amse
école d'économie d'aix-marseille
aix-marseille school of economics

Working Papers / Documents de travail

The Welfare Economics of Oil Exploration

Renaud Coulomb
France d'Agrain
Fanny Henriet

WP 2025 - Nr 39

The Welfare Economics of Oil Exploration*

Renaud Coulomb [a] France d'Agrain [b] Fanny Henriet [c]

December 2025

Abstract

Despite growing calls to phase it out, oil exploration persists, often justified by the natural decline of existing fields and potential efficiency gains from discoveries. This paper quantifies the global welfare and environmental impacts of restricting oil exploration. We develop a global dynamic model calibrated to a granular dataset of 14,637 proven oilfields, accounting for heterogeneity in private extraction costs, capacity constraints, life-cycle carbon intensities of oil barrels, along with exploration dynamics and basin-specific estimates of yet-to-find resources. We find that exploration restrictions are an effective second-best climate policy: in the absence of a global carbon tax, a universal ban increases global welfare by $12.5 trillion due to lower social costs of oil production and use (assuming a social cost of carbon of $200/t$CO_2$eq). A partial ban by OECD and BRICS countries alone captures 66% of these gains. Under optimal carbon pricing, however, a global ban yields a modest $0.3 trillion welfare loss, as it precludes access to lower-social-cost deposits and prevents the easing of short-run capacity constraints.

Keywords: Oil exploration, Climate change, Carbon tax, Ban, Second-best, Stranded assets.

JEL Codes: Q58, Q54, Q31, Q35, Q41, Q38, H23.

*The authors are grateful to Mathias Abitbol, Lassi Ahlvik, Killian Dengreville, Simon Dietz, Bryan Leonard, François Lévêque, Shengyu Li, Leslie Martin, Pol Morvan, Katheline Schubert, Joel Stiebale, Frank Venmans, Axel Verrier, Yanos Zylberberg, as well as Léo Jean and Léo Reitzmann for research assistance. The authors also thank audiences at the 2025 ES World Congress, the 8th LSE/Imperial Workshop in environmental economics, the Manchester MWEET Workshop, the 46th IAEE conference, the F-AEE conference, the 13th ZEW-MCEE Conference, the 73rd AFSE Congress, the 30th EAERE conference and seminar audiences at CentraleSupelec, Mines Paris - PSL University, Nanterre University, World Bank (C3A group), Dauphine University and Yonsei University. The project leading to this publication has received funding from the French government under the France 2030 investment plan managed by the French National Research Agency (ANR-17-EURE-0020) and from Excellence Initiative of Aix-Marseille University—A*MIDEX. Financial support from the Faculty of Business and Economics of the University of Melbourne, the ENG Chair (Mines Paris–PSL University, TSE, Dauphine University, CentraleSupélec), and the DDX Chair (Ecole Polytechnique) is also gratefully acknowledged. This project was provided with computing HPC and storage resources by GENCI at TGCC thanks to the grant 2024-AD011015758 on the supercomputer Joliot Curie's ROME partition. Author names follow alphabetical order. The usual disclaimer applies.

[a]Mines Paris–PSL University, CEDP, France; ORCID iD: 0000-0003-3213-8891; renaud.coulomb@minesparis.psl.eu (corresponding author)

[b]Mines Paris–PSL University, CEDP, France; france.dagrain@minesparis.psl.eu

[c]Aix Marseille University-CNRS-AMSE, France; fanny.henriet@univ-amu.fr

1 Introduction

Oil remains a cornerstone of the global energy system, contributing about 30% of global energy supply in 2022 (IEA, 2025). Although greenhouse gas emissions from the extraction, refining, and combustion of proven reserves already exceed the carbon budgets compatible with limiting warming to 1.5°C or even 2°C (McGlade and Ekins, 2014; Welsby et al., 2021; Trout et al., 2022), substantial investments in oil exploration and development persist. Projections suggest that annual oil and gas exploration expenditures will reach approximately $22 billion through 2027 (Wood Mackenzie, 2023).[1]

Continued exploration carries two principal risks. First, additional discoveries could lead to an oversupply of oil, exacerbating greenhouse gas (GHG) emissions and undermining decarbonization efforts. Second, tightening climate policies in the future may render newly developed reserves unexploitable, resulting in more stranded assets. These risks have fueled growing policy interest toward restricting fossil fuel expansion through bans. In 2016, the United States banned oil and gas drilling across most U.S.-owned Atlantic and Arctic waters, with a contemporaneous move by Canada covering parts of its Arctic waters. More recently, the UK pledged in 2024 to halt new oil and gas licensing within its territory.[2] At the international level, calls for a fossil fuel non-proliferation treaty have gained momentum (Asheim et al., 2019; Newell and Simms, 2020; Van Asselt and Newell, 2022; Green et al., 2024), with endorsements from high-profile figures such as UN Secretary-General António Guterres (Guterres, 2023). In parallel, various NGO campaigns, including Just Stop Oil, 350.org, and People & Planet's Fossil Free initiative, have actively pressured governments to implement legal restrictions on oil resource expansion (Olsen and Patel, 2020).

Against this backdrop, the oil industry justifies ongoing exploration and development in the context of a carbon-constrained future on two primary grounds: (i) potentially uncovering lower-cost or lower-carbon resources to displace existing reserves, and (ii) counteracting production declines from existing fields. First, oil resources are heterogeneous, not only in private extraction costs but also in life-cycle GHG emissions per barrel (Masnadi et al., 2018).[3] Industry stakeholders suggest that emissions from new developments can be signif-

[1] This figure is conservative as it excludes post-discovery appraisal expenses.

[2] Belize, Costa Rica, Denmark, France, Greenland, Ireland, Marshall Islands, Portugal, New Zealand, Samoa, Spain, Sweden, Tuvalu, and Vanuatu have implemented similar measures to prohibit future oil exploration within their national boundaries. Several of these countries are part of the "Beyond Oil and Gas Alliance" that aim at phasing out fossil fuels. Appendix F elaborates on country measures to reduce licensing of new oil and gas projects.

[3] For instance, the life-cycle emissions of a barrel of crude from Canada are approximately 20% higher than those of a barrel from Saudi Arabia. Differences in the carbon intensity of barrels primarily stem from variations in oil type, flaring levels, and venting emissions during extraction, as well as differences in oil composition affecting refining emissions.

icantly reduced through selective field development, lower flaring and venting, and oilfield electrification using renewable energy (Gilbert, 2011; Beck et al., 2020; Reuters, 2024).[4] In principle, new discoveries could thus lower the average GHG emission intensity of the global oil supply if they displace higher-carbon reserves (Shell, 2021; Worland, 2021; TotalEnergies, 2023a).

Second, declining production resulting from well depletion and reservoir-specific factors such as pressure drops and increasing water content, as documented by IEA (2025), may outpace the rate of demand reduction in some low-carbon transition pathways, leading to short-term saturated extraction capacity and potentially elevated prices and significant losses in consumer surplus. This reinforces the incentives for exploration and the development of new fields to ease these annual production bottlenecks even though the pollution potential of proven reserves already exceeds carbon budgets (Latham and Bandal, 2024; TotalEnergies, 2024; Moore, 2025).

These two justifications have substance from a global social planner perspective engaged in mitigating climate change. Yet, they must be scrutinized against the risks of excessive GHG emissions and stranded assets.

This paper addresses a fundamental question: What are the social costs of continuing oil exploration in a world where carbon regulation remains far from optimal? To answer this, we quantify the social costs and benefits of oil exploration under two contrasting carbon pricing scenarios: (i) a climate-conscious world, where life-cycle GHG emissions are fully internalized through optimal carbon pricing, and (ii) a carbon-ignorant world, where exploration and extraction decisions are driven solely by private extraction costs and market profitability. Since pricing of oil-related GHG emissions remains systematically inadequate in practice (World Bank, 2025), the latter scenario provides a more realistic depiction of near-term developments. Our goal is therefore to quantify the external costs associated with continued exploration under insufficient climate regulation, and to compare their magnitude to the private gains from exploration accruing to producers and consumers. This comparison informs a central policy question: In the absence of optimal carbon taxation, should oil exploration be banned outright? Focusing on an exploration ban is interesting as it offers a unique policy lever. In contrast to downstream carbon taxation, which partially transfers rents to consumer countries, a ban can be more politically palatable to some producers by preserving their economic surplus. Despite its potential, this policy margin has been

[4]Electrification, flaring and venting reductions can also be implemented in existing fields. Flaring is the controlled burning of natural gas released during oil extraction. Venting, by contrast, is the direct release of uncombusted natural gas, primarily methane, into the atmosphere. It constitutes a major source of methane emissions. Recent literature highlights the importance of reducing methane emissions as a key lever for near-term climate mitigation. See Stoerk et al. (2025) for recent estimates of the social cost of methane.

understudied.

This paper is the first to quantify the welfare and environmental implications of fully or partially restricting oil exploration using a global field-level model, under various carbon taxation scenarios. The model captures the heterogeneity in private costs and carbon intensities across 14,637 documented oil fields, covering the world's proven reserves.[5] Our approach to estimating basin-specific YTF involves fitting statistical distributions to observed field sizes and characteristics in proven reserves to infer the likely properties—such as cost, carbon intensity, and volume—of future discoveries. Each field's production is capped by capacity constraints calibrated on observed data. Deposit-level production, resource discoveries, and global oil demand adjust in response to carbon taxes. By explicitly modeling deposits' heterogeneity, we capture how exploration influences not just aggregate oil supply but also the deposit-level supply structure. The model also represents the trade-offs between consumer surplus, producer surplus, and the environmental costs of oil production and use; exploration affects these metrics through an efficiency channel, potentially improving supply characteristics in terms of private cost and carbon footprint per barrel, as well as through aggregate effects on equilibrium demand size and total environmental costs. This modeling framework enables us to project future supply and demand trajectories that maximize social welfare under four policy regimes: (i) an optimal carbon tax with exploration allowed, (ii) an optimal carbon tax with an exploration ban, (iii) no carbon tax with exploration allowed, and (iv) no carbon tax with an exploration ban. We then contrast welfare outcomes of a global exploration ban across carbon-tax regimes.

In a first-best setting where life-cycle carbon emissions are priced at their full social cost (set at \$200/tCO_2eq for most of the analysis), the net social benefits of exploration are inherently positive but remain small, estimated at approximately \$0.3 trillion. For new deposits to be beneficial, they must replace higher-cost, higher-emission deposits in current proven reserves. Achieving this requires proper carbon pricing to ensure that only low-social-cost new deposits are exploited, while high-social-cost deposits in proven reserves are left untapped. Yet, due to the abundance of low-private cost, low-emission deposits within existing proven reserves,—even when accounting for annual production limits from these fields—, further exploration generates only small welfare gains, underscoring the limited justification for expanding resource pools in a climate-wise perspective.

In contrast, without adequate carbon taxation, further exploration results in substantial

[5]We use data on deposit extraction costs, production and proven reserves from the Rystad Energy database (Rystad Energy, 2022) and derive deposit carbon intensity measures from OPGEE (Oil Production Greenhouse Gas Emissions Estimator, Brandt et al. 2025), PRELIM (Petroleum Refining Life Cycle Inventory Model), and OPEM (Oil Products Emissions Model, Gordon et al. 2025) from the Oil-Climate Index.

welfare losses due to increased GHG emissions. Notably, a global exploration ban in the absence of carbon pricing moves welfare and emission levels closer to the first-best scenario, a world with exploration allowed but under an optimal carbon tax. A ban would result in a welfare increase of \$12.5 trillion and a pollution reduction of 114 $GtCO_2eq$ compared to an unregulated scenario (no carbon taxation, exploration allowed). This reduction accounts for roughly 28% of the emission savings that would be achieved by implementing the first-best policy compared to *laissez-faire*.

In summary, the welfare implications of exploration depend critically on the existence of a robust carbon price. Allowing exploration yields social benefits only when carbon taxes are set close to the social cost of carbon. With a social carbon cost of $\$200/tCO_2eq$, our analysis finds that the tax must reach at least $\$144/tCO_2eq$ for exploration to be welfare-improving. This required tax level, applied over the life-cycle emissions of an oil barrel from exploration to combustion, is significantly higher than current carbon prices, which often only partially cover GHG emissions from oil production, refining, and consumption. Furthermore, starting from an unregulated situation that describes well the present context, we show that introducing an exploration ban would bring similar welfare gains to what a global tax set around $\$69/tCO_2eq$ would. More conservative bans targeting the development of both YTF resources and undeveloped proven reserves would yield even larger welfare gains when carbon price remains substantially below its optimal level. Finally, we show that even under the assumption of a costless technological progress that would drive the GHG emissions of extraction from YTF fields to zero, our core findings remain robust.

Our findings bring three main policy insights: First, in the absence of an optimal global carbon tax, banning exploration emerges as a clear welfare-improving mitigation strategy. By limiting resource availability, it curbs oil demand growth, reduces global emissions, and moves the economy closer to its social optimum. A ban is not only an effective second-best instrument; it also has appealing features in terms of policymaking and political acceptance. By directly restricting supply, such bans put upward pressure on prices and therefore tend to preserve producer surplus. This can make them more attractive to key stakeholders than globally harmonized carbon taxes.

Second, restricting oil exploration solely in OECD and BRICS countries can substantially improve global welfare and reduce emissions, in a context of weak climate policies. A partial ban limited to these countries captures 66% of the potential gains of a universal ban, amounting to \$8.2 trillion. These results support asymmetric regulations that account for countries' unequal responsibilities and capacities for climate mitigation, in light of difficulties associated with implementing effective international transfer mechanisms.

Finally, bans on exploration and development represent a pragmatic first step toward

phasing out active fields.[6] However, timing is crucial, early enactment is paramount to avoid the creation of new stranded reserves: were a tax to start only in 2030 (set to keep cumulative emissions post-2021 below the first-best carbon budget) following a *laissez-faire* policy with authorized exploration would strand about 82% of post-2021 discovered resources.

Our primary contribution to the literature on energy policy and climate economics is to quantify the welfare gains and costs of allowing or banning oil exploration under different carbon taxation scenarios, using granular deposit-level data and incorporating basin-specific estimates of yet-to-find resources. By contrasting scenarios with and without carbon taxation, we offer novel insights into the trade-offs between exploration-induced supply increases and the potential for social gains rooted in deposits' heterogeneity in carbon intensity and private extraction cost.

Our work contributes to the growing literature on supply-side climate policy and fossil-fuel supply restrictions, including theoretical analyses (e.g., Harstad 2012; Eichner and Pethig 2017; Eichner et al. 2023) and quantitative modeling (e.g., Fæhn et al. 2017; Prest 2022; Ahlvik et al. 2022). In particular, Fæhn et al. (2017) and Prest (2022) quantify the impacts of unilateral supply-side policies aimed at curbing oil production in the contexts of Norway and the United States, respectively. Using historical variation in production taxes in the global oil industry over the 2000—2019 period, Ahlvik et al. (2022) show that higher production taxes reduce long-run output primarily through lower exploration expenditures and diminished discoveries. Looking forward, they estimate that an additional global 20-percentage-point royalty (holding other production taxes fixed) would reduce annual global emissions by 9—20 percent by the end of the century. Our approach differs in three key respects. First, we focus on quantifying the welfare costs of oil exploration across alternative carbon-tax scenarios from a normative perspective. Second, we study outright exploration bans rather than adjustments in production tax rates that treat all barrels in proven or new fields in a similar way. Third, we incorporate heterogeneity in production-related emissions across oil barrels—an argument central to the industry's defense of continued exploration—allowing us to study how exploration bans interact with life-cycle carbon pricing and to assess the potential for resource displacement resulting from new discoveries and developments.

Related to this last point, our paper builds on Coulomb et al. (2025a), who quantify historical and future carbon misallocation resulting from the neglect of heterogeneity in the carbon intensity of oil barrels in extraction decisions.[7] This previous work does not incorporate

[6] Production from active fields naturally declines (about 4–5% per year with maintenance, and up to 8% without, as documented in IEA 2025), reducing emissions. In the medium to long run, however, this natural decline is insufficient to keep emissions within carbon budgets compatible with 1.5–2°C, making additional measures to curb production from proven reserves necessary.

[7] This literature also connects to broader research on production cost inefficiencies, which examines the

oil exploration and consider aggregate supply to be fixed across alternative policy/market structure scenarios to pin down misallocation sources in the supply. In contrast, the current paper captures the impact of oil exploration on both the supply structure (deposit-level productions) and the aggregate supply to be able to quantify the welfare implications of oil exploration.

A growing literature focuses on OPEC's market power in the context of carbon mitigation. Market power is often viewed as beneficial for resource conservation (Hotelling, 1931; Solow, 1974) and thus for pollution mitigation. Asker et al. (2024) quantify the environmental benefits of OPEC's production restrictions, demonstrating that, despite the welfare losses from market distortions, restricting supply has reduced emissions over the 1970—2022 period. In contrast, Benchekroun et al. (2020) use a two-resource model to demonstrate that OPEC's strategies may inadvertently accelerate pollution by facilitating the earlier market entry of higher-cost, higher-emission producers. In the same line of research, De Cannière (2025) analyzes how carbon taxes make the cartel shift production forward in time to avoid asset devaluation, in a cartel-fringe setting with granular data on deposits. Unlike these studies, our analysis focuses on aggregate supply changes driven by exploration rather than by market organization or its interplay with carbon taxes. Although our baseline model does not explicitly formalize market power, we account for OPEC supply restrictions by calibrating field-level capacity constraints using observed production data. This approach captures the systematic under-utilization of technical capacity among certain OPEC members and thus provides a plausible approximation of market equilibrium and welfare implications arising from producer market power. We further demonstrate the robustness of our results to alternative assumptions about industry structure. Critically, even under the polar case of perfect collusion among all producers, the estimated social costs and benefits of an exploration ban remain quantitatively robust.[8]

Our paper relates to theoretical literature on the social desirability of oil exploration (Bialek and Weichenrieder, 2022) and development of new fields (Van den Bijgaart and Rodriguez, 2023). More generally, it builds on the resource literature that models the extraction of heterogeneous resources (Chakravorty et al., 2008; Van der Ploeg and Withagen, 2012;

under-utilization of lower-cost production units within sectors without pollution concern (Asker et al., 2019; Stefanski et al., 2025).

[8]This result is anchored in resource finiteness and in our global carbon budget framework, where only cumulative emissions matter and not their timing. In a carbon-ignorant world, market power does not affect cumulative emissions, as the short-term conservation effect from restricted supply is perfectly offset over time as all profitable resources (i.e., those with a private extraction cost below the clean backstop alternative) will eventually be exploited. In such a world, an exploration ban impacts welfare through changes in producer and consumer surplus and not through changes in environmental costs associated with oil production and use.

Michielsen, 2014; Fischer and Salant, 2017; Coulomb and Henriet, 2018). We depart from that theoretical literature by aiming at quantifying the welfare cost of exploration using field-level data that cover proven reserves and estimates of YTF.

Our paper also contributes to the literature on stranded assets and unburnable fossil fuels (Meinshausen et al., 2009; McCollum et al., 2014; McGlade and Ekins, 2015; Welsby et al., 2021). These studies do not consider the implications of exploration, nor enclose extraction decisions into a social cost-benefit framework. Our findings suggest that, in a world with optimal carbon taxation, one possible benefit of continued exploration is its capacity to displace dirtier existing fields—effectively creating more stranded assets. However, we also show that many newly discovered deposits in an unregulated world would offer limited social value and exacerbate environmental risks if exploited, and thus most discoveries could qualify as "carbon bombs" (Kühne et al., 2022). Our analysis indicates that delaying carbon taxation or exploration ban leads to significantly larger stranded reserves, underscoring the urgent need to slash oil supply now rather than later. Our paper relates to Kellogg (2024) and Van der Ploeg and Rezai (2020), who study how expectations about future demand shape exploration and investment decisions but do not consider exploration bans or their social costs. In our first-best scenario, only 14% of the yet-to-find resources exploited in the no-carbon-tax baseline with exploration allowed are produced by 2100. Introducing a carbon tax, and thus inducing expectations of a long-term decline in demand, depresses current investment in discovering new resources, consistent with Kellogg (2024).

Finally, our paper connects with the literature on energy system modeling that examines market equilibrium in oil supply. Many conventional models (e.g., the UCL-TIAM-BUEGO module, McGlade 2014) incorporate exploration only through an exogenous discovery rate, assuming that exploration is unresponsive to carbon policies and neglecting intertemporal trade-offs in production decisions. The IEA's World Energy Model (IEA, 2023, 2024) provides a detailed representation of current assets and project developments but does not explicitly model YTF fields or account for carbon intensity heterogeneity. While Wood Mackenzie's "Lens Upstream" model (Latham and Bandal, 2024) incorporates emission heterogeneity, it does not capture changes in aggregate supply driven by exploration decisions. In contrast, our analysis employs granular deposit-level data and basin-specific estimates of YTF resources to explicitly quantify the welfare and environmental impacts of exploration bans under different carbon pricing regimes. Our approach to estimate basin-specific YTF fits statistical distributions to observed field sizes and characteristics to infer the likely properties—such as cost, carbon intensity, and volume—of future discoveries. This modeling framework improves transparency, ensures internal consistency across basins, and enables robustness checks against external datasets like Rystad's, thereby providing a credible foun-

dation for scenario analysis and policy evaluation. Specifically, we allow exploration activity to respond to market incentives and analyze the welfare effects of combining policy instruments—carbon taxation and an exploration ban.

The remainder of the paper is organized as follows. Section 2 describes the data and estimation of deposit-level carbon intensities, capacity constraints and extraction costs of proven reserves. Section 3 discusses the industry's primary justifications to continue exploration. Section 4 describes the modeling of yet-to-find deposits. Section 5 details the numerical model and its calibration. Section 6 presents our results on the welfare impacts of banning oil exploration under alternative carbon pricing regimes. Section 7 develops robustness checks, and Section 8 offers concluding remarks and discusses the associated policy implications and challenges.

2 Oil Proven Reserves

Proven reserves refer to quantities of oil that geological and engineering data demonstrate with reasonable certainty to be recoverable under specified economic and operational conditions, while yet-to-find resources represent the volumes expected to be discovered in the future through exploration, but which have not yet been located or confirmed.

This section briefly presents the oilfield data on proven reserves and explains how we calculate field-level private extraction costs, capacity constraints, and life-cycle carbon intensities of oil barrels.

2.1 Oil-deposit Data

The Rystad Upstream dataset Our empirical analysis uses data from the Rystad UCube Database (Rystad, hereafter), one of the most comprehensive datasets of oil fields. This covers most of World oil production, with 14,637 active deposits in 2021.[9] It includes precise field-level data on oil production, exploitable reserves, discoveries, capital and operational expenditures from exploration to field decommissioning as well as production taxes, current governance (e.g., ownership and operators), field-development dates (license, discovery, approval, start-up, and production end), oil characteristics (e.g., oil type, density and sulfur content), and reservoir information (e.g., water depth, basin and location). We restrict the sample to oilfields discovered before 2022, with positive reserves in 2022 and production starting before 2080 in Rystad's forward scenario.

[9] In certain cases, a field might refer to a single oil well. In the majority of cases, especially for onshore oil fields, a field consists of multiple oil wells.

The Rystad dataset does not contain information on fields' upstream carbon intensities but records key variables that influence emissions from extraction or refining, such as oil type (e.g., bitumen or light), API gravity, gas-to-oil ratio, and sulfur content.

Additional data Two extraction techniques that affect emissions from extraction—methane flaring and steam injection— are not recorded precisely in Rystad. Flaring consists of the on-site burning of the co-produced gas. As only a minority of countries and companies collect and publish data on flared gas, this information is missing in Rystad data. We complement the data using the geocoded flaring volumes calculated by the Earth Observation Group (EOG), within the Payne Institute for Public Policy at the Colorado School of Mines, in collaboration with the National Oceanic and Atmospheric Administration (NOAA). Satellite data is collected from the Visible Infrared Imaging Radiometer Suite (VIIRS) embarked on satellites launched in 2012 and 2017. Heat emitted by gas flares are observed at night and corresponding flared volumes are computed using the VIIRS Nightfire algorithm (VNF) developed by the EOG (Elvidge et al., 2016; Zhizhin et al., 2021). We use the 95,617 positive annual flaring volumes recorded by the EOG between 2012 and 2020 and match it geographically to Rystad oil and gas fields.

Steam injection is a thermal Oil Enhanced Recovery (EOR) technique employed in some fields—mostly those producing heavy oil—to facilitate extraction. Rystad data identify the use of steam injection only for bitumen fields. We add steam-injection data from the International Energy Agency (IEA, 2018b).

Online Appendix A provides a detailed description of our data.

2.2 Modeling Deposit Private Extraction Costs and Extraction Capacities of Proven Reserves

Marginal costs Extraction costs are calculated as average costs. Specifically, for each asset that is already developed, we divide the sum of annual costs—"selling, general and administrative", "transportation", "production" OPEX, and "well" and "facility" CAPEX (after the initial development phase that lasts two years)—by the sum of annual productions. For yet-to-develop fields in proven reserves, development costs—defined in line with the literature as the "well" and "facility" CAPEX up to two years after production startup—are incorporated in the computation of average cost. All annual costs are adjusted for inflation using deflators from Rystad's database.

Field-specific capacity constraints The amount of resources extracted annually from a deposit is constrained by its extraction capacities. The extraction process at the field level typically unfolds in two stages: an initial steady output phase, the plateau, followed by a

decline phase (Fetkovich, 1980; Robelius, 2007; Höök and Aleklett, 2008; Höök et al., 2014; Jackson and Smith, 2014). During the latter phase, the decline rate is constant over time and production in a given year is thus limited to a certain fixed percentage of the remaining reserves. Field-specific capacity constraints are calibrated using observed productions recorded in Rystad database.

Online Appendix C provides detailed information on the data, our main approach for estimating field-level private extraction costs and alternative cost measures used in robustness checks, and on the calibration of capacity constraints.

2.3 Measuring Deposit Life-cycle Carbon Intensities

Upstream carbon intensity Before oil extraction starts, GHG emissions are generated from field exploration and the setting up of injecting and extracting wells.[10] After production begins, activities such as well maintenance, oil extraction and surface processing, emit greenhouse gases.

Field-level carbon intensities are assumed to be exogenous to carbon policies and time-invariant in our main approach: this reflects the role of exogenous factors, such as oil viscosity, density and content, in emissions from extraction. Emissions are also linked to extraction techniques, which are largely tied to oil type. For example, lifting heavy oils requires a more intensive use of Enhanced Oil Recovery (EOR) techniques, such as thermal EOR or Gas-EOR. Another example is flaring: methane flaring is typically associated with high carbon intensity, and is partly determined by exogenous factors such as the reservoir's gas-to-oil ratio and the distance to a significant consumer market for natural gas.

We use the *Oil Production Greenhouse Gas Emissions Estimator* (OPGEE3, Brandt et al. 2025) of the Oil-Climate Index (OCI) from the Carnegie Endowment for International Peace to estimate upstream emissions together with a regression-based approach. OPGEE3 is a bottom-up, engineering-based life-cycle assessment tool that estimates upstream greenhouse-gas emissions from oil production by modeling field geology, extraction and processing operations, and transport to refinery gate.

Midstream carbon intensity When reaching a refinery, crude oil from different fields is blended and refined into petroleum products, such as gasoline and other fuels. Additionally, a given field production can be split up between different crude streams in order to be blended with a variety of other crudes. We employ the OCI *Petroleum Refinery Life-Cycle Inventory Model* (PRELIM) to assess midstream emissions that mostly consist of CO_2, CH_4

[10] We will use the terms carbon emissions and pollution to refer to GHG emissions, similarly carbon intensity must be understood as GHG intensity. Local (air/water/soil) pollution from oil extraction, transportation, refining and use is ignored.

and N_2O. PRELIM is a crude-oil refining process model that estimates midstream emissions and product yields by simulating refinery configurations (hydroskimming, medium conversion, deep conversion) based on crude properties such as API gravity and sulfur content.

Downstream carbon intensity The downstream sector represents the transport of refined products to end consumers, and their use (mostly combustion). We use the *Oil Products Emissions Module* (OPEM) model from the Oil-Climate Index to compute downstream carbon intensity of petroleum products for each crude assay (Gordon et al., 2025). OPEM calculates downstream carbon intensity by applying emission factors and product conversion assumptions to petroleum products from each crude assay, covering transport, distribution and end-use combustion.

Online Appendix B describes OPGEE, PRELIM, OPEM, the estimation models, and the robustness checks.[11]

2.4 Descriptive Evidence

Pollution potential from proven reserves largely exceeds 1.5°C-compatible carbon budget. Oil proven reserves amount to about 1,235 Gbbl. Extracting, refining and burning all proven reserves of oil would generate about 582 $GtCO_2eq$, which is about 1.2 times the average *total* remaining carbon budget (that encompasses emissions from oil, gas and coal, and land-use change) required to keep the temperature increase below 1.5°C (IEA, 2021).

Carbon intensity varies across oil deposits. Based on our estimation, the average upstream CI of oil in 2021 was 59.7 $kgCO_2eq/bbl$, while average midstream and downstream carbon intensities were 31.9 $kgCO_2eq/bbl$ and 381.5 $kgCO_2eq/bbl$, respectively. Unconventional oil, such as bitumen, is about 25% more polluting than conventional oil, such as light oil when accounting for life-cycle emissions from exploration to combustion (Appendix Figure G1, top panel). Most of the heterogeneity comes from the extraction and refining stages of the oil life-cycle, oil type here matters a lot: bitumen upstream-midstream average carbon intensity is for instance twice as large as the one of light oil. Flaring and steam injection also play a role, and partly explain the large variation in carbon intensity within oil categories.

There is little correlation between carbon intensities and private extraction costs. At the fine level of disaggregation in our data, there is little correlation between upstream (or life-cycle) carbon intensity and private extraction costs for proven reserves or yet-to-find resources, as shown in Appendix Figures G2 and G3. This suggests that if

[11] Online Appendix of Coulomb et al. (2025a) (that uses a previous version of OPGEE) also provides more detail on CI estimation and comparison with CI estimates from other sources.

Figure 1: Life-cycle GHG Emissions, Private Extraction and Exploration Costs per Barrel by Field Discovery Year

(a) Life-cycle GHG emissions per barrel

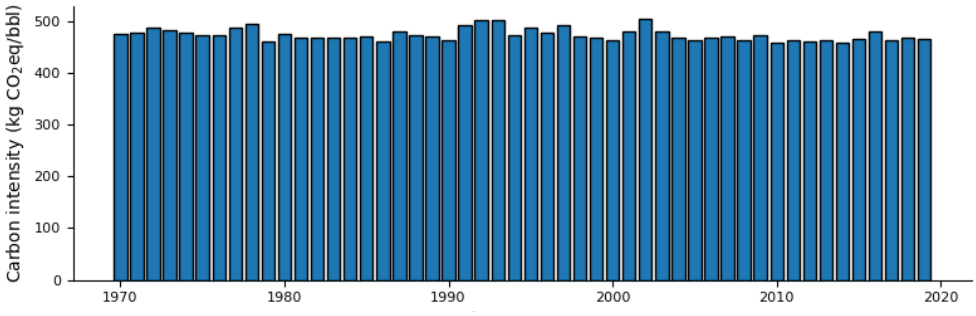

(b) Private extraction cost per barrel

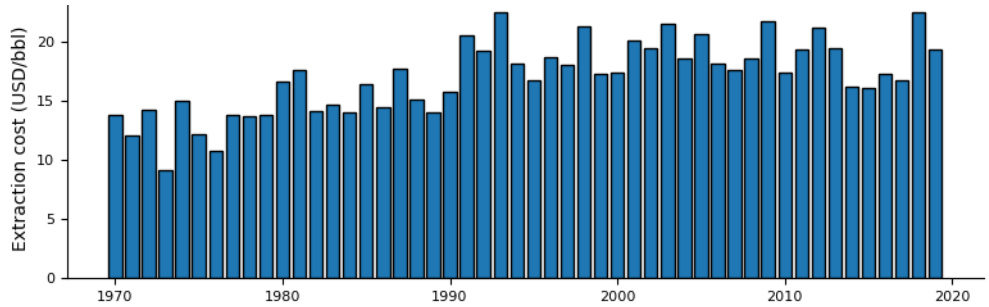

(c) Private exploration cost per barrel

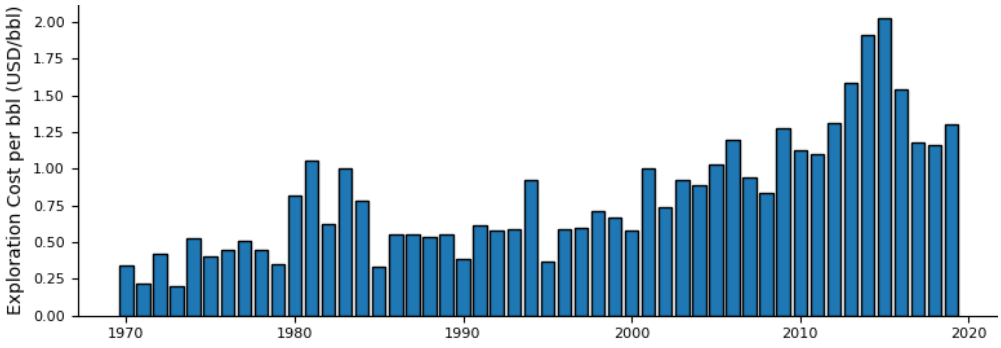

Note: Figure (a) represents the weighted average life-cycle GHG emissions per barrel by discovery year, with the initial reserves as weights. Figures (b) and (c) represent the weighted average private extraction cost and exploration cost per barrel by discovery year, with the initial reserves as weights.

exploration focuses on low private-cost fields, it will not bring on average low-carbon oil online.

The average social cost of barrels does not vary much by field's discovery year. Examining how barrel characteristics vary with the field discovery year, we find no economically significant pattern, as illustrated in Figure 1. On average, more recent discoveries tend to be associated with slightly lower emissions and higher extraction costs compared to older fields. We observe an increase in per barrel exploration costs through time, but this cost remaining small compared to development and extraction costs.

3 Preliminary Insights: Rationales for Oil Exploration

We now examine the two main arguments commonly invoked by the oil industry to justify continued oil exploration, which, on their face, may also appear relevant from the perspective of a climate-conscious social planner: 1). discovering better resources that could displace high–social cost fields in proven reserves, and 2). alleviating annual production constraints related to field capacity. Without modeling yet-to-find (YTF) resources, we find that, when examined against granular data on proven reserves, both justifications have limited empirical support.

3.1 Displacement of Existing Resources by YTF Resources

This subsection shows evidence that the characteristics of YTF resources in terms of private costs or CI must be outstanding from a social point of view to displace a significant part of the proven reserves that would be optimally used in the net-zero transition absent exploration.

We first compute the optimal future supply under the assumption of no exploration efforts and an exogenous future demand to match. Post-2021 annual oil demand follows the Net-Zero 2050 pathway (NZE), which targets a 1.5°C temperature limit with a 50% probability (IEA, 2021). We label this supply, that minimizes the social cost to match this exogenous demand path, the *Clean future without exploration*. As in the rest of our analysis, the social cost of carbon is set at \$200/tCO_2eq. Starting from the *Clean future without exploration* scenario, we then allow exploration to assess its potential to further reduce social costs by enacting proven reserves *displacement* or stranding while keeping aggregate supply unchanged.

Given that demand is exogenous in that exercise and downstream emissions are uniform across barrels, the social benefits of exploration manifest as minimizing social production costs under an exogenous demand path. Social cost reductions can only result from utilizing

YTF deposits with lower upstream-midstream carbon intensity to displace high social-cost resources.

Next, we construct a theoretical pool of YTF resources, assigning them a uniform private cost set at the median value of private costs across proven reserves and a hypothetical constant upstream-midstream carbon intensity.

Varying the hypothetical upstream-midstream carbon intensity of YTF assets between 0 and 90 $kgCO_2eq/bbl$, we compute the share of reserves utilized in the *Clean future without exploration* that should be optimally stranded by these hypothetical YTF assets under the exogenous NZE demand. The solid blue curve in Figure 2 illustrates the relationship between the theoretical upstream-midstream CI of YTF and resource displacement induced by exploration. By construction, the curve is monotonically decreasing: the greater the share of existing reserves displaced by new discoveries, the lower the CI of these newly discovered resources must be. The figure also includes a box plot depicting the observed upstream-midstream CI distribution for proven reserves and the red dotted line represents the average midstream carbon intensity of proven reserves.

Replacing even the first percentage point of global supply along the NZE pathway with YTF requires discovering fields with a CI at or below 95 $kgCO_2eq/bbl$—already below the third quartile of observed values. Replacing 10% of the global supply would require YTF fields with an upstream-midstream carbon intensity of 74 $kgCO_2eq/bbl$, approximately the first quartile of observed carbon intensity for proven reserves. Replacing 60% of the global supply or more would require a totally carbon-free extraction.

Replacing even the first percentage point of global supply along the NZE pathway with YTF requires discovering fields with a CI at or below 70 $kgCO_2eq/bbl$—already below the median of observed values. Replacing 10% of the global supply would require YTF fields with an upstream-midstream carbon intensity of 57 $kgCO_2eq/bbl$, well below the first quartile of observed carbon intensity for proven reserves. Replacing 50% of the global supply or more would require a totally carbon-free extraction.[12]

These preliminary findings, largely agnostic to true YTF characteristics, have a clear implication: the likelihood that new discoveries will be sufficiently low-carbon to significantly displace proven resources in the NZE transition is very limited. This is particularly the case given, as shown in Figure 1, the absence of any significant improvement in field CI over discovery years.[13]

[12] Any displacement target above 50 % would also require decarbonizing refining activities for oil from YTF fields; however, there is no reason to believe that such abatement would not also reduce refining emissions from proven reserves, since crudes from both proven oil and YTF are processed in the same existing refineries.

[13] Furthermore, abandoned wells may continue to emit methane in the absence of stringent environmental closure standards (IEA, 2025), weakening the climate benefits of resource reallocation.

Figure 2: Displacing Proven Reserves by YTF Resources Along an Exogenous NZE Pathway

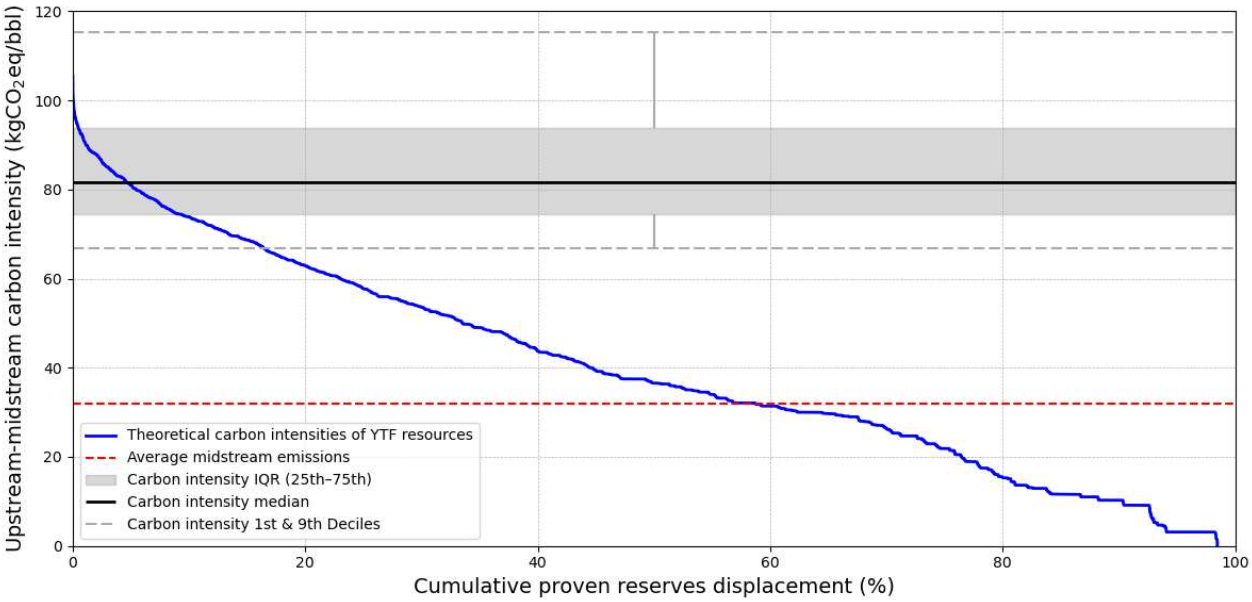

Note: The plain blue curve shows the relationship between the theoretical upstream-midstream CI of the YTF pool (y-axis) to optimally replace a given share of global supply used in the *Clean future without exploration* (x-axis), assuming a uniform private extraction cost of the theoretical pool set at the median value of that of existing reserves. For context, the box plot illustrates the observed upstream-midstream CI distribution of proven reserves, while the red dotted line indicates average midstream CI per barrel of proven reserves.

3.2 Relaxing Short-term Capacity Constraints

Industry stakeholders often argue that short-term capacity constraints necessitate continued exploration and development of new oil fields, as well as further development of existing fields, to adequately meet demand due to constraints on the supply stemming from field extraction capacities—even if the exploitation of existing, developed reserves already exceeds carbon budgets aligned with limiting global temperature rise to 1.5 or 2°C.

In the following analysis, we examine whether projected future oil demand under Net Zero Emissions (NZE) pathways defined by the International Energy Agency (IEA) can be fulfilled solely using currently proven reserves or two subsets of proven reserves (already producing fields, or already developed fields as of 2021). As in the rest of our analysis, we first narrow oil resource to oil contained in fields with a gas-to-oil (GOR) below $10,200$ scf/bbl. Our approach relies on production profiles exhibiting plateau-decline patterns, calibrated using historical data, as outlined in Section 2.2. We then extend this analysis to alternative demand projections corresponding to the IEA's Announced Pledges Scenario (APS), Stated Policies Scenario (STEPS) as well as oil demand scenarios produced by the IPCC and four private oil companies (BP, Equinor, Shell, TotalEnergies). Appendix G1 describes these scenarios in more detail. For all demand scenarios, we use the projected oil demand until

2050 and, from 2050 we assume a linear decline to zero by 2060.[14]

Table 1 records our findings, with a green check mark for achievable demand scenarios and a red cross from unfeasible scenarios. Our findings demonstrate that, even after considering short-term production constraints, developed fields (with continued maintenance investments) alone are sufficient to meet the IEA's NZE scenario oil demand. Adding oil resources from high-GOR fields ($GOR > 10,200$ scf/bbl) increases the number of feasible scenarios. The marks shown in brackets indicate a change in the feasibility of each counterfactual when the GOR constraint is discarded. Finally, considering the broader subset of proven reserves (last column), most oil demand scenarios become feasible, including scenarios significantly diverging from the net-zero objective. Thus, short-term capacity constraints alone cannot justify further exploration if future economic planning fully integrates climate targets.

This feasibility exercise tells us whether projected demand can be met without new exploration under limited capacity assumptions, but it does not quantify the welfare cost of over-exploring, i.e., discovering more resources than is socially efficient given the social cost of carbon, or under-exploring, when inefficiencies arise because resources that would be used in the first best cannot be accessed once exploration is banned. Doing so would require a genuine characterization of yet-to-find resources, as well as private and GHG–related costs along the entire production path. In the next sections, we carry out this exercise. In this analysis, we incorporate demand elasticity so that exploration decisions can be evaluated within a complete welfare framework.

4 Yet-to-find Resources

We need a modeling for yet-to-find (YTF) resources because, contrary to proven reserves, which are already discovered and well-documented, the volumes and economic viability of YTF resources are uncertain. While Rystad provides estimates for these resources, the proprietary nature of their methodology precludes a full decomposition of the underlying drivers. Consequently, we develop a bottom-up estimation strategy based on explicit hypotheses, which we then compare with Rystad's estimates to assess robustness.

To estimate the characteristics of YTF fields, such as private costs, carbon intensity, and reserves, we start from what is already known. In the oil sector, a basin is a geographically and geologically defined area where oil and gas accumulations occur. All major geological basins have already been identified and mapped through decades of exploration. YTF re-

[14]Most IEA scenarios provide demand estimates only up to 2050, with no publicly available projections beyond that date.

Table 1: Feasibility of Oil-demand Scenarios Across Different Field Pools

Demand scenario	**Demand (Gbbl)**	**Producing fields only**	**Developed fields only**	**All proven fields**
IEA Net Zero	550	✗(✓)	✓	✓
IEA APS	782	✗	✗	✓
IEA STEPS	1034	✗	✗	✗
IPCC Below-1.5°C and 1.5°C low-OS	666	✗	✗	✓
IPCC 1.5°C high-OS	901	✗	✗	✗(✓)
Shell sky	800	✗	✗	✓
Shell archipelagos	1066	✗	✗	✗
Equinor walls	1102	✗	✗	✗
Equinor bridges	689	✗	✗	✓
BP accelerated	849	✗	✗	✓
BP NZE	693	✗	✗	✓
BP New Momentum	1050	✗	✗	✗
TE Momentum	912	✗	✗	✓
TE Rupture	793	✗	✗	✓

Note: Appendix G1 describes these scenarios in more detail. *Demand*: The cumulative demand from 2022 to 2060 in each scenario, expressed in Gbbl. *Producing fields only*: Only oilfields ($GOR < 10,200$ scf/bbl) with observed historical production are allowed to produce, and only up to their historical production maximum (all the development capex are already spent). *Developed fields only*: Only oilfields ($GOR < 10,200$ scf/bbl) where development has already occurred—even partially—are allowed to produce, with production capped at the historical or modeled plateau multiplied by the share of development CAPEX already spent. *All proven fields*: Development is allowed across all proven oil fields ($GOR < 10,200$ scf/bbl), regardless of whether production or development has started. Marks in brackets indicate the feasibility of the counterfactual without restricting fields to oil fields ($GOR < 10,200$ scf/bbl); marks in brackets are not displayed if relaxing the GOR restriction has no impact on the feasibility of a counterfactual.

sources are therefore not expected to come from the discovery of new basins, but rather from additional fields within these known basins, typically in underexplored, deepwater, or technically challenging areas. For each of the 366 known basins, available data provide information on already discovered (proven) fields within these basins, including their production profiles and reserve estimates. YTF resources refer to new fields that are expected to be discovered in the future within these same basins.

The core idea of our procedure is simple: assuming an underlying size distribution and that discovery probabilities are proportional to field size, we estimate the parameters that best reproduce the observed sequence of past discoveries. More precisely, our procedure to estimate the characteristics (private cost, carbon intensity, reserves) of YTF fields follows two steps. First, we fit a log-normal population to discovered field sizes—correcting for size-biased discovery. Second, we draw the YTF fields from the fitted log-normal distribution to fill the unobserved ranks; their total volume, combined with past discoveries, constitutes the basin's Ultimately Recoverable Resources (URR). Third, we assign oil type, carbon intensity, and private cost to these simulated fields by sampling from the attributes of already discovered fields in the same basin.

4.1 Modeling

YTF field-size's distribution To model the sizes of yet-to-find oil fields, we assume that field sizes follow a log-normal distribution, consistent with empirical evidence that their distribution is heavily right-skewed, with a few large fields accounting for most of the resources (Arps and Roberts, 1958; Kaufman, 1964; McCrossan, 1969; Attanasi and Drew, 1985). Given this log-normal population, we also account for the empirical tendency of exploration to favor larger fields. Specifically, we treat observed discoveries as size-biased draws: if $S \sim \text{LogNormal}(\mu_b, \sigma_b^2)$ and the chance of discovery is roughly proportional to size (Barouch and Kaufman, 1976; Ratnaparkhi and Naik-Nimbalkar, 2012), the induced discovery distribution is the same log-normal with a location shift of $+\sigma_b^2$, i.e. $S_{\text{obs}} \sim \text{LogNormal}(\mu_b+\sigma_b^2, \sigma_b^2)$.

We first estimate (μ_b, σ_b) for each basin using the size-biased lognormal likelihood fitted to the observed discoveries, initialized from decile-based seeds. Based on these parameters, we obtain an initial estimate of the total number of fields, $N_b^{(0)}$, by matching the empirical size ranks to their expected quantiles under the fitted log-normal distribution. Following Hyndman and Fan (1996), we compute $N_i^{(0)} = (i - \delta)/[1 - F_{\mu_b,\sigma_b}(s_{(i)})]$ with $\delta \simeq 0.35$, and take the trimmed median over ranks to obtain $\hat{N}_b^{(0)}$.

Starting from these initial estimates (μ_b, σ_b) and $\hat{N}_b^{(0)}$, we simulate a large candidate pool $S_i \sim \text{LogNormal}(\mu_b, \sigma_b^2)$ for $i = 1, \ldots, \hat{N}_b^{(0)}$, convert each candidate to a score $U_i =$

$\log S_i + \varepsilon_i$ with $\varepsilon_i \sim \text{Gumbel}(0, 1)$, and retain the same number of top-scoring candidates as in the observed discoveries. The Gumbel term introduces controlled randomness: larger fields typically score higher, but smaller ones can occasionally outrank them, yielding an approximately size-proportional selection that reflects imperfect information in exploration. We sort the selected candidates, compute their cumulative discovery curve, and align it with the empirical curve by penalizing rank-by-rank relative gaps.

Parameter estimation combines two components: a shifted-lognormal likelihood for the observed magnitudes (to account for size bias) and the curve penalty for the ordering (simulation uses the unshifted lognormal). These are fitted jointly, and we use the posterior medians $(\hat{\mu}_b, \hat{\sigma}_b)$ for the next steps. Frontier basins (fewer than ten discoveries) borrow pooled medians from mature basins before applying the same procedure.

Ultimate Recoverable Resources (URR) per basin Using the fitted $(\hat{\mu}_b, \hat{\sigma}_b)$ for the *population*, we first sort the discoveries from largest to smallest. For each rank, we convert the observed size into an implied total field count by inverting the fitted lognormal CDF. Taking a trimmed median over these per-rank totals yields a new stable estimate $\hat{N}_b$.

We run up to ten refinement rounds that repeatedly (1) update the estimated total number of fields from the quantile-fill step and (2) refit the size-distribution parameters using the size-biased likelihood plus the curve-alignment penalty. Finally, we build a complete population of $\hat{N}_b$ model-implied sizes from evenly spaced percentiles of the fitted lognormal, sort them, and *overwrite* the entries at the observed ranks with the actual discoveries.

Summing this completed list gives the basin's URR. This preserves every observed size at its rank and fills only the unobserved ranks with model-consistent values.

In Appendix D.5, we assess the sensitivity of basin-level YTF estimates to the pool of observed discoveries by re-estimating YTF under a leave-one-out design, where up to 20 randomly selected discoveries per basin are removed one at a time and the resulting parameter variation is propagated to 1,000 simulated YTF asset scenarios. As is apparent from Online Appendix Figure D2, most of adjusted global YTF estimates (89%) falls in a range of ±25% of our main estimate. In a set of robustness checks, we assess the sensitivity of our findings to YTF estimates.

YTF fields' characteristics We assign properties to YTF fields based on observed fields within the same basin. Within each basin, we preserve the empirical distribution of oil types, which generally includes only two or three distinct categories. For each oil type and size bin, we apply kernel density estimation (KDE) to the observed data to assign private costs and carbon intensity, thereby capturing intra-basin geological and operational heterogeneity. Onshore/offshore status and supply segment classifications are likewise assigned to YTF fields by sampling from the corresponding observed distributions. For production dynamics,

we impute each YTF asset's production plateau and decline rate from observed fields using the same basin–oil type–size hierarchy. We then apply the same flooring rules as for observed assets, distinguishing between the top 5% largest fields and the remaining 95% (the full procedure is detailed in Online Appendix C.2 on extraction capacity).

Exploration costs To estimate the full exploration cost associated with each discovered field, we account for the fact that available data (e.g., from Rystad) only record expenditures linked to commercially successful wells. However, operators also incur costs from dry holes and non-commercial discoveries that are not observed in the dataset. To reconstruct the true average cost of exploration, we use the probability of commercial success in basin b, P_b^{CS}, based on basin maturity (Rudolph, 2016).

We assign each supply segment (e.g., onshore, shallow offshore, deepwater) a mean observed drilling cost CE_s, computed from successful exploration wells only. The expected exploration CAPEX required to discover a given field i is then inferred as:

$$E_i = \frac{CE_s}{P_b^{CS}},$$

where s is the supply segment of field i and b is its basin. This expression corresponds to the average number of wells required to achieve a successful discovery (i.e., $1/P_b^{CS}$), each costing CE_s. We use this adjusted cost as the deterministic exploration cost associated with YTF field i.

We could allow for more uncertainty in the discovery process. To the extent that identifying specific resource types requires greater effort, this feature could be included as higher exploration costs. Because observed exploration costs per barrel are small relative to private extraction costs and the social cost of oil, we expect such changes to have limited implications for our assessment of the welfare effects of oil exploration.

4.2 Descriptive Evidence

We estimate YTF assets to amount to about 248 Gbbl. According to IEA (2025), around 10 Gbbl of new oil discoveries per year would be needed through 2050 to maintain production. Our estimate of 248 Gbbl of YTF resources would therefore be sufficient to sustain production at this pace. Online Appendix D elaborates on URR estimates[15] and details our estimations of YTF assets' characteristics and alternative estimations used in robustness checks.

[15]Our calculated YTF values are roughly 25% greater than Rystad's corresponding estimates. Looking at a subsample of the world basins for which USGS data are available (50 basins matched between Rystad and the USGS), URR are 136 Gbbl for the USGS (after adjusting USGS estimates due to differences in years of estimation), 115 Gbbl for Rystad, and 97 Gbbl using our in-house method.

5 The Model of Welfare Maximization

This section introduces the social-planner welfare maximization framework and model calibration.

5.1 The Social-planner's Objective

The social planner's goal is to maximize discounted social surplus from extraction, considering pollution costs. The gross surplus from energy consumption, $u_t()$, is twice continuously differentiable, strictly increasing and concave, with $\lim_{e\to\infty} u_t'(e) = 0$ and $\lim_{e\to 0} u_t'(e) = +\infty$. We write $p(t)$ the endogenous energy price, that corresponds to marginal utility of energy consumption at each date.[16]

Energy can be provided by oil deposits and a costly, carbon-free backstop that provides a perfect substitute for oil at a constant marginal cost, c_c, per barrel-equivalent energy amount, with no capacity constraint, available from the beginning of the period. We call I, the universe of available oil deposits, representing proven oil reserves as of 2022, and Υ the universe of YTF deposits. We call Ω, the universe of proven and YTF oilfields. The exploration cost of YTF field d is E_d (nil for already found resources). Exploration costs are calibrated using historical basin-level data.[17] Each barrel has a private extraction cost, c_i. Furthermore, each deposit has an external cost that varies with its life-cycle carbon intensity, θ_i, and the social cost of carbon (SCC). The SCC is assumed exogenous and grows at the social discount rate, r, consistent with a carbon-budget approach. Writing μ the 2022 marginal carbon cost with year 2022 the year zero, the current-value SCC reads:

$$\mu(t) \quad = \mu e^{rt} \quad \text{for all } t. \tag{1}$$

Environmental costs are only a function of accumulated emissions, and the timing of pollution does not matter.[18]

[16] We assume uniformity in transportation, refining, and end-use costs across all oil barrels. Additionally, we posit that downstream emissions are homogeneous due to the minimal cost associated with adjustments in crude slates at refineries. We assume sufficient global refining capacities across our scenarios. Furthermore, our modeling presupposes that refineries currently configured for processing heavier crude can be reconfigured to handle lighter crude at minimal additional cost. This assumption is consistent with the fact that complex refinery configurations, required for processing heavier crudes, inherently include the capability to refine lighter crudes as well (U.S. Energy Information Administration, 2020a).

[17] Our approach partly captures heterogeneity across world regions regarding institutional context suitability when this pertains in observed investment and operational costs. Cust and Harding (2019) and Arezki et al. (2019) provide evidence that institutional context helps explain oil exploration activity and discovery patterns.

[18] The scientific literature shows that temperature increase is proportional to cumulative carbon emissions at a given time horizon, and independent of emission pathways (Allen et al., 2009; Matthews et al., 2009;

Oil extraction is subject to feasibility constraints: each deposit's extraction is capped by reserves, delayed until its availability year, and restricted by field-specific extraction capacities, calibrated on historical data. Most fields follow a two-phase extraction with plateau and decline stages (see Section 2.2). Throughout most of our analysis, field-level capacity constraints are taken as exogenous. In Section 7, we also consider the case in which plateau capacity arises endogenously from firms' investment decisions.

We include development lags after a field discovery before production can start as in Arezki et al. (2017); Bornstein et al. (2017). We use the median values of the number of years between a field's discovery and its first production, as observed from historical data, to calibrate development lags for onshore, offshore, and shale fields. Onshore (offshore) extraction can begin as early as 7 (12) years post-discovery or their observed startup year, whichever is sooner; for shale and tight oil, a two-year lag applies. For simplicity, we assume that oil-and-gas leases have long maturities, so potential distortions in production decisions arising from the risk of losing leases if drilling does not occur by a specified date are ignored (Herrnstadt et al., 2024).

Let $x_c(t)$, the energy flow from the carbon-free backstop in year t, $x_i(t)$ denote the annual production of deposit i in year t, $k_i(t)$ its capacity during plateau, α_i its maximal extraction rate during decline, τ_i the discovery year of YTF, H_i the discovery year of proven reserves, l_i field-specific development lag between discovery and possible year of first production, T_i the observed first-production year of already-producing proven reserves ($T_i < 2022$), R_i its initial reserves, and $R_i(t)$ its remaining reserves at date t. The social planner chooses exploration activity, $\tau = (\tau_i)_{i\in\Upsilon}$, and extraction from each oilfield and the clean backstop production, $x \equiv (x_j(\cdot))_{j\in\Omega\cup\{c\}}$, to maximize global welfare accounting for environmental costs associated with oil production and use. Taking 2022 as the starting year, the social

Meinshausen et al., 2009; Gillett et al., 2013; Stocker, T. and Xie, 2014; Allen, 2016). This evidence serves as a rationale for mitigation targets formulated in terms of carbon budgets (Rogelj et al., 2019).

welfare maximization problem is:

$$\mathcal{P}_1(\mu(t)): \quad \max_{\substack{x_i(t)\in\mathbb{R}_+ \\ \tau_i\in[2022,\infty)}} \left\{ \int_{2022}^{\infty} e^{-r(t-2022)} \left(u_t\Big(\sum_{i\in\Omega\cup\{c\}} x_i(t) \Big) - \sum_{i\in\Omega\cup\{c\}} (c_i + \theta_i \mu(t))\, x_i(t) \right) dt \right.$$
$$\left. - \sum_{i\in\Upsilon} E_i\, e^{-r(\tau_i-2022)} \right\}$$

s.t.

$$R_i(t) \equiv R_i - \int_0^t x_i(u)\, du \geq 0 \quad \text{for all } t,\ i \in \Omega \tag{2}$$
$$0 \leq x_i(t) \leq k_i \quad \text{for all } t,\ i \in \Omega \tag{3}$$
$$x_i(t) \leq \alpha_i R_i(t) \quad \text{for all } t,\ i \in \Omega \tag{4}$$
$$x_i(t) = 0 \quad \text{for all } t < \min(H_i + l_i,\ T_i),\ i \in I \tag{5}$$
$$x_i(t) = 0 \quad \text{for all } t < \tau_i + l_i,\ i \in \Upsilon \tag{6}$$

where (2) ensures cumulative production does not exceed initial reserves, (3) caps annual production by plateau capacity, (4) limits decline-phase production, and (5)–(6) prevent extraction from starting before the year the field is available for extraction. The time path of $\mu(t)$ is set by (1).

The solution to $\mathcal{P}_1(\mu(t))$ is not trivial because of the presence of capacity constraints (Holland, 2003). We thus directly solve the optimization program using Gurobi linear optimization solver.

5.2 Calibration

Utility is coherent with a crude demand function that is assumed isoelastic and calibrated to 2019.[19] The (absolute) price elasticity is set to 0.1 in 2022 (Uria Martinez et al., 2018) and increases by 0.005 per year to reflect greater long-run substitutability as consumers adjust to price changes.

Private extraction costs, carbon intensities, field-specific capacity constraints, and reserves estimates of proven reserves, are specified in Section 2. YTF field assessments by basin, along with their characteristics (CI, exploration costs, extraction costs, and volumes), are described in Section 4. The social discount rate is 3%. In our main exercise, the social cost of carbon (SCC) in 2022 is set to \$200 per ton of CO_2eq, in accordance with recent literature (Lemoine, 2021; Rennert et al., 2022).

[19]The year 2019 was selected for calibrating oil demand to avoid incorporating the market distortions associated with the COVID-19 pandemic in 2021.

We set the clean-backstop price to \$180 per barrel-equivalent so that, when GHG emissions are fully internalized (μ =\$200/tCO_2eq), the calibrated model forecasts cumulative oil consumption of 526 gigabarrels over 2022–2060, consistent with the IEA's Net Zero by 2050 scenario. Optimal deployment of the clean-energy backstop begins in 2039 and then ramps up, fully displacing oil by the end of 2044. See Appendix E.4.

5.3 Model Performance and Projections

How Well Does the Market Model Perform? We begin by validating that our model, which incorporates capacity constraints calibrated on observed production levels, can replicate the historical production split across countries. Specifically, we solve the program $\mathcal{P}_1(0)$ (our *laissez-faire* scenario with no carbon tax and unrestricted exploration), with the optimization process initialized in 2017. We first verify that predicted World production over the 2017–2019 period corresponds to observed production: we found that the predicted production is about 5% smaller than the observed one. We then predict production by country over the 2017–2019 period and compare these productions with observed data. Appendix Figure G4 shows that our model reproduces well the historical distribution of production across countries. However, some discrepancies remain, stemming primarily from two sources. First, external shocks, such as wars and sanctions, have affected production in Iraq, Iran, and Libya. While these events are partially captured in the calibration of capacity constraints, they are not fully addressed. Second, deviations of some major producers, such as Saudi Arabia, stem partially from our model's limited ability to capture the market power effects of OPEC. OPEC's market power over the past years led to reduced productions in low-cost OPEC countries such as Saudi Arabia. As capacity constraints are calibrated over observed productions, they partly account for distortions induced by market power.

Projections Looking ahead, we verify that our model's projections under a *laissez-faire* (business-as-usual) scenario,—defined as the solution to $\mathcal{P}_1(0)$, with no carbon tax, and unrestricted exploration—, are consistent with the existing literature on future oil demand. Specifically, by 2080, the endogenous cumulative demand reaches 1298.4 Gbbl to be compared to 1505.8 Gbbl projected by Rystad over the 2022–2080 period.

Our model indirectly captures market power effects through capacity constraints calibrated to observed production levels. In particular, OPEC producers are subject to tighter constraints, limiting their output relative to non-OPEC counterparts. Although these constraints are exogenous, they function as structural limitations that approximate the effects of market power. While our main framework does not permit explicit strategic behavior, the calibrated constraints reflect the systematic underutilization of OPEC true capacity ob-

served among some producers. This approach offers a reasonable first-order approximation of market power by ensuring that persistent capacity underuse by OPEC is embedded in the model's equilibrium outcomes.

To further account for market power, in one part of our analysis, we project future welfare and aggregate emissions under alternative taxation scenarios with and without authorized exploration, assuming perfect collusion among oil producers as an extreme case of a global cartel.

6 Welfare Gains and Costs of Oil Exploration

This section estimates the social gains and costs, and GHG emissions of continuing exploration activities in the future, and thus, conversely, the social benefits of exploration bans, depending on whether the social carbon cost is accounted for.

6.1 Main Results

Depending on the carbon tax policy in place, exploration will bring significantly different outcomes in terms of social welfare and additional emissions since carbon taxation will impact the selection of supplied resources and the demand size. To quantify the welfare impact of exploration, we compare two worlds: one where optimal carbon taxation is implemented and one where it is absent. Within each case, we evaluate outcomes with and without exploration. By allowing exploration under optimal carbon taxation, where life-cycle GHG emission intensities are factored into extraction decisions, we estimate the upper bound of exploration's welfare gains. Conversely, by allowing exploration without pricing external costs related to oil production and use, we estimate its lower bound. Table 2 records the main quantities of interest (private, environmental and social cost associated with supply of oil, cumulative GHG emissions and oil productions, and variation in social welfare compared to the first-best) across four policy scenarios.

Welfare gains generated by exploration under optimal carbon taxation Exploration can only yield social gains in this case, as social cost is factored in exploration and production decisions. Indeed, YTF are used to satisfy future demand only if their social cost is below the marginal utility consumers get from consuming energy. Comparing the welfare and emissions between the first-best scenario and the exploration ban scenario reveals that allowing exploration yields benefits of about \$0.3 trillion (last column, row 2). Exploration increases pollution by 5 $GtCO_2eq$ (column 4, row 1 − row 2) as it pushes production up by 12 Gbbl (column 5, row 1 − row 2). This modest welfare gain is because high-quality

Table 2: Welfare Gains and Costs of an Exploration Ban

Policy scenario	**Private cost** (TUSD)	**Env. cost** (TUSD)	**Social cost** (TUSD)	**GHG** ($GtCO_2eq$)	**Prod.** (Gbbl)	**ΔWelfare/1st best** (TUSD)
Carbon tax, with explo.	70.2	48.6	118.8	243	526	0.0
Carbon tax, no explo.	71.3	47.6	118.8	238	514	-0.3
No carbon tax, no explo.	36.8	107.5	144.3	537	1138	-23.0
No carbon tax, with explo.	29.4	130.2	159.6	651	1379	-35.5

Note: Each row corresponds to a distinct policy scenario. Scenarios are the solutions of $\mathcal{P}_1(\mu(t))$, assuming either a carbon tax of $\mu = 0$ (rows 3 and 4) or $\mu = \$200/tCO_2eq$ (rows 1 and 2) and restricting the sample to already discovered deposits (rows 2 and 3) or allowing YTF resources to enter the optimization program (rows 1 and 4). All costs are expressed in trillion USD (TUSD) and expressed in 2022 value. *Private cost* is the present value of the sum of future exploration, development and extraction costs of oil, and the clean backstop production cost. *Environmental cost* is calculated as the sum of future GHG emissions in Gigatons multiplied by $\$0.2/GtCO_2eq$. *Social cost* is the sum of private and environmental costs. *GHG* are cumulative life-cycle GHG emissions in Gigatons of CO_2eq. *Production* is measured in gigabarrels (Gbbl). *ΔWelfare/1st best* represents the absolute variation in the present value of social welfare in TUSD, compared to the first-best scenario which corresponds to the first row (*Carbon tax, with exploration*). Carbon taxation here refers to "optimal" carbon taxation.

assets are already abundant in proven reserves, as shown in Section 3. With a carbon tax of \$200 per ton of CO_2eq, demand remains low irrespective of exploration status, leaving only a small number of undiscovered assets capable of enhancing global social welfare.[20] Furthermore, annual extraction limits do not impose significant consumer losses or welfare costs at the global level.

Welfare losses brought by exploration under no carbon taxation In a laissez-faire world, that is a more plausible near future given the current mispricing of GHG emissions (Appendix Table F), exploration is driven towards searching and using low-private costs deposits.[21] Welfare impacts of such process are less a priori predictable and depend on two key factors: the correlation between carbon intensities and private costs in YTF resources that determines whether an exploration process oriented towards cheap resources brings deposits that have low social costs, and the demand response to newly discovered resources. Since there is little correlation between fields' private costs and carbon intensities, exploration oriented towards private interest does not bring better resources from a social point of view. Meanwhile, there is additional demand from the discovery of low-cost fields in a world that disregards environmental costs. This second aspect with no doubt deteriorates the welfare as it brings to the market resources whose externality is not priced. Eventually, exploration reduces social welfare if pollution cost is factored in ex post only to value environmental damages. An exploration ban would result in a welfare increase of \$12.5 trillion

[20] A carbon tax substantially curbs production from proven reserves and discoveries, revealing significant supply elasticity to carbon taxation. As shown in Appendix Table G2, the effect of a tax increase on oil production is nonlinear, and there is a clear asymmetry between how exploration responds compared to production from proven reserves.

[21] Our laissez-faire scenario includes, in some specifications presented in Section 7, private costs augmented by production taxes (e.g., royalties, opex taxes). These taxes can be interpreted as indirect uniform carbon taxation (Ahlvik et al., 2022). On this issue, see also Agnolucci et al. (2024); Belfiori and Rezai (2024).

(last column, row 3 − row 4) and a 114 $GtCO_2eq$ reduction in pollution (column 4, row 3 − 4). This abatement corresponds to a decline in global warming of about 0.05°C over the century.[22]

Exploration bias toward low-social cost or low-private cost deposits In contrast to a world in which GHG emissions are correctly priced, in an unregulated world, exploration is oriented toward low private costs deposits. Figure 3(a) displays the supply curves in an unregulated World based on proven reserves only (blue) or proven reserves augmented with YTF (red). Vertical dash lines represent cumulative demand in equilibrium with (red) and without (blue) exploration and show the increase in production due to exploration if GHG emissions are left unchecked. As the externality is not accounted for, all assets with a private cost below that of the clean backstop must be fully depleted in the long run. Exploration significantly lowers the supply curve, resulting in higher equilibrium demand. Furthermore, as apparent from Figure 3(b) that plots the environmental cost associated to the private supply curves, there is no significant reduction in the environmental cost per barrel thanks to exploration. This outcome arises from the weak correlation between private extraction costs and CI, implying that newly discovered resources are not environmentally superior.

[22]This estimate assumes that for each additional 1,000 $GtCO_2$ the temperature increases by about 0.45°C (Masson-Delmotte et al., 2021). Note that this back-of-the-envelope calculation incorporates methane emissions expressed in CO_2-equivalent terms, whereas the IPCC's temperature–carbon equivalence applies exclusively to CO_2.

Figure 3: Private Supply Curves With/Without Exploration and Their Environmental Costs

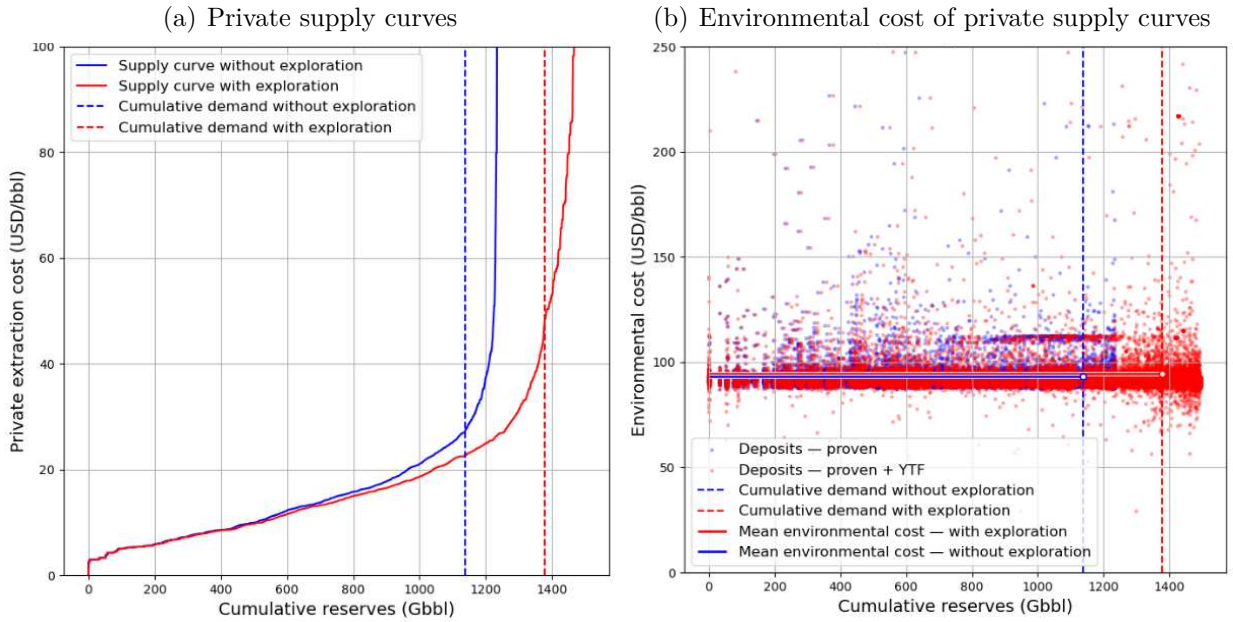

Note: Panel (a): the blue solid line represents the supply curve under the *No carbon tax, no exploration* scenario, with the blue vertical dashed line indicating the corresponding global equilibrium demand. The red solid line represents the supply curve for the *No carbon tax, with exploration* scenario, while the red vertical dashed line marks the total equilibrium demand in this scenario. Reserves are ranked by private extraction cost. Panel (b): Blue dots correspond to deposits in proven reserves, red dots to deposits of proven and YTF reserves assuming exploration is driven by laissez-faire. Dashed blue lines correspond to the cumulative demand in scenarios without exploration, red dashed lines to scenarios with exploration. The blue solid horizontal line represents the average environmental cost of deposits whose cumulative reserves is below the aggregated demand level without exploration. The red horizontal line represents the average environmental cost of assets whose cumulative reserves are below the equilibrium aggregate demand with exploration.

6.2 Minimal Carbon Tax to Make Exploration Socially Desirable

In this section, we identify the minimal carbon tax required for exploration to have a neutral effect on social welfare (assuming a social cost of carbon of \$200 per ton of CO_2eq as in the rest of the analysis). This tax value serves as a threshold for the minimum carbon tax needed to make exploration welfare-improving.

Figure 4 shows variations in welfare differences, expressed in trillions of US dollars, between the first-best supply (optimal taxation, exploration allowed) and each of the two following scenarios: sub-optimal taxation with exploration (green square), suboptimal taxation without exploration (green crosses), as the suboptimal tax rises from \$0 to \$200/tCO_2eq. The figure also compares the welfare gap between the first-best and the no-tax scenario with exploration (red squares) or without (red crosses).

Figure 4: Welfare Difference Across Policies: Suboptimal Taxes

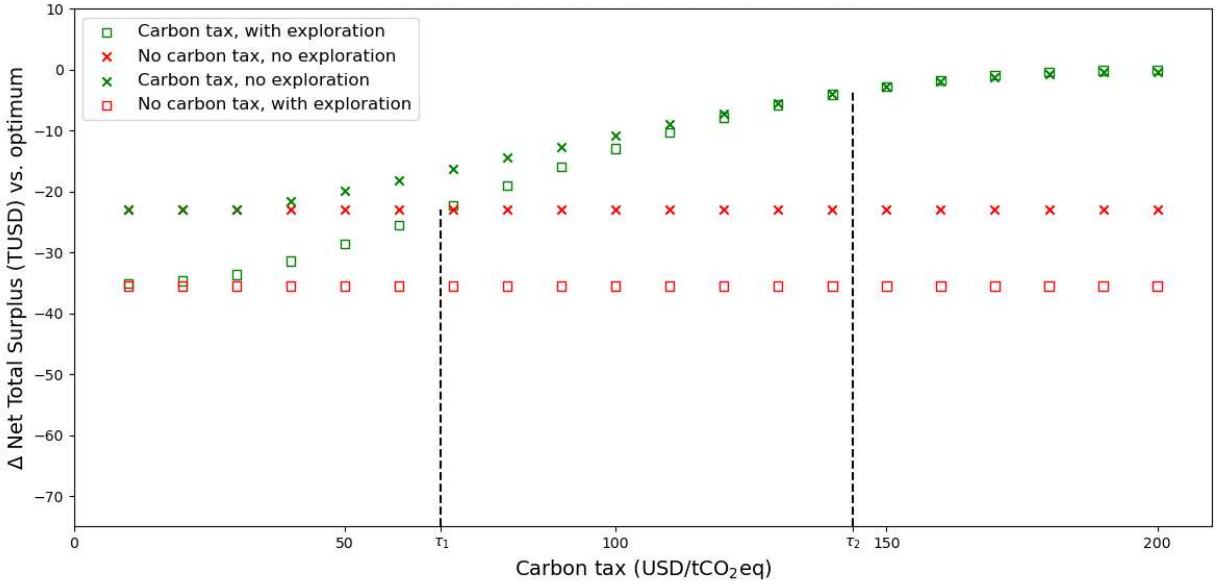

Note: This figure displays the differences in social welfare in trillion USD between the first-best (carbon tax at \$200/tCO$_2$eq and exploration allowed) and scenarios in which the carbon tax varies from \$0 to \$200/tCO$_2$eq (in 2022 US dollars) and either exploration is allowed (green square) or banned (green cross). Red crosses (squares) mark the difference in social welfare between the first-best and a no-carbon-tax scenario without (with) exploration.

Allowing exploration yields social benefits only when carbon taxes are set close to the social cost of carbon. With a social carbon cost of \$200 per tCO$_2$eq, the tax must reach at least \$144/tCO$_2$eq for exploration to be welfare-improving (the τ_2 value depicted on the graph). This tax threshold—applied over the life-cycle emissions of an oil barrel—is significantly higher than current carbon pricing schemes, which often only partially account for GHG emissions from oil production, refining, and use (Appendix Table F).

Furthermore, starting from an unregulated situation (no carbon taxation, exploration allowed) that describes well the present situation, introducing an exploration ban brings a welfare increase similar to what a carbon tax set around \$69/tCO$_2$eq (the τ_1 value depicted on the graph) would generate.

Figure 5 displays the cumulative GHG emissions across scenarios with or without exploration allowed, and with a carbon tax varying between \$0 and \$200/tCO$_2$eq. For any tax level below \$200, GHG emissions are higher when exploration is allowed. An exploration ban with no carbon taxation results in lower GHG emissions than a carbon tax below \$81 (the τ_3 value on the graph given by the intersection of the green square curve and the red cross curve) with exploration allowed.

Figure 5: GHG Emissions Across Policies: Suboptimal Taxes

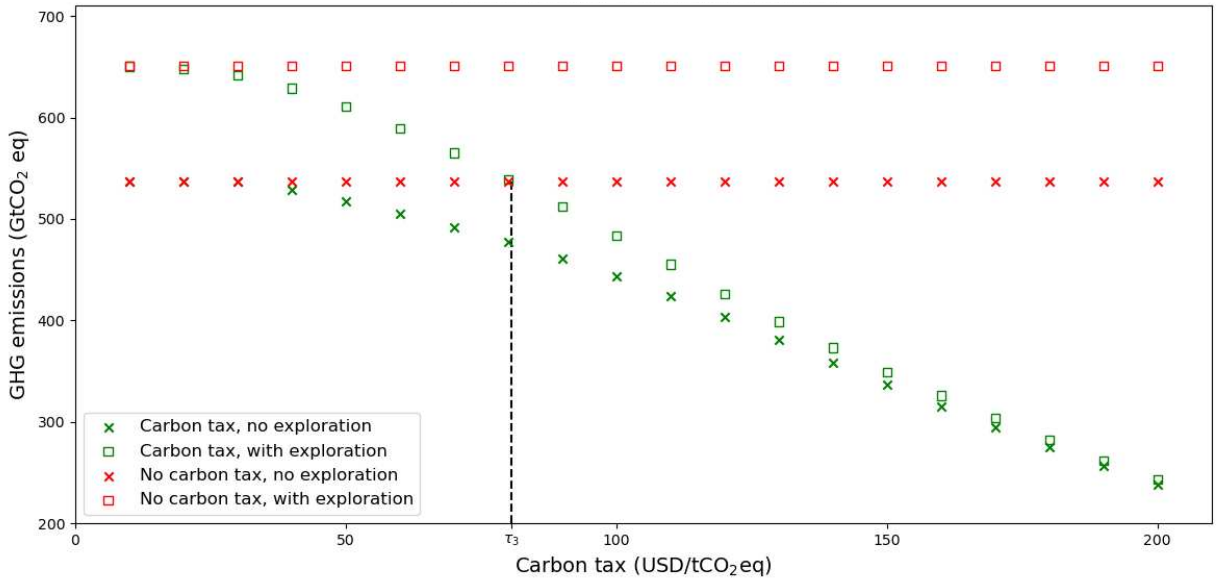

Note: This figure displays cumulative GHG emissions in the four scenarios described in Table 2. In the scenarios that include a carbon tax (green crosses and squares), it varies from \$0 to \$200/tCO_2eq (in 2022 US dollars).

6.3 Producers' surplus: Bans vs. taxes

Exploration bans may be more politically feasible than alternative policy instruments. By imposing a direct restriction on supply, such bans exert upward pressure on prices and thereby tend to preserve producer surplus. This feature can make them more appealing to certain key stakeholders relative to globally harmonized carbon taxes. Figure 6 displays, by world region, the change in profits relative to the laissez-faire benchmark in a setting where exploration is permitted and no carbon taxes are levied.[23] We consider four scenarios: (i) an optimal carbon tax with exploration allowed; (ii) a global exploration ban; and (iii) two suboptimal carbon-tax scenarios in which exploration is allowed. In these latter two cases, the tax is set so that either welfare (τ_1 in previous section) or cumulative emissions (τ_3) match the levels obtained under an exploration ban with no carbon taxation.

As is apparent, an exploration ban significantly increases global surplus by reducing the overall supply of oil, thereby pushing prices upward. Carbon taxes, by contrast, have a negative impact on producers: the higher the tax, the stronger the adverse effect on profits. This pattern holds consistently across all world regions.[24]

[23]Profits are defined as the net present value of before-tax profits (i.e., prior to production taxes) but *net* of carbon taxes. We do not report profits net of production taxes because fiscal regimes for new discoveries are difficult to anticipate.

[24]See Fischer and Salant (2017) for a comparison of the distributional impacts of carbon taxes, clean backstop cost reductions, and coalition expansion in a two-region model.

Figure 6: Discounted Profits Compared to *Laissez-faire* Across Mitigation Scenarios

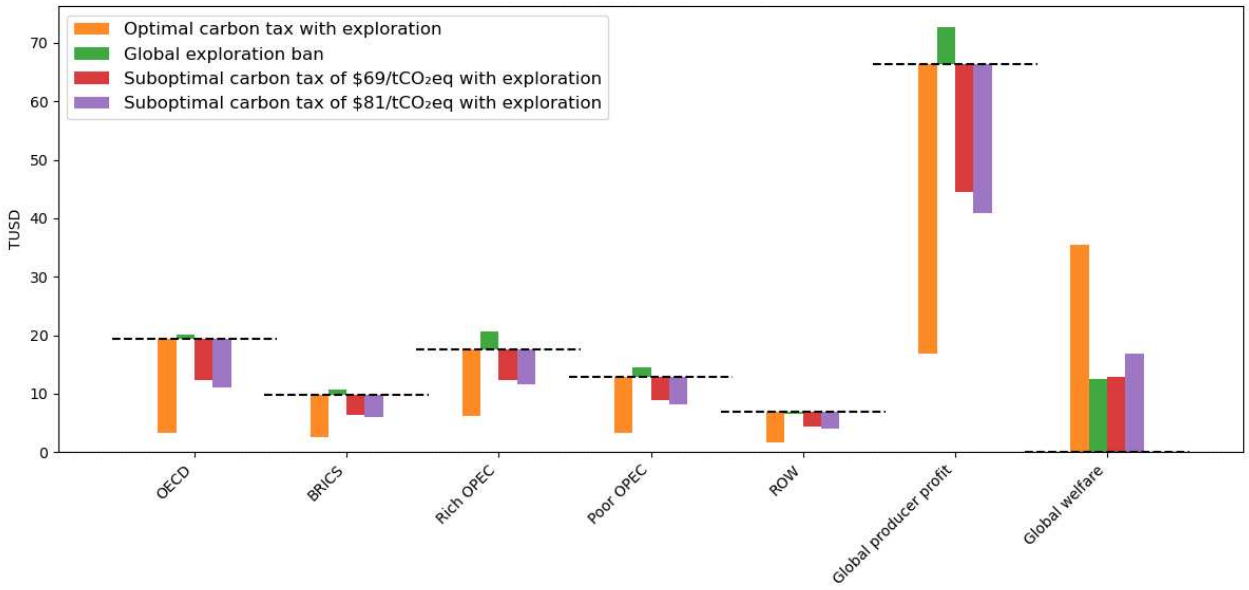

Note: The figure reports, by world region, the change in profits relative to the laissez-faire benchmark (defined as the case with exploration permitted and no carbon taxes). Profits are measured as the net present value of before-tax profits (i.e., prior to production taxes) but net of carbon taxes. The figure compares the laissez-faire (benchmark) to four scenarios: (i) an optimal carbon tax with exploration allowed; (ii) a global exploration ban; and (iii) two suboptimal carbon-tax scenarios with exploration allowed. In these latter two scenarios, the tax is calibrated so that either welfare (τ_1=\$69/tCO$_2$eq) or cumulative emissions (τ_3=\$81/tCO$_2$eq) equal the levels achieved under an exploration ban without carbon taxation.

6.4 Alternative Bans

A Ban on New Developments Our prior analyses focused on the welfare implications of exploration bans. However, approximately 25% to 32% of proven reserves remain yet-to-develop (YTD), when delineating YTD fields as those that, as of 2022, have either not been developed at all or had only been partially developed (including fields with no production). A more stringent policy option involves prohibiting all new developments—both in YTF or proven reserves. We quantify the welfare impacts of such a development ban under scenarios with and without optimal carbon taxation.

Appendix Table G3 records our findings. In a world without carbon taxation, banning new developments delivers substantially larger welfare gains compared to solely prohibiting exploration (\$19.5 trillion compared to 12.5). Yet, under an optimal carbon taxation regime, a development ban imposes larger costs compared to an exploration ban (\$2.4 trillion compared to 0.3). Overall, these findings are robust to alternative definitions of developed fields.

A Partial Exploration Ban Extraction bans can be arguably difficult to implement,

especially in developing nations that heavily depend on new oil resources.[25] In what follows, we first look into exploration bans enacted only in OECD and BRICS.[26] Appendix Table G4 records our findings. A partial ban combined with global optimal carbon taxation has little impact in terms of welfare or emission reductions. This is different if emissions are not properly priced. When a partial exploration ban is enacted with no carbon taxation, oil emissions reach 584 $GtCO_2eq$, down from 651 $GtCO_2eq$ in the *laissez-faire* scenario. This partial ban provides about 59% of the emission reduction a full ban would provide. A partial ban increases welfare by $8.2 trillion compared to the unregulated scenario; in contrast, a full ban increases welfare by $12.5 trillion. The performance of this partial ban is due to the fact that a significant proportion of YTF reserves is located in OECD and BRICS. Furthermore, these resources are on average similar to those of the rest of the world, as apparent from Appendix Figure G5. Naturally, advanced economies may resist such a ban if concerns over supply risks and the increasing leverage of OPEC in global oil markets rise to the level of first-order political considerations. Extending the partial ban to encompass high-income OPEC members—specifically the United Arab Emirates, Kuwait, and Saudi Arabia—would generate an additional reduction of 9 $GtCO_2eq$ in the absence of a carbon tax, and of 2 $GtCO_2eq$ under an optimal carbon tax.

A ban implemented in only some world regions would likely induce leakage in production and discoveries, through adjustments in global oil prices or through changes in extraction and exploration costs.[27] We find that a ban imposed in OECD, BRICS, and high-income OPEC countries, for instance, increases production from YTF resources in the remaining regions by roughly 27% relative to laissez-faire (i.e., no ban anywhere) over the 2022–2052 period. However, viewed over an infinite horizon, this spatial leakage vanishes. The reason

[25] We do not study the fiscal implications of banning oil exploration. Coulomb et al. (2025b) explore the changes in countries' fiscal revenues across various carbon taxation scenarios but exclude oil exploration from their analysis.

[26] An alternative to partial bans could be a system of compensations for low-income countries in exchange for their participation in the ban. Since our framework does not separate demand by country, we cannot compute the transfer that would leave each country welfare-indifferent between (i) joining the ban while receiving compensation and (ii) continuing to explore and produce oil. Nevertheless, by examining producer profits across scenarios, we find that producer profits in countries outside the OECD, BRICS, and high-income OPEC (i.e., UAE, Kuwait and Saudi Arabia) would actually increase under a full ban relative to the absence of a ban due to upward price pressure (see Appendix Figure G6). This suggests that compensations are not required from an equity perspective if the sole objective is to offset losses in oil producer profits. Yet, country participation condition to a global ban coalition—that implies no incentive to unilaterally quit it—cannot be satisfied without transfers or additional mechanisms: indeed, low-income countries would benefit in terms of producer surplus if they decide to unilaterally quit the coalition.

[27] Ahlvik et al. (2022) provide evidence of substantial leakage arising from an increase in oil royalties applied solely to OECD producers. On leakage from supply-side policies, one can also refer to Prest (2022) and Fæhn et al. (2017) for quantifications, as well as Van der Meijden et al. (2015); Fischer and Salant (2014) for theoretical contributions on spatial and/or intertemporal leakage.

is that the long-run stock of profitable resources in the unregulated region—defined by their cost relative to the clean backstop alternative—remains unchanged and is not impacted by a partial ban. Consequently, the ban has no effect on long-run environmental damages in regions where exploration continues, since these damages remain constant in discounted value under our carbon-budget approach.

Since equity is a first-order concern for policymakers, we examine the emissions impact of progressively extending an exploration ban, beginning with the wealthiest nations. Figure 7 plots cumulative global emissions as countries—ordered by GDP per capita—sequentially ban domestic production. The trajectory reveals that significant abatement can be achieved through a limited coalition of high-income economies. Notably, a partial ban that exempts nations with a GDP per capita below $10,000 achieves more than two-thirds of the total emission reductions realized under a full global ban.

Figure 7: Global GHG Emissions as Exploration Bans Extend by Country Income Level

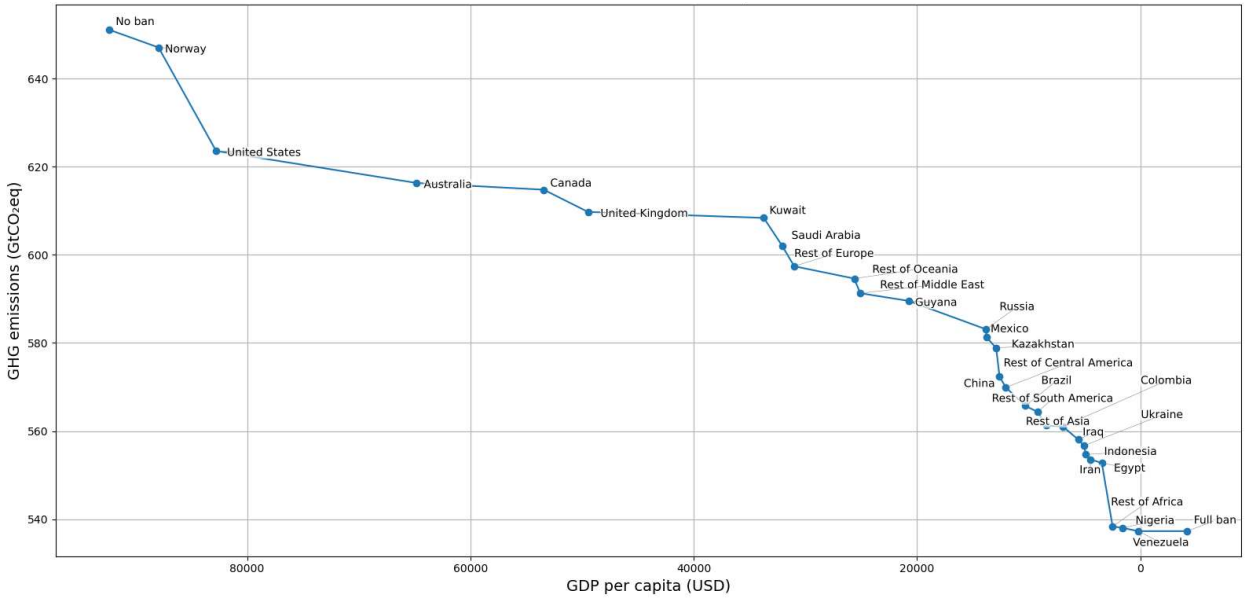

Note: Producers are removed in descending order of GDP per capita ("No ban" → "Full ban"). The x-axis is inverted so richer groups appear on the left. When a country joins the ban policy, its impact on aggregate emissions (in $GtCO_2eq$) can be read from the difference in the blue dot marked with the country name and the blue dot immediately on its left. Appendix Figure G7 replicates this figure adding other ways—used in our sensitivity analysis—to compute basin-level YTF. Appendix D3 provides some evidence of the potential variation in YTF per country based on changes in the pool of observed discoveries used in basin YTF estimation, as described in Section D.5.

6.5 Market Power

Our baseline model indirectly captures market power effects through calibrated capacity constraints reflecting observed production levels. Specifically, tighter constraints imposed on OPEC producers relative to non-OPEC ones approximate the influence of market power.

Although these constraints are exogenous, they reflect observed systematic under-utilization of actual production capacity among certain OPEC members. This approach provides a plausible first-order approximation of the welfare and market equilibrium implications of market power. Of course, under policies that imply declining future oil demand and that affect available oil resources worldwide, OPEC may have incentives to strategically adjust production and either make fuller use of its capacity to expand market share or, conversely, to lower utilization, to aim at putting upward pressure on prices. In robustness checks reported in Section 7, we therefore verify that our results are robust to varying usage level of OPEC capacity constraints.

Extreme case of perfect collusion of oil producers To further assess the implications of market power, we replicate our counterfactual constructions of oil exploration across taxation regimes, assuming perfect collusion among oil producers operating as a global cartel. In Table 3, we replicate earlier analyses under the assumption of joint profit maximization by collusive oil producers. Across all scenarios, the global cartel opts for a limit pricing strategy, coherently with findings from Andrade de Sá and Daubanes (2016). Our main results still hold: exploration generates small welfare gains if GHG emissions are correctly priced, but very large costs if they are left unchecked.[28] We explain below why market power in our framework has a negligible impact on our quantification of the costs and benefits of oil exploration ban.

Table 3: Welfare Gains and Costs of an Exploration Ban: Oil Producers Collude

Policy scenario	**Private cost** (TUSD)	**Env. cost** (TUSD)	**Social cost** (TUSD)	**GHG** ($GtCO_2eq$)	**Prod.** (Gbbl)	**ΔWelfare/1st best** (TUSD)
Carbon tax, with explo.	69.9	47.5	117.4	238	515	0.0
Carbon tax, no explo.	71.0	46.7	117.7	234	505	-0.3
No carbon tax, no explo.	33.9	107.4	141.3	537	1138	-23.9
No carbon tax, with explo.	24.6	130.0	154.6	650	1376	-37.2

Note: See the note to Table 2.

Environmental neutrality of market power in a carbon-ignorant world In the absence of carbon taxation, provided the extraction cost of an oilfield remains lower than the clean backstop cost, the field will ultimately be exploited irrespective of the prevailing market structure—be it perfect competition, cartel-fringe, oligopoly, or full collusion. Here, the conservation effect induced by market power emerges only temporarily through short-

[28] We also compare welfare changes across identical policy scenarios varying the market structure in unreported exercises. We confirm that the monopoly decreases global welfare when there is no carbon taxation compared to a competitive market, as there is no conservation effect in the long-run, therefore the welfare decreases due to the reduction in consumer surplus. Finally, we observe that perfect collusion leads to a lower welfare even when carbon emissions are correctly priced: the conservation effect is welfare-improving but small compared to the detrimental impact of market power on consumer surplus.

term supply reductions, eventually offset by future extraction, thus market power exerts no impact on cumulative emissions.[29] Under a carbon-budget framework—where the marginal social cost of carbon rises at the discount rate—the environmental cost depends exclusively on cumulative emissions. Hence, in our *laissez-faire* scenario, market power influences welfare solely through reductions in global surplus, primarily via diminished consumer surplus.

In contrast, under carbon taxation, the short-term conservation arising from market power translates into lasting conservation. Slower production rates induced by market power render certain oil reserves economically nonviable relative to the clean alternative, as oil prices increase through time due to the tax increase. In the long-term, escalating social carbon costs ensure that the transition to clean alternatives occurs in finite time, making delayed extraction increasingly unprofitable due to rising carbon taxes.

6.6 Delayed Mitigation Policies and Stranded Assets

The following analysis examines the welfare implications and stranded assets—measured as forgone reserves—when society delays climate action until 2030. Table 4 records capex expenses, social costs, welfare variation compared to first-best, total future production including that coming from post-2021 discoveries, total reserves left in 2100, as well as reserves left in 2100 coming from post-2021 discoveries.

Under the first-best policy (row 1), about 786 Gbbl of oil reserves are stranded, representing roughly 64% of 2022 proven reserves that amount to 1,235 Gbbl.[30] Only 33 Gbbl are produced using post-2021 discoveries which indicates again the limited impact discoveries can have in displacing existing reserves: about 3% of reserves used in the first-best scenario without exploration (row 2) are displaced by new discoveries if exploration is allowed.

We consider a delayed policy scenario in which society initially prepares for a laissez-faire scenario but abruptly changes course in 2030 by imposing a tax that increases at the social discount rate such as cumulative emissions from 2022 onwards equal cumulative emissions of the first-best scenario. The society must implement a carbon tax corresponding to \$205/tCO_2eq (in 2021 value) for a late carbon pricing policy starting in 2030, in order to obtain the same cumulative emissions over the 2022—2100 period as in the first-best policy. With late carbon taxation starting in 2030, exploration allowed (row 4), 853 Gbbl

[29] This is of course due to some assumptions of our setting. In particular, exploration efforts and outcomes are not tied to the industry past profits, and there is no demand destruction. In our sensitivity analysis, we study the case where the clean backstop price decreases exogenously through time regardless of the state of the carbon tax policy. Appendix Table G5 examines such a case together with full collusion of producers. The conservation effect of market power remains limited.

[30] Newly found reserves end up being exhausted in the long run, but due to decline in field extraction capacity, full exhaustion is asymptotic. Therefore, the 786 Gbbl figure is an approximation.

of reserves are stranded by 2100, including 118 Gbbl from resources discovered post-2021. Around 82% of the discovered resources post-2021 are ultimately stranded (= 118/144). Figure 8 illustrates the distribution of these stranded post-2021 discoveries across countries, expressed as a percentage of the total new stranded reserves. The US stands out with a large contribution to global new stranded assets.[31]

If a carbon tax is never implemented, an extraction ban introduced in 2030 is considerably more costly than one enacted in 2022, with an additional welfare loss of approximately $6.3 trillion (row 6 – row 2).

Table 4: Stranded Assets and Welfare Costs of Delayed Mitigation Policies

Policy scenario	**E. Capex** (TUSD)	**GHG** ($GtCO_2eq$)	**ΔWelfare** (TUSD)	**Prod (YTF)** (Gbbl)	**2100 Reserves (YTF)** (Gbbl)
CT, with explo.	0.01	243	0	526 (33)	786 (45*)
No CT, no explo.	0	537	-23.0	1,139 (0)	96* (0)
No CT, with explo.	0.7	651	-35.5	1,379 (241)	98* (2*)
CT 2030, with explo.	0.09	245	-0.9	526 (26)	853 (118)
CT 2030, no explo.	0	240	-1.0	515 (0)	720 (0)
No CT, no explo. 2030	0.08	604	-29.2	1,279 (141)	97* (1*)

Note: See the note to Table 2 for the first three rows. *CT 2030, with explo.* follows a laissez-faire scenario before 2030, without anticipation of policy change, and applies a post-2030 carbon tax increasing at the social discount rate, starting at a level ensuring cumulative 2022—2100 emissions match the first-best scenario. *CT 2030, no explo.* follows a no-taxation, no-exploration scenario before 2030, without anticipation of policy change, and applies a post-2030 carbon tax increasing at the social discount rate, starting at a level ensuring cumulative 2022—2100 emissions match the first-best scenario. *No CT, no explo. 2030.* follows a laissez-faire scenario until 2030 (without anticipation of a policy change), after which exploration and development are banned, with no carbon taxation throughout the period. *E. Capex* represents the discounted sum of all Exploration capex. *GHG* are cumulative life-cycle GHG emissions in Gigatons of CO_2eq. Δ*Welfare* represents the variation in social welfare in TUSD, compared to the first-best scenario which corresponds to the first row (*Carbon tax, with exploration*). *Prod (YTF)* represents total productions over the 2022—2100 period, and in parenthesis production coming from post-2021 discoveries. *2100 Reserves (YTF)* represents total reserves remaining in 2100 and in parenthesis remaining reserves coming from post-2021 discoveries. *: these values tend to zero as the time horizon increases to infinity. Indeed, all profitable reserves end up being exhausted in the long run, but due to decline in field extraction capacity, full exhaustion is asymptotic.

[31] Appendix Figure G9 illustrates the share of discovered reserves in each country that remain unexploited under the late carbon price scenario (*CT 2030, with explo.*). Overall, while most discovered assets end up being stranded if a carbon tax is implemented in 2030, there is significant variation across countries due to differences in their oil characteristics.

Figure 8: Distribution of Stranded Post-2021 Discoveries: Carbon Tax Set in 2030

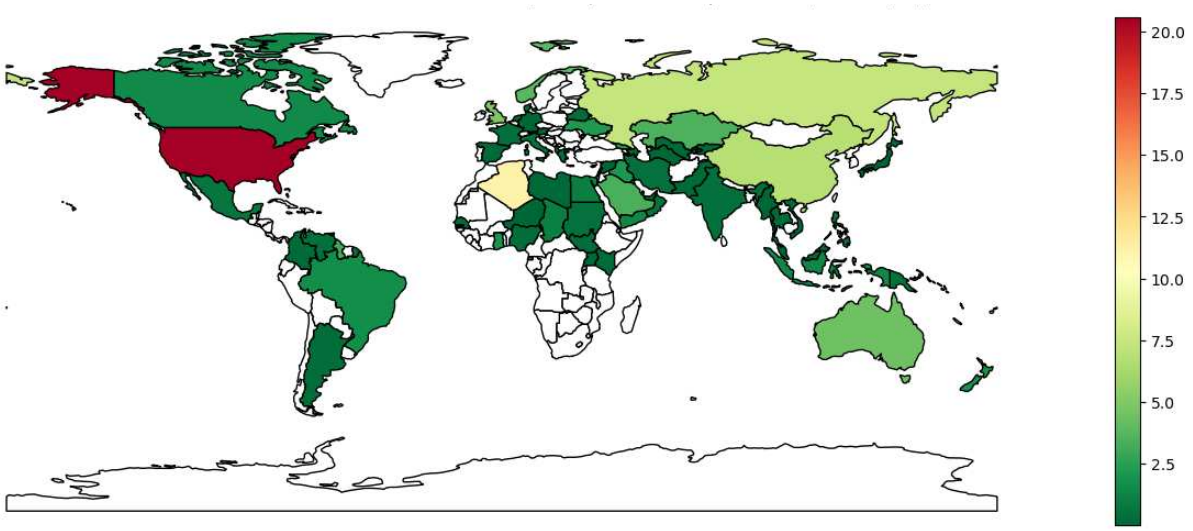

Note: The map illustrates how total stranded post-2021 discovered reserves over the scenario *CT in 2030, with explo.*, in which a carbon tax is unexpectedly set in 2030 following a laissez-faire period, are distributed across countries, showing the percentage share of each nation in the total. Countries shown in white are not assigned positive YTF values in our analysis; see Appendix D for details on the construction of YTF.

7 Sensitivity Analysis

We conduct a comprehensive sensitivity analysis on YTF resource modeling, carbon intensity estimates, private cost measures, resource availability, extraction capacity constraints, demand elasticity, and clean backstop prices. Appendices B—E provide more detail. Unless specified otherwise, the changes listed above are incorporated into all four counterfactual supply scenarios.

YTF fields We first analyze the sensitivity of our main results to changes in basin-level estimates of YTF. We reproduce our estimation of YTF in each basin by modifying the pool of observed discoveries used to estimate basin-level ultimate recoverable reserves. We use a leave-one-out resampling procedure, as described in Appendix D.5, to generate alternative YTF asset lists. Appendix Figures D2 and D3 display the global and country-level YTF distributions obtained from the resampling procedure, respectively. For each basin-level YTF estimate, there exists a specific distribution of YTF assets with characteristics (private cost, CI, reserves) attributed using the procedure described in Section D.5. We then randomly draw 1,000 basin-level YTF estimates from the pool of potential basin-level YTF distributions and re-estimate the gains and costs of exploration across taxation scenarios for each draw. Figures 9 and 10 plot the distribution of welfare variations and GHG emissions

due to exploration restriction, respectively: panels (a) show results when GHG emissions are perfectly priced (sensitivity of row 1 − row 2, columns 4 and 6, Table 2), while panels (b) show results when GHG emissions are not priced (sensitivity of row 3 − row 4, columns 4 and 6, Table 2). Overall, these figures provide evidence that our main conclusions on the gains or costs of exploration across taxation scenarios are robust to changes in YTF estimates.

We also explore how our results respond to alternative assumptions about the characteristics of YTF fields. We report these robustness checks, as well as others listed below, in Figures 11(a) and 11(b). First, we allow the estimated URR per basin to vary by ±10% or ±25% for all basins (in the previous exercise, 89% of the distribution of global YTF falls within the ±25% interval around the benchmark value). Then, private costs and carbon intensities are assigned based on the average value per *basin* × *oil type* × *size bin* rather than based on random draw from observed conditional distributions. Third, we discard all assets in frontier basins except those in offshore areas, acknowledging that technical progress may lead to more frequent discoveries in these regions. Fourth, we assume that exploration costs are zero or double estimated ones to address uncertainty in future exploration costs. Finally, we estimate the URR thanks to a creaming-curve method and assume that field size distribution within a basin follows a Pareto distribution or a log-normal distribution.

Carbon Intensities We first explore variation in upstream carbon intensities. We allow for upstream emissions to reach zero at no cost for all resources. Then, we assume that upstream CI goes to zero only for YTF fields to capture any potential heterogeneity in abatement costs between already developed and yet-to-develop fields. Another robustness check uses the Co-Product Displacement (CPD) approach in OPGEE to account for emission credits when natural gas lifted with oil is sold, reducing total CO_2eq emissions. Finally, we estimate life-cycle carbon intensities using a 20-year Global Warming Potential (GWP) instead of the standard 100-year metric.

Private Extraction Costs and Reserves Our estimates for proven reserves are consistent with estimates from the literature (British Petroleum, 2019; IEA, 2019; OPEC, 2019; BGR, 2020; EIA, 2021). Yet we first verify robustness by reducing each field's initial reserves by 10% for all fields (including YTF), to incorporate possible strong non-linearity in extraction costs when a field approaches full depletion that would make part of reserves not profitable to extract. We also use the Levelized Cost Of Energy (breakeven price) instead of average cost. Finally, we increase private extraction costs from distorsive production taxes such as royalties and other production taxes.

Extraction Capacities We consider two alternative approaches. First, we assume constant extraction capacities by relaxing constraint (4) in $\mathcal{P}_1(\mu(t))$. Second, we endogenize capacity constraints through investments in field development CAPEX and assume that by

spending x% of the total development costs recorded in Rystad forecasts, the plateau is set to x% of the plateau production forecast in Rystad, that is the average of the three highest observed annual productions. We do not allow for extra development that would increase the plateau-level production above that recorded in Rystad. Finally, we capture the potential impact of OPEC's market power by reducing OPEC potentially-used field-level capacities by 10 or 20% compared to their benchmark values.

Demand and Clean Backstop Prices We vary the 2021 demand price elasticity (0.05, 0.15, 0.2, or 0.3 instead of 0.1) and also test robustness without time variation (instead of a 0.005 exogenous annual increase). Then, we keep production for regular and light oil types above or equal to their market shares of 2019, an indirect way to ensure that key petroleum products would not decrease in market share if refineries' product slates were left unchanged. Finally, we either increase or decrease the clean backstop price by 20% and, in a separate exercise, following the 2023 version of DICE (Barrage and Nordhaus, 2024), we assume the backstop cost falls exogenously at 1% yr^{-1} over 2022–2050 and 0.1% yr^{-1} from 2051 onward.

Social Cost of Carbon We change our value for the social cost of carbon to equal either \$150 or \$250 instead of \$200/tCO_2eq. Across these SCC scenarios, the optimal tax equals the SCC. Recent evidence (Bilal and Känzig, 2024) suggests that the SCC may be even higher than assumed here, which would increase the benefits of a ban in the absence of carbon pricing.

Discount rates We finally change the discount rate, used in the social planner problem, that also governs the dynamics of the SCC to either 1.5% or 4.5 %.

Summary Overall, our main results are stable across these robustness checks as illustrated by Figures 11(a) and 11(b) that report values recorded in columns 6 and 4 of Table 2, respectively, across the specification changes described above. Exploration generates small social gains when optimal carbon taxation is in place (gap between green squares and green crosses), but large losses when no optimal carbon taxation is factored in exploration and extraction decisions (gap between red squares and red crosses).

Figure 9: Changes in Welfare due to Exploration: Sensitivity to YTF Estimates

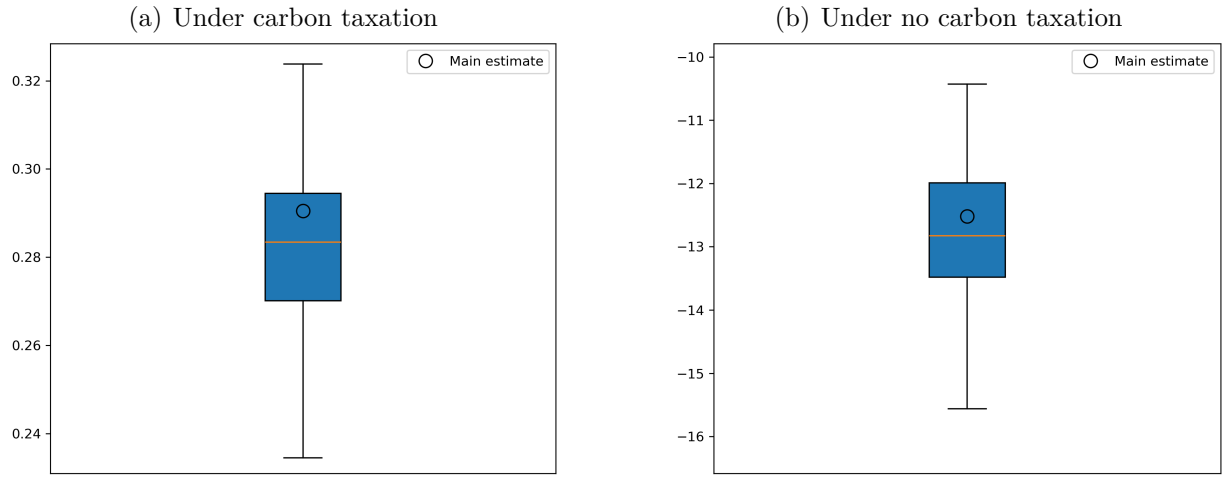

Note: The panels show the box plots of the distributions of the welfare impact of exploration (in TUSD) across carbon tax policies using the global YTF distribution obtained from the procedure described in Appendix D.5. The light-orange horizontal line denotes the median. Whiskers extend to the lower and upper decile cutoffs. The blue rectangle represents the interquantile range, spanning the 25th to the 75th percentiles. The black circle indicates the main estimate derived using the full set of observed discoveries.

Figure 10: Changes in GHG Emissions due to Exploration: Sensitivity to YTF estimates

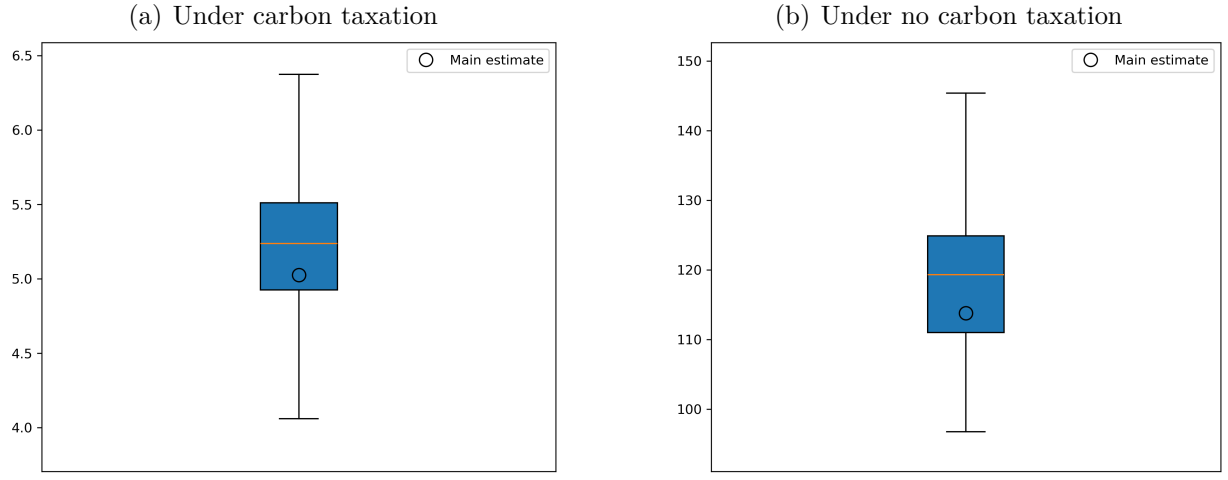

Note: The panels show the box plots of the distributions of the cumulative GHG emissions due to exploration (in $GtCO_2eq$) across carbon tax policies using the global YTF distribution obtained from the procedure described in Appendix D.5. The light-orange horizontal line denotes the median. Whiskers extend to the lower and upper decile cutoffs. The blue rectangle represents the interquantile range, spanning the 25th to the 75th percentiles. The black circle indicates the main estimate derived using the full set of observed discoveries.

Figure 11: Sensitivity Analysis

(a) Social welfare differences

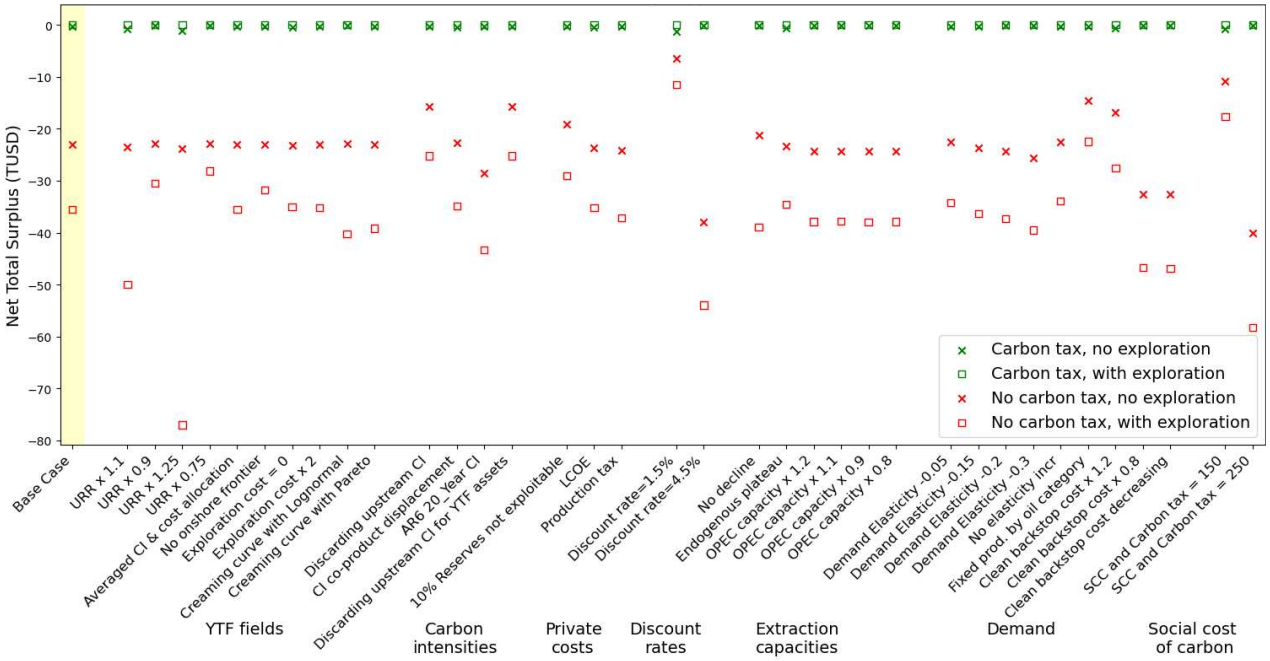

(b) Cumulative GHG emissions

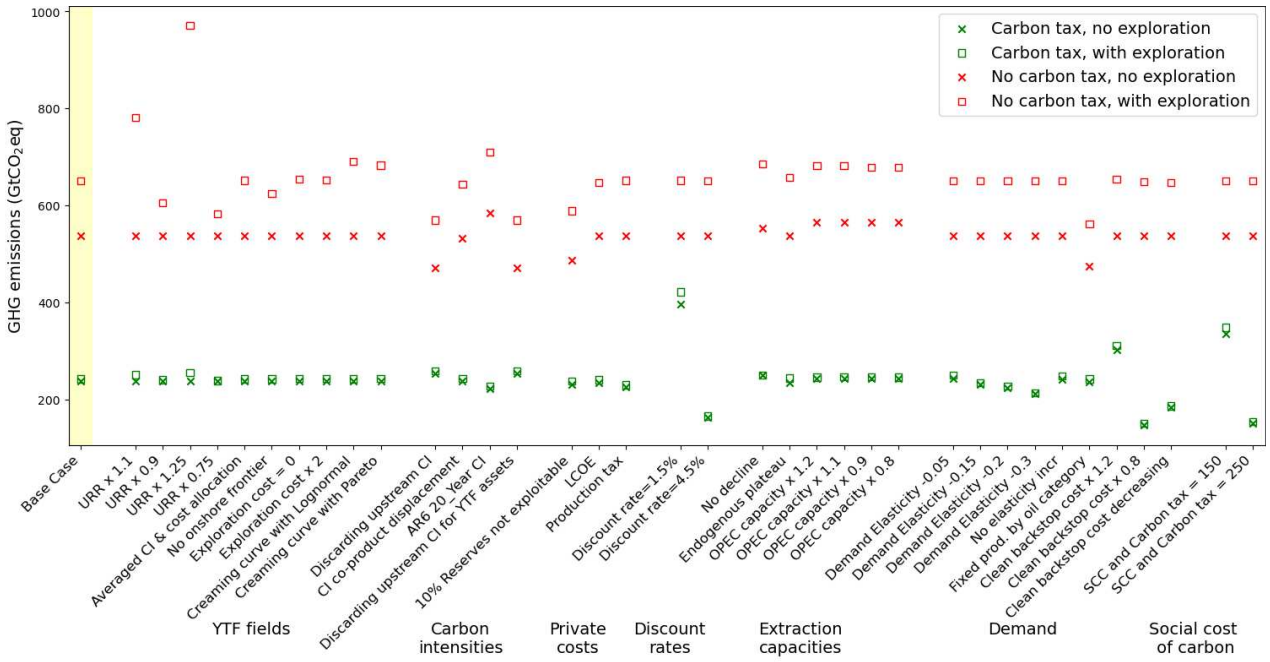

Note: Panel (a) illustrates the social welfare differences (in 2021 Trillion USD) between the optimal policy scenario *Carbon tax with exploration* and the three alternative counterfactuals: *Carbon tax, no exploration* (green cross), *No carbon tax, no exploration* (red cross) and *No carbon tax, with exploration* (red square), for each robustness check presented in Section 7. The bar highlighted in yellow shows the base case that corresponds to the last column of Table 2. Panel (b) displays the total GHG emissions (in $GtCO_2eq$) in the four counterfactuals (i) *Carbon tax, with exploration* (green square), (ii) *Carbon tax, no exploration* (green cross), (iii) *No carbon tax, no exploration* (red cross) and (iv) *No carbon tax, with exploration* (red square), for each robustness check presented in Section 7. The bar highlighted in yellow shows the base case that corresponds to Column 4 of Table 2.

8 Concluding Remarks

Our analysis demonstrates that in the realistic scenario where GHG emissions are significantly underpriced, banning oil exploration largely increases global welfare. Continued oil exploration is only socially beneficial under stringent carbon pricing that almost fully internalizes GHG emissions. In a hypothetical world in which oil-related GHG emissions are perfectly priced, exploration can yield welfare gains by replacing high-private cost, high-emission fields with lower-private cost, lower-emission fields. However, given the already abundant supply of cheap and relatively low-carbon oil among existing reserves, the net benefits of further exploration even in such a world remain limited. Furthermore, our analysis reveals that supply bottlenecks arising from field-level constraints on oil flows do not justify extended exploration and development once greenhouse gas emissions are correctly priced.

A key policy implication of our findings is that, in the absence of an optimal global carbon tax, banning exploration is an effective second-best policy that moves the economy closer to a socially optimal outcome. The welfare gains of an exploration ban are comparable to those achieved by a carbon tax on all life-cycle emissions of oil set at approximately \$69/tCO_2eq, a tax level substantially higher than current global carbon pricing regimes. Overall, policies that restrict resource expansion, alongside efforts to strengthen carbon pricing mechanisms, are crucial to aligning the energy sector with long-term climate objectives.

Exploration bans may be more politically feasible than alternative policy instruments. As they amount to a forced restriction of supply, they exert upward pressure on prices and thereby tend to preserve producer surplus, making them more attractive to some key stakeholders compared to global carbon taxes. Furthermore restricting exploration may be more politically viable than outright *extraction* bans.

Significant challenges remain. First, enforcing exploration bans is particularly challenging for low-income countries, where economic and political constraints may undermine ban viability. Advanced economies could finance reserves stranding in low-income countries with buying mechanisms such as the one described in Harstad (2012). Policies aimed at purchasing and preserving oil reserves should not be limited to producing fields but should also encompass basins where YTF resources are located, which makes them costlier from a public-budget perspective. The Yasuní case—an attempt at North–South compensation for leaving oil resources unexploited—illustrates the challenges of this approach (Einhorn et al., 2023). In contrast, partial bans allowing exploration in low-income countries while restricting it elsewhere can substantially reduce emissions, account for differences in countries' capacity to pay for mitigation, and avoid dependence on direct international transfers.

Second, each country imposing an exploration ban contributes to tightening global supply, thereby exerting upward pressure on oil prices. Higher prices, in turn, weaken the incentives for other producing countries to adopt similar bans, thus limiting their overall effectiveness. Non-cooperative countries that do not impose exploration bans—or, more generally, oil producers that benefit from bans implemented elsewhere—could be asked to contribute to a compensation fund that would support countries choosing to restrict exploration.

Third, exploration bans may lead to greater market concentration, which in turn can affect market power and the security of supply for oil-consuming countries. Yet, since oil is largely homogeneous from the demand side and transportation costs are likely relatively limited, supply risks for oil remain more contained than for certain other resources, such as critical minerals. Moreover, this concern could be factored in the design of partial bans.

Fourth, the legal feasibility of an exploration ban varies across institutional settings. In most countries, governments retain subsurface mineral rights, making such bans a straightforward policy instrument. The United States is a notable exception. While bans on federal lands (including offshore areas) would not raise ownership concerns, restrictions on privately owned onshore reserves could entail compensation requirements, increasing legal and political complexity.

Finally, the impermanence of exploration bans remains a critical concern, as future political majorities may repeal them. The same concern applies to many other policy instruments—such as carbon taxes, subsidies for pollution abatement, or emission quotas—which influence long-term investment decisions only when they are perceived as credible and durable. It is therefore essential to design financing mechanisms that compensate countries for durably leaving oil reserves unexploited and halting exploration. Since the permanence of such bans is crucial, one promising approach—drawing inspiration from forest preservation mechanisms (Harstad and Storesletten, 2023)—is to establish conditional loans with infinite maturity, where principal and interest become repayable only if exploration activities resume.

Despite its contributions, three main limitations remain. First, our model does not explicitly capture OPEC's strategic response to exploration bans or carbon pricing. Since our primary focus is on the welfare implications of exploration across policy regimes, rather than the industrial structure of the oil market, we adopt a simplified treatment of market power. While market power is approximated through capacity constraints calibrated to historical production levels, future strategic behavior remains uncertain. Robustness checks confirm that our main quantification of the social costs and gains of exploration remains unchanged under the assumption of perfect collusion of world oil producers, which clearly overstates OPEC's actual influence.

Second, estimating the future availability and characteristics of yet-to-find resources is inherently complex and subject to many assumptions. To make the analysis tractable while retaining key insights from the energy economics literature, we have represented exploration uncertainty through an augmented exploration cost framework that accounts for the risk of unsuccessful discoveries. While this approach allows us to capture the key trade-offs of exploration under different policy settings, future work could refine these estimates by incorporating more detailed geological and technological uncertainties. We could introduce greater uncertainty in the timing of resource discoveries with specific characteristics and define the decision as whether to exploit a discovered resource, conditional on its characteristics. This extension would be more computationally intensive, but we expect the results to remain quantitatively similar. If identifying particular resource types requires more intensive exploration effort, this would essentially translate into higher exploration costs. In robustness checks, we increase exploration costs and find that our main results are largely unaffected, because exploration costs are small relative to private extraction costs or to the social cost of oil. Moreover, rapid discovery has little effect on environmental costs when there is no carbon tax, because profitable YTF resources are not sensitive to discovery timing in that case, but rather to the size of the pool of profitable YTF resources. By contrast, discovery timing matters for the welfare gains from exploration when an optimal carbon tax is available. In that setting, our framework therefore tends to overestimate the gains from exploration. In the same spirit, uncertainties about future demand and prices—and their implications for oil production and investment decisions—could also be explored (Van der Ploeg, 2010; Kellogg, 2014).

Finally, our partial-equilibrium approach has left aside important macroeconomic phenomena attached to carbon policy (Golosov et al., 2014) and to resource discoveries e.g., production factors move across sectors, changes in exchange rates, trade balance and inflation, see Arezki et al. 2017; Perez-Sebastian et al. 2021. Incorporating these macroeconomic channels into an analysis of the welfare implications of oil exploration is an important avenue for future research.

References

Agnolucci, Paolo, Carolyn Fischer, Dirk Heine, Mariza Montes de Oca Leon, Joseph Pryor, Kathleen Patroni, and Stéphane Hallegatte. "Measuring total carbon pricing." *The World Bank Research Observer* 39, 2: (2024) 227–258.

Ahlvik, Lassi, Jørgen Juel Andersen, Jonas Hveding Hamang, and Torfinn Harding. "Quantifying supply-side climate policies." Technical report, Working Paper No 01/2022, Centre for Applied Macro- and Petroleum economics (CAMP), BI Norwegian Business School, 2022.

Allen, Myles R. "Drivers of peak warming in a consumption-maximizing world." *Nature Climate Change* 6, 7: (2016) 684–686.

Allen, Myles R., David J. Frame, Chris Huntingford, Chris D. Jones, Jason A. Lowe, Malte Meinshausen, and Nicolai Meinshausen. "Warming caused by cumulative carbon emissions towards the trillionth tonne." *Nature* 458, 7242: (2009) 1163–1166.

Andrade de Sá, Saraly, and Julien Daubanes. "Limit pricing and the (in)effectiveness of the carbon tax." *Journal of Public Economics* 139: (2016) 28–39.

Anthoff, David, and Richard S.J. Tol. "The uncertainty about the social cost of carbon: A decomposition analysis using fund." *Climatic Change* 117, 3: (2013) 515–530.

Arezki, Rabah, Valerie A. Ramey, and Liugang Sheng. "News shocks in open economies: Evidence from giant oil discoveries." *The Quarterly Journal of Economics* 132, 1: (2017) 103–155.

Arezki, Rabah, Frederick van der Ploeg, and Frederik Toscani. "The shifting natural wealth of nations: The role of market orientation." *Journal of Development Economics* 138: (2019) 228–245.

Arps, John J, and Thomas G Roberts. "Economics of drilling for Cretaceous oil on east flank of Denver-Julesburg basin." *AAPG Bulletin* 42, 11: (1958) 2549–2566.

Asheim, G. B., T. Fæhn, K. Nyborg, M. Greaker, C. Hagem, B. Harstad, M. O. Hoel, D. Lund, and K. E. Rosendahl. "The case for a supply-side climate treaty." *Science* 365, 6451: (2019) 325–327.

Asker, John, Allan Collard-Wexler, Charlotte De Canniere, Jan De Loecker, and Christopher R Knittel. "Two Wrongs Can Sometimes Make a Right: The Environmental Benefits of Market Power in Oil." Technical report, National Bureau of Economic Research, 2024.

Asker, John, Allan Collard-Wexler, and Jan De Loecker. "(Mis)Allocation, Market Power, and Global Oil Extraction." *American Economic Review* 109, 4: (2019) 1568–1615.

Associated Press. "Oregon approves temporary ban on 'fracking'", 2019. https://apnews.com/general-news-103b23719d85402c9a49bb1dc310d5a7.

Attanasi, E. D., and L. J. Drew. "Lognormal field size distributions as a consequence of economic truncation." *Journal of the International Association for Mathematical Geology* 17, 4: (1985) 335–351.

Barouch, Eytan, and Gordon M. Kaufman. "Oil and Gas Discovery Modeled as Sampling Proportional to Random Size." Working paper, MIT Alfred P. Sloan School of Management, Cambridge, MA, 1976.

Barrage, Lint, and William Nordhaus. "Policies, projections, and the social cost of carbon: Results from the DICE-2023 model." *Proceedings of the National Academy of Sciences* 121, 13: (2024) e2312030,121.

BBC News. "Obama bans oil drilling 'permanently' in millions of acres of ocean.", 2016. https://www.bbc.com/news/world-us-canada-38387525. Accessed: 2025-12-18.

Beck, Chantal, Sahar Rashidbeigi, Occo Roelofsen, and Eveline Speelman. "The Future Is Now: How Oil and Gas Companies Can Decarbonize." *McKinsey & Company* https://www.mckinsey.com/industries/oil-and-gas/our-insights/the-future-is-now-how-oil-and-gas-companies-can-decarbonize.

Belfiori, Elisa, and Armon Rezai. "Implicit carbon prices: Making do with the taxes we have." *Journal of Environmental Economics and Management* 125: (2024) 102,950.

Benchekroun, Hassan, Gerard van der Meijden, and Cees Withagen. "OPEC, unconventional oil and climate change - On the importance of the order of extraction." *Journal of Environmental Economics and Management* 104: (2020) 102,384.

Beyond Oil and Gas Alliance. "BOGA announces support fund for fossil phase out.", 2022a. https://beyondoilandgasalliance.org/wp-content/uploads/2022/12/BOGA_COP27_Press_Release_updated.pdf.

———. "Who We Are.", 2022b. https://beyondoilandgasalliance.org/who-we-are/.

———. "Beyond Oil and Gas Alliance adds new members and announces funding for Global South producer countries to catalyse a just transition away from oil and gas.", 2023a. https://beyondoilandgasalliance.org/wp-content/uploads/2023/12/BOGA-Press-Release-COP28.pdf.

———. "Marshall Islands Joins Beyond Oil and Gas Alliance.", 2023b. https://beyondoilandgasalliance.org/wp-content/uploads/2023/10/BOGA_Press_Release_RMI.pdf.
———. "Who We Are.", 2023c. https://beyondoilandgasalliance.org/who-we-are/.
BGR. "BGR Energy Study 2019: Data and Developments Concerning German and Global Energy Supplies.", 2020.
Bialek, Sylwia, and Alfons J. Weichenrieder. "Should the global community welcome new oil discoveries?" *Journal of Economics* 137, 3: (2022) 255–278.
Van den Bijgaart, Inge, and Mauricio Rodriguez. "Closing wells: Fossil development and abandonment in the energy transition." *Resource and Energy Economics* 74: (2023) 101,387.
Bilal, Adrien, and Diego R Känzig. "The Macroeconomic Impact of Climate Change: Global vs. Local Temperature." Working Paper 32450, National Bureau of Economic Research, 2024.
Bornstein, Gideon, Per Krusell, and Sergio T. Rebelo. "Lags, Costs and Shocks: An Equilibrium Model of the Oil Industry." Technical report, CEPR Discussion Paper No. DP12047, 2017.
Brandt, Adam R, and Alexander E Farrell. "Scraping the bottom of the barrel: greenhouse gas emission consequences of a transition to low-quality and synthetic petroleum resources." *Climatic Change* 84, 3: (2007) 241–263.
Brandt, Adam R., et al. "OPGEE: The Oil Production Greenhouse Gas Emissions Estimator." https://eao.stanford.edu/research-project/opgee-oil-production-greenhouse-gas-emissions-estimator, 2025. Accessed: 2025-02-12.
British Petroleum. "BP Statistical Review of World Energy.", 2019.
———. "BP Energy Outlook 2023.", 2023. https://www.bp.com/en/global/corporate/energy-economics/energy-outlook/energy-outlook-downloads.html.
Calel, Raphael, and Paasha Mahdavi. "Opinion: The unintended consequences of antiflaring policies–and measures for mitigation." *Proceedings of the National Academy of Sciences* 117, 23: (2020) 12,503–12,507.
Calel, Raphael, and David A. Stainforth. "On the physics of three integrated assessment models." *Bulletin of the American Meteorological Society* 98, 6: (2017) 1199–1216.
Calel, Raphael, David A. Stainforth, and Simon Dietz. "Tall tales and fat tails: the science and economics of extreme warming." *Climatic Change* 132, 1: (2015) 127–141.
Chakravorty, Ujjayant, Michel Moreaux, and Mabel Tidball. "Ordering the Extraction of Polluting Nonrenewable Resources." *American Economic Review* 98, 3: (2008) 1128–1144.
Consejo Nacional Electoral. "CNE proclama resultados definitivos de binomio presidencial y Consulta Popular del Yasuní.", 2023. https://www.cne.gob.ec/cne-proclama-resultados-definitivos-de-binomio-presidencial-y-consulta-popular-del-yasuni/.
Coulomb, Renaud, and Fanny Henriet. "The Grey Paradox: How fossil-fuel owners can benefit from carbon taxation." *Journal of Environmental Economics and Management* 87: (2018) 206–223.
Coulomb, Renaud, Fanny Henriet, and Léo Reitzmann. ""Bad" Oil, "Worse" Oil, and Carbon Misallocation." *The Review of Economic Studies* rdaf018.
Coulomb, Renaud, Léo Jean, and Fanny Henriet. "The Climate-wise Values of Oil." Mimeo, 2025b.
Cust, James, and Torfinn Harding. "Institutions and the Location of Oil Exploration." *Journal of the European Economic Association* 18, 3: (2019) 1321–1350.
Danish Ministry of Climate, Energy and Utilities. "Denmark introduces cutoff date of 2050 for oil and gas extraction in the North Sea, cancels all future licensing rounds.", 2020. https://www.en.kefm.dk/news/news-archive/2020/dec/denmark-introduces-cutoff-date-of-2050-for-oil-and-gas-extraction-in-the-north-sea-cancels-all-future-licensing-rounds.
De Cannière, Charlotte. "Pump it? Market Power and the Energy Transition in the Global Oil Market." Mimeo, 2025.
Dietz, Simon, and Nicholas Stern. "Endogenous growth, convexity of damage and climate risk: how Nordhaus' framework supports deep cuts in carbon emissions." *Economic Journal* 125, 583: (2015) 574–620.
EIA. "Annual crude and lease condensate reserves.", 2021.
Eichner, Thomas, Gernot Kollenbach, and Martin Schopf. "Demand- versus Supply-Side Climate Policies with a Carbon Dioxide Ceiling." *The Economic Journal* 133, 652: (2023) 1371–1406.
Eichner, Thomas, and Rüdiger Pethig. "Trade in Fossil Fuel Deposits for Preservation and Strategic Action." *Journal of Public Economics* 147: (2017) 50–61.
Einhorn, Catrin, Manuela Andreoni, and Erin Schaff. "Ecuador Tried to Curb Drilling and Protect the

Amazon. The Opposite Happened." The New York Times, 2023-04-14, 2023. https://www.nytimes.com/2023/01/14/world/americas/ecuador-oil-drilling-amazon.html.

Elvidge, Christopher, Mikhail Zhizhin, Kimberly Baugh, Feng-Chi Hsu, and Tilottama Ghosh. "Methods for global survey of natural gas flaring from visible infrared imaging radiometer suite data." *Energies* 9, 1: (2016) 14.

Equinor. *Energy Perspectives 2022*. 2022.

Fæhn, Taran, Cathrine Hagem, Lars Lindholt, Ståle Mœland, and Knut Einar Rosendahl. "Climate Policies in a Fossil Fuel Producing Country: Demand versus Supply Side Policies." *The Energy Journal* 38, 1: (2017) 77–102.

Fetkovich, M. "Decline curve analysis using type curves: SPE, v. 4629." Technical report, 1980.

Financial Accountability Office of Ontario. "Cap and Trade: A Financial Review of the Decision to Cancel the Cap and Trade Program." Technical report, Financial Accountability Office of Ontario, 2018.

Fischer, Carolyn, and Stephen Salant. "Quantifying Intertemporal Emissions Leakage." In *Climate Policy and Nonrenewable Resources: The Green Paradox and Beyond*, The MIT Press, 2014.

Fischer, Carolyn, and Stephen W. Salant. "Balancing the carbon budget for oil: The distributive effects of alternative policies." *European Economic Review* 99: (2017) 191–215.

Gilbert, Daniel. "Drilling for Crude Goes Solar." *The Wall Street Journal* https://www.wsj.com/articles/SB10001424052970203405504576602891728012656.

Gillett, Nathan P., Vivek K. Arora, Damon Matthews, and Myles R. Allen. "Constraining the Ratio of Global Warming to Cumulative CO2 Emissions Using CMIP5 Simulations." *Journal of Climate* 26, 18: (2013) 6844–6858.

Gillingham, Kenneth, William Nordhaus, David Anthoff, Geoffrey Blanford, Valentina Bosetti, Peter Christensen, Haewon McJeon, and John Reilly. "Modeling uncertainty in integrated assessment of climate change: a multimodel comparison." *Journal of the Association of Environmental and Resource Economists* 5, 4: (2018) 791–826.

Golosov, Mikhail, John Hassler, Per Krusell, and Aleh Tsyvinski. "Optimal Taxes on Fossil Fuel in General Equilibrium." *Econometrica* 82, 1: (2014) 41–88.

Gordon, Deborah, Adam R. Brandt, Joule Bergerson, and Jon Koomey. "Know your oil: creating a global oil-climate index.", 2015.

Gordon, Deborah, et al. "OPEM: Oil Pipeline Emissions Model." https://oci.carnegieendowment.org/, 2025. Accessed: 2025-02-12.

Government of Belize. "Petroleum Operations (Maritime Zone Moratorium) Act.", 2017. https://faolex.fao.org/docs/pdf/blz175462.pdf.

Government of Costa Rica. "Executive Decree No. 36693: Moratorium on Oil and Gas Exploration.", 2011. https://www.pgrweb.go.cr/scij/Busqueda/Normativa/Normas/nrm_texto_completo.aspx?param1=NRTC&nValor1=1&nValor2=70858&nValor3=106736&strTipM=TC.

Government of France. "French Law No. 2017-1839: Ban on Hydrocarbon Exploration.", 2017. https://www.legifrance.gouv.fr/jorf/id/JORFTEXT000036339396.

Government of New Zealand. "Planning for the future - no new offshore oil and gas exploration permits.", 2018. https://www.beehive.govt.nz/release/planning-future-no-new-offshore-oil-and-gas-exploration-permits.

Government of Portugal. "Portugal Ends Fossil Fuel Licensing.", 2018. https://www.portugal.gov.pt.

Government of Québec. "Bill 21 - Québec Hydrocarbon Ban Law.", 2022. https://www.assnat.qc.ca/fr/travaux-parlementaires/projets-loi/projet-loi-21-42-1.html.

Government of Scotland. "Scottish Government says no to fracking.", 2017.

Government of Spain. "Climate Change and Energy Transition Law 7/2021.", 2021. https://www.boe.es/eli/es/l/2021/05/20/7.

Green, Fergus, Olivier Bois von Kursk, Greg Muttitt, and Steve Pye. "No new fossil fuel projects: The norm we need." *Science* 384, 6699: (2024) 954–957.

Greene, Suzanne, Haiying Jia, and Gabriela Rubio-Domingo. "Well-to-tank carbon emissions from crude oil maritime transportation." *Transportation Research Part D: Transport and Environment* 88: (2020) 102,587.

Guterres, A. "Secretary-General's Video Message for Press Conference to Launch the Synthesis Report of the Intergovernmental Panel on Climate Change.", 2023.

Harstad, Bård. "Buy coal! A case for supply-side environmental policy." *Journal of Political Economy* 120, 1: (2012) 77–115.

Harstad, Bård, and Kjetil Storesletten. "Conservation by Lending." Technical report, Cesifo Working Paper, 2023.

Herrnstadt, Evan, Ryan Kellogg, and Eric Lewis. "Drilling Deadlines and Oil and Gas Development." *Econometrica* 92, 1: (2024) 29–60.

Holland, Stephen P. "Extraction capacity and the optimal order of extraction." *Journal of Environmental Economics and Management* 45, 3: (2003) 569–588.

Höök, Mikael, and Kjell Aleklett. "A decline rate study of Norwegian oil production." *Energy Policy* 36, 11: (2008) 4262–4271.

Höök, Mikael, Simon Davidsson, Sheshti Johansson, and Xu Tang. "Decline and depletion rates of oil production: a comprehensive investigation." *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 372, 2006: (2014) 20120,448.

Hotelling, Harold. "The Economics of Exhaustible Resources." *Journal of Political Economy* 39, 2: (1931) 137–175.

Hyndman, Rob J., and Yanan Fan. "Sample Quantiles in Statistical Packages." *The American Statistician* 50, 4: (1996) 361–365.

IEA. "The Future of Petrochemicals." Technical report, International Energy Agency, Paris, 2018a.

———. "World Energy Outlook 2019." Technical report, International Energy Agency, Paris, 2019.

———. "World Energy Outlook 2021." Technical report, International Energy Agency, Paris, 2021.

———. *The Oil and Gas Industry in Net Zero Transitions*. OECD, 2023. https://www.oecd.org/en/publications/the-oil-and-gas-industry-in-net-zero-transitions_fd522f59-en.html.

———. "Global Energy and Climate Model Documentation 2024." https://iea.blob.core.windows.net/assets/89a1aa9a-e1bd-4803-b37b-59d6e7fba1e9/GlobalEnergyandClimateModelDocumentation2024.pdf.

IEA. "Global Methane Tracker 2025.", 2025. https://www.iea.org/reports/global-methane-tracker-2025. Licence: CC BY 4.0.

IEA. "The Implications of Oil and Gas Field Decline Rates.", 2025. https://www.iea.org/reports/the-implications-of-oil-and-gas-field-decline-rates.

IEA. "Where does the world get its oil?" https://www.iea.org/world/oil, 2025. Accessed: 2025-02-17.

IEA, Christophe McGlade, Glenn Sondak, Mei Han. "Whatever happened to enhanced oil recovery?" International Energy Agency, Paris, November 28, 2018b.

IPCC. "Climate Change 2014: Synthesis Report. Contribution of Working Groups I, II and III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change [Core Writing Team, R.K. Pachauri and L.A. Meyer (eds.)].", 2014.

———. *Global Warming of 1.5°C: IPCC Special Report on Impacts of Global Warming of 1.5°C above Pre-industrial Levels in Context of Strengthening Response to Climate Change, Sustainable Development, and Efforts to Eradicate Poverty*. Cambridge University Press, 2022, 1 edition. https://www.cambridge.org/core/product/identifier/9781009157940/type/book.

Irish Government. "Climate Action and Low Carbon Development (Amendment) Act 2021.", 2021. https://www.oireachtas.ie/en/bills/bill/2021/39/.

Jackson, Peter M., and Leta K. Smith. "Exploring the undulating plateau: the future of global oil supply." *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 372, 2006: (2014) 20120,491.

Kaufman, Gordon. "Size and distribution of oil and gas fields." *AAPG Bulletin* 48, 4: (1964) 534–534.

Kellogg, Ryan. "The Effect of Uncertainty on Investment: Evidence from Texas Oil Drilling." *American Economic Review* 104, 6: (2014) 1698–1734.

———. "The End of Oil." Working Paper 33207, National Bureau of Economic Research, 2024.

Kühne, Kjell, Nils Bartsch, Ryan Driskell Tate, Julia Higson, and André Habet. ""Carbon Bombs" - Mapping key fossil fuel projects." *Energy Policy* 166: (2022) 112,950.

Latham, Andrew, and Simran Bandal. "No country for old fields: why high-impact oil and gas exploration is still needed." Technical report, Wood Mackenzie, 2024. https://www.woodmac.com/horizons/high-impact-oil-gas-exploration/?__FormGuid=9cc01a3f-5784-4719-a268-9e82b1c21a12&__FormLanguage=en&__FormSubmissionId=7e9e5ad6-7e26-44f8-a5d2-8e527320d88b.

Lemoine, Derek. "The Climate Risk Premium: How Uncertainty Affects the Social Cost of Carbon." *Journal of the Association of Environmental and Resource Economists* 8, 1: (2021) 27–57.

Lenton, Timothy M., and Juan-Carlos Ciscar. "Integrating tipping points into climate impact assessments." *Climatic Change* 117, 3: (2013) 585–597.

Maryland General Assembly. "Oil and Natural Gas - Hydraulic Fracturing - Prohibition.", 2017. https://mgaleg.maryland.gov/mgawebsite/legislation/details/hb1325?ys=2017rs.

Masnadi, Mohammad S., Hassan M. El-Houjeiri, Dominik Schunack, Yunpo Li, Jacob G. Englander, Alhassan Badahdah, Jean-Christophe Monfort, James E. Anderson, Timothy J. Wallington, Joule A. Bergerson, Deborah Gordon, Jonathan Koomey, Steven Przesmitzki, Inês L. Azevedo, Xiaotao T. Bi, James E. Duffy, Garvin A. Heath, Gregory A. Keoleian, Christophe McGlade, D. Nathan Meehan, Sonia Yeh, Fengqi You, Michael Wang, and Adam R. Brandt. "Global carbon intensity of crude oil production." *Science* 361, 6405: (2018) 851–853.

Masson-Delmotte, Valérie, Panmao Zhai, Anna Pirani, Sarah L Connors, Clotilde Péan, Sophie Berger, Nada Caud, Y Chen, L Goldfarb, MI Gomis, et al. "Climate change 2021: the physical science basis." Geneva, Switzerland, 2021, volume 2, 2391.

Matthews, H. Damon, Nathan P. Gillett, Peter A. Stott, and Kirsten Zickfeld. "The proportionality of global warming to cumulative carbon emissions." *Nature* 459, 7248: (2009) 829–832.

McCollum, David, Nico Bauer, Katherine Calvin, Alban Kitous, and Keywan Riahi. "Fossil resource and energy security dynamics in conventional and carbon-constrained worlds." *Climatic change* 123, 3-4: (2014) 413–426.

McCrossan, R. G. "An analysis of size frequency distribution of oil and gas reserves of Western Canada." *Canadian Journal of Earth Sciences* 6, 2: (1969) 201–211.

McGlade, Christophe. "A review of the uncertainties in estimates of global oil resources." *Energy* 47, 1: (2012) 262–270. Asia-Pacific Forum on Renewable Energy 2011.

McGlade, Christophe, and Paul Ekins. "Un-burnable oil: An examination of oil resource utilisation in a decarbonised energy system." *Energy Policy* 64: (2014) 102—112.

———. "The geographical distribution of fossil fuels unused when limiting global warming to 2°C." *Nature* 517, 7533: (2015) 187–190.

McGlade, Christophe E. *Uncertainties in the outlook for oil and gas*. Ph.D. thesis, UCL (University College London), 2014.

Meinshausen, Malte, Nicolai Meinshausen, William Hare, Sarah C.B. Raper, Katja Frieler, Reto Knutti, David J. Frame, and Myles R. Allen. "Greenhouse-gas emission targets for limiting global warming to 2°C." *Nature* 458, 7242: (2009) 1158–1162.

Michielsen, Thomas O. "Brown backstops versus the green paradox." *Journal of Environmental Economics and Management* 68, 1: (2014) 87–110.

Ministry of Mineral Resources and Justice of Greenland. "Stop oil exploration in Greenland.", 2021. https://naalakkersuisut.gl/Nyheder/2021/07/1507_oliestop?sc_lang=da.

Moore, Malcom. "Saudi Aramco chief warns of global oil shortage if industry fails to invest." *Financial Times* https://www.ft.com/content/698b826a-8d14-4eef-bca6-f84634965080. Accessed: 2025-10-20.

Mui, Simon, Luke Tonachel, and E. Shope. "GHG emission factors for high carbon intensity crude oils." *Natural Resources Defense Council* 2.

Murray, Brian, and Nicholas Rivers. "British Columbia's revenue-neutral carbon tax: A review of the latest "grand experiment" in environmental policy." *Energy Policy* 86: (2015) 674–683.

New York Assembly. "Assembly Passes Legislation to Prohibit Use of Carbon Dioxide to Drill and Extract Natural Gas and Oil.", 2024. https://nyassembly.gov/Press/?sec=story&story=109485.

Newell, Peter, and Andrew Simms. "Towards a fossil fuel non-proliferation treaty." *Climate Policy* 20, 8: (2020) 1043–1054.

Nordhaus, William D. "The 'DICE' model: Background and structure of a dynamic integrated climate-economy model of the economics of global warming." Technical report, Cowles Foundation for Research in Economics, Yale University, 1992.

———. *Managing the global commons: the economics of climate change*, volume 31. MIT Press Cambridge, MA, 1994.

———. "A review of the Stern review on the economics of climate change." *Journal of Economic Literature* 45, 3: (2007) 686–702.

———. "Revisiting the social cost of carbon." *Proceedings of the National Academy of Sciences* 114, 7: (2017) 1518–1523.

Nystad, Arild N. "Petroleum taxes and optimal resource recovery." *Energy Policy* 13, 4: (1985) 381–401.

———. "Rate sensitivity and the optimal choice of production capacity of petroleum reservoirs." *Energy Economics* 9, 1: (1987) 37–45.

Office of the Governor of California. "Governor Newsom Takes Action to Phase Out Oil Extraction in California.", 2021. https://www.gov.ca.gov/2021/04/23/governor-newsom-takes-action-to-phase-out-oil-extraction-in-california/.

Olsen, R., and S. Patel. "Government Policies and NGO Advocacy: Pushing for a Ban on New Fossil Fuel Exploration." *Global Environmental Politics* 20, 2: (2020) 120–139.

OPEC. "OPEC Annual Statistical Bulletin 2019.", 2019.

Parliament of Sweden. "Ban on extraction of coal, oil and natural gas and stricter rules for extraction in alum shale.", 2022. https://www.riksdagen.se/sv/dokument-och-lagar/dokument/betankande/forbud-mot-utvinning-av-kol-olja-och-naturgas-och_h901nu23/html/.

Perez-Sebastian, Fidel, Ohad Raveh, and Frederick van der Ploeg. "Oil discoveries and protectionism: Role of news effects." *Journal of Environmental Economics and Management* 107: (2021) 102,425.

Pindyck, Robert S. "Climate change policy: What do the models tell us?" *Journal of Economic Literature* 51, 3: (2013) 860–72.

———. "The social cost of carbon revisited." *Journal of Environmental Economics and Management* 94: (2019) 140–160.

Prest, Brian C. "Supply-Side Reforms to Oil and Gas Production on Federal Lands: Modeling the Implications for CO_2 Emissions, Federal Revenues, and Leakage." *Journal of the Association of Environmental and Resource Economists* 9, 4: (2022) 681–720.

Prime Minister of Canada. "United States-Canada Joint Arctic Leaders' Statement.", 2016. https://www.pm.gc.ca/en/news/statements/2016/12/20/united-states-canada-joint-arctic-leaders-statement. Accessed: 2025-12-18.

Ratnaparkhi, Makarand, and Uttara Naik-Nimbalkar. "The Length-Biased Lognormal Distribution and Its Application in the Analysis of Data from Oil Field Exploration Studies." *Journal of Modern Applied Statistical Methods* 11, 1.

Rennert, Kevin, Frank Errickson, Brian C. Prest, Lisa Rennels, Richard G. Newell, William Pizer, Cora Kingdon, Jordan Wingenroth, Roger Cooke, Bryan Parthum, et al. "Comprehensive evidence implies a higher social cost of CO2." *Nature* 610, 7933: (2022) 687–692.

Reuters. "Upstream Electrification Can Cut Oil and Gas Production Emissions by More Than 80%, Report Says." *Reuters* https://www.reuters.com/business/environment/upstream-electrification-can-cut-oil-gas-production-emissions-by-more-than-80-2024-09-19/.

Revesz, Richard L., Peter H. Howard, Kenneth Arrow, Lawrence H. Goulder, Robert E. Kopp, Michael A. Livermore, Michael Oppenheimer, and Thomas Sterner. "Global warming: Improve economic models of climate change." *Nature News* 508, 7495: (2014) 173.

Robelius, F. "Giant oil fields—the highway to oil: giant oil fields and their importance for future oil production." Doctoral thesis, Uppsala University, 2007.

Rogelj, Joeri, Piers M. Forster, Elmar Kriegler, Christopher J. Smith, and Roland Séférian. "Estimating and tracking the remaining carbon budget for stringent climate targets." *Nature* 571: (2019) 335–34.

Rudolph, Kurt. "Benchmarking Exploration Predictions and Performance Using 20+ Years of Drilling Results: One Company's Experience." *AAPG Bulletin* IN PRESS.

Rystad Energy. "UCube Global Asset-Level Oil and Gas Database." https://www.rystadenergy.com/services/energytype/oil-gas/ucube/, 2022.

Shell. "Energy Transition Strategy 2021.", 2021. https://www.shell.com/promos/energy-and-innovation/energy-transition-strategy/_jcr_content.stream/1617270377735/e260cd1e47a7f226e5e6b39d35187ee06529904a/shell-energy-transition-strategy-2021.pdf.

———. *The Energy Security Scenarios.* 2023.

Solow, Robert M. "The Economics of Resources or the Resources of Economics." *The American Economic Review* 64, 2: (1974) 1–14.

Stefanski, Radek, Lassi Ahlvik, Jørgen Juel Andersen, Torfinn Harding, and Alex Trew. "Extracting Wedges: Misallocation and Taxation in the Oil Industry." Technical Report 11915, CESifo Working Paper, 2025.

Stern, Nicholas. "The Economics of Climate Change." *American Economic Review* 98, 2: (2008) 1–37.

Stocker, D. Qin G.-K. Plattner L. Alexander S. Allen N. Bindoff F.-M. Bréon J. Church U. Cubasch S. Emori P. Forster P. Friedlingstein N. Gillett J. Gregory D. Hartmann E. Jansen B. Kirtman R. Knutti K. K. Kumar P. Lemke J. Marotzke V. Masson-Delmotte G. Meehl I. Mokhov S. Piao V. Ramaswamy D. Randall M. Rhein M. Rojas C. Sabine D. Shindell L. Talley D. Vaughan, T., and S.-P. Xie. "Technical summary. In T. Stocker, D. Qin, G.-K. Plattner, M. Tignor, S. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex, and P. Midgley (Eds.), Climate Change 2013: the physical science basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change.", 2014.

Stoerk, Thomas, James Rising, Drew Shindell, and Simon Dietz. "Global methane action pays for itself at least six times over.", 2025. Mimeo.

Tol, Richard S.J., Kenneth J. Arrow, Maureen L. Cropper, Christian Gollier, Ben Groom, Geoffrey Heal, Richard G. Newell, William D. Nordhaus, Robert S. Pindyck, William A. Pizer, Paul R. Portney, Thomas Nils Samuel Sterner, and Martin L. Weitzman. "How Should Benefits and Costs Be Discounted in an Intergenerational Context?" Working papers, Department of Economics, University of Sussex, 2013.

TotalEnergies. "Sustainability and Climate Progress Report 2023.", 2023a. https://corporate.totalenergies.us/news/totalenergies-publishes-its-sustainability-climate-2025-progress-report-and-further.

———. "TotalEnergies Energy Outlook 2023.", 2023b. https://totalenergies.com/fr/actualites/communiques-de-presse/totalenergies-energy-outlook-2023.

TotalEnergies. "2024 Strategy & Outlook Presentation (Transcript).", 2024. https://totalenergies.com/sites/g/files/nytnzq121/files/documents/totalenergies_TotalEnergies-Strategy-and-Outlook-transcript_2024_pdf.pdf.

Trout, Kelly, Greg Muttitt, Dimitri Lafleur, Thijs Van de Graaf, Roman Mendelevitch, Lan Mei, and Malte Meinshausen. "Existing fossil fuel extraction would warm the world beyond 1.5 C." *Environmental Research Letters* 17, 6: (2022) 064,010.

UK Department for the Economy. "Petroleum licensing.", 2024. https://www.economy-ni.gov.uk/articles/petroleum-licensing.

Uria Martinez, Rocio, Paul Newsome Leiby, Gbadebo A Oladosu, David Bowman, and Megan Johnson. "Using Meta-Analysis to Estimate World Oil Demand Elasticity." Technical report, Oak Ridge National Lab.(ORNL), Oak Ridge, TN (United States), 2018.

U.S. Energy Information Administration. "Oil and petroleum products explained: Refining crude oil." U.S. Energy Information Administration, 2020a.

———. "What countries are the top producers and consumers of oil?" U.S. Energy Information Administration, 2020b.

USGS. "Heavy Oil and Natural Bitumen." USGS, Online, August 2003; accessed on June 1st, 2023, 2003.

Van Asselt, Harro, and Peter Newell. "Pathways to an international agreement to leave fossil fuels in the ground." *Global Environmental Politics* 22, 4: (2022) 28–47.

Van den Bergh, Jeroen CJM, and Wouter JW Botzen. "A lower bound to the social cost of CO2 emissions." *Nature climate change* 4, 4: (2014) 253–258.

Van der Meijden, Gerard, Frederick van der Ploeg, and Cees Withagen. "International capital markets, oil producers and the Green Paradox." *European Economic Review* 76: (2015) 275–297.

Van der Ploeg, Frederick. "Aggressive oil extraction and precautionary saving: Coping with volatility." *Journal of Public Economics* 94, 5: (2010) 421–433.

Van der Ploeg, Frederick, and Armon Rezai. "The risk of policy tipping and stranded carbon assets." *Journal of Environmental Economics and Management* 100: (2020) 102,258.

Van der Ploeg, Frederick, and Cees Withagen. "Too much coal, too little oil." *Journal of Public Economics* 96, 1: (2012) 62–77.

Venables, Anthony. "Depletion and Development: Natural Resource Supply with Endogenous Field Opening." *Journal of the Association of Environmental and Resource Economists* 1, 3: (2014) 313 – 336.

Vermont General Assembly. "Hydraulic fracturing; prohibition.", 2012. https://legislature.vermont.gov/statutes/section/29/014/00571.

Victorian Legislation. "Constitution Amendment (Fracking Ban) Act 2021.", 2021. https://www.legislation.vic.gov.au/as-made/acts/constitution-amendment-fracking-ban-act-2021.

Weitzman, Martin L. "On modeling and interpreting the economics of catastrophic climate change." *Review*

of Economics and Statistics 91, 1: (2009) 1–19.

———. "Fat-Tailed Uncertainty in the Economics of Catastrophic Climate Change." *Review of Environmental Economics and Policy* 5, 2: (2011) 275–292.

———. "GHG targets as insurance against catastrophic climate damages." *Journal of Public Economic Theory* 14, 2: (2012) 221–244.

Welsby, Dan, James Price, Steve Pye, and Paul Ekins. "Unextractable fossil fuels in a 1.5°C world." *Nature* 597, 7875: (2021) 230–234.

Welsh Government. "Drilling down: the Welsh Government proposes policy to ban petroleum extraction.", 2018. https://research.senedd.wales/research-articles/drilling-down-the-welsh-government-proposes-policy-to-ban-petroleum-extraction/.

Wood Mackenzie. "Oil and gas exploration spending to recover from historic lows, average $22B per annum through 2027, Wood Mackenzie.", 2023.

Worland, Justin. "BP's CEO Bernard Looney Believes He Can Help Save the Planet—and Still Make Money Doing It." *Time Magazine* https://time.com/6125315/bernard-looney-bp-ceo-interview/.

World Bank. "Carbon Pricing Dashboard." World Bank, 2020.

———. "Carbon Pricing Dashboard." World Bank, 2025.

Zhizhin, Mikhail, Alexey Matveev, Tilottama Ghosh, Feng-Chi Hsu, Martyn Howells, and Christopher Elvidge. "Measuring gas flaring in russia with multispectral VIIRS nightfire." *Remote Sensing* 13, 16: (2021) 3078.

Online Appendix – *Not for Publication*

A Field-level Data **55**
A.1 Rystad Upstream Database 55
A.2 Flaring 55

B Estimating Carbon Intensities **58**
B.1 Upstream 58
B.2 Midstream 65
B.3 Downstream 70

C Private Extraction Costs and Extraction Capacities **71**
C.1 Data on Private Extraction Costs 71
C.2 Extraction Capacities 71

D Modeling of Yet-to-find Resources **74**
D.1 Data Scope 74
D.2 Field-size Distribution and Discovery Bias 74
D.3 Joint Estimation of Field Size Distribution and Total Count . 77
D.4 Assigning Oil Type, Private Costs and Carbon Intensity to YTF Deposits 79
D.5 Sensitivity Analysis in YTF Estimation to the Pool of Observed Discoveries 79
D.6 Alternative YTF Constructions for Robustness Checks 81

E Other Modeling and Parameter Choices **82**
E.1 The Selection of Deposits 82
E.2 Reserves 82
E.3 Resource Availability 83
E.4 The Social Cost of Carbon 83
E.5 Demand Function 85
E.6 Social Welfare Maximization 85

F The Global Carbon Policy Context and The Oil Industry **85**
F.1 Carbon-pricing Initiatives in the Oil Industry 85
F.2 National-Level Bans on New Fossil Fuel Exploration/Development 86
F.3 Subnational and Regional Bans/Moratoria 87

G Additional Figures and Tables **89**

Online Appendix

A Field-level Data

A.1 Rystad Upstream Database

Our empirical investigation employs one of the most comprehensive datasets on oil and gas fields, the Rystad Upstream Database (UCube). This database covers World oil production since 1900, with over 90,000 oil and gas assets. It brings together precise field-level data on oil and gas production,[1] exploitable reserves, capital and operational expenditures from exploration to field decommission as well as production taxes, current governance (ownership and operating companies), field development dates (discovery, license, start-up, and production end), oil types, and reservoir characteristics (water depth, basin, and location), among others. For some parts of our analysis, we use geocoded boundaries of each field obtained from Rystad GIS team, as well as data on the annual number of drilled and active wells per asset from the Rystad WellCube. The Rystad data sources include governments and companies' operation reports.

A.2 Flaring

Data Since only a minority of countries and companies collect and publish data on flared gas, nearly 95% of the fields in Rystad have missing flaring data. We complement the data using the geocoded flaring volumes calculated by the Earth Observation Group (EOG), within the Payne Institute for Public Policy at the Colorado School of Mines, in collaboration with the National Oceanic and Atmospheric Administration (NOAA). Satellite data is collected from the Visible Infrared Imaging Radiometer Suite (VIIRS) embarked on satellites launched in 2012 and 2017. Heat emitted by gas flares are observed at night and corresponding flared volumes are computed using the VIIRS Nightfire algorithm (VNF) developed by the EOG (Elvidge et al., 2016; Zhizhin et al., 2021). We use the 95,617 positive annual flaring volumes recorded by the EOG between 2012 and 2020 and match it geographically to Rystad oil and gas fields.

Computing field- flaring-to-oil ratios (FORs) Among the 28,939 geocoded boundaries of assets obtained from the Rystad GIS team, 10,154 were generic shapes (single points). The geocoded data was merged with data on active wells per year from the WellCube as well as production data from the UCube. A first correction was applied to the generic shapes:

[1] All productions are expressed in energy-equivalent barrels, using Brent Crude as the benchmark.

for 498 assets with generic shapes that were contained in a non-generic one, the non-generic shape was assigned to the corresponding asset. These assets generally correspond to the re-development or secondary phases of an already existing asset. In the geocoded data, 17,320 assets had positive production and/or active producing wells in the 2012–2020 period.[2] 4,871 of these assets had a generic shape, with a highly heterogeneous geographical distribution. 21 countries with more than 90% of shapes being generic were excluded from the matching, among which China and Ukraine. An additional restriction was made: all onshore US and Canadian assets were also excluded, leaving only offshore ones. The remaining assets with positive production or number of active wells amounted to 13,636 (of which 1,192 were generic). An additional correction was performed to better account for generic shapes: an average number of wells per km^2 was computed for non-generic assets, divided in 3 categories (offshore, onshore shale and other onshore), allowing us to estimate the size of assets missing a shape. A circle of the corresponding size was attributed to these 1,192 assets.

95,617 positive flaring volumes were recorded by the EOG in the 2012–2020 period, among which 94,964 were identified as relating to oil and gas production (thus excluding refineries and oil and gas processing plants). Applying the countries exclusion explained above, 51,224 flaring observations were left to be matched at the asset-level, corresponding to 90.6% of total flared volumes worldwide during the period.

The flare-field matching is implemented through the following sequential steps: First, we match a flare to a field if the flare is located within the field's geographic boundaries or less than 800 meters from it.[3] This allows us to match 31,658 flares out of 51,224, representing about 65.6% of total volumes considered. Second, if a flare does not lie within a field's boundaries, we associate the flare with the closest cluster of fields. We compute for that flare its distance to all fields located less than 40 km away from it. We then rank these distances by increasing order. Denote I_x the number of fields located less than 40 km from the flare x, $(d_{i,x})$ the vector of distances between fields $i \in [|1, I_x|]$ and the flare x, such as $d_{i,x} \leq d_{i+1,x}$ for all $i \in [|1, I_x - 1|]$. We define the closest cluster of fields to a flare x as the subgroup of fields $\{1, ..., N\}$ with N the smallest integer such as either $(d_{N+1,x} - d_{N,x} > 2km$ and $d_{N+1,x} > 1.5d_{N,x})$ or $d_{N+1,X} - d_{N,X} > 10km$. When a cluster of fields is associated to a flare, the flared volume of the flare is split across the fields in the cluster based on their production of hydrocarbons. This step allows us to match 10,892 flares, that represent 29.2% of the total volume flared. Third, we apply a minor correction of previous matches to account for possible misreporting in Rystad data and for the fact that VIIRS may exhibit inaccuracies in flare detection. We correct matches obtained from step 2 as follows. If the

[2]Producer, appraisal and wildcat wells in the Rystad data

[3]The resolution of a pixel in the data collected by the VIIRS is 800m × 800m at nadir.

closest field matched to a flare is at least 10 km away from the flare, we verify whether the flare is located within the boundaries of a field producing in the 4-year window around the year of observation of the flare: if such a field exists, we associate the flare to this latter field instead. This correction concerns 310 flares, that represent together 0.28% of the flared volume over that period. Fourth, for flares left unmatched through steps 1 through 3, we associate all fields located less than 40 km away and split flared volumes based on their production. This matching concerns 1,210 flares representing 1.5% of total flared volumes. Finally, 69 unmatched flares representing 0.03% of the total volume flared are matched to assets within 40 km that produced in the 4-year window around the observation year of the flare. Overall, 96.6% of the total volume of flaring from the upstream oil and gas industry observed over the 2012–2020 period in countries with fields' precise geographic boundaries is matched to Rystad fields that produced some hydrocarbons over the 2012–2020 period.

Total flaring volume in countries with imprecise asset boundaries represent 9.4% of the total flared volume of the upstream Oil & Gas industry over the 2012–2020 period. We attribute a regional average corresponding to the "Area" level to each field located in these countries. These Areas are geographic units defined in the Rystad dataset: 1,259 Areas in the World produced oil between 2012 and 2020, with an average of 13.9 fields producing oil in the period per Area. Areas flaring to oil ratios were computed by matching worldwide flaring observations to the nearest asset having produced between 2011 and 2021 in Rystad data (with a 100 km limit) and dividing the total flared volumes by the area's total oil production. Flaring was observed in 937 areas, and 322 had no flaring observed, which can be explained by complete integration of assets in the gas market or flared volumes being low or intermittent enough not to be captured by the VIIRS instruments.

Overall, among assets having produced some oil in the period, 12,982 assets were assigned a FOR through field-level matching (with 6,610 of them being equal to 0), and 4,520 were assigned an area average (226 equal to 0). Regarding the 20,458 assets with no recorded oil production in the 2012–2020 period, their area average was attributed, when available (86.3% of them). The rest was attributed the most precise average available: province $\times$ oil type (4.9%), province (0.1%), country $\times$ oil type (3.1%), and finally oil type (5.3%). Finally, 48 assets with no oil production recorded in the period but having positive oil reserves were attributed a FOR equal to 0 on the account of being mainly gas fields with no recorded flaring while producing gas in the period.

Estimated field-level FORs translate into an average of 156.5 scf/bbl when weighting by 2021 oil production, and a weighted median of 31.4 scf/bbl. This translates into a total of 143.97 bcm flared in 2021, remarkably close to the 144 bcm announced by the World Bank in its 2022 Global Gas Flaring Report. Large heterogeneity exists across countries: Norway

8.6 scf/bbl, Saudi Arabia 14.5 scf/bbl, US 93.1 scf/bbl, Iran 280.8 scf/bbl, Russia 310.2 scf/bbl, Nigeria 327.3 scf/bbl, Venezuela 347.1 scf/bbl (production-weighted averages, 2021 productions).

B Estimating Carbon Intensities

B.1 Upstream

The OPGEE. To estimate upstream emissions, we rely on the *Oil Production Greenhouse Gas Emissions Estimator* (OPGEE3) of the Oil-Climate Index.[4] The OPGEE measures GHG emissions in grams of CO_2 equivalent (gCO_2eq) in a field and then divides it per megajoule (MJ) or energy extracted or oil barrel (bbl).

Input data for OPGEE. Masnadi et al. (2018) provide a sample of 958 large fields, with data formatted to be used as inputs in OPGEE (values of 2015). We update these data using Rystad data for the following variables: 2015 production, FOR, GOR, number of producing and injection wells, field age, fraction of remaining natural gas injected. The oil type is defined as the main oil type in terms of initial reserves. NGL and condensate are merged as a single oil type.

Running OPGEE. We run the OPGEE model for both the EBA and CPD configuration. In most oil fields, associated gas is extracted together with oil and is sometimes sold on the market. It is therefore necessary to decide how to account (or not) for the associated gas production. The OPGEE model was run with the two following approaches.

- *Energy-based allocation* (EBA): Oil carbon intensity is $CI_{EBA} = \frac{E_{field}}{Prod_{oil}+Prod_{gas}}$. Total carbon emissions of the field, E_{field}, are calculated and then assigned to oil and gas proportionally to their respective productions in megajoules, $Prod_{oil}$ and $Prod_{gas}$.
- *Co-product displacement* (CPD): Oil carbon intensity is $CI_{CPD} = \frac{E_{field}-Prod_{gas}.\overline{CI}_{gas}}{Prod_{oil}}$, with $\overline{CI}_{gas}$ the carbon intensity of the displaced system producing gas, and E_{field} the total pollution of the field, i.e. the pollution of the joint production of oil and gas, $Prod_{oil}$ and $Prod_{gas}$. With the CPD approach, sold gas gives rise to carbon-emission credits that are equal to the amount of gas sold multiplied by an estimate of the average carbon intensity of gas production in the World, $\overline{CI}_{gas}$, calculated from the CA-GREET model.[5]

[4] To access the model and its documentation, see http://oci.carnegieendowment.org/#models. For a detailed description of the model, see Gordon et al. (2015).

[5] The model is available at https://greet.es.anl.gov.

Matching OPGEE and Rystad fields. We match the 958 publicly-available OPGEE fields with those from the Rystad dataset using field name and location. 200 fields out of 958, mostly Californian fields that represent only a marginal share of US production (jointly producing about 0.16% of World production in 2015) were unmatched. The 758 OPGEE fields were matched with 1,196 fields from the Rystad dataset. The difference in these figures is due to some fields in one dataset (either Rystad or OPGEE dataset) being represented by more than one field in the other dataset. Despite representing only 8% of producing fields, these 1,196 fields account for over 53% of 2015 World oil production.

This subsample includes most oil-producing countries, although fields located in the UK and the US are over-represented. The OPGEE sample is representative of the rest of the Rystad oil assets in the dimensions that matter for the estimation of upstream and mid-stream carbon intensity. There are however some differences. The fields operated by Major companies are over-represented in the OPGEE sample, as are fields extracting Bitumen and Synthetic crudes, those employing steam injection, and Offshore fields. The GOR tends to be larger for fields in the OPGEE sample, but the difference is not statistically significant. The OPGEE sample consists of relatively larger fields: we have verified that adding field size to the set of explanatory variables in our regression model below has no significant impact on our estimation. This indicates that the over-representation of large fields in the OPGEE sample is not expected to bias our estimation of CIs.

Estimation model. After attributing the OPGEE-calculated carbon intensities to the corresponding Rystad fields, we specify a regression model to explain the across-field variation in carbon intensity using the variables from the Rystad dataset. The existing literature (Brandt and Farrell, 2007; Mui et al., 2010; Masnadi et al., 2018) has highlighted a number of deposit characteristics that help explain upstream GHG emissions: The American Petroleum Institute (API) gravity or the Oil type (light, regular, extra heavy), the gas-to-oil ratio (GOR), the flaring-to-oil ratio (FOR), and the steam-to-oil ratio (SOR). We briefly discuss below how these deposit characteristics relate to GHG emissions.

- The American Petroleum Institute (API) gravity is a measure of oil density. Together with oil viscosity (a measure of the fluid's resistance to flow), this is used to characterize the main oil types. The U.S. Geological Survey uses the following classification: *light oil* has an API gravity over 22° and viscosity under 100 centipoise (cP); *Heavy oil* is an asphaltic, denser oil type with an API gravity under 22° and a viscosity of at least 100 cP; *Extra heavy oil* is a type of heavy oil with an API gravity below 10°. Heavier oil requires more energy to be brought up to the surface by traditional well-based extraction methods (USGS, 2003). Their extraction thus generates more carbon emissions, all else equal.

- Gas-to-oil ratio (GOR): when oil is extracted, natural gas is also brought up to the surface. Deposits differ in terms of the quantity of gas that comes with each oil barrel. The extraction of high-GOR deposits can generate large emissions depending on how this gas is handled. Gas can be either sold on the gas market, reinjected underground, flared, or directly vented into the atmosphere.

- Flaring-to-oil ratio (FOR): gas flaring, i.e. the on-site combustion of gas (or part of it), helps control the pressure in the well and the plant equipment. This technology is largely tied to the oil reservoir's specificity (e.g. GOR) and the distance to the gas market (field localization). Flaring is a major contributor to upstream carbon intensity.[6]

- Steam-to-oil ratio (SOR): steam injection concerns mostly heavy oils and bitumen. Injecting steam heats the oil and thereby reduces its viscosity, enabling oil to flow toward the extraction wells. Steam injection requires vast amounts of energy, and therefore produces substantial GHG emissions.

Only fields with a GOR $<$ 40,000 scf/bbl, a WOR $\leq$ 40 bbl water/bbl oil, and a production of at least 10 bbl/d were retained for the estimation. Finally, 1144 fields from Rystad are used for the model estimation (717 OPGEE fields). We refer to these fields as the "OPGEE sample".

Models with all potential explanatory variables, at different orders (up to cubic) with multiple interactions were tested. Variables with little explanatory power were then dropped one by one to avoid data over-fitting. The EBA approach was selected for the baseline carbon-intensity estimation. The selected model[7] includes the main explanatory variables found in the literature: oil-type dummies, GOR and FOR, and SOR, and a dummy for Offshore fields. The selected model for the EBA CI is:[8]

$$\begin{aligned} CI_f^{EBA} \times (ED) \quad &= \textstyle\sum_0^8 \beta_i OilType_{i,f} + \beta_9 GOR_f + \beta_{10} FOR_f \\ &+ \beta_{11} SOR_f + \beta_{12} Offshore_f + \epsilon_f, \end{aligned} \tag{7}$$

[6] It is still however preferable to venting—the direct release of methane into the atmosphere without burning. Burning one ton of methane results in the generation of approximately 2.75 tons of carbon dioxide, and a ton of methane has a 30-times larger global-warming potential than that of a ton of carbon dioxide (IPCC, 2014). Unfortunately, very little data is available on the amount of vented gas in the World, as this is not systematically and truthfully reported by corporations or governments, and is difficult to detect at a very granular level using remote sensors (Calel and Mahdavi, 2020).

[7] To account for some OPGEE fields being matched to more fields in the main database (due to differences in field definitions between the OPGEE and Rystad datasets), weights were added in the regression so that all OPGEE fields have equal weight.

[8] The selected model for the CPD CIs is: $CI_f^{CPD} = \sum_0^8 \beta_i OilType_{i,f} + \beta_9 GOR_f + \beta_{11} SOR + \epsilon_f$.

where f denotes a field, CI_f^{OPGEE} (EBA, AR6, 110y) in kgCO_2e/bbl, ED the energy denominator OPGEE ($ED = 1 + GOR - FOR$), $OilType_i$ oil-type dummies (e.g., Extra Heavy, Light, Regular, Bitumen), GOR the ratio of the gas quantity (in thousands standard cubic feet, kscf) to the oil quantity (in barrels, bbl) in the reservoir, FOR the ratio between flared gas and extracted oil, also in kscf/bbl, SOR the ratio between injected steam and extracted oil, in bbl of water/bbl of oil, *Offshore* a dummy for the asset being located offshore.

Regression results. The regression result for CI calculated with the EBA approach appear in Table B1. The estimation model explains most of the CI variation (an R-squared of 98%). All coefficients are significant at the 1% level. The final field-level upstream dataset of carbon intensities contains the CIs calculated using OPGEE model for fields common to the OPGEE and Rystad datasets, and the CI predicted using the regression above for the remaining Rystad fields (using estimated βs coefficients and Rystad data)

Robustness checks. We consider the sensitivity of the estimated coefficients of specification (7) to specific observations, and re-estimate it multiple times removing the fields in our main estimation sample one by one. Figure B1 displays the estimated coefficients before each explanatory variable. Overall, these robustness checks indicate that estimated parameters of equation (7) are robust to changes in the estimation sample of deposits.

We then look at the robustness of our estimates to the exclusion of the most influential observations. For a given variable, we define the n most-influential observations as those with the n largest Cook's distances (CD), with $CD_i = \sum_{j=1}^{N} \left(\widehat{y}_{\,j} - \widehat{y}_{\,j(i)}\right)^2 / ps^2$, where $\widehat{y}_{\,j}$ and $\widehat{y}_{\,j(i)}$ are the fitted response values from using the full sample and after excluding field i, respectively; s^2 is the mean squared error of the regression model, p the number of coefficients, and N the number of observations. We re-estimate equation (7) after dropping the 5, or 10 most-influential observations by this definition, and record the new estimates in Table B2. Overall, the coefficients of the explanatory variables are stable.

Finally, we compare EBA and CPD CIs. Figures B2 illustrate that the EBA and CPD approaches produce very similar carbon-intensity estimates. In a robustness check, we also reproduce the main outcomes of our analysis using carbon intensities estimated via the CPD approach.

Table B1: Upstream CI estimations

Variables	Estimates
Synthetic	146.332***
	(5.868)
Bitumen	60.823***
	(4.327)
Extra Heavy	57.418***
	(2.987)
Heavy (15-19)	54.599***
	(5.159)
Heavy (20-23)	52.627***
	(3.215)
Regular	48.547***
	(1.016)
Sour (> 3%)	44.476***
	(4.829)
Light	48.363***
	(1.641)
Condensate	44.799***
	(2.834)
Offshore	-6.609***
	(1.17)
SOR (bbl water/bbl oil)	23.837***
	(0.902)
GOR (kscf/bbl)	4.609***
	(1.26e-04)
FOR (kscf/bbl)	81.698***
	(4.08e-04)
R-squared	0.986
Adjusted R-squared	0.986

Note: *p<0.1; **p<0.05; ***p<0.01. The dependent variable is the upstream carbon intensity in $kgCO_2eq/bbl$ calculated in OPGEE (EBA, AR6, 100y). Column 1 displays the estimated coefficients from equation (7) with carbon intensities calculated in OPGEE. *Synthetic*, ..., *Sour (> 3%)* are dummy variables for each oil type. *Sour (> 3%)* represents oil with a sulfur content of over 3%. *Offshore* is a dummy for the field being offshore. *GOR* is the ratio between the quantity of gas in Thousands of cubic feet (kscf) and the quantity of oil in barrels in the reservoir. *FOR* is the ratio between the quantity of gas flared and the quantity of oil extracted, also in kscf/bbl. *SOR* is the ratio between injected water and the quantity of oil extracted, in bbl of water/bbl of oil. See Appendix B for the definition of these variables.

Figure B1: Upstream CI estimation: Estimates of the coefficients on the explanatory variables removing the OPGEE fields one-by-one.

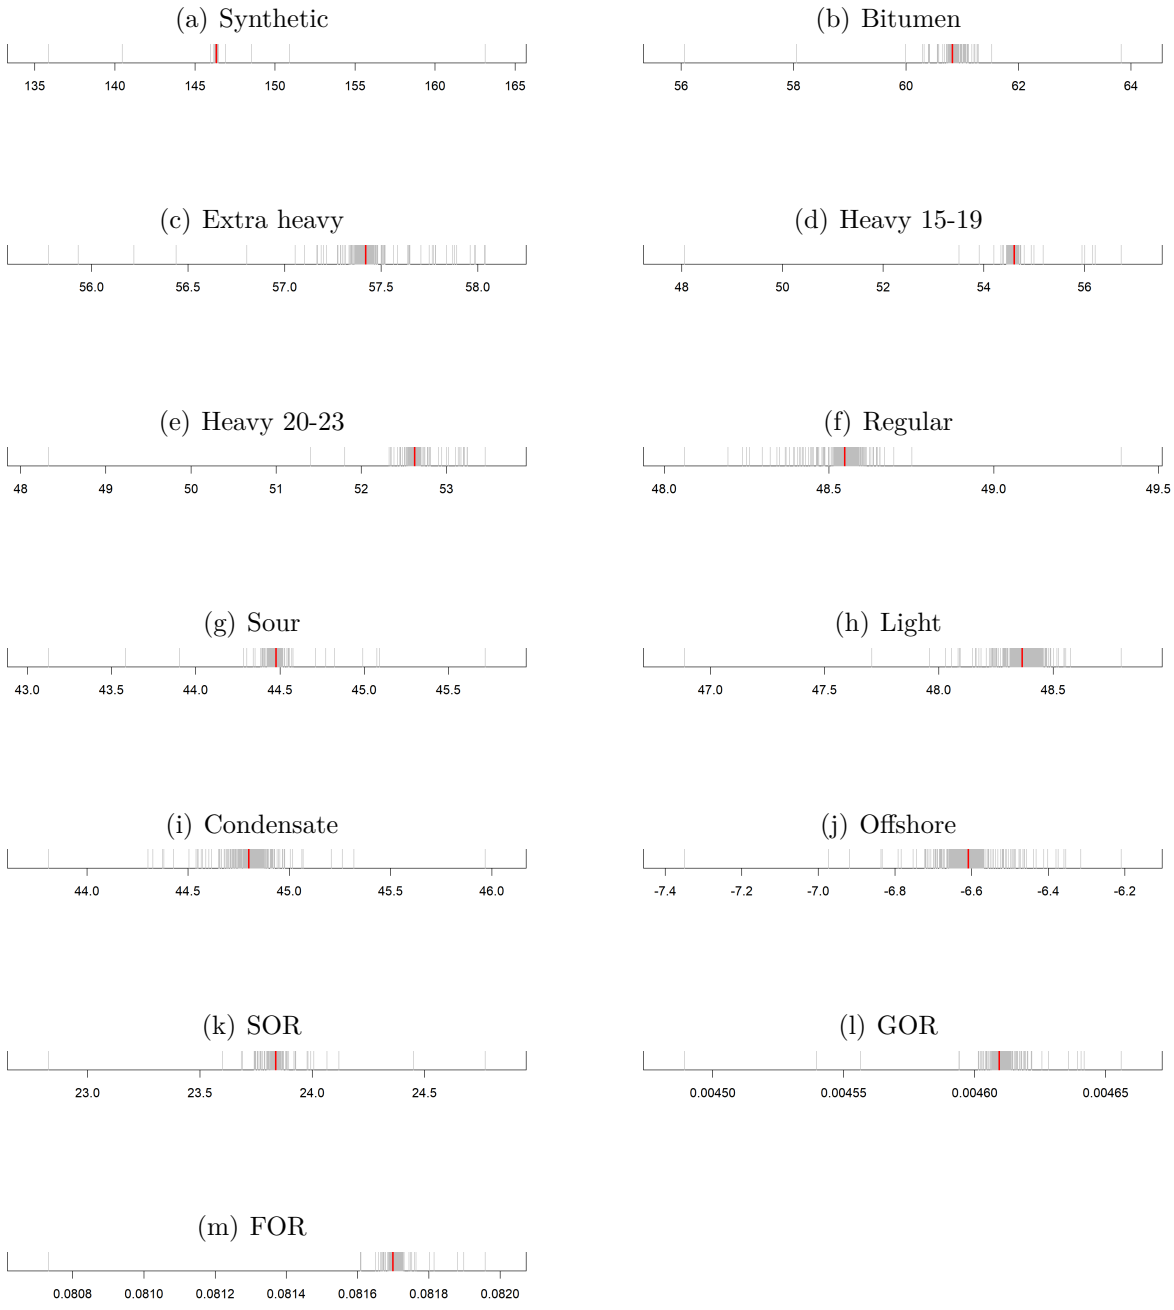

Note: These panels show the estimated coefficients on each of the explanatory variables in equation (7) removing OPGEE fields one-by-one. Each gray line corresponds to a different estimation with the same number of observations: $N(OPGEE_{sample}) - 1$. The thick red lines show the original estimates obtained with the full sample of OPGEE publicly-available fields (OPGEE sample). The coefficients before the GOR and FOR variables are for units in million of scf per bbl, thus they need to be multiplied by 1000 to be comparable to coefficients in Table B1 where units are in kscf/bbl for these variables.

Table B2: Upstream CI estimation: Excluding influential observations (Cook's distance).

N obs.	Synthetic	Bitumen	Condensate	Extra-Heavy	Heavy 15-19
5	140.259*** (7.379)	61.118*** (3.995)	45.776*** (2.594)	54.708*** (2.747)	55.087*** (4.717)
10	127.934*** (7.473)	60.929*** (3.505)	45.821*** (2.282)	54.784*** (2.410)	48.687*** (4.406)
	Heavy 20-23	Light	Regular	Sour	Offshore
5	53.119*** (2.940)	48.114*** (1.501)	49.265*** (0.933)	45.030*** (4.415)	-7.169*** (1.073)
10	49.076*** (2.633)	46.828*** (1.326)	49.624*** (0.822)	45.211*** (3.874)	-7.932*** (0.944)
	SOR	GOR	FOR		
5	23.738*** (0.846)	4.617*** (0.115)	80.744*** (0.386)		
10	23.801*** (0.743)	4.664*** (0.101)	80.783*** (0.382)		

Note: *p<0.1; **p<0.05; ***p<0.01. Each line represents the coefficient on the variable indicated at the column head when estimating equation (7) excluding the N most-influential observations from the OPGEE sample. The number N appears in the first column. We define the n most-influential observations as those with the n-largest Cook's distances. The Cook's distance of observation i is $CD_i = (1/ps^2)\sum_{j=1}^{N}\left(\widehat{y}_j - \widehat{y}_{j(i)}\right)^2$, where $\widehat{y}_j$ is the fitted response value obtained with the full sample and $\widehat{y}_{j(i)}$ that after excluding i, s^2 the mean squared error of the regression model, p the number of coefficients, and N the number of observations in the OPGEE sample. See Appendix B for the definition of these variables and the note to Table B1.

Figure B2: The distribution of upstream carbon intensities: EBA and CPD.

(a) OPGEE-sample fields (calculated CIs) (b) All producing fields 1992–2018 (Rystad)

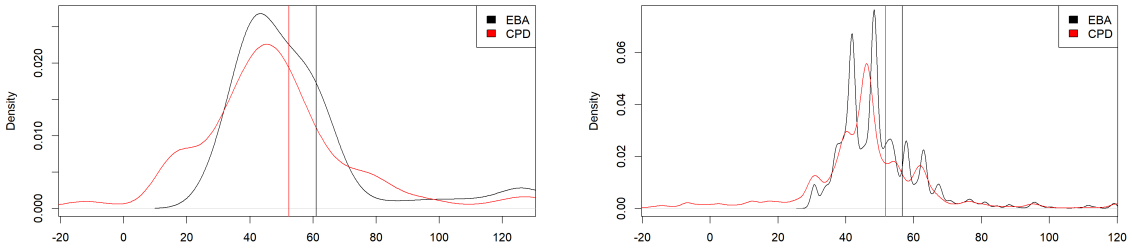

Note: Panel (a) plots the production-weighted distributions of upstream CIs (2015 production, $GOR < 10,200$) calculated in OPGEE for fields in the OPGEE sample using the EBA (black) and CPD (red) methods. On the horizontal axis, CIs are in kilograms of CO_2eq per barrel (kgCO_2eq/bbl). The vertical bars are the CI means weighted by 2015 production. Panel (b) plots the production-weighted distributions of predicted CIs (using (7) and Rystad data) and those calculated in OPGEE when possible with the EBA (black) and CPD (red) methods for all fields producing oil between 1992 and 2018 ($GOR < 10,200$). On the horizontal axis, CIs are in kilograms of CO_2eq per barrel (kgCO_2eq/bbl). The vertical bars are the CI means weighted by 1992–2018 productions. We exclude fields with CI > 125 or $CI < -20$ (CI > 120 or $CI < -20$) for visibility reasons in Panel (a) (Panel (b)). To smooth the curves, fields with $CI > 150$ are dropped as well as CI < -30 for the CPD approach.

B.2 Midstream

PRELIM. The *Petroleum Refinery Life-Cycle Inventory Model* (PRELIM),[9] from the Oil-Climate Index, is an engineering-based model to estimate energy use, refining yields and GHG emissions from crude-oil refining. Three types of refinery (hydroskimming, medium conversion, and deep conversion) appear in the model. PRELIM selects the refining type based on the crude's API gravity and sulfur content as follows: deep conversion for heavy crudes (under 22 API) with any sulfur content; medium conversion for medium crudes (22-32 API) and for light sour crude (over 32 API and 0.5% sulfur content by weight); and hydroskimming for light sweet crude (over 32 API and below 0.5% sulfur content by weight).

The Rystad dataset contains only a subset of the crude properties required to run PRELIM. To the best of our knowledge, there is no systematic collection of crude assays for global oil production. We thus run PRELIM with 657 assays — from companies, specialized websites and past research — that are publicly available in the PRELIM 1.6 Assay Inventory. We run PRELIM and obtain the carbon intensities for these 657 assays.

Linking PRELIM's crude assays and Rystad's fields. PRELIM crude assays were linked to the Rystad crudes assays using operator name, crude name, and location information (fuzzy matching on name, with manual checks). 513 PRELIM assays out of 657 were matched to Rystad fields (a few old assays were dropped). For some fields (mostly in Canada), the extracted oil was split into different refined crudes, depending on the company. For each field, Rystad gives the percentage of production allocated to each crude assay. Refining emissions were allocated to a field based on the split of its oil production across crude assays. For these fields, the weighted averages were calculated when data on the use of each field's oil was available.[10] Otherwise, weighted averages using the companies' shares in the fields were calculated.[11]

Estimation model. After matching the PRELIM crude assays to the corresponding fields in the Rystad dataset, we regress midstream carbon intensities on key variables from the Rystad dataset that can impact the refining process in this sample of 953 fields.[12] As

[9]The model and documentation is available at http://oci.carnegieendowment.org/#models. PRELIM and OPGEE use the same units of energy, and are designed so that their carbon intensities can be summed to track emissions from oil exploration to refining.

[10]For example, 85% of the oil from Christina Lake (Canada) is refined as "Christina Dilbit Blend", while the remaining 15% is blended with other oils and refined as "Western Canadian Select". Data on the use of each field's oil is available at https://www.oilsandsmagazine.com/projects/.

[11]For instance, there are two crudes in the American field "Thunder Horse" (one from BP and another from ExxonMobil). It was assumed that as ExxonMobil owns 75% of the field, 75% of the production was refined by ExxonMobil (and so 25% by BP). Data on companies' shares come from Rystad.

[12]This sample is representative of the remainder of the Rystad oil assets in the dimensions that matter for the estimation of midstream carbon intensity. For the Rystad assets matched to the PRELIM assays, mean API is 37.0, 18% contain over 1% of sulfur, and 12% are operated by Major companies. The Rystad dataset

discussed above, the PRELIM model automatically selects the best refinery configuration for a crude based on its API gravity and sulfur content: heavy and sour crudes require complex deep conversion refineries that emit more GHG emissions. We also add a Major company dummy, as Majors blend oil crudes to create specific oil crudes with stable chemical and physical properties for refining. Note that we are interested in explaining only the part of midstream carbon intensity that relates to oil characteristics, so that a change in deposit extraction could reduce the carbon footprint of refining. We consider refineries as fixed, and abstract from heterogeneity in midstream emissions that relate to refinery particularities independent of the oil type. We estimate the following model:

$$\begin{aligned} CI_f^{PRELIM} &= \textstyle\sum_0^4 \beta_i Config_{i,f} + \beta_5 API_f + \beta_6 API_f^2 + \beta_7 API_f^3 \\ &+ \beta_8 ln(Sulfur_f) + \beta_9 ln(Sulfur_f)^2 + \epsilon_f, \end{aligned} \tag{8}$$

where CI^{PRELIM} is the midstream carbon intensity of crude f, API is the crude API gravity, *Sulfur* its Sulfur content of the crude and $Config_i$ a set of dummies for the refinery configuration. The allocation of an oil to a specific configuration is based on the oil API and sulfur content following the PRELIM criteria explained above. For each field in the Rystad dataset, we then use the coefficient estimates from 8 to predict the refining carbon emissions of its oil.

Regression results. The estimates of equation (8) appear in Table B3. Our model explains over half of the variance in the midstream carbon intensities of crudes (for both the "coking" and "hydrocracking" refining configurations). This reflects the link between the refining process (deep or medium conversion, and hydroskimming) and the API gravity and sulfur content. There is significant CI variation between heavy and light oils, with bitumen refining being the most polluting. The Sour dummy also captures the particularities of the refining process. The Major dummy coefficient is not significant at conventional levels when considering hydrocracking refining.

The final field-level midstream dataset of carbon intensities contains the CIs estimated using PRELIM for crudes that are common to the PRELIM and Rystad datasets, and the CIs predicted using estimates from regression (8) and Rystad data for the rest of the fields.

Robustness checks. Panel (a) shows the strong correlation between midstream CIs computed in PRELIM and those predicted (8) for the sample of Rystad field that could be matched to a PRELIM crude assay. In Panel (b) of Figure B3, we plot the distribution of the CIs calculated in PRELIM for the Rystad fields matched to the PRELIM crude assays and those predicted using (8) for all other producing fields over the 2012—2021 period.

assets that are unmatched to PRELIM oil assays, with positive reserves in 1992, have analogous figures of 33.5, 24% and 8%.

Table B3: Midstream CI estimation.

Variables	Estimates
Deep conversion	56.552***
	(1.374)
Medium conversion	51.361***
	(1.604)
Hydroskimming	40.093***
	(1.608)
API	-0.933***
	(0.124)
API^2	0.016***
	(0.003)
API^3	-9.62e-05***
	(2.74e-05)
ln(Sulfur)	-0.988***
	(0.103)
$ln(Sulfur)^2$	-0.135***
	(0.025)
R-squared	0.997
Adjusted R-squared	0.997

Note: *p<0.1; **p<0.05; ***p<0.01. This table shows the estimates from equation (8). The dependent variable is the midstream carbon intensity in $kgCO_2eq/bbl$ computed in PRELIM (AR6 100y). *Deep conversion*, *Hydroskimming*, *Medium conversion* are refinery configurations attached to each crude. *API* is the American Petroleum Institute (API) gravity index, a measure of oil density. *Sulfur* is the sulfur content in percent. See Appendix B for the definition of these variables.

We then consider the sensitivity of the estimated coefficients of specification (8) to specific observations, and re-estimate it multiple times removing the fields in our main estimation sample one by one. Figure B4 displays the estimated coefficients before each explanatory variable. Each gray line corresponds to a different estimation, and the thick red line to the original correlation coefficient in the full OPGEE sample. Overall, these robustness checks indicate that estimated parameters of equation (8) are robust to changes in the estimation sample of deposits.

Figure B3: Midstream carbon intensities.

(a) Prelim and Predicted midstream CI

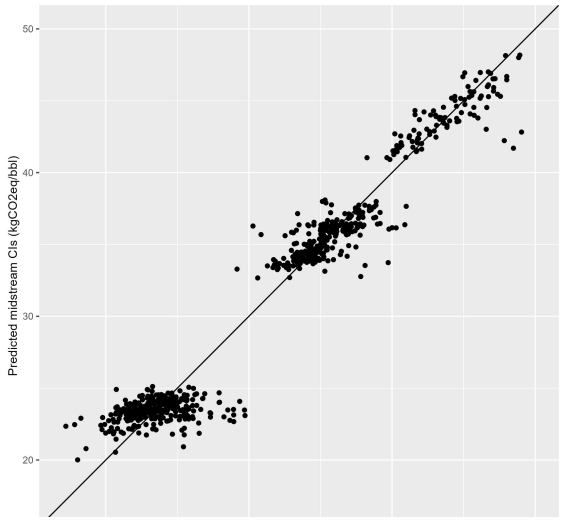

(b) Distribution of midstream CI

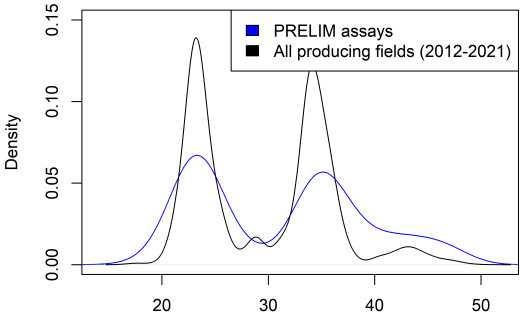

Note: The figure plots the distributions of refining carbon intensities for PRELIM crude assays calculated in PRELIM and those predicted by our regression model. On the horizontal axis, CIs are in kg of CO_2eq per barrel ($kgCO_2eq/bbl$). We exclude the crudes with CI > 17 for visibility reasons. The figure plots the distributions of refining carbon intensities for PRELIM crude assays for all fields producing oil between 2012 and 2018. On the horizontal axis, CIs are $kgCO_2eq/bbl$. We exclude the crudes with CI > 50 (Panel (a)) and fields with $CI > 14$ or CI > 52 (Panel (b)) for visibility reasons.

We then look at the robustness of our estimates to the exclusion of the most influential observations. For a given variable, we define the n most-influential observations as those with the n largest Cook's distances (CD), with $CD_i = \sum_{j=1}^{N} \left(\widehat{y}_j - \widehat{y}_{j(i)}\right)^2 / ps^2$, where $\widehat{y}_j$ and $\widehat{y}_{j(i)}$ are the fitted response values from using the full sample and after excluding field i, respectively; s^2 is the mean squared error of the regression model, p the number of coefficients, and N the number of observations. We re-estimate equation (8) after dropping the 5, or 10 most-influential observations by this definition, and list the new estimates in Table B4. Overall, the coefficients of the explanatory variables are stable.

Figure B4: Midstream: Estimates of the coefficients on the explanatory variables removing the OPGEE fields one-by-one.

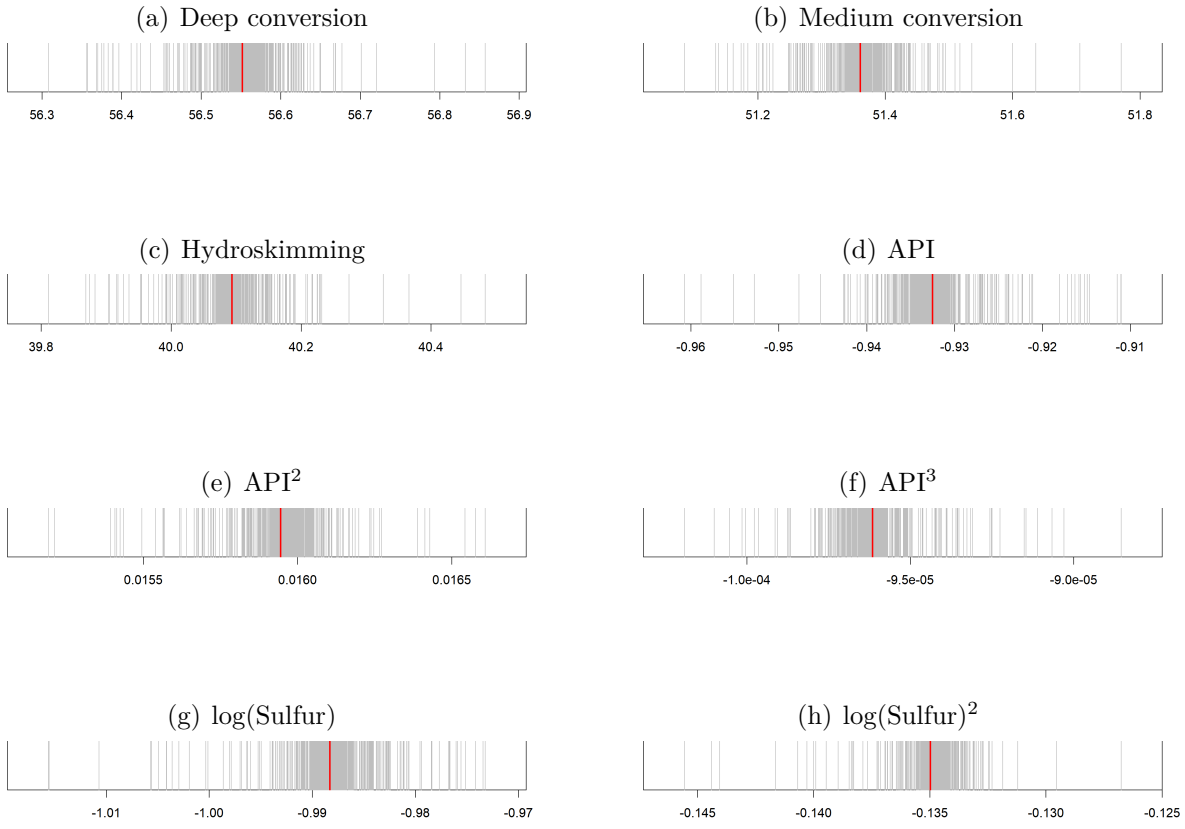

Note: These panels show the estimated coefficients on each of the explanatory variables in equation (8) removing OPGEE fields one-by-one. Each gray line corresponds to a different estimation with the same number of observations: $N(OPGEE_{sample}) - 1$. The thick red lines show the original estimates obtained with the full sample of OPGEE publicly-available fields (OPGEE sample).

Table B4: Midstream CI estimation: Excluding influential observations (Cook's distance).

N obs.	Deep conversion	Hydroskimming	Medium conversion	API
5	56.74*** (1.26)	40.565*** (1.473)	51.707*** (1.47)	-9.37e-01*** (1.14e-01)
10	56.702*** (1.234)	40.593*** (1.44)	51.744*** (1.438)	-9.33e-01*** (1.12e-01)
	API^2	API^3	ln(Sulfur)	$ln(Sulfur)^2$
5	1.56e-02*** (2.99e-03)	-9.05e-05*** (2.52e-05)	-1.07e+00*** (9.45e-02)	-1.70e-01*** (2.37e-02)
10	1.53e-02*** (2.93e-03)	-8.77e-05*** (2.47e-05)	-1.10e+00*** (9.31e-02)	-1.79e-01*** (2.45e-02)

Note: *p<0.1; **p<0.05; ***p<0.01. Each line represents the coefficient on the variable indicated at the column head when estimating equation (8) excluding the N most-influential observations from the OPGEE sample. The number N appears in the first column. We define the n most-influential observations as those with the n-largest Cook's distances. The Cook's distance of observation i is $CD_i = \frac{1}{p\,s^2}\sum_{j=1}^{N}\left(\widehat{y}_j - \widehat{y}_{j(i)}\right)^2$, where $\widehat{y}_j$ is the fitted response value obtained with the full sample and $\widehat{y}_{j(i)}$ is that after excluding i, s^2 is the mean squared error of the regression model, p is the number of coefficients, and N is the number of observations in the OPGEE sample. See Appendix B for the definition of these variables and the note to Table B3.

B.3 Downstream

The downstream sector represents the transport of refined products to end consumers, and their use (mostly combustion). Downstream emissions do not vary much across the original crudes. Oil-product combustion represents most of their life-cycle carbon emissions (downstream emissions average 382 $kgCO_2eq/bbl$, out of 473 $kgCO_2eq/bbl$ total life-cycle emissions), but combustion emissions do not vary across oil crudes for a given composition of final products (e.g. gasoline, kerosene etc.) that we take as given, as our focus is on supply rather than demand. In a robustness check, we allow the composition of petroleum products to vary with crude origin: this generates variation in combustion-related emissions per barrel of crude oil. We use the *Oil Products Emissions Module* (OPEM) model from the Oil-Climate Index (https://oci.carnegieendowment.org/) to compute downstream carbon intensity of petroleum products for each crude assay, we then match these crude assays with Rystad fields using a matching approach identical to that used for midstream emissions.

Greene et al. (2020) show that the carbon emissions from well-to-tank seaborne transportation vary between 5 and 27 gCO_2eq per liter (i.e. between 0.78 and 4.2 $kgCO_2eq/bbl$), with an average figure of 10 gCO_2eq per liter (1.6 $kgCO_2eq/bbl$). Even the difference between the two extreme values, 3.4 $kgCO_2eq/bbl$, is small compared with the standard deviation of life-cycle carbon intensity in our dataset (43 $kgCO_2eq/bbl$). Therefore, we abstract from the consequences of a change in the structure of the oil supply, in terms of deposit-level productions, on total GHG emissions.

Finally, we also take as given the quantity of oil used in the petrochemical industry (about 12% according to IEA 2018a), for which downstream emissions are small.

C Private Extraction Costs and Extraction Capacities

C.1 Data on Private Extraction Costs

In our main specification, we measure annual costs as the sum of "well" and "facility" CAPEX, and the "selling, general and administrative", "transportation", "production" OPEX.

According to the Rystad definitions, "well" CAPEX is capitalized costs related to well construction, including drilling costs, rig lease, well completion, well stimulation, steel costs and materials. "Facility" CAPEX includes development CAPEX ("costs to develop, install, maintain and modify surface installations and infrastructure"), exploration CAPEX ("costs incurred to find and prove hydrocarbons"), and abandonment costs ("costs associated with shutting down and dismantling the surface and subsea facilities"). Production OPEX is "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 OPEX includes "the costs of bring the oil and gas from the production site/processing plant to the pricing point (only upstream transportation)". We use costs and productions recorded in Rystad for each field to compute field-level extraction costs per barrel: these costs and productions are observed for producing fields before 2022, and productions and costs are modeled for post-2021.[13] Most of the fields in proven reserves have been producing before 2022, thus our cost measure for prove fields is partly computed ex post based on *realized* productions and costs. We assume that the marginal extraction is constant and equal to average cost as long as extraction is below capacity constraint.

We calculate average costs as: $c_d = (\sum_{1900}^{2099} Exp_{dt})/(\sum_{1900}^{2099} x_{dt})$, where x_{dt} and Exp_{dt} denote the annual deposit production and *total* cost.

C.2 Extraction Capacities

The literature usually describes field-level extraction as having two sequential phases: a plateau, and then a phase of decline in extraction (Fetkovich, 1980; Robelius, 2007; Höök and Aleklett, 2008; Höök et al., 2014; Jackson and Smith, 2014). The "plateau–decline" profile is captured by two simultaneous flow limits. For every deposit d and year t we require

$$x_{dt} \leq \min(k_d, \alpha_d R_{dt}),$$

[13] Following Rystad's guidelines, if a field is split into several assets for fiscal reasons, cost data are aggregated at the block level.

Where x_{dt} is production from deposit d at date t in any scenario and $R_{d,t}$ remaining reserves from deposit d at date t in any scenario. The first term k_d fixes the maximum plateau rate ; the second, $\alpha_d R_{d,t}$, lets the cap fall as the field depletes. We now explain the two right-hand-side terms of the inequality above: first the plateau bound k_d, then the reserves-based decline bound $\alpha_d R_{dt}$ –under two alternative capacity specifications: an exogenous version calibrated to historical production data and an endogenous version that adjusts with the field's development spending.

Exogenous capacity constraints In our main specification, a deposit's plateau is exogenous to firms' decision. For deposit d and year t, $x_{dt} \leq k_d$, where

$$k_d = \max \left(\max_{j \in [1970,2018]} (x_{d,j,observed}), m_d \times R_{d,1970,observed} \right)$$

where m_d is 1.9% if d is in the top 5% of deposits, 3.3% otherwise (remaining 95%), x_{dt} is production from deposit d at date t in any scenario, $x_{d,j,observed}$ is observed production from deposit d at date j and $R_{d,1970,observed}$ is the deposit's 1970 reserve.

The "plateau" constraint forces a field's production to be below either 1.9% or 3.3% of the field's reserves as of 1970 or to be below the maximal observed production of the field since 1970 if the latter is larger. For the bottom 95% of fields that start producing after 1970, the maximal observed production since 1970 is larger than 3.3% of the field's 1970 reserves for 95% of them. For the top 5% largest fields, the same ratio exceeds 1.9% for 95% of the fields. For a typical field with no (or little) observed past production, forcing production to stay below these size-specific thresholds (3.3% or 1.9% of the field's 1970 reserves) remains very conservative.

Endogenous capacity constraints In another specification, we introduce endogenous capacity constraints, separating development and extraction costs. Following Rystad's definition, we define development costs as the sum of well" and facility" CAPEX up to 2 years after production startup: $C_d = \sum_{1901}^{Start-upYear+1} FacilityCapex_{dt} + WellCapex_{dt}$. We allow for partial development of assets and assume that by spending x% of the total development costs recorded in Rystad, the plateau is set to x% of the plateau k_d defined above. We do not allow for extra development that would increase the plateau-level k_d. Extraction costs are then computed as averages costs purged from the capital expenses included in development costs for each asset.

Decline constraints The "decline" part reads for deposit d and year t, $x_{dt} \leq \alpha_d R_{d,t}$, where

$$\alpha_d = \max \left(\max_{j \in [1970,2018]} \left(\frac{x_{d,j,observed}}{R_{d,j,observed}} \right), n_d \right)$$

Where $R_{d,j,observed}$ is observed reserves of deposit d at date j, $R_{d,t}$ remaining reserves from deposit d at date t in any scenario and n_d is 3.8% for the top 5% fields and 6.6% for the remaining 95% fields.

The "decline" constraint forces a field's extraction-to-current-reserves to be below n_d or to its maximal historical extraction-to-current-reserves, if the latter is larger. The n_d decline rate is chosen because on average the maximal observed extraction-to-current-reserves rate during the decline phase (when production is less than 95% of the maximal production, as defined in Höök et al. 2014) is twice the maximal observed extraction-to-1970-reserves rate. The n_d decline rate is twice the m_d plateau coefficient defined previously.

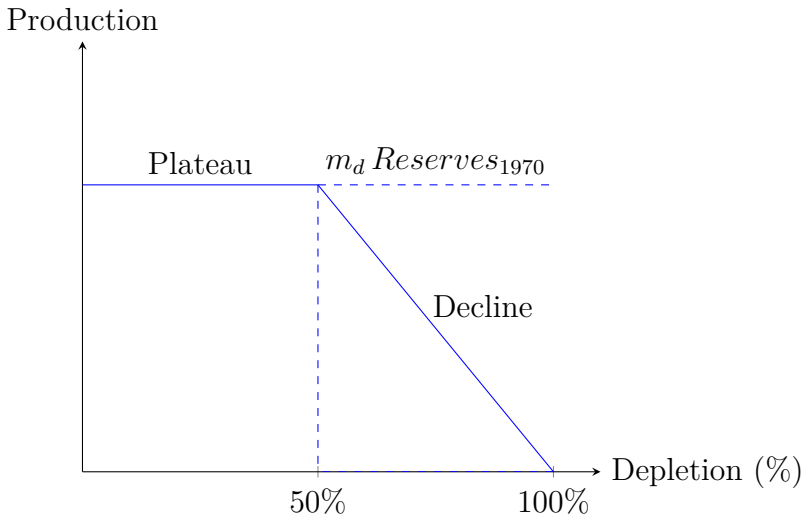

Figure C1: The "Plateau-decline" constraint for a virtual field

The "exogenous plateau–decline" specification allows for field-specific extraction. There is significant variation in extraction rates across field sizes: numerous small fields (e.g., shale oil) have high extraction rates and can be exhausted over the course of a few years, whereas larger fields tend to have lower extraction rates and longer plateau durations. Our capacity constraint reflects this heterogeneity.

To illustrate, suppose the plateau cap is determined by $k_d = m_d R_{d,1970}$ (i.e., historical production does not bind) and the decline cap is governed by $x_{dt} \leq \alpha_d R_{d,t}$. The binding constraint switches from plateau to decline when $k_d = \alpha_d R_{d,t}$, i.e.

$$\frac{R_{d,t}}{R_{d,1970}} = \frac{m_d}{\alpha_d}.$$

Under our baseline calibration for fields with little (or no) observed past production, $\alpha_d = n_d = 2m_d$, so the plateau lasts until remaining reserves fall to $R_{d,t} = 0.5\, R_{d,1970}$ (50%

depletion), after which production declines proportionally to remaining reserves at rate n_d, as illustrated in Figure C1. When the historical maximum extraction-to-current-reserves exceeds n_d, we have $\alpha_d > n_d$ and the transition to decline can occur earlier.

We perform a capacity-constraint sensitivity check with two variants: (i) an *endogenous-plateau* specification, in which both the plateau and the decline cap adjust with investment; and (ii) a *plateau-only* specification, in which the plateau is fixed (exogenous) and the decline cap is removed.

Finally, in Venables (2014), a faster rate of depletion can reduce the size of recoverable reserves in a field. This rate sensitivity is based on Nystad (1985, 1987) who categorize fields as "Hotelling" (recoverable reserves are not rate sensitive), "intermediate", and "geosensitive" (recoverable reserves are strongly rate sensitive) and consider the theoretical implications of reservoirs' rate-sensitivity on optimal extraction rate. Our conservative assumption regarding capacity constraint ensures that extraction rates in our counterfactuals are below what is observed in the data and likely below levels of depletion rates that would significantly reduce the size of recoverable reserves. We have included robustness checks on reserves' recoverability that indirectly address the issue of highly rate-sensitive reservoirs such as dropping 10% of each field reserves.

D Modeling of Yet-to-find Resources

We model yet-to-find (YTF) resources within *known* geological basins. Unlike proven reserves, YTF are uncertain in both quantity and characteristics. While the Rystad database provides YTF estimates, the underlying assumptions are not fully transparent. We therefore construct YTF assets under explicit, testable hypotheses and compare our results to Rystad's ones for robustness.

D.1 Data Scope

We use field-level discoveries from Rystad across 366 sedimentary basins. A basin enters the model if it has at least 10 discoveries ; frontier basins (< 10) are handled separately with pooled information.

D.2 Field-size Distribution and Discovery Bias

In modeling the size distribution of oil fields within a region, several statistical distributions have been proposed in the geology and oil engineering literatures, each with distinct

characteristics reflecting different underlying assumptions about field sizes. The most commonly considered distributions are the lognormal distribution, the Pareto distribution, and the parabolic fractal model. The lognormal distribution has been traditionally favored in the oil and gas industry following early observations by Arps and Roberts (1958) that the logarithm of discovered field sizes tends to resemble a normal distribution. This implies that field sizes themselves are lognormally distributed, characterized by a right-skewed distribution where a small number of large fields contain the majority of the resources. Empirical studies have reinforced the applicability of the log-normal model, making it a conventional choice for resource assessment (McCrossan, 1969; Kaufman, 1964).

Figure D1 illustrates the lognormality of observed field size distribution in the context of the Volga-Urals basin.

Figure D1: Distribution of the natural logarithm of field initial reserves, Volga-Urals basin

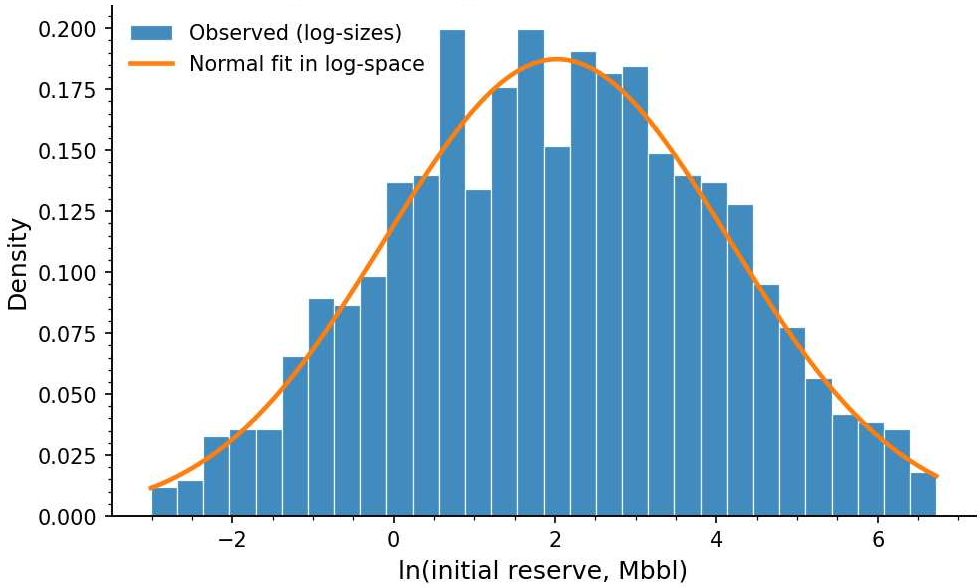

Note: Histogram shows the empirical distribution of the natural logarithm of initial reserves (Mbbl) for 1057 observed initial reserves. The solid curve is a Normal distribution fitted to log-sizes (equivalent to a lognormal model in original reserve units), with parameters estimated by maximum-likelihood (μ and $\sigma$of ln-size). For visual robustness, the plotted range is restricted to the 1st—99th percentiles of ln-sizes while densities are computed with a fixed binning scheme (30 bins).

To model the sizes of yet-to-find (YTF) assets, we employ a log-normal distribution, which is suitable due to the right-skewed nature of field size distributions in the oil and gas industry, where a few large fields dominate the total reserves. For each basin with at least 10 discoveries, we estimate the location (μ_b) and scale (σ_b) from the observed sizes, using a size-biased likelihood and the curve-alignment penalty.

Let S denote the true field size in a basin. We assume the population follows a lognormal

distribution,

$$S \sim \text{LogNormal}(\mu_b, \sigma_b^2), \qquad f_S(s) = \frac{1}{s\,\sigma_b\sqrt{2\pi}} \exp\!\Big[-\tfrac{(\ln s-\mu_b)^2}{2\sigma_b^2}\Big],\ s>0,$$

with mean $\mathbb{E}[S] = \exp(\mu_b + \frac{\sigma_b^2}{2})$.

Because larger fields are easier to find (size-biased discovery), an observed discovery $S^\star$ is drawn with probability proportional to size. The observed density is therefore the *size-biased* transform of f_S:

$$f_{S^\star}(s) = \frac{s\, f_S(s)}{\mathbb{E}[S]}\,, \qquad s>0.$$

(full derivation below):

$$S^\star \sim \text{LogNormal}(\mu_b + \sigma_b^2,\ \sigma_b^2),$$

i.e. the *observed* sizes follow a lognormal with the same scale but a location shifted by $+\sigma_b^2$. This shift is a property of the discovery mechanism and must be included in the likelihood for the observed discoveries and is not applied when simulating the underlying population.

Proof of the $+\sigma^2$ shift (size-biased lognormal). Let $S \sim \text{LogNormal}(\mu, \sigma^2)$ with density

$$f_S(s) = \frac{1}{s\,\sigma\sqrt{2\pi}} \exp\!\Big[-\tfrac{(\ln s-\mu)^2}{2\sigma^2}\Big], \qquad s>0,$$

Under size-biased sampling, the discovered size $S^\star$ has density

$$f_{S^\star}(s) \;=\; \frac{s\, f_S(s)}{\mathbb{E}[S]} = \frac{1}{s\,\sigma\sqrt{2\pi}}\ \exp\!\Big[-\tfrac{(\ln s-\mu)^2}{2\sigma^2} - \mu - \tfrac{\sigma^2}{2} + \ln s\Big].$$

Completing the square in $\ln s$ using

$$-\frac{(\ln s-\mu)^2}{2\sigma^2} + \ln s - \mu - \frac{\sigma^2}{2} = -\frac{(\ln s-(\mu+\sigma^2))^2}{2\sigma^2},$$

gives the shifted lognormal form

$$f_{S^\star}(s) = \frac{1}{s\,\sigma\sqrt{2\pi}} \exp\!\Big[-\tfrac{(\ln s-(\mu+\sigma^2))^2}{2\sigma^2}\Big], \quad s>0\,.$$

Hence $S^\star \sim \text{LogNormal}(\mu+\sigma^2, \sigma^2)$.

D.3 Joint Estimation of Field Size Distribution and Total Count

We jointly estimate the parameters of the basin-level field-size distribution and the total number of fields. The estimation proceeds in four steps: Bayesian inference of the lognormal parameters, refinement of these estimates through the modeling of the discovery process, estimation of the total field count, and iterative convergence of all parameters.

D.3.1 Bayesian Inference of (μ_b, σ_b) from Discoveries

We first estimate the parameters (μ_b, σ_b) of the *population* field size distribution in basin b from the observed discoveries, using the size–biased lognormal likelihood:

$$S^\star \sim \text{LogNormal}(\mu_b + \sigma_b^2,\ \sigma_b^2).$$

In `PyMC` we implement this as a `Lognormal` likelihood with parameters $(\mu_b + \sigma_b^2,\ \sigma_b)$ applied to the observed sizes.

Priors: we use weakly informative basin-level priors for μ_b and σ_b, anchored to decile-specific pooled medians from mature basins, providing gentle regularization without dominating the data.

From these estimates, we compute an initial basin-level estimate of the total number of fields, $N_b^{(0)}$. Following Hyndman and Fan (1996), we equate the empirical size rank of each observed discovery i with its expected quantile in the fitted population distribution:

$$N_i^{(0)} = \frac{i - \delta}{1 - F_{\mu_b,\sigma_b}(s_{(i)})}, \qquad \delta \simeq 0.35,$$

and take the trimmed median of $\{N_i^{(0)}\}$ to obtain $\hat{N}_b^{(0)}$. This $N_b^{(0)}$ serves as the initial population size for the subsequent simulation of the discovery process.

D.3.2 Discovery mechanism and Gumbel noise (curve alignment)

In a second step, we refine the estimates of (μ_b, σ_b) by exploiting the known discovery order, modeling the probabilistic ranking of fields during exploration rather than relying solely on the size distribution of the discovered fields.

Exploration generally targets larger prospects first, though not perfectly in size order due to imperfect information, access constraints, timing, and chance. We capture this process by simulating latent population sizes $S_i^{\text{sim}} \sim \text{LogNormal}(\mu_b, \sigma_b^2)$ for $i = 1, \ldots, \hat{N}_b^{(0)}$, and

assigning each a *score*

$$U_i = \log S_i^{\text{sim}} + \varepsilon_i, \qquad \varepsilon_i \sim \text{Gumbel}(0, 1).$$

The Gumbel term introduces "choice noise," a standard device in discrete choice models, turning the strict size ranking into a probabilistic one where larger fields are more likely—but not certain—to be discovered first. We then select the top k_b fields by U_i, corresponding to the number of observed discoveries, and compare the resulting cumulative size–rank curve to the empirical one. The discrepancy between both curves is penalized at each rank. The shifted-lognormal likelihood and the curve penalty for the ordering are fitted jointly, and provide estimates of the posterior medians $(\hat{\mu}_b, \hat{\sigma}_b)$.

D.3.3 Estimating Total Field Count and URR per Basin

Given $(\hat{\mu}_b, \hat{\sigma}_b)$ for the population, we sort the observed sizes in decreasing order $s_{(1)} \geq \cdots \geq s_{(k_b)}$ and evaluate the fitted CDF $F_{\hat{\mu}_b, \hat{\sigma}_b}$. For the i-th observed size, the *population* fraction expected to be larger is

$$p_i \;:=\; 1 - F_{\hat{\mu}_b, \hat{\sigma}_b}(s_{(i)}).$$

Under a largest-first discovery pattern, the number of larger fields is approximately $i - 1$. To reduce small-sample bias, we instead use $i - \delta$ with $\delta \approx 0.35$, a pragmatic adjustment consistent with practices in the quantile literature (see Hyndman and Fan, 1996). Equating expected exceedances to the corrected rank and inverting gives a per-observation estimate of the *total* field count:

$$N_i \;=\; \frac{i - \delta}{p_i} \;=\; \frac{i - \delta}{1 - F_{\hat{\mu}_b, \hat{\sigma}_b}(s_{(i)})}.$$

We aggregate $\{N_i\}$ robustly (trimmed median) to obtain $\hat{N}_b$.

D.3.4 Fixed-point refinement

We then perform up to ten refinement rounds that alternate between: (a) updating the effective total number of fields $\hat{N}_b$ using the quantile-fill procedure; and (b) refitting the size-distribution parameters (μ_b, σ_b) via the size-biased likelihood combined with the curve-alignment penalty at the updated $\hat{N}_b$. Iterations continue until convergence—defined as changes below 5 fields in $\hat{N}_b$, 0.1 in parameter estimates, and 1% in URR. If convergence is not reached after ten rounds, we retain the iteration with the best sampler diagnostics, this occurs in fewer than 20% of basins.

Finally, we use a simple rank-aware fill: build a full list of $\hat{N}_b$ model-implied sizes from the

fitted lognormal (using evenly spaced percentiles), sort it, replace the entries at the observed ranks with the actual observed sizes, and sum the completed list. This keeps every observed size at its rank and fills only the unobserved ranks with model-implied sizes from the fitted lognormal.

In basins with fewer than 10 existing fields, the limited data necessitates an alternative approach. We utilize global parameters derived from well-represented basins to define the log-normal distribution for these under-represented frontier basins.

D.4 Assigning Oil Type, Private Costs and Carbon Intensity to YTF Deposits

Inside each basin, we ensure that the oil type distribution is replicated, as most basins typically contain only two or three oil types. Private costs and carbon intensity for YTF deposits are assigned based on observed distributions within the same basin, oil type, and size bin, using Kernel Density Estimation (KDE). Missing values are filled using basin-wide averages. In frontier basins—defined as those with fewer than ten discovered deposits—we impose a minimum cost threshold: only newly created YTF assets with private costs above the 3rd quartile of observed values are retained. This reflects the assumption that exploration in such basins is economically viable only for higher-cost resources.

D.5 Sensitivity Analysis in YTF Estimation to the Pool of Observed Discoveries

For each basin, we begin by fitting our primary Bayesian model using all available data. From this fit, we store the point-by-point log-likelihood contributions (that is, how much each individual discovery supports the fitted model).

We then assess how sensitive the results are to individual discoveries using a leave-one-out procedure. Let k denote the number of observed discoveries in the basin, and set $K = min(20, k)$. We randomly select K discoveries and, in turn, consider analyses in which each selected discovery is removed from the dataset.

Rather than refitting the Bayesian model from scratch for every removal (which would be computationally expensive), we approximate the "discovery-removed" posterior distribution by reweighting the posterior draws from the original full-data fit. The reweighting is based on the stored log-likelihood contributions and uses a stabilized importance-weighting scheme, so that draws that better represent the dataset without the removed discovery receive higher weight.

For each removal, we then draw a single posterior sample from this reweighted set, yielding one representative pair of parameter values (μ, σ). Repeating this for all K removals produces a leave-one-out output file containing K such (μ, σ) parameter pairs for each basin.

Finally, to generate 1000 future scenarios, we construct a smooth approximation to the empirical joint distribution of (μ, $log(\sigma)$) implied by the leave-one-out parameter pairs, using a multivariate Gaussian model. We draw 1000 joint samples of μ and σ from this fitted distribution and, for each sampled pair, generate a full "yet-to-find" asset list and the corresponding total "yet-to-find" volume using the same quantile-based population-filling procedure as in the main approach.

Figure D2: Distribution of Global YTF: Sensitivity to Pool of Observed Discoveries

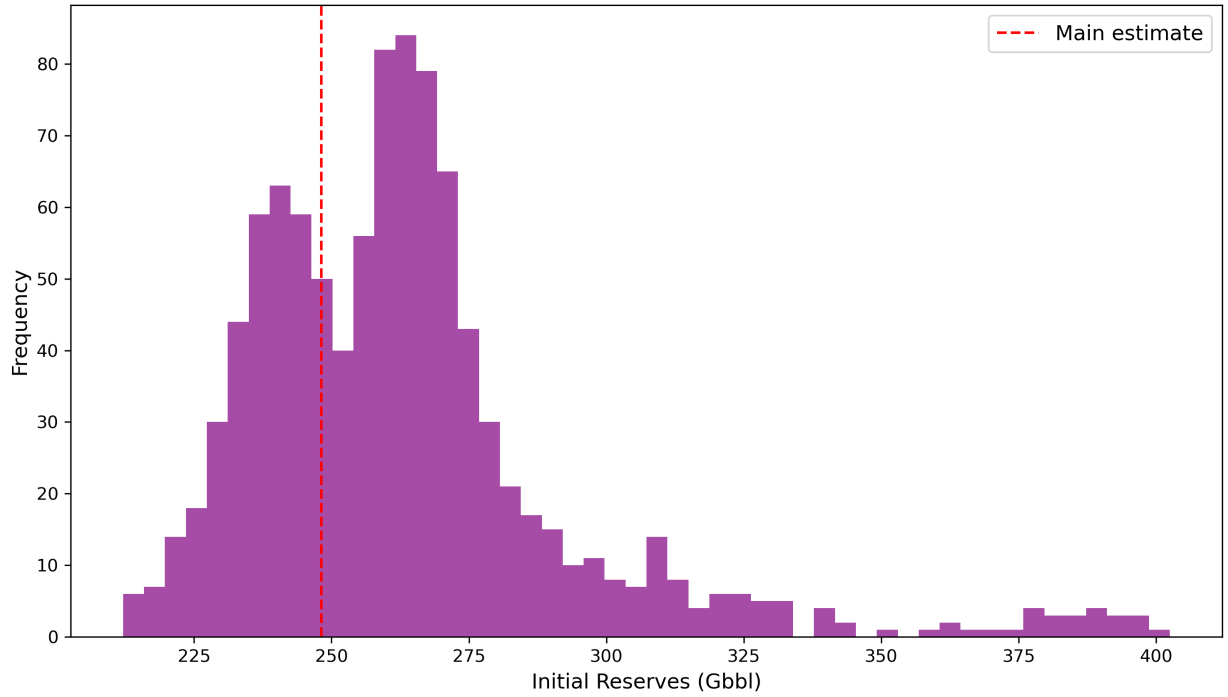

Note: The figure plots the distribution of global YTF in Gbbl estimated using the approach described in D.5. The dashed red line indicates our main estimate when all observed discoveries are used for YTF estimations.

Figure D3: Distribution of Country-level YTF: Sensitivity to Pool of Observed Discoveries

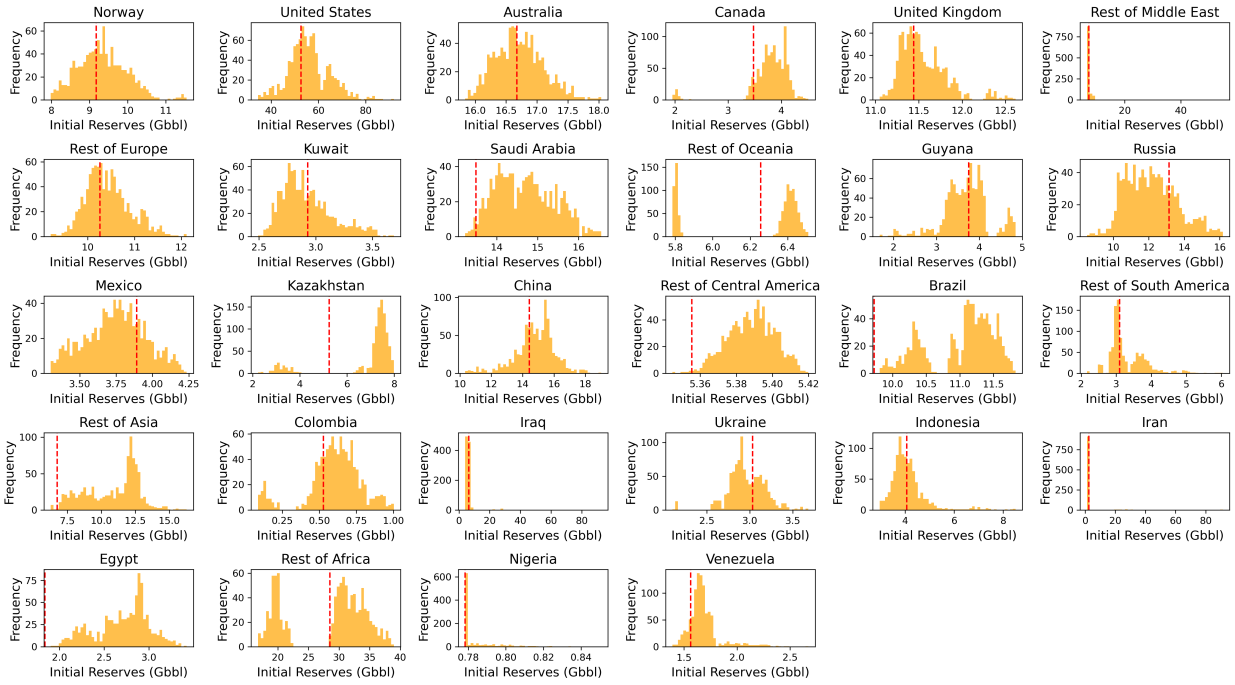

Note: The figure plots the distribution of country-level YTF in Gbbl estimated using the approach described in D.5. The dashed red line indicates our main estimate when all observed discoveries are used for YTF estimations.

D.6 Alternative YTF Constructions for Robustness Checks

In other specifications, we first assign private costs and carbon intensities by using the average per basin × oil type × size bin instead of reconstructing the entire distribution. Concretely, for each combination of basin, oil type, and size bin, we calculate the mean private cost and carbon intensity from the observed fields and apply these averages to all new YTF assets in that category. Given that most exploration effort in frontier basins focuses on offshore fields, we discard all non-offshore frontier-basin assets in another variant. We also test the robustness of our results to the possibility that newly found fields could have zero upstream carbon emissions. Specifically, we re-create the YTF assets using an alternative carbon intensity estimation that excludes upstream emissions from the observed fields, thereby reflecting a scenario in which future discoveries generate no additional upstream emissions. Additionally, we estimate the URR in other analysis thanks to a creaming curve method (we fit the cumulative sum of discoveries with a concave "creaming curve," and take its horizontal asymptote as the URR estimate) and assume that field size distribution within a basin follows a Pareto distribution or a log-normal distribution. Finally, we test the sen-

sitivity of our results to exploration cost assumptions by either discarding exploration costs entirely (reflecting a zero-cost exploration scenario) or doubling the original Rystad-based exploration cost estimates, which are aggregated at the supply segment level.

E Other Modeling and Parameter Choices

E.1 The Selection of Deposits

Our analysis is restricted to Rystad assets that have been discovered before 2022, with positive reserves at the beginning of 2022 of at least one of the following oil types: "Condensate", "Light", "Regular", "Sour (> 3%)", "Extra Heavy Oil", "Heavy Oil 15-19", "Heavy Oil 20-23", "Bitumen", and "Synthetic crude" (the numbers for heavy oils represent the API gravity range). The sample is composed of 14,637 oil fields.

E.2 Reserves

The reserve estimates are from the Rystad UCube dataset. A deposit reserve is defined as its economically-recoverable volume, assuming an oil price of US$120 per barrel (scenario "Resources High Case"). Current World reserves amount to 1,485 billion barrels in the Rystad dataset, as against 1,729 billion barrels for World proven reserves according to British Petroleum (2019). The difference in these numbers mostly reflects that the Rystad dataset does not record the proven reserves for some bitumen and extra heavy oils that are expected to be too expensive to be profitable, even in the long run. For instance, in Rystad, Canada and Venezuela have reserves estimated respectively at 101 and 34 billion barrels (Gbbl) at the end of 2018; the analogous figures in British Petroleum (2019) are 167 and 303 billion barrels. The resources not in the Rystad dataset are in general not only too expensive to be extracted but also very polluting. As a consequence, they would not be used in counterfactual scenarios in which social or private extraction costs are minimized. Overall, using a less-restrictive definition of reserves in our analysis would have little impact on our findings. There exist various definitions of oil reserves, as noted by McGlade (2012). Our reserve estimates are close to that of OPEC (2019) (1,498 Gbbl) and below those of British Petroleum 2019; EIA 2021; BGR 2020. The EIA estimates global reserves to amount to 1,659 Gbbl (EIA, 2021), the BGR to 1,790 Gbbl (BGR, 2020), the IEA to 1,700 Gbbl (IEA, 2019). In contrast, global reserves are lower at 1,276 Gbbl (2018 figure) in Welsby et al. (2021). In our main specification, we retain only deposits with a gas-to-oil ratio below 10,200 scf/bbl, which brings our reserve estimate of current reserves to 1,341 billion barrels. Rystad data include

yet-to-find resources within already-discovered fields from unconventional assets. These yet-to-find resources that correspond to resource extension represent only a small part (about 8%) of total reserves available as of 2019 (fields discovered before 2019). In a robustness check, we reduce reserves by 10% to account for the possibility of that overestimation in field reserve size.

E.3 Resource Availability

All fields in our sample have been discovered before 2021. Yet, for fields that do not produce by 2022, we include development lags after a resource discovery before production can start as in Arezki et al. (2017); Bornstein et al. (2017). Bornstein et al. (2017), in particular, highlight that development lags vary with oil type, a piece of evidence we do verify with our data. We assume that for all resource types,—except shale oil and tight oil—, the production can only start seven years after the discovery date (twelve years for offshore resources) (always anterior to 2022), or after the first year of observed production if the latter is sooner. For tight and shale oil, we use a two-year lag.

E.4 The Social Cost of Carbon

The value of the SCC in 2022 The development of Integrated Assessment Models (IAMs) in the 1990s—e.g., the Dynamic Integrated Climate and Economy (DICE), developed by William Nordhaus and used in Nordhaus (1994, 1992), the Policy Analysis of the Greenhouse Effect (PAGE), developed by Chris Hope and used in Stern (2008), and the Climate Framework for Uncertainty, Negotiation, and Distribution (FUND), developed by Richard Tol and used in Anthoff and Tol (2013)—constituted a turning point in the estimates of the social cost of carbon (SCC).[14] These models describe interactions between carbon-dioxide concentration, the climate, damage from climate change, and human activities that produce carbon-dioxide emissions. Although criticisms about the limitations of IAM remain (e.g., Pindyck, 2013, 2019, which stress arbitrary modeling choices, the deterministic nature of most IAMs, and the exclusion of catastrophes), the IAM estimates have been widely discussed in the public debate.

The SCC estimates from IAMs vary greatly, in particular depending on the choice of the social discount rate (Nordhaus, 2007; Stern, 2008; Tol et al., 2013), the coverage of damage, and the modeling (if any) of uncertainty (Weitzman, 2009, 2011; Calel et al., 2015; Gillingham et al., 2018). A growing body of literature argues that social costs of carbon, such as those used by the U.S. Government Interagency Working Group, that are based on

[14] See Calel and Stainforth (2017) for a comparison of these IAMs.

simulations from the DICE, PAGE and FUND models underestimate the true social costs, as they ignore key uncertainties, some climate-change related damage, irreversibilities and acceleration factors, such as the possibility of tipping points, and usually keep the valuation of ecosystems constant despite their rarefaction (for a summary of these criticisms, see e.g., Revesz et al., 2014; Van den Bergh and Botzen, 2014).

Nordhaus (2017) finds a SCC of $37 in 2020 (2010 US dollars) along the optimized emission path, which jumps to $87 with a discount rate of 3%. In the baseline parametrization of the DICE2016R, constraining the increase in temperature over the next 100 years as compared to 1900 to be under 2.5°C increases the SCC to $229 per tCO_2 (2010 US dollars). Dietz and Stern (2015) bring three main changes to the 2010 version of the DICE model: (i) Climate change negatively affecting the accumulation of physical, technological and intellectual capital; (ii) Modifying the function translating temperature increases into GDP loss to account for possible tipping points (Weitzman, 2012; Lenton and Ciscar, 2013); (iii) The "climate-sensitivity" parameter, which links atmospheric concentrations of greenhouse gases to (expected) temperature increases, being updated to account for new climate-science knowledge, although Nordhaus' discount rate is left unchanged (at twice the level of that preferred by Stern). They find a carbon price in the $32–103/$tCO_2$ range (2012 prices) in 2015. Lemoine (2021) estimates a 200-year social cost of carbon in 2014 of $362 per tCO_2 in a calibration exercise accounting for uncertainty about both warming and the impact of warming on consumption, and including stochastic shocks to consumption growth. Rennert et al. (2022), using the new open-source Greenhouse Gas Impact Value Estimator (GIVE) model with up-to-date knowledge regarding socioeconomic projections, climate models, damage functions, and discounting methods, found a central value of $185 per tCO_2 (2020 US dollars).

Overall, our choice of a SCC of US$200 per tCO_2 in 2022 is consistent with the SCC in DICE2016R when the temperature increase is kept strictly below 2.5°C over the next 100 years.

The rate of increase of the SCC The social cost of carbon in current value increases at the rate of the discount rate, thus the discounted cost of emissions is not time-dependent. The reason for this modeling choice is to be found in the scientific literature that shows that temperature increase is proportional to cumulative carbon emissions at a given time horizon, and independent of the carbon emission pathways (Allen et al., 2009; Matthews et al., 2009; Meinshausen et al., 2009; Gillett et al., 2013). This evidence serves as a rationale behind mitigation targets formulated in terms of carbon budgets (Rogelj et al., 2019).

E.5 Demand Function

World oil demand is assumed to be isoelastic (constant price-elasticity):

$$D_t \;=\; D_{\text{ref}} \left(\frac{P_t}{P_{\text{ref}}} \right)^{\varepsilon_t},$$

where D_t is total demand in year t, P_t is the oil price in year t, $(D_{\text{ref}}, P_{\text{ref}})$ are the observed demand and price in the reference year 2019 and $\varepsilon_t < 0$ is the absolute price-elasticity in year t.

Calibration of demand parameters Utility is coherent with a demand function that is assumed isoelastic and calibrated to 2019[15]. We assume price elasticity to be 0.1 in 2022. To reflect greater long-run adaptability, the absolute price elasticity is increased by 0.005 per year. In other specification, we set the demand elasticity to 0.05, 0.15, 0.2 and 0.3, with or without elasticity increment.

E.6 Social Welfare Maximization

We linearize the consumer-surplus term by sampling the demand curve on a fine grid, computing the exact surplus at each point, and linking these points with a piece-wise-linear constraint. All computations are implemented in *Python*, and the mixed-integer linear programs are solved via the Gurobi Optimizer (Python API), see https://www.gurobi.com/products/gurobi-optimizer/.

F The Global Carbon Policy Context and The Oil Industry

F.1 Carbon-pricing Initiatives in the Oil Industry

The 10 top producers of oil and gas over the 1992—2018 period (Saudi Arabia, Russia, U.S., Iran, China, Mexico, UAE, Canada, Venezuela, Iraq) have limited carbon pricing initiatives. More importantly, with the exception of some attempts in Canada, these initiatives do not cover GHG emissions related to oil extraction.

[15]The year 2019 was chosen instead of 2021 in order to avoid capturing market effects of the COVID pandemic.

Table F1: Carbon-pricing Initiatives in the 10 Largest Oil Producers as of 2018.

Country	Year (start-end)	Sectors
US ETS		
RGGI†	2009-	power
Washington	2017-	industry, power, transport, waste, buildings
Massachusetts	2018-	power
California	2012-	power, road fuel distribution
Canada		
Alberta	2007-17	industry, power
Alberta CCIR	2018-	industry, power, large oil-sands mines
Quebec ETS	2013-	power, industry, distribution, fossil-fuel imports
BC tax	2008-	all sectors except agriculture (from 2013)
Ontario CaT	2017-19	all sectors except agriculture, waste, aviation, marine transport
China ETS		
Shanghai	2013-	power, petrochemicals, aviation, heavy industry
Shenzhen	2013-	power, manufacturing
Tianjin	2013-	petrochemicals, power, oil & gas, heavy industry
Guangdong	2013-	power, cement, steel, petrochemicals
Chongqing	2014-	power, heavy industry
Hubei	2014-	power, heavy industry, petrochemicals
Beijing	2013-	power, heavy industry, petrochemicals
Fujian	2016-	petrochemical, heavy industry
Schenyang	2021-	power
Hubei	2021-	all sectors
Saudi Arabia GCOM‡	2024-	all sectors
Other top-10 producers	Nil	Nil

Note: Production includes the domestic production of crude oil, all other petroleum liquids, biofuels, and refinery processing gain. Data from U.S. Energy Information Administration (2020b) and World Bank (2020). For the British Columbia (BC) carbon tax and Ontario CaT, data are from Murray and Rivers (2015) and Financial Accountability Office of Ontario (2018) respectively. †: The Regional Greenhouse Gas Initiative (RGGI) covers Connecticut, Delaware, Maine, Maryland, Massachusetts, New Hampshire, New York, Rhode Island, and Vermont. ‡: GCOM stands for Greenhouse Gas Crediting and Offsetting Mechanism.

F.2 National-Level Bans on New Fossil Fuel Exploration/Development

Many countries have moved to prohibit new oil, gas, and coal extraction projects. Table F2 summarizes national policies that ban or impose moratoria on new fossil fuel exploration and/or development.

Table F2: National-Level Fossil Fuel Exploration/Development Bans

Country	Ban/Moratorium Details	Date & Status	Source
Canada	Arctic Canadian waters as indefinitely off limits to future offshore Arctic oil and gas licensing	Dec 2016 (Policy)	Prime Minister of Canada (2016)
USA	No new permits for oil and gas in parts of U.S.-owned waters in the Arctic and Atlantic Oceans	Dec 2016 (Law)	BBC News (2016)
France	No new permits for oil, gas and coal exploration; existing production to end by 2040. Fracking banned since 2011.	Dec 2017 (Law)	Government of France (2017)
Spain	Ban on new coal, oil, and gas permits; existing operations end by 2042. Fracking banned.	May 2021 (Law)	Government of Spain (2021)
Denmark	End of oil/gas licensing; fossil extraction ends by 2050.	Dec 2020 (Parliamentary Agreement)	Danish Ministry of Climate, Energy and Utilities (2020)
Ireland	No new oil/gas licenses.	Feb 2021 (Climate Act)	Irish Government (2021)
New Zealand	No new offshore oil/gas permits; existing permits honored.	Apr 2018 (Policy pledge)	Government of New Zealand (2018)
Belize	Permanent offshore oil exploration/drilling moratorium.	Dec 2017 (Law)	Government of Belize (2017)
Greenland	Halt on new oil licenses;	June 2021 (Policy)	Ministry of Mineral Resources and Justice of Greenland (2021)
Portugal	Ended new oil/gas licenses (offshore/onshore).	2018–2020 (Policy)	Government of Portugal (2018)
Sweden	No permits for coal, oil, gas extraction.	July 2022 (Law)	Parliament of Sweden (2022)
Costa Rica	Moratorium on oil/gas exploration and production.	2011–2050 (Executive Decree)	Government of Costa Rica (2011)
Marshall Islands	BOGA pledge, no new fossil fuels permits for extraction or exploration.	Sep 2023 (Pledge)	Beyond Oil and Gas Alliance (2023b)
Samoa	BOGA pledge, no new fossil fuels permits for extraction or exploration.	Nov 2021 (Pledge)	Beyond Oil and Gas Alliance (2023a)
Tuvalu	BOGA pledge, no new fossil fuels permits for extraction or exploration.	Nov 2021 (Pledge)	Beyond Oil and Gas Alliance (2023c)
Vanuatu	BOGA pledge, no new fossil fuels permits for extraction or exploration.	Nov 2022 (Pledge)	Beyond Oil and Gas Alliance (2022b)

F.3 Subnational and Regional Bans/Moratoria

Several subnational jurisdictions have also banned or pledged against fossil fuel exploration as reported in Table F3.

Table F3: Subnational and Regional Fossil Fuel Exploration/Development Bans

Jurisdiction	Ban/Moratorium Details	Date & Status	Source
Quebec (Canada)	Complete ban on oil/gas exploration and extraction, revoked licenses.	Aug 2022 (Law)	Government of Québec (2022)
Wales (UK)	Proposal to ban oil/gas exploration and extraction including fracking.	Sep 2018 (Policy)	Welsh Government (2018)
Scotland (UK)	Ban on fracking	2018 (Policy)	Government of Scotland (2017)
Northern Ireland (UK)	Active ban oil and gas exploration and fracking, planned legislation to ban extraction.	2024 (Legislation)	UK Department for the Economy (2024)
California (USA)	No new fracking permits after 2024, oil phase-out by 2045	2021–2024 (Executive Order)	Office of the Governor of California (2021)
New York (USA)	Legal ban on fracking	2024 (Law)	New York Assembly (2024)
Maryland (USA)	Legal ban on fracking	2017 (Law)	Maryland General Assembly (2017)
Vermont (USA)	Legal ban on fracking	2012 (Law)	Vermont General Assembly (2012)
Washington State (USA)	Fracking moratorium, joined BOGA	2022 (Moratorium & Pledge)	Beyond Oil and Gas Alliance (2022a)
Oregon (USA)	Moratorium on fracking until 2025	2019 (Law)	Associated Press (2019)
Victoria (Australia)	Constitutional ban on fracking and coal seam gas	2021 (Law)	Victorian Legislation (2021)
Yasuní National Park (Ecuador)	Voter-approved ban on oil extraction	2023 (Referendum)	Consejo Nacional Electoral (2023)

G Additional Figures and Tables

Figure G1: Life-cycle GHG Emissions and Private Extraction Cost Per Barrel by Oil Type

(a) Life-cycle GHG emissions per barrel by oil type

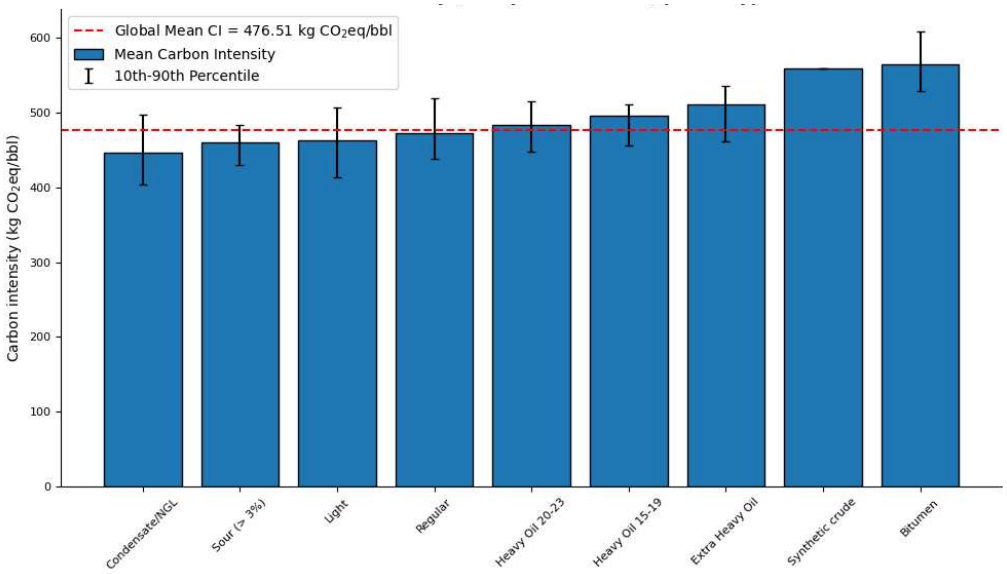

(b) Private extraction cost by oil type

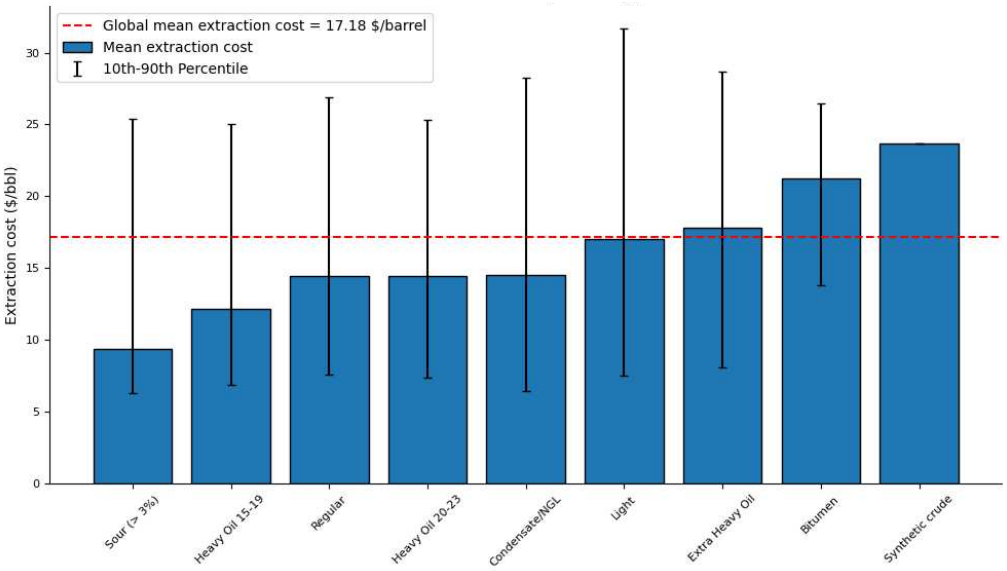

Note: Panel (a) represents the life-cycle GHG emissions per barrel of 2022 reserves by oil type. The bar height represents the average (weighted by 2022 reserves), and the extremities of the lines the 10% and 90% deciles. Panel (b) represents the private extraction cost per barrel of 2022 reserves by oil type. The bar height represents the average (weighted by 2022 reserves), and the extremities of the lines the 10% and 90% deciles.

Figure G2: Life-cycle GHG Emissions Per Barrel from Discovered and Modeled YTF Fields

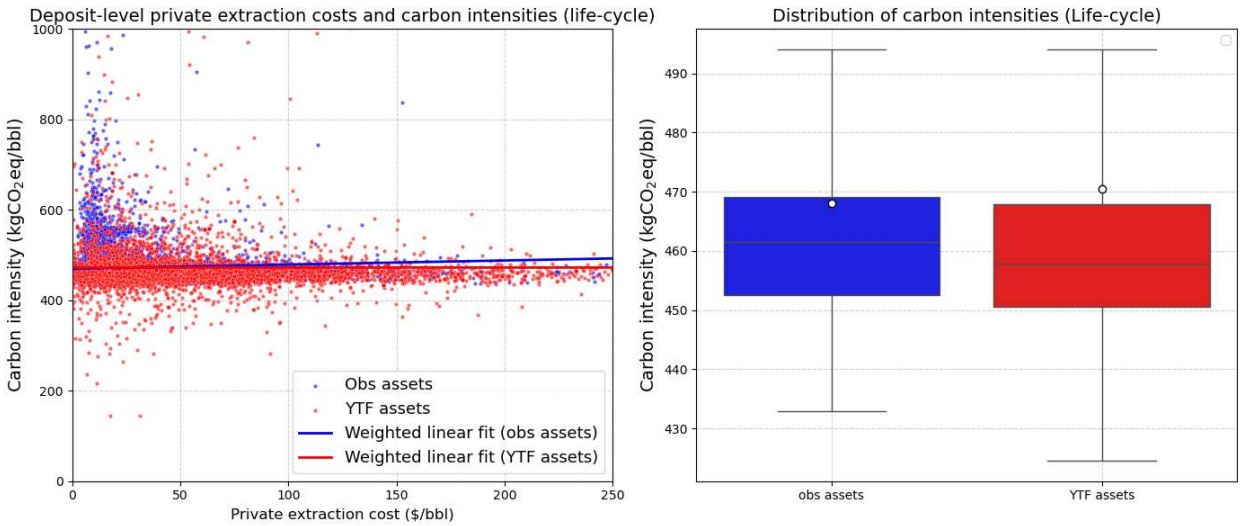

Note: Left Figure: each data point is a deposit with positive reserves in 2021. The red dots are the modeled YTF assets whereas the blue dots are the observed assets from Rystad. The blue line represents the best linear fit for observed assets when deposits are weighted by their 2021 reserves and the red line represent the same linear fit for YTF assets. The correlation coefficients between carbon intensity and private extraction cost are 0.090 and 0.00 when deposits are weighted by their initial reserves and 2021 reserves, respectively. Section 2.3 describes the estimation of carbon intensities, and Appendix E the reserves data and the estimation of the private extraction costs. Fields with either carbon intensities above 1000 or private extraction costs above 250 are excluded from the graph for visibility reasons. Right Figure: the distribution of carbon intensities of the observed and YTF deposits. The white dot represent the average carbon intensity in each category (observed or YTF). The extremities of the lines represent the 10% and 90% deciles and the box itself stands for the interquartile range.

Figure G3: Deposit-level Private Extraction Costs and Upstream Carbon Intensities

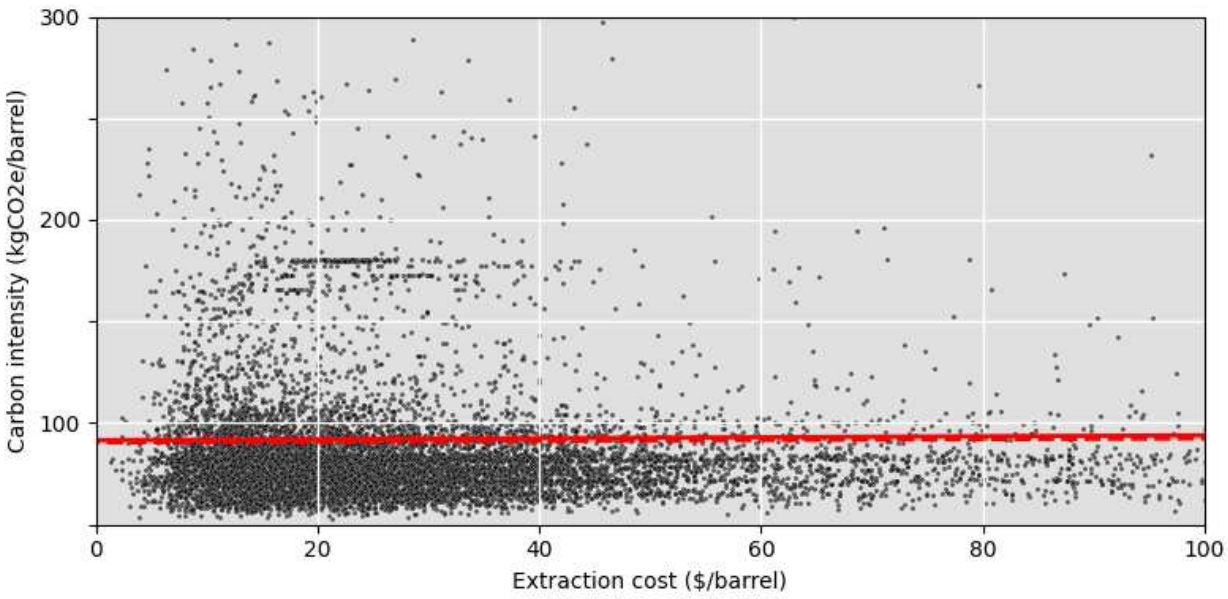

Note: Each data point is a deposit with positive reserves in 2021. The unbroken red line represents the best linear fit when deposits are weighted by their initial reserves (sum of all production), and the dashed red line the best linear fit when deposits are weighted by their 2021 reserves. The correlation coefficients between carbon intensity and private extraction cost are 0.024 and 0.020 when deposits are weighted by their initial reserves and 2021 reserves, respectively. Section 2.3 describes the estimation of carbon intensities, and Appendix E the reserves data and the estimation of the private extraction costs. The figure displays the average costs measure, so that one dot represents one field. Fields with either carbon intensities above 300 or private extraction costs above 100 are excluded from the graph for visibility reasons.

Figure G4: Top-30 Producing Countries from 2017 to 2019: Predicted vs. Observed

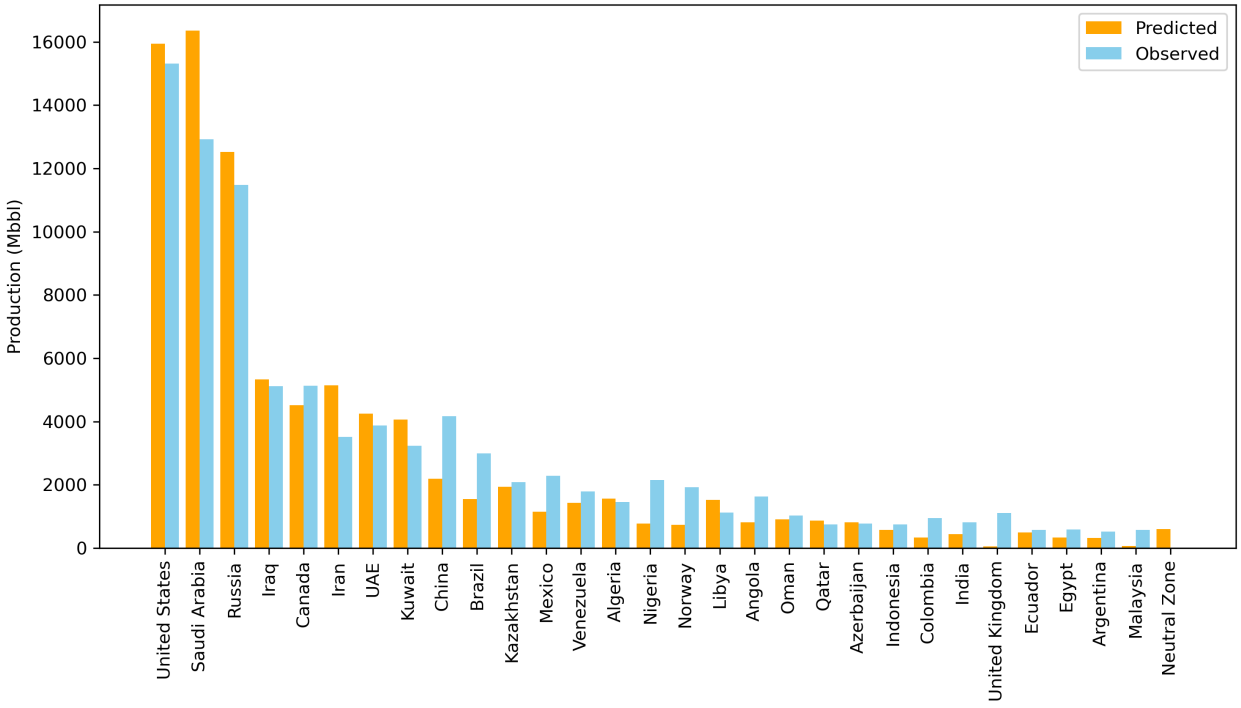

Note: This figure represents predicted and observed country-level total oil production over 2017–2019, measured in million barrels (Mbbl), for the 30 largest oil-producing countries. Predicted productions are the solution of our optimization program initialized in 2017, with no carbon taxes, while retaining distortionary production taxes, preventing a field's production to start before its observed production year in the predictions' counterfactual. As in our main analysis, the demand is calibrated over the year 2019.

Figure G5: Reserves, Carbon Intensity and Private Extraction Cost by World Region

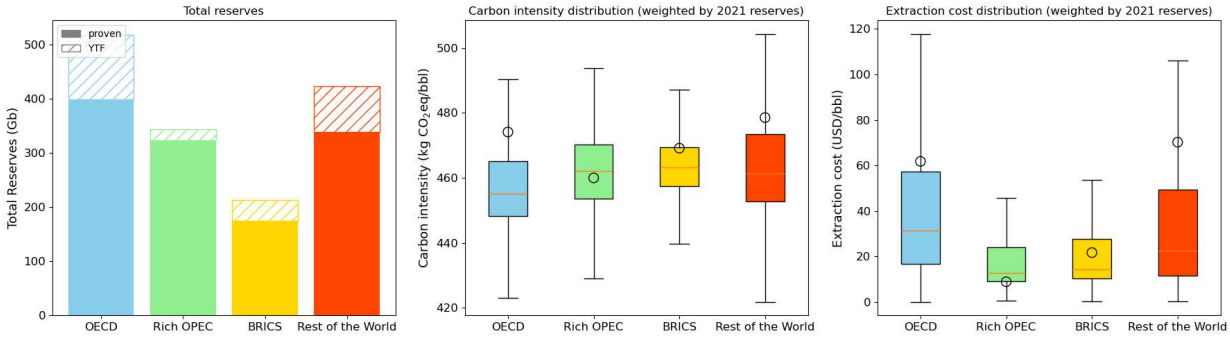

Note: Left panel: Proven and YTF reserves for OECD, rich OPEC countries, BRICS, and the rest of the world. Middle panel: Distribution of carbon intensity per barrel ($kgCO_2eq/bbl$) weighted by 2021-reserves for OECD, rich OPEC countries, BRICS, and the rest of the world. Right panel: Distribution of private extraction costs per barrel (USD/bbl) weighted by 2021-reserves for OECD, rich OPEC countries, BRICS, and the rest of the world. Black circles in box plots represent weighted means.

Figure G6: Discounted Profits by Country Income Level Under Different Bans

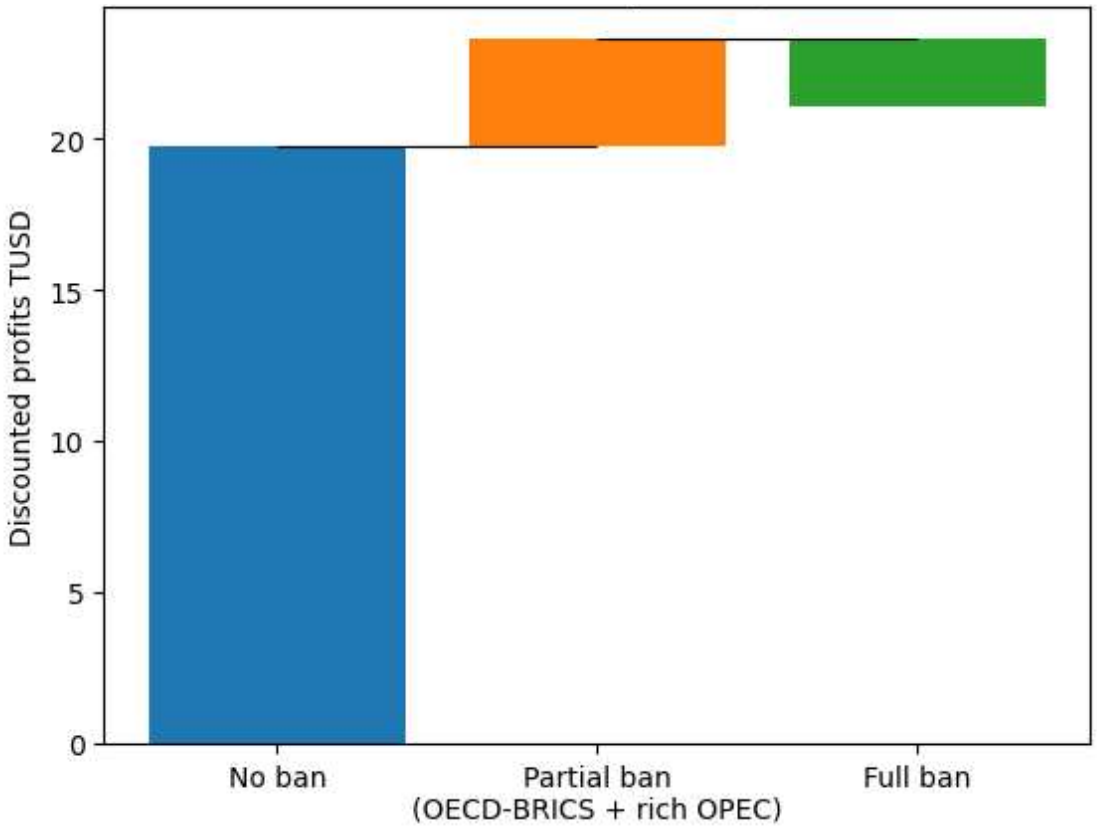

Note: Waterfall bars show changes relative to the immediate scenario on the left; the first bar ("No ban") is the absolute level. Profits are discounted at 3% and reported in trillion USD (TUSD). The figure reports discounted profits for producers outside the OECD and BRICS, excluding rich-OPEC countries that are Saudi Arabia, the UAE, and Kuwait.

Figure G7: Global GHG Emissions as Exploration Bans Extend by Country Income Level under Different YTF Modeling Approaches

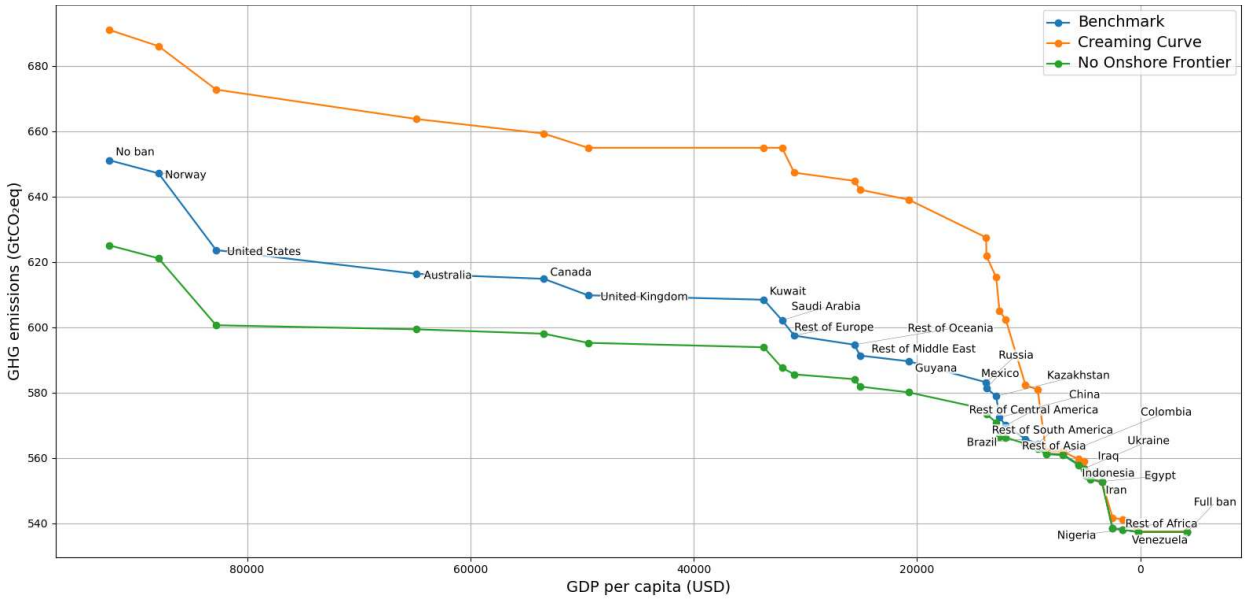

Note: Producers are removed in descending order of GDP per capita ("No ban" → "Full ban"). The x-axis is inverted so richer groups appear on the left. Each point shows total modeled CO_2 emissions ($GtCO_2$) for the corresponding ban. The orange curve estimates basin-level URR by fitting a concave creaming (cumulative-discoveries) curve and using its asymptote as the URR. The green curve follows the same Bayesian setup as the blue benchmark, with the additional assumption that frontier discoveries are restricted to offshore basins.

Figure G8: Country Pollution Potentials: Sensitivity to Pool of Observed Discoveries

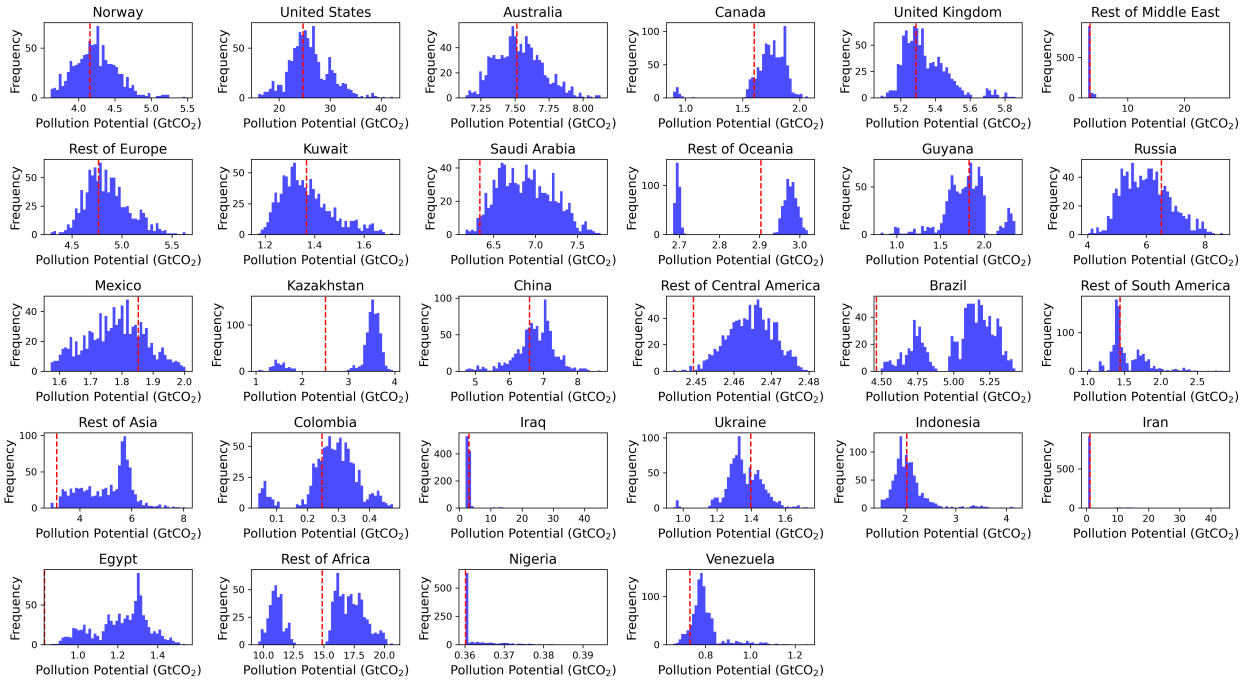

Note: The figure plots the distribution of country-level pollution potentials (in $GtCO_2eq$) from the burning of YTF, using YTF re-estimated using the the approach described in D.5. CI of YTF field are then estimated using the procedure descrived in the main text. The dashed red line indicates our main estimate when all observed discoveries are used for YTF estimations.

Figure G9: Shares of Post-2021 Discoveries That End Up Being Stranded by Country: Carbon Tax Set in 2030

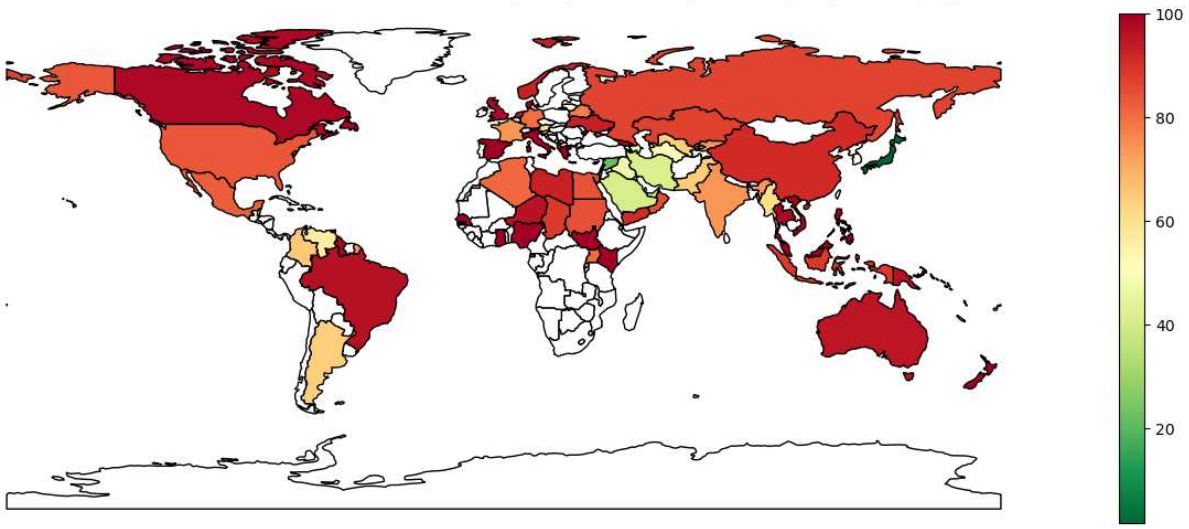

Note: The map illustrates the percentage of post-2021 discovered resources that remain stranded if a carbon tax is introduced in 2030 following a laissez-faire period in which the tax was not anticipated, for each oil-producing country. Countries shown in white are not assigned positive YTF values in our analysis; see Appendix D for details on the construction of YTF.

Table G1: Oil Demand Scenarios

Organization	Scenario	Short description	Source
IEA	NZE (Net Zero Emissions)	Describes a pathway that limits global warming to 1.5°C through immediate, deep emissions cuts and a massive clean energy transition.	IEA (2023)
IEA	APS (Announced Pledges Scenario)	Describes a world where all announced government climate and energy commitments are fully implemented, including net-zero targets not yet backed by concrete policy.	IEA (2023)
IEA	STEPS (Stated Policy Scenario)	Describes a future based on current policies and measures already in place, assuming no additional action.	IEA (2023)
IPCC	Below-1.5°C and 1.5°C low-OS pathways	Describes scenarios that either stay below 1.5°C throughout the century (50–66 % chance) or slightly overshoot (<0.1°C) before returning by 2100 (50–67 % chance).	IPCC (2022)
IPCC	1.5°C high-OS	Describes scenarios where warming overshoots by 0.1–0.4°C before returning to 1.5°C by 2100 with >67 % probability.	IPCC (2022)
Shell	Sky	Describes a rapid, citizen-driven global decarbonization that achieves net-zero by 2050 and reduces warming below 1.5°C after 2075.	Shell (2023)
Shell	Archipelagos	Describes a fragmented, security-first transition where delayed emissions cuts lead to net-zero beyond 2100 and warming near 2.2°C.	Shell (2023)
Equinor	Bridges	Describes a cooperative global transition where structural barriers are overcome through coordinated policies, innovation, and behavioral change, aligning with the 1.5°C goal.	Equinor (2022)
Equinor	Walls	Describes a fragmented world where geopolitical tensions, short-term economic priorities, and limited cooperation delay the energy transition and miss climate targets.	Equinor (2022)
BP	NZE	Describes a highly ambitious global transition with strong policy and societal change, reducing CO_2e emissions by around 95 % by 2050 to align with net-zero goals.	British Petroleum (2023)
BP	Accelerated	Describes a future with significantly tighter global climate policy, reducing CO_2e emissions by 75 % by 2050 through rapid low-carbon adoption.	British Petroleum (2023)
BP	New Momentum	Describes the current policy trajectory, incorporating recent developments and resulting in emissions falling by around 30 % by 2050 compared to 2019.	British Petroleum (2023)
TotalEnergies	Momentum	Describes an uneven global transition led by NZ50 countries and China, where fossil fuel use declines slowly, leading to 2.1–2.2°C warming by 2100.	TotalEnergies (2023b)
TotalEnergies	Rupture	Describes a deeper global decarbonization supported by rich-country aid to the Global South, keeping warming between 1.7–1.8°C by 2100.	TotalEnergies (2023b)

Table G2: Supply Elasticity to Carbon Tax

Tax $/tCO$_2$eq	Total prod. (Gbbl)	Prod. in developed proven (Gbbl)	Prod. in undeveloped proven (Gbbl)	Prod. in post-2021 discoveries (Gbbl)	Discoveries (Gbbl)
0	1379	841	297	241	243
50	1297	806	289	202	218
100	1031	684	238	109	150
150	752	534	155	63	117
200	526	398	95	33	78
250	337	271	53	14	60

Note: Developed and *undeveloped* proven reserves are defined as in Section 3. Production and discoveries represent cumulative volumes over the simulation horizon.

Table G3: Welfare Gains and Costs of a Ban on New Oil Developments

Policy scenario	Private cost (TUSD)	Env. cost (TUSD)	Social cost (TUSD)	GHG ($GtCO_2eq$)	Prod. (Gbbl)	ΔWelfare/1st best (TUSD)
Carbon tax, with explo./dev.	70.2	48.6	118.8	243	526	0.0
Carbon tax, no explo./dev.	80.3	39.8	120.1	199	426	-2.4
No carbon tax, no explo./dev.	59.7	74.0	133.7	370	789	-16.0
No carbon tax, with explo./dev.	29.4	130.2	159.6	651	1379	-35.5

Note: See the Note to Table 2.

Table G4: Welfare Gains and Costs of a Partial Exploration Ban

Policy scenario	Private cost (TUSD)	Env. cost (TUSD)	Social cost (TUSD)	GHG ($GtCO_2eq$)	Prod. (Gbbl)	ΔWelfare/1st best (TUSD)
CT, with explo.	70.2	48.6	118.8	243	526	0.0
CT, no explo. in OECD/BRICS	70.5	48.3	118.8	242	523	-0.1
CT, no explo. in OECD/BRICS/Rich OPEC	70.8	48.1	118.9	240	520	-0.2
No CT, no explo. in OECD/BRICS	33.0	116.8	149.8	584	1235	-27.3
No CT, no explo. in OECD/BRICS/Rich OPEC	33.8	115.0	148.8	575	1216	-26.7
No CT, with explo.	29.4	130.2	159.6	651	1379	-35.5

Note: See the Note to Table 2. Rich OPEC refers to UAE, Kuwait and Saudi Arabia.

Table G5: Welfare Gains and Costs of an Exploration Ban: Oil Producers Collude and the Clean Backstop Cost Decreases

Policy scenario	Private cost (TUSD)	Env. cost (TUSD)	Social cost (TUSD)	GHG ($GtCO_2eq$)	Prod. (Gbbl)	ΔWelfare/1st best (TUSD)
Carbon tax, with explo.	69.2	36.6	105.8	183	397	0.0
Carbon tax, no explo.	69.9	36.1	106.0	180	390	-0.2
No carbon tax, no explo.	32.0	107.4	139.4	537	1138	-33.6
No carbon tax, with explo.	25.3	129.3	154.6	646	1369	-40.0

Note: See the Note to Table 2, and Sections 6.5 and 7.