969	1971	1921	2006
Startup year	1976	1978	1931	2010
Off-shore	0.20	0.00	0.00	1.00
Shale Tightoil	0.04	0.00	0.00	0.00
Oil Sands	0.00	0.00	0.00	0.00
Expenditures in millions:				
Exploration Capex	0.44	0.00	0.00	0.00
Well Capex	5.49	0.00	0.00	21.79
Facility Capex	2.75	0.09	0.00	8.55
Abandonment cost	0.01	0.00	0.00	0.00
Production Opex	5.73	0.35	0.00	22.57
Transportation Opex	1.64	0.09	0.00	5.00
SGA Opex	1.76	0.18	0.00	6.48
Government Profit Oil	10.14	0.00	0.00	9.55
Royalty	12.13	0.35	0.00	32.44
Taxes Opex	0.98	0.00	0.00	1.45

A.2 Emission cost: country, technology and crudes

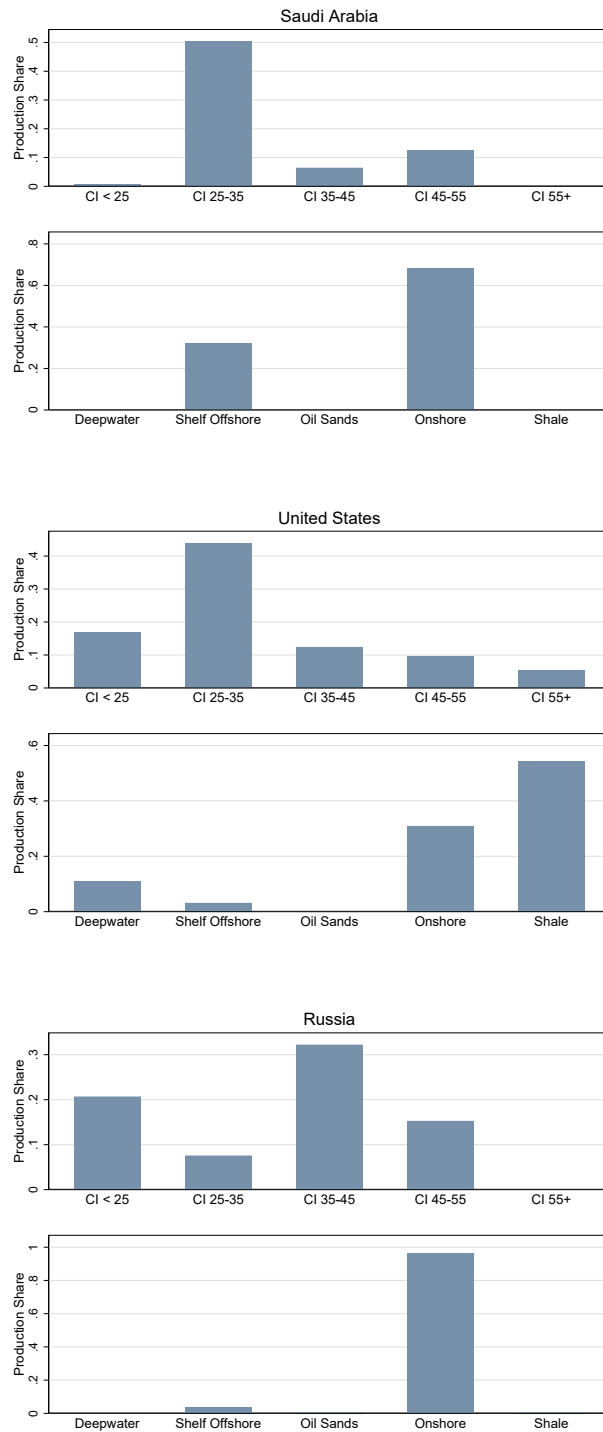
Figure A1 illustrates the distribution of crude oil production in 2015 by midstream carbon intensity (CI) categories (top panel) and production technology (bottom panel) for Saudi Arabia, the United States, and Russia.

For midstream emissions, the type of crude is a main determinant of carbon intensity (Jing et al., 2020). The top panel compares the carbon intensity across different crudes produced within each country. Since fields within countries tend to produce similar types of crude, a production-weighted average of carbon intensity across crude types effectively captures midstream emission costs. For example, Saudi Arabia’s production is concentrated in the lower carbon intensity range (25–35 $kg\ CO_2eq/bbl$), whereas most of Russia’s production falls within the 35–45 $kg\ CO_2eq/bbl$ category. This consistency suggests that a production-weighted average effectively represents midstream carbon intensity without significant loss of detail.

Upstream emissions are largely driven by extraction technology (Masnadi et al., 2018). The bottom panel displays the distribution of production across extraction technologies for each country, typically dominated by one or two main technologies. For instance, over 68% of Saudi Arabia’s production is from onshore technology, while more than 50% of U.S. production is from shale oil. These patterns imply that a country-level average can reliably capture upstream emissions as well.

This distribution underscores that variations in emission costs across countries are primarily due to their unique technological mix and crude types. Therefore, aggregating to the country level adequately reflects the heterogeneity in carbon intensity across regions within data constraints.

Figure A1: Carbon Intensity (CI) and Production Technology Distribution within Saudi Arabia, US, and Russia



Note: For each country, the top panel displays the share of crude oil production categorized by carbon intensity (CI) in $kg\ CO_2eq/bbl$ in 2015, based on data from Jing et al. (2020). The bottom panel illustrates the distribution of production by extraction technology.

Table A2: Definitions of cost components

Exploration Capital Expenditures:	Costs incurred to find and prove hydrocarbons: seismic, wildcat and appraisal wells, and general engineering costs.
Well Capital Expenditures	Capitalized costs related to well construction, including drilling costs, rig lease, well completion, well stimulation, steel costs and materials.
Facility Capital Expenditures	Costs to develop, install, maintain and modify surface installations and infrastructure.
Abandonment Cost	Costs for decommissioning a field.
Production Operating Expenditures	Operational expenses directly related to the production activity. The category includes materials, tools, maintenance, equipment lease costs and operation-related salaries. Depreciation and other non-cash items are not included.
Transportation Operating Expenditures	Represents the costs of bringing the oil and gas from the production site/processing plant to the pricing point (only upstream transportation). The category includes transport fees and blending costs.
SGA Operating Expenditures	Operating expenses not directly associated with field operations. The category includes administrative staff costs, office leases, related benefits (stocks and stock option plans) and professional expenses (legal, consulting, insurance). Only exploration and production-related SG&A are included.
Taxes Operating Expenditure	Local US taxes that are directly related to production. The category includes ad valorem taxes (county-based) and severance taxes (state-based).
Royalties	The sum of all gross taxes, including royalties and oil and export duties.
Government Profit Oil	The production-sharing agreement equivalent to petroleum taxes, but paid in kind (that is, the government contracts with a company to develop and operate the field, but retains rights to a proportion of the production). Government Profit Oil reduces the company's entitlement production and is treated as a royalty effect in company reports.

Source: Rystad U-Cube External Use Documentation.

B Demand Estimation

We follow Asker, Collard-Wexler and De Loecker (2022), where the main objective is to recover a (long-run) price coefficient. An Annual (t) global oil demand is estimated as

$$Q_t = \alpha + \beta P_t + \gamma GDP_t + g(t) + \epsilon_t, \quad (10)$$

where Q_t and P_t denote the observed annual global quantity and the price, respectively.²³ GDP_t denotes world GDP., $g(t)$ is a fourth order polynomial time trend and this allows for population growth and related macro trends to be controlled for.

To account for the endogeneity of price in this regression, a supply shifter is typically used that is uncorrelated with the error term (ϵ_t). The literature on oil demand estimation (and demand for commodities in general) acknowledges that finding suitable instruments is often challenging.²⁴

The field-level data set allows us to construct an instrumental variable that tracks shifts in the aggregate supply curve. In every year, order all barrels produced from the lowest to highest unit cost of production (c_{it}). Select the Q th lowest barrel in this ordering. The cost of this Q th barrel, is the instrumental variable used to track shifts in the supply curve.²⁵ We set Q at 15 billion bbl per year, where the lowest level of annual consumption in a year in our data is strictly below this threshold. The estimated coefficients are robust to the choice of this underlying threshold value.

