How does the cost of capital affect oil production?*

Helena Cordt[a] Julien Daubanes[b] Yiding Ma[c]

April 30, 2026

Abstract

Major investors are divesting from the oil-producing industry. In the spirit of green finance taxonomies, reducing carbon-generating businesses' access to financial capital is hoped to contribute to addressing the climate problem. We build a model of oil production driven by project-level net-present-value decisions and calibrate it to all US and Canadian oil projects producible over 2000-2024, capturing projects' earliest production dates, fixed development costs, variable operating costs, and production trajectories. On the one hand, a higher cost of financial capital—isolated from the cost of physical capital—reduces the profitability of oil projects. On the other hand, it generates short-termism, leading the industry to exploit profitable projects at earlier dates. Our counterfactual simulations show that modest increases in the cost of capital are counterproductive, leading to greater production over 25 years. Effective increases in the cost of capital, equivalent to a $25 carbon tax, are unrealistically large. The effect of the cost of capital is inconsistent with individual projects' internal rates of return. This is because, at the industry level, a higher cost of capital generates equilibrium price adjustments that increase oil profitability.

JEL codes: G1; H21; Q31
Keywords: Oil divestment; Green finance; Short-termism; Unintended policy impact; Internal rates of return

*We have benefited from valuable suggestions by Léo Jean, Chris Knittel, Pedro Linares, Gib Metcalf, John Parsons, Lasse Heje Pedersen, as well as by participants at various seminars and conferences: Toulouse Business School, UPPA, NAERE, MIT CEEPR, OECD, Paris School of Economics, University of Copenhagen. The research leading to these results has been supported by the Carlsberg Foundation (2023 Semper Ardens Accelerate grant number $CF23-0609$).

[a]Department of Technology, Management and Economics, Technical University of Denmark (DTU).
[b]Department of Technology, Management and Economics, Technical University of Denmark (DTU), Center for Energy and Environmental Policy Research, Massachusetts Institute of Technology (MIT), and CESifo. Corresponding author: jxada@dtu.dk.
[c]Department of Technology, Management and Economics, Technical University of Denmark (DTU).

1 Introduction

In the past decade, a rapidly growing number of investors—currently representing more than $40 trillion in assets, about one third of all financial assets under management worldwide—have committed to ceasing their financing of the production of oil (Global Fossil Fuel Divestment Commitments Database, 2025). These investors mainly include pension funds, banks, and climate-themed asset managers. Now more than ever, in the spirit of burgeoning green finance taxonomies, it is hoped that less funding for oil production will increase producers' cost of capital, thus partly substituting for carbon pricing in addressing the climate problem. The goal of this paper is to explore how an augmented cost of capital affects oil exploitation decisions and aggregate oil production.

Our main question of interest concerns how the assessment of real oil production projects and producers' profit-maximizing decisions are affected by the cost of capital. This cost, as finance theory explains, is profit-maximizing firms' opportunity cost of funds, which determines the rate at which their projects' future cash flows are discounted.[1] The first and obvious effect of the cost of capital is to reduce the net present value of firms' oil projects. The second effect is that the cost of capital induces producers to give priority to projects that generate cash in the short run.

The resource economics literature has long warned that higher discount rates may increase oil production in the short run—see, e.g., Long and Sinn (1985), Sinn (2008), van der Ploeg and Withagen (2012), among many others. Interest in this short-termist response to bad future prospects was recently revived by Kellogg's (2024) study of the anticipated effect of the long-term destruction of oil demand.[2]

The financial approach to the climate problem—restricting the availability of financial resources for carbon-producing industries—sharply contrasts with economists' traditional call that the carbon externality should be priced directly. Yet, policymakers and financial regulators have rapidly embraced it as a new paradigm. For example, following the EU-ETS carbon pricing scheme, the second pillar of the EU climate action is its "green finance taxonomy"—followed by many similar

[1] Oil projects typically involve both upfront development costs and future benefits that depend on the oil price.

[2] For example, anecdotal evidence suggests that the 2022-2023 rise in interest rates and, thus, in investors' expected returns not only pressured oil majors into reducing their long-term investment in the development of future projects, but also induced them to prioritize "short-cycle" projects—see *The Economist*, February 6, 2023, available at: `https://www.economist.com/business/2023/02/06/where-on-earth-is-big-oil-spending-its-150bn-profit-bonanza`. This short-termist response is consistent with the fact that oil companies' profits were higher in 2022-2023 than in the period 2010-2014, despite lower prices.

initiatives worldwide—aimed to make financial capital more costly for climate-harmful activities.

The literature in finance has recently paid attention to the effect of the cost of capital on carbon-intensive industries. Pedersen (2025) is the first to point to a theoretical equivalence between carbon pricing and green finance, modeled as an increase in firms' cost of capital. Quantifying this equivalence—for recent attempts, see Pedersen (2025), Chittaro et al. (2025), and Gollier (2025)—yields industry-level costs of capital benchmarks, indicating the financial returns that markets should forego to produce incentives in line with climate objectives. To our knowledge, existing attempts have focused on calibrated static marginal cost formulas relating the price of carbon and the cost of the capital input, leaving unexplored the effect of the cost of financial capital on the assessment of real projects' generating cash flows over time, the gap that we seek to fill.

In this paper, we focus on the US and Canadian oil production industry. Oil accounts for approximately one third of global greenhouse gas emissions worldwide and is expected to remain an essential component of the energy mix for decades. At the same time, meeting climate targets implies that a substantial share of currently known and profitable reserves must remain unexploited (McGlade and Ekins, 2015; Welsby et al., 2021; Trout et al., 2022)—e.g., about one third of these reserves, should the objective of a two-degree temperature rise be met. US shale and Canadian assets, among the least profitable, should be among the most impacted ones.

We examine how oil production responds to (i) a carbon price and (ii) an increase in the cost of capital. Since a persistent increase in the cost of capital is without historical precedent, the prediction of industry and market responses must rely on economic theory supplemented by calibrated simulations.