In our preferred specification the estimated price coefficient is -69.66 with a standard error of 12.22.²⁶ Given that there are only 52 observations of annual price and quantity data, this seems a reasonably precise estimate.²⁷

²³The global price is measured at the annual level and therefore be seen as a quantity-weighted average price of a particular calendar year.

²⁴A popular approach is to specify a VAR and rely on specific aggregate supply-side instruments. See Caldaraa, Cavalloa, and Iacoviello (2018) for a recent study and an overview of studies and surveys reporting estimates of short-run demand elasticities.

²⁵The plant-level analogue to our approach is adopted by Foster Haltiwanger and Syverson (2008). They estimate linear demand systems for a set of (physically) homogeneous products across US plants and instrument plant-level prices with a plant-level measure of physical efficiency (TFPQ in their terminology) reflecting cost differences across plants that are plausibly excluded from the demand system.

²⁶Full details of the estimates from this baseline specification and alternative specifications are provided in Asker, Collard-Wexler and De Loecker (2022).

²⁷The standard errors are obtained using the Newey-West HAC correction with one lag, with a small-sample correction of the degrees of freedom.

C Alternative Specifications

C.1 Alternative Social Cost of Carbon

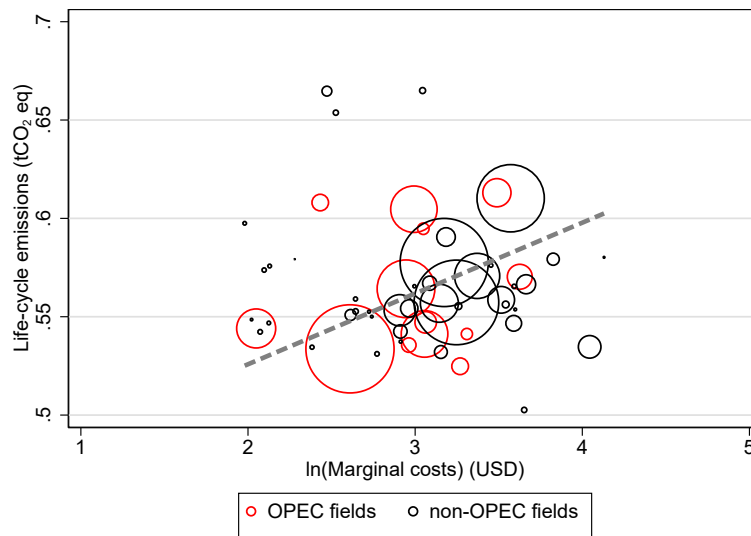
Table A3: Net present value of environmental damages, assuming a $SCC = \$190/tCO_2$

	NPV emissions (\$B)	Difference (\$B)
<u>Panel A: Homogenous emissions</u>		
Observed production	42 309	
SP production	46 417	4108
<u>Panel B: Heterogenous emissions</u>		
Observed production	44 632	
SP production	47 728	3095
<i>SP(a): Composition</i>	<i>43 482</i>	<i>- 1149</i>
<i>SP(b): Volume</i>	<i>4245</i>	<i>4245</i>

Notes: Net Present Values of environmental damages, valued at a social cost of carbon of $\$190/tCO_2$, following the estimates of Environmental Protection Agency (2023). All values are reported in billion 2021 USD.

C.2 Alternative Refinery Emission Allocation

Figure A2: Marginal costs and life-cycle emission intensity: midstream emissions allocated to refinery country



Note: Correlation of marginal extraction costs and carbon intensity of the midstream and upstream production process. In this alternative specification, midstream emissions are allocated to the country where the refinery takes place, rather than the oil production. Observations are weighted by reserves in 2018.

Table A4: Net present value of environmental damages: alternative midstream emission allocation

	NPV emissions (\$B)	Difference (\$B)
<u>Heterogenous emissions</u>		
Observed production	56 434	
SP production	60 724	4290
<i>SP(a): Composition</i>	<i>55 510</i>	<i>- 923</i>
<i>SP(b): Volume</i>	<i>5213</i>	<i>5213</i>

Notes: Net Present Values of environmental damages, valued at a social cost of carbon of \$250/ tCO_2 . In this specification, refinery emissions are allocated to the country of the refinery, rather than the country of oil production.

D The DICE model

Emissions We apply the DICE framework on the **difference** in emissions between the observed production path, and the counterfactual production path. We follow the parametrization of Nordhaus (2014).

For each 5 year period, the difference in emissions between the actual world and the counterfactual scenario can be computed. Thus, there are $E_{diff,t}$ more carbon emissions in the counterfactual world.

$$E_{diff}(t) = E_{Counterfactual}(t) - E_{observed}(t) \quad (11)$$

Atmospheric concentrations We compute the *changes* in concentrations in the atmosphere, upper ocean and lower oceans due to the additional emissions of the counterfactual scenario. To this end, the evolution in atmospheric and ocean concentration of carbon is computed as a function of E_{diff}

The carbon concentration in the atmosphere follows the cycle as described below:

$$M_{AT,diff}(t) = E_{diff}(t) + b_{11}M_{AT,diff}(t) + b_{21}M_{UP,diff}(t-1) \quad (12)$$

The upper oceans follows the following cycle:

$$M_{UP,diff}(t) = b_{12}M_{AT,diff}(t-1) + b_{22}M_{UP,diff}(t-1) + b_{32}M_{LO,diff}(t-1) \quad (13)$$

The cycle for the lower oceans is displayed below:

$$M_{LO,diff}(t) = b_{23}M_{UP,diff}(t-1) + b_{33}M_{LO,diff}(t-1) \quad (14)$$

Forcing parameter For computing the difference in radiative forcing, we include the observed level of atmospheric concentration as well. For information on the observed atmospheric concentration of carbon, satellite data (NOAA)²⁸ is used.

Following Nordhaus (2014), evolution in actual radiative forcing can be described as follows

$$F(t) = \eta \left(\log_2 \left[\frac{M_{AT}(t)}{M_{AT}(1750)} \right] \right) + F_{EX}(t) \quad (15)$$

, where $F(t)$ is the radiative forcing, and $F_{EX}(t)$ is exogenous radiative forcing.

Therefore, the *difference* in radiative forcing, compared to our actual amount, is the following:

$$F_{diff}(t) = \eta [\log_2(M_{AT}(t) + M_{AT,diff}(t)) - \log_2(M_{AT}(t))] \quad (16)$$

Where $M_{AT}(t)$ is satellite data input from NOAA.

²⁸The data can be found at gml.noaa.gov.

Temperature There's a feedback effect between atmospheric temperature and lower ocean temperature. Note that one 'period' in the DICE model equals 5 years.

Atmospheric temperature evolution is modeled as follows:

$$T_{AT,diff} = T_{AT,diff}(t-1) + \xi_1 (F_{diff}(t) - \xi_2 T_{AT,diff}(t-1) - \xi_3 [T_{AT,diff}(t-1) - T_{LO,diff}(t-1)]) \quad (17)$$

Lower ocean temperature progresses as follows:

$$T_{LO,diff} = T_{LO,diff}(t-1) + \xi_4 (T_{AT,diff}(t-1) - T_{LO,diff}(t-1)) \quad (18)$$

Contact.

MIT CEEPR Working Paper Series

is published by the MIT Center for Energy and Environmental Policy Research from submissions by affiliated researchers.

For inquiries and/or for permission to reproduce material in this working paper, please contact:

General inquiries: ceepr@mit.edu

Media inquiries: ceepr-media@mit.edu

Copyright © 2025

Massachusetts Institute of Technology



MIT CEEPR

Center for Energy and
Environmental Policy Research

**MIT Center for Energy and
Environmental Policy Research**
Massachusetts Institute of Technology
77 Massachusetts Avenue, E19-411
Cambridge, MA 02139-4307
USA

ceepr.mit.edu



MASSACHUSETTS INSTITUTE OF TECHNOLOGY