Model. We model intertemporal oil supply decisions and characterize the dynamic equilibrium of the market for US and Canadian oil as a residual demand net of the rest-of-the-world supply. Our model's specification is guided by facts about oil extraction projects which we present in Section 2. Our model focuses on producers' choices of when to undertake their extraction projects, capturing rigidities and production constraints. Each project is characterized by (i) an earliest feasible production date, (ii) fixed development costs incurred before production begins, (iii) an operational cost of extraction per barrel, (iv) finite recoverable reserves, and (v) an exogenously specified production profile once drilling occurs. The key endogenous choice is the initial drilling

date for each asset. Once drilled, the asset produces according to a predetermined extraction path (reflecting technological and geological constraints) until its recoverable reserves are depleted. Given a price trajectory, each field's drilling date, if any, is determined so as to maximize the net present value of its exploitation. Prices are established endogenously on a dynamic competitive market, balancing supply and demand at all times. Our specifications focus the analysis on two main margins determining the impact of policies on oil production: At the extensive margin, producers decide whether or not to exploit an asset; at the intensive margin, they decide whether each produced asset should be started earlier or later.

Besides the dynamic dimension of oil supply and its calibration based on real projects, a noticeable novelty of our model is that it distinguishes the cost of financial capital and the cost of physical capital (e.g., pipelines, platforms, and other required infrastructure) whose costs are determined by the price of material rather than by the cost of funds. Apart from these characteristics, the main assumption of our theory is that supply decisions are solely guided by the net present value of oil extraction projects.[3]

We then calibrate model parameters using production and economic data covering all US and Canadian oil assets, presented in Section 3: We collect, infer, or estimate projects' recoverable reserve, first production year, production path, as well as fixed development and per-barrel operational costs. We consider all oil projects that have been producible since 2000 as candidates for production over 2000-2024.[4] We consider projects started prior to 2000 as exogenous, to account for short-term rigidities, and allow for projects unfinished or untapped by 2024 to pay off over a longer horizon.

Main results. We use our calibrated model to simulate the dynamic competitive equilibrium under various counterfactual policies for the period 2000-2024, including carbon pricing and "green finance," modeled as an augmented cost of financial capital. In our setting, carbon pricing effectively discourages oil production. As expected, US and Canadian oil supply is very sensitive to the

[3]As a matter of fact, major oil companies require that their projects meet so-called "hurdle" rates of return; profitable projects that don't meet hurdle rates are sold to other operators taking over production. See, for a recent example, Total Energy's sale as part of a strategy to focus on more profitable assets (*Reuters*, October 1, 2025, available at `https://www.reuters.com/business/energy/totalenergies-agrees-sell-oil-assets-off-norway-2025-10-01/`). We hold that projects with a positive NPV are systematically undertaken.

[4]Our focus on known oil means that we take the oil industry "technology" as fixed. In the very long-run, technological change may affect the industry's production capacities, an effect that is out of the scope of this analysis.

producer price because a lot of these assets lie relatively high along the worldwide oil marginal cost curve. For example, a tax of $25 per ton of CO2—causing a wedge $14 between the final price of oil and the price earned by producers—reduces total oil production over the 2000-2024 period by around 17%.

By contrast, the cost of capital has little effect on oil supply even when we consider spreads above those observed in financial markets. Perhaps surprisingly, increasing the cost of capital by up to a few percentage points is counterproductive because it generates industry short-termism. Costs of capital that are effective at deterring oil production are unrealistically large. For example, the cost of capital that has the same effect on total production as a $25 ($50) carbon tax is of 23% (86%). The limited effect is partly due to the fact that oil projects' internal rates of return (IRRs) are high in general. Yet, paradoxically, the simulated effects of the cost of capital are at odds both with current estimates of oil projects' internal rates of returns and our own, based on the baseline equilibrium price trajectory. We show that this is due to equilibrium price adjustments.

Oil projects' IRRs rightly indicates the cost of capital that individually turn them unprofitable. By construction, our own estimates do indicate the exact cost of capital that makes a project's NPV nonpositive assuming the price trajectory that would prevail in the baseline. Yet, existing IRR estimates and our baseline IRRs fail to capture the effect of the cost of capital at the aggregate level. For example, IRRs estimated under baseline market conditions indicate that a cost of capital of 14% would discourage as much oil as a $25 carbon price would do. After taking into account the price increase generated by such penalties, this rate becomes 23%. The discrepancy between baseline IRRs and the aggregate effect of the cost of capital becomes larger as we consider higher carbon pricing levels. For example, based on IRRs estimated under baseline market conditions, the cost of capital required to deter oil production as much oil as a $50 carbon price would do is 32%. After incorporating equilibrium price adjustments into the IRR calculation, the policy-equivalent cost of capital rises sharply to 86%.

Related literature and contribution. Our analysis mainly contributes to the recent finance literature comparing green finance with carbon pricing (Pedersen, 2025; Chittaro et al., 2025; Gollier, 2025). The novelty of our study is to assess the effect of the cost of capital using real projects' data in a setting that allows us both to distinguish the cost of financial capital from the

cost of physical capital inputs and to take equilibrium price adjustments into account.

Our paper is also related to a set of recent papers modeling and simulating the dynamic equilibrium of the oil market based on the same oil data as we use here. This literature studies various sources of inefficiency in the actual allocation of oil—see, e.g., Asker et al. (2019), Coulomb et al. (2025), and Asker et al. (2024). Our objective is very different: Our simulation exercises do not seek to assess the real allocation of oil but to compare various policies under an unchanged setting featuring realistic economic responses. Besides, we depart from Hotelling's (1931) conventional model of nonrenewable resource management, in the spirit of Anderson et al.'s (2018) criticism of the Hotellian model for oil supply, recently substantiated by Ahlvik et al. (2026): We consider that oil supply is determined decisions to tap oil assets under production-trajectory constraints. In line with Ahlvik et al. (2026), we also address the peculiarity of shale oil by defining a project at a more granular level allowing more flexibility in the production of tapped resources.

At the intersection between finance and the economics of oil, Erickson et al. (2017) have estimated internal rate of returns of oil projects, based on the same data as we use in this paper. IRRs are often-used to establish projects' profitability according to the cost of capital. Yet we find that, perhaps surprisingly, these rates do not indicate the true effect of the cost of capital at the industry level. The first reason is that an increased cost of capital modifies producers' decision about when to undertake their projects. The second and main reason is that projects' abandonment cause the equilibrium price to increase, leading to an increase in infra-marginal oil projects' IRRs. We find that the IRR of the marginal oil project may be more than doubled once this equilibrium adjustments are taken into account.

Finally, our paper extends the study of supply-side climate policies initiated by Harstad (2012)—see also, e.g., Asheim et al. (2019) and Ahlvik et al. (2026)—to the analysis of financial penalties.

2 Model

Facts about oil projects. Our model's specification and calibration are guided by the following facts about oil extraction projects.

(1) In reality, the economic commitment to undertake an oil exploitation project, including its financing, is typically made at the "asset" level as defined in Rystad Energy data. In the sequel,

we will use interchangeably asset or project to refer to an asset. Assets may include a single well or multiple wells, typically aimed at exploiting the same geological resource. Recoverable reserves are assessed at the asset level, and a project seeks to exploit all of them. Assets are first discovered, then developed, and exploited.

(2) Once an asset is discovered, it may be developed and exploited as soon as possible thereafter or may be delayed until prices and financing conditions warrant exploitation. For example, most oil projects undertaken in the US and Canada over 2000-2024 started to produce with a delay of a few years with respect to the first possible production year we estimate based on their discovery year and type of oil.

(3) After an asset is developed and tapped, its exploitation is virtually uninterrupted; exceptions may be explained by unanticipated demand (e.g., COVID-19 crisis) or supply (e.g., sanctions) shocks.[5] The asset-level production trajectory is governed both by wells' production and by the opening of additional wells, if any. Well-level production is known to be determined by physical factors (Anderson et al., 2018); in general, the dispersion of wells is constrained by pre-determined capacity constraints (e.g., transportation) and by inter-wells competition for pressure.

(4) Production from shale assets feature more flexibility. In an extension, we verify that our results are qualitatively unchanged when we consider that a shale projects are batches of wells (about ten time more granular than an asset on average) instead of assets.[6]

(5) US and Canadian's oil projects do not exactly sell at the WTI oil price, the North-American benchmark, mainly because of assets' heterogeneity in output quality and access to transportation infrastructure. Many sell at a discounted price—e.g., heavy oil from Canadian resources.

Oil supply. Our model of oil supply is based on project-level decisions. We focus our model on the price-taking oil industry.

The oil industry consists of many projects indexed by $i \in \mathcal{I} = \{0, 1, ..., I\}$ that may produce oil at a set of dates $\mathcal{T} = \{0, 1, ..., T\}$. Each project i may only start producing after its earliest production date $\bar{t}^i$, exogenous. For example, there are minimum development delays following discovery. The first production date, denoted $\tau^i (\geq \bar{t}^i)$, generates a fixed development cost $C_F^i > 0$

[5] Our baseline calibration and our counterfactual simulations abstract from these shocks, assuming regular market conditions. We have verified that our main results hold when we calibrate demand at the year level, capturing actual changes in the demand for US and Canadian oil. This extension will appear in an Appendix by September 2026.

[6] This robustness check is discussed in Section 6. It will be presented in detail in an appendix by September 2026.

incurred at the same date. The production cost per barrel extracted, inclusive of quality- and location-discounts is $c_V^i > 0$. The amount of recoverable reserves from project i is $\overline{R}_0^i > 0$.

The novelty of our model is to assume that each project's recoverable oil reserves $\overline{R}_0^i$ can only be produced through an exogenous production path, so that the producers' sole decision is the date of first production τ^i, if any.

Project i's flow of production q_t^i is 0, for all $t < \tau^i$. For all $t \geq \tau^i$, $q_t^i = q_t^i(\tau^i) \equiv \min\left(\bar{q}^i(t-\tau^i), R_t^i\right)$, where $\bar{q}^i(t-\tau^i)$ is an exogenous function of time since first production and $R_t^i \equiv \overline{R}_0^i - \sum_{s \leq t} q_t^i$ is the remaining amount of oil to be recovered from the project.

For a given price trajectory $(p_t)_{t \in \mathcal{T}}$, and a cost of capital r, the industry's problem is to choose first production dates τ^i, one for each field i, were this asset be tapped, such that:[7]

$$\forall i \in \mathcal{I}, \quad \max_{\tau^i \in \mathcal{T} \cup \infty} NPV^i(\tau^i) \equiv \sum_{\mathcal{T}} \left(p_t - c_V^i\right) q_t^i(\tau^i) \left(\frac{1}{1+r}\right)^t - C_F^i \left(\frac{1}{1+r}\right)^{\tau^i}, \tag{1}$$

where $\forall t < \tau^i$, $q_t^i = 0$, and $\forall t \geq \tau^i$, $q_t^i(\tau^i) = \min(\overline{q}^i(t-\tau^i), R_t^i)$. We denote by $\tau^i = \infty$ the decision not to exploit an asset at all, which occurs when the maximum NPV is nonpositive. Therefore: $NPV^i(\infty) = 0$.

The solution to this problem determines the path of industry supplied quantities $(q_t)_{\mathcal{T}} = \sum_{\mathcal{IT}} q_t^i$.

To sum up, taking into account physical production rigidities at the project level—projects being oil assets for conventional oil and oil wells for shale oil—leads us to formulate the oil-supply problem as a choice of projects' start dates, if any. Accordingly, our model captures that oil production responds to policies through two margins. At the extensive margin, producers decide to launch profitable oil projects; at the intensive margin, they decide whether profitable projects should start earlier or later.

A simple example. In order to illustrate the the intensive and extensive margins through which the cost of capital affects supply, consider a single asset with a simple production and cost structure. Assume that the project may only be started at dates $\tau = 0, 1$, against a fixed cost of development C_F and that its production consists of a flow q at the date of development and at the following one ($q = R_0/2$). Given prices, the asset will be tapped at a date that depends on its cost of capital

[7] We disregard major companies constraints such as their hurdle rates of returns, under the rationale that when assets are not deemed sufficiently profitable, they are sold to other operators.

if at all. The graphical example of Figure 1 reflects that (1) high prices at later dates warrant a late development when the cost of capital is low; (2) a higher cost of capital commands that development be advanced (short-termism); (3) a sufficiently high cost of capital makes the project unprofitable (NPV destruction).

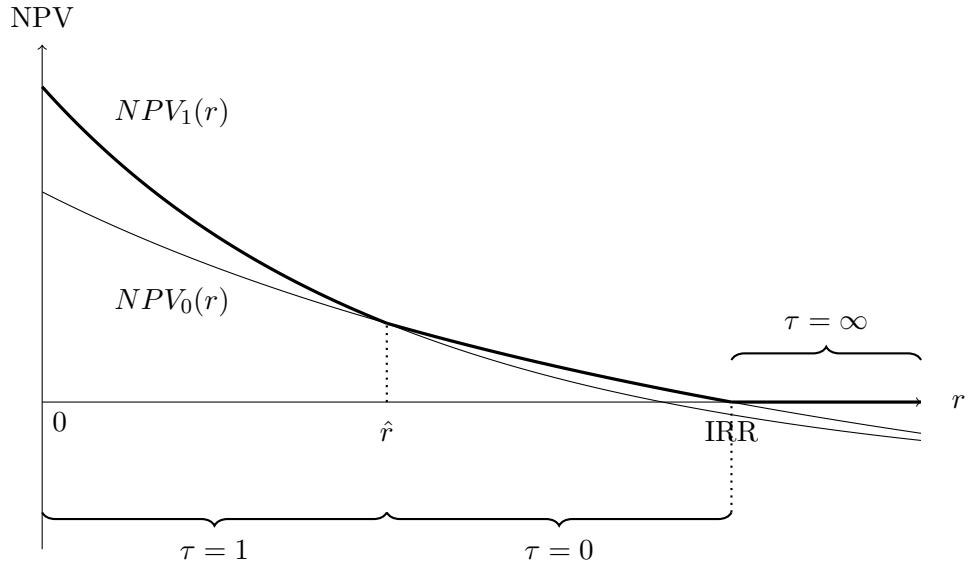


Figure 1: Development timing as a function of the discount rate.

Notes: The downward-sloping curves depicts NPV conditional on the date of development. The thick upper envelope shows the optimized NPV, determining supply decisions. This illustrative example assumes the following parameter values: $C_F = 100$, $c_V = 0$, $q = 1$, $p_0 = 70$, $p_1 = 70$, $p_2 = 80$.

Demand and dynamic equilibrium. Assuming that the demand for oil at each date is given by $D_t(p_t)$, we define a competitive equilibrium as a sequence $\left(p_t, (q_t^i)_{\mathcal{I}}\right)_{\mathcal{T}}$ such that (1) $\sum_{\mathcal{I}} q_t^i = D_t(p_t)$, $\forall t \in \mathcal{T}$, and (2) $(q_t^i)_{\mathcal{I}\mathcal{T}}$ solves problem (1).

Equilibrium selection. In general, for a set price trajectory $p \equiv (p_t)_{\mathcal{T}}$, NPV-maximizing production start dates τ^i*, $i \in \mathcal{I}$, are unique, guaranteeing the existence of an aggregate supply function associating a single vector of supplied quantities $\sum_{\mathcal{I}} q_t^i$, $t \in \mathcal{T}$, to price trajectory p.

This property of supply, however, does not guarantee, the uniqueness of the dynamic equilibrium. Our algorithm of resolution, explained in the appendix, selects the fist equilibrium reached by tâtonnement starting from the constant price trajectory at level \$85 per barrel. Due to discounting, anticipation of a constant price leads the industry to start projects as soon as possible, i.e., to corner timing decisions $\tau^i* = \bar{t}^i$, generating larger excess supply at (exogenous) dates when

more projects becomes producible. Our algorithm adjusts every year's global price accordingly, increasing the price more at dates when less projects become producible.

We have verified that, starting from the assumed $85 price, the resolution algorithm changes the price at each date in a monotonic way, selecting *the* first dynamic equilibrium it meets. We will verify that the baseline equilibrium is robust to alternative starting price levels.

3 Data and calibration

We seek to simulate the effect of policies on US and Canadian equilibrium oil production over the historical period 2000-2024. To take into account short-term production rigidities and to allow for projects unfinished or untapped by 2024 to pay off over a longer horizon, the set of projects that we consider is larger than those producible over our period of interest.

To calibrate our model, we use Rystad Energy data covering all oil fields in the US and Canada. Besides discovery dates, reserves, and production, we estimate the levelized cost of extraction (LCOE for short) of all variable costs associated with production and operation of oil assets, which we use as a proxy for marginal production costs. Following Coulomb et al. (2025), pre-production investment expenditures (well and facility CapEx) are treated as a stream of development costs. We allow for different development-cost structures across shale and non-shale fields: Shale developments involve primarily well-level CapEx, whereas non-shale fields also incur substantial multi-year facility investments. The LCOE and the pre-production costs capture the costs of physical inputs required for extraction—e.g., pipelines, platforms, etc—which are independent of the cost of financial capital.

Projects candidates for endogenous production over this period are all projects untapped by 2000, with earliest estimated production date prior to 2025. In order to take into account that projects may be started years after their earliest production year, our baseline model assumes an additional 10-year "buffer" period, allowing the industry to start projects with a delay, in competition with projects with earliest production year in the 2025-2034 period.

The assets' annual production profile is modeled as a discrete Weibull trajectory, whose parameters are estimated at the asset level by maximum likelihood based on historical production data and Rystad Energy's projections. Additionally, our model captures short-term rigidities by taking as exogenous the continued production, at its actual level, from all projects that were started before

2000 with some production afterwards.

For simplicity, we assume that all assets are equally emission-intensive, generating lifetime emissions of 0.57 tCO2 per barrel, in line with Asker et al. (2024).

Finally, we treat demand for US and Canadian oil explicitly as a residual demand, net of foreign supply, and assume that both the global demand for oil and foreign supply are iso-elastic. Following Kellogg (2024), we assume that the elasticity of global demand is 0.5.[8] Moreover, we assume an elasticity of oil supply of 0.5—for a recent long-run estimate, see, e.g., Vatter et al. (2022). We calibrate global demand and foreign supply based on 2000 and 2024 point observations and Rystad Energy's 2035 forecast, and intrapolate their levels linearly to years 2001-2023 and 2025-2034.[9]

4 Simulation of policy effects

Based on the above calibrated model, we build a resolution algorithm that solves for the dynamic competitive equilibrium of the market for US and Canadian oil over the 2000-2024 period. The equilibrium establishes drilling dates of all projects and, therefore, US and Canada's oil production over time.

In the baseline scenario, we set the cost of capital to 3%.[10] We consider two counterfactual scenarios, both assuming perfect competition and perfect foresight for comparison purposes.

Carbon pricing. In the first scenario, oil traded is subject to a carbon price of $25 per ton of CO2—in 2024 real terms, constant in real present value—while the cost of capital remains unchanged at the baseline level. As in Figure 3, our simulation shows that a $25 carbon price—a wedge of $14.25 per barrel between producer and final prices—reduces total production (and emissions) by 17% over the model period 2000-2024. As expected, the main mechanism of carbon pricing effectiveness is that it generates a lower producer price, depressing the net present value of all assets, and making a number of them unprofitable.

[8] See, for example, Hamilton's (2009) survey of long-run demand elasticity estimates ranging from 0.2 and 0.9, and Krupnick et al.'s (2017) and Balke and Brown's (2018) own estimates close to 0.5.

[9] Our main results survive when we calibrate global demand and foreign supply at the year level and when we take Saudi oil production as a price insensitive component of foreign supply, fixed at its observed level every year, capturing both actual demand variations and OPEC policy.

[10] This baseline level is purposely assumed lower than the industry's WACC because there is no risk in our model.

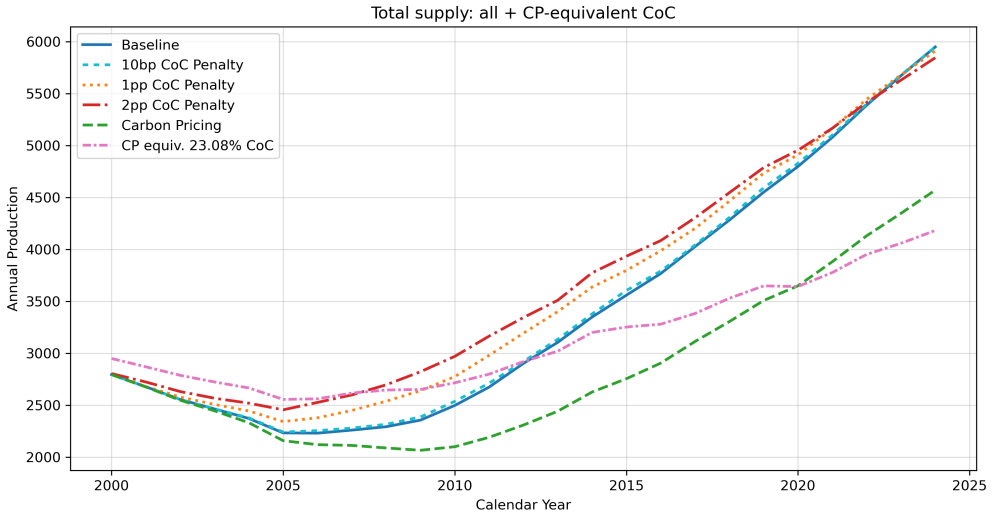


Figure 2: Oil production equilibrium response to cost-of-capital and carbon pricing (2000-2024)

Notes: The figure shows oil production trajectories under alternative policy scenarios. The blue solid line represents the baseline scenario with a 3% cost of capital and no carbon pricing. The teal short-dashed, orange dotted, and red dash–dot lines represent cost-of-capital (CoC) penalty scenarios of 10 basis points, 1 percentage point, and 2 percentage points, respectively, increasing the cost of capital from 3% to 3.1%, 4%, and 5% in the absence of carbon pricing. The green dashed line represents a carbon pricing scenario with a $25 carbon tax (in 2024 present value) and a 3% cost of capital as in the baseline. The pink dash–dot–dot line represents the policy-equivalent cost of capital that discourages the same cumulative production as the $25 carbon tax, corresponding to a 23.08% increase in the cost of capital relative to the baseline. Annual production is measured in million barrels.

Cost of capital. In the second scenario, we examine green finance by increasing the cost of financial capital while imposing no carbon price. We first consider increases in the cost of capital from 50 basis points to 2 percentage points (see Figure 2), which are relatively modest although larger than currently observed green premiums.

In general, our simulations indicate that these costs of capital have a weak effect on production reduction. Moreover, perhaps surprisingly, aggregated production over the model horizon becomes higher than the one in baseline (see Figure 3), which is due to the higher production during the first years (see Figure 2).

While oil asset NPVs are negatively impacted by the cost of capital,[11] firms limit this impact by advancing the dates when they start exploiting their assets. In other words, the rise in the cost of capital makes the industry more short-termist, inducing it to exploit more oil fields immediately. The effect will perversely increase the overall production over the period of interest, resulting in a counterproductive outcome.

[11] As the cost of capital increases, systematically less assets are tapped in total, i.e., over the entire horizon including the 2025-2035 buffer period.

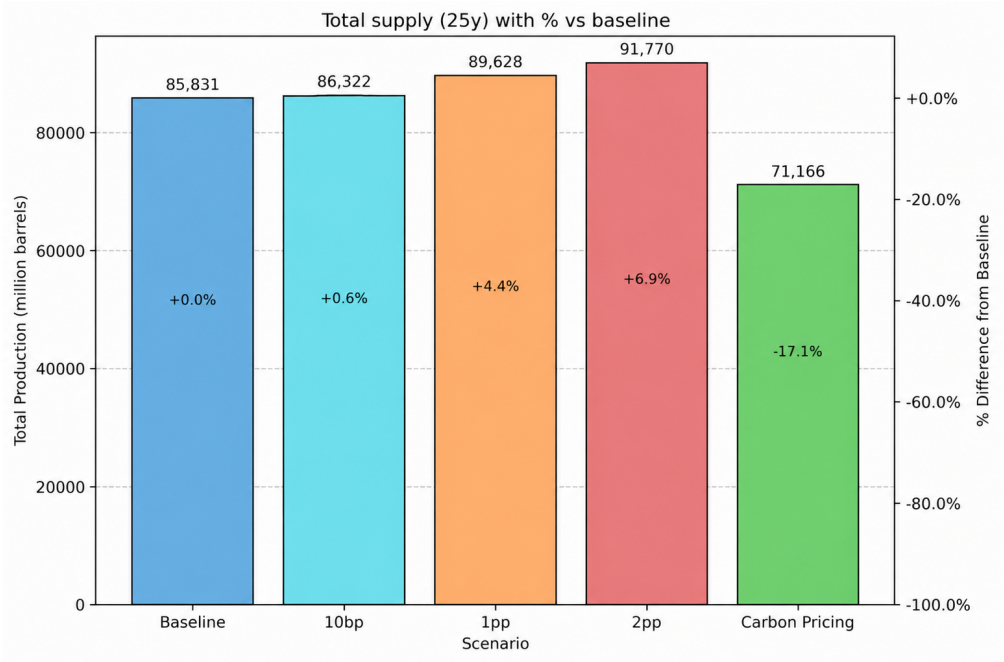


Figure 3: Total oil production response to cost-of-capital and carbon pricing (2000-2024)

Notes: The figure shows a bar plot of total oil production modeled from 2000 to 2024 under alternative policy scenarios. The baseline scenario is with a 3% cost of capital and no carbon pricing. The cost-of-capital (CoC) penalty scenarios of 10 basis points, 1 percentage point, and 2 percentage points represent increasing the cost of capital from 3% to 3.1%, 4%, and 5% in the absence of carbon pricing. The carbon pricing scenario is with a $25 carbon tax (in 2024 present value) and a 3% cost of capital that is same as in the baseline. Total production is measured in million barrels. The second vertical axis represents the percentage difference of scenario total production from the one in baseline.

Green finance becomes effective at discouraging production over the horizon when the cost of capital increases sufficiently to make a substantial number of oil fields unprofitable. Effective green finance calls for high costs of capital, significantly higher than those resulting from the common magnitude of green premiums. For instance, achieving an impact on oil production equivalent to a $25 carbon price would require increasing the cost of capital by 20 percentage points. Similarly, achieving the same effect as a $50 carbon price would require increasing the cost of capital by 83 percentage points. These rates are obviously unrealistically high, which leads us to conclude that an augmented cost of capital is unable to match the effect of oil production on carbon pricing.

5 Why oil projects' IRRs do not inform the effect of the cost of capital

The effect of the cost of capital simulated by our model is at odds with oil projects' conventional IRRs. Oil companies notoriously characterize extraction projects' profitability by their IRRs, defined as the cost of capital that turns a project's NPV negative. They seldom make these calculations public, except for isolated announcements. Analysts' and academic specialists' estimates reflect that returns to oil extraction projects are generally high.[12] High returns to oil extraction projects *partly* explain why an extraordinarily high cost of capital is required to significantly discourage oil production. Yet, they don't fully explain our results.

Conventional IRRs assume fixed producer prices, typically assumed constant in corporate assessments. Because they assume given prices, conventional IRRs indicate the effect of the cost of capital on whether a single, *small* project is profitable or not. However, these metrics fail to indicate the effect of the cost of capital at the industry level. This is because the financial approach to discourage oil production raises the cost of capital at the sector level, causing an adjustment in equilibrium prices.

A simple theoretical example. To illustrate how IRRs may be augmented by equilibrium price adjustments, consider a single *large* project with a purposely simple production and cost structure.

[12] For example, based on its 2024 price assumption, Rystad Energy estimated the IRR of US tight oil to around 35% and the IRR of oil sands to approximately 12%—see *Rystad Energy News*, October 1, 2024, available at `https://www.rystadenergy.com/news/upstream-breakeven-shale-oil-inflation`. See also, Erickson et al. (2017).

Assume that the project is developed at date 0 and that its production consists of a single flow q at date 1. If the project is developed, the market's demand-supply balance establishes the date-1 price $\underline{p}$. When it is not, the market price jumps to $\overline{p} > \underline{p}$. Under price $\underline{p}$, the project's NPV is denoted $\overline{NPV}(r)$. Its root, $\underline{IRR}$, is the threshold cost of capital that would turn the project unprofitable would the price remain $\underline{p}$. Figure 4 depicts how the project's NPV is shifted upward to $\overline{NPV}$ when the cost of capital reaches $\underline{IRR}$, so that the threshold cost of capital that actually turns the project unprofitable is $\overline{IRR}$ instead of the baseline $\underline{IRR}$.

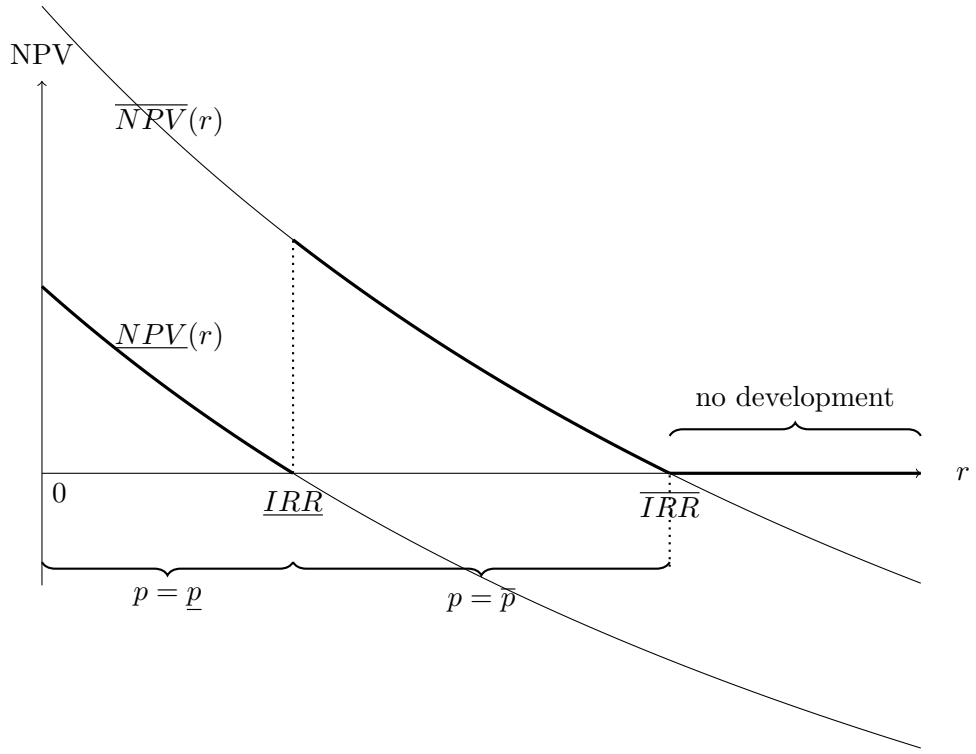


Figure 4: NPV and IRR with price adjustment.

Notes: The upper NPV envelope captures that the large project's abandonment affects the price upward. The baseline IRR does not indicate the effect of the cost of capital on the project's development. Instead, the adjusted IRR indicates the cost of capital that effectively discourages the project. This illustrative example assumes the following parameter values: $C_F = 100$, $c_V = 0$, $q = 1$, $\underline{p} = 120$, $\overline{p} = 150$.

Real oil projects' IRRs with equilibrium price adjustment. To illustrate the discrepancy between real projects' IRRs and the effect of the cost of capital, we calculate project-level IRRs of all candidate oil projects in our simulations, based on fixed baseline market conditions, in line with

the conventional practice.[13] These rates range from 3%, the baseline cost of capital, until more than 100%. They are depicted in Figure 5's blue curve as their distribution over the produced assets' reserves, where we have limited rates to 100% for readability.

This distribution of baseline IRRs alone does not fully account for the effect of the cost of capital indicated by our simulations. For example, based on this distribution, one would conclude that a cost of capital of 14% would make enough assets unprofitable to generate the same effect as a $25 carbon price. However, the simulations presented in Section 4 show that achieving the same effect as a $25 carbon price requires increasing the cost of capital to 23%, rather than 14%. In the case of a $50 carbon price, a similar exercise shows that achieving the same effect through the cost of capital requires increasing its rate to 86%, rather than the 32% suggested by the IRRs' distribution.

This manifest discrepancy between baseline IRRs and the effect of the cost of capital reflects that as we impose higher costs of capital that effectively discourage oil production, the global oil price increases on average, making most oil projects more profitable than when assessed under baseline conditions. Although IRRs adequately indicate the cost of capital that makes an individual project unprofitable, they do not indicate the effect of green finance. To reconcile oil IRRs with the effect of the cost of capital at the industry level, one should adjust their calculation to take into account price equilibrium adjustments caused by the cost of capital and the industry's optimized drilling dates. For example, Figure 5 shows how oil projects' IRRs are modified when they are calculated under the price trajectory resulting from the 23% cost of capital, the rate that impacts oil production as much as a $25 carbon price does. These updated IRRs are indicated by the dots in Figure 5, one for each asset. We verify that they are consistent with the effect of a 23% cost of capital. The abandonment of all projects with updated IRRs below 23% indeed has the same effect as a $25 carbon price.

6 Robustness: Sensitivity and shale-production flexibility

Our main results survive qualitatively and are not significantly affected quantitatively when we modify the parameters of the main analysis or account for additional sources of sensitivity.

[13] One difference with the standard practice, however, is that the IRRs we estimate capture that projects' start production dates are re-optimized as the cost of capital changes.

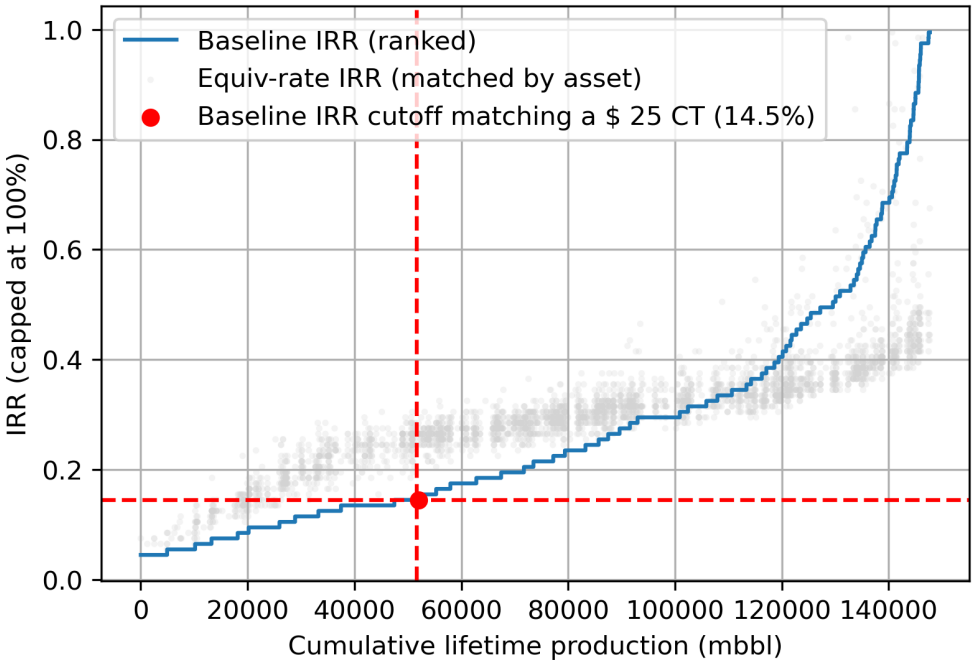


Figure 5: Asset-level IRR and cumulative lifetime oil produced (2000–2024)

Notes: The solid step curve plots cumulative lifetime oil production (in million barrel) ranked by baseline asset-level internal rate of return (IRR) under the baseline equilibrium price and optimal production timing choice τ^i. Grey markers report the IRR for the same assets evaluated under the cost-of-capital that generate an equivalent lifetime supply reduction as a $25 carbon tax. Each marker is positioned at the midpoint of the asset's lifetime production contribution along the baseline cumulative supply axis, ensuring that baseline and counterfactual profitability are compared for identical projects holding lifetime production weights fixed. The sample includes all assets entering baseline equilibrium production between 2000 and 2024 with valid production timing τ^i and positive lifetime output.

Basic changes in the parameters of the main analysis. Our robustness exercises include:

- An increase in the baseline cost of capital to 4%;
- A post-2025 buffer period extended to 20 years;
- A change in the (residual) demand elasticity by $-/+10\%$, regardless of whether the change in responsiveness is due to global demand or other producers' supply;
- An endogenous income taxation scheme calibrated to the US and Canadian tax systems.

Well-level production decisions within shale assets. Finally, we have extended our model to account for the fact that production is more flexible in shale assets than in conventional oil assets. This extension is more complex than the asset-level model of the main text because we assume that shale wells belonging to the same asset are linked in two ways. First, while each group of wells has its own fixed development cost, asset-level fixed development costs (e.g., transportation infrastructure) are attributed to the first group of wells within the asset. The development of the asset, necessary for the drilling of wells therein, is warranted only by the overall profitability of the drilled wells. Second, we impose and calibrate an asset-level production-flow constraint limiting, at each date, the sum of wells' production within each asset.

This sensitivity analysis, including the shale-wells extension, will be presented in the Appendix of the next version of our paper, by September 2026.

7 Conclusion

Concerned investors' sacrifice is hoped to curb the production of carbon resources by raising oil firms' cost of capital. Our analysis shows, however, that such capital penalties have limited effects, and can potentially turn counterproductive in the short run, with unintended effects even over a relatively long period of time.

Our main results can be summarized as follows. First, a relatively modest increase in the cost of capital, even larger than currently estimated green premiums, is counterproductive, increasing oil production over the 2000-2024 horizon. This is because an augmented cost of capital induces industry short-termism. Second, effective increases in the cost of capital that reduce oil production

by the same magnitude as a relatively modest carbon price would be unrealistically large. Third, oil projects' internal rates of return do not account for the effect of the cost of capital on oil supply. This is because effective green finance, unlike carbon pricing, calls for higher producer oil prices, thereby increasing the profitability of oil projects.

Appendix

Construction of the US and Canada Asset Sample

This section documents the data processing steps used to construct the asset-level sample for the United States and Canada. Our analysis uses production and economic data from Rystad Energy, aggregated at the asset-year level from 1970 to 2100. We include static asset characteristics such as discovery year, asset supply group and country. We restrict the country of assets to be the US and Canada.

Both production and economic data are from Rystad's default scenario. For economic variables, we include cost components that associated with asset development and operations. Specifically, we consider well CapEx, facility CapEx, production OpEx, SG&A OpEx, transportation OpEx, and abandonment cost. In addition, we include goverment profit oil, royalty effects, taxes opex, and income tax to capture the effect of fiscal regimes. Nominal economic variables are deflated using the U.S. GDP deflator from the World Bank with base year as 2024, the last year of realized data entries from Rystad Energy as the time of data access. For years on and before 2024, nominal values are converted to real terms using the GDP deflator. For 2025 onward, values from Rystad's default scenario are retained as real values.

To focus on upstream crude oil production, we restrict assets to a predefined set of oil product categories, including light, regular, sour, heavy oil grades, bitumen, condensate, and synthetic crude. Gas, NGL, LNG, biofuels, refinery gains, and other non-oil products are excluded.

For each asset, the first production year is defined as the earliest year with positive production. Observations with zero or missing production throughout 1970 and 2100 are excluded from the sample.

Description of Model Calibration

We define assets for model candidates based on their first production year relative to the model start year (2000). Endogenous assets are those whose first production year is greater than or equal to 2000. We explicitly model the development decision of endogenous assets within the designated horizon. Exogenous assets are those that began production before 2000 and continue to produce after 2000. We treat their post-2000 production as exogenous supply.

The algorithm solves for equilibrium oil production by iterating between asset-level production timing decisions and market-level price adjustment to align demand and supply. Starting from an initial price path, each asset chooses its optimal production start year from the set of feasible years. For each candidate start year, the asset's predetermined production profile and development cost stream are shifted accordingly, and the asset's net present value (NPV) is computed as the sum of discounted operating profits net of variable costs and discounted development costs. The asset selects the start year that maximizes NPV. If the resulting NPV is non-positive, the asset remains undeveloped and contributes no endogenous supply within the model horizon.

Given these asset-level choices, the algorithm aggregates endogenous production across all developed assets and adds exogenous supply to obtain total supply in each period. Demand is then evaluated using the residual demand equation, where the effective price incorporates any carbon-tax component, if applicable. Excess demand is computed as the difference between demand and supply at each model period. If the maximum absolute excess demand across periods falls below the convergence tolerance, the algorithm terminates, returning asset-level development choice and equilibrium price path. If convergence is not achieved, the model computes the inverse-demand price implied by current supply and updates the price path using log-space relaxation. The algorithm then takes the updated price path to start another round of development timing optimization at asset level. This process repeats until the market-clearing condition is satisfied or the maximum number of iterations is reached. The iterative process is represented by the pseudo-code outlined in 1.

Algorithm 1 Equilibrium Oil Production with Asset-Level Timing Choice

Initialize price path $\{p_t^{(0)}\}_{t=1}^{T}$, tolerance ε, and iteration counter $k = 0$
repeat
  for each asset $i \in \mathcal{I}$ **do**
    for each feasible start year $\tau^i \geq \bar{t}^i$ **do**
      Align production profile and development cost stream with τ^i
      Compute net present value $NPV^i(\tau^i)$
    end for
    Choose optimal timing $\tau_i^{(k)}$ that maximizes the NPV
    If $NPV^i(\tau_i^{(k)}) \leq 0$, set $\tau_i^{(k)} = \infty$
    Record production path for assets with a valid $\tau_i^{(k)}$
  end for
  Aggregate endogenous production based on the determined production paths of assets
  Add endogenous supply and exogenous supply to obtain total supply
  Compute demand and excess demand
  if maximum excess demand $< \varepsilon$ **then**
    Stop
  end if
    Update prices using inverse-demand and log-space relaxation
    $k \leftarrow k + 1$
  until convergence within tolerance ε

References

Ahlvik, L., Andersen, J. J., Hamang, J. H., and Harding, T. (2026). Quantifying supply-side climate policies. *Review of Economic Studies*. Forthcoming.

Anderson, S. T., Kellogg, R., and Salant, S. W. (2018). Hotelling under pressure. *Journal of Political Economy*, 126(3):984–1026.

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

Asker, J. W., Collard-Wexler, A., De Canniere, C., De Loecker, J., and Knittel, C. R. (2024). Two wrongs can sometimes make a right: The environmental benefits of market power in oil. NBER Working Paper w33115, National Bureau of Economic Research.

Asker, J. W., Collard-Wexler, A., and De Loecker, J. (2019). (mis)allocation, market power, and global oil extraction. *American Economic Review*, 109(4):1568–1615.

Balke, N. S. and Brown, S. P. (2018). Oil supply shocks and the us economy: An estimated dsge model. *Energy policy*, 116:357–372.

Chittaro, L., Piazzesi, M., Sena, M., and Schneider, M. (2025). Asset returns as carbon taxes. NBER Working Paper 34342, National Bureau of Economic Research.

Coulomb, R., Henriet, F., and Reitzmann, L. (2025). "bad" oil, "worse" oil, and carbon misallocation. *Review of Economic Studies*, page rdaf018.

Erickson, P., Down, A., Lazarus, M., and Koplow, D. (2017). Effect of subsidies to fossil fuel companies on united states crude oil production. *Nature Energy*, 2:891–898.

Global Fossil Fuel Divestment Commitments Database (2025). Global fossil fuel divestment commitments. Accessed April 28, 2025.

Gollier, C. (2025). *Économie de l'(in)action climatique*. Presses Universitaires de France, Paris.

Hamilton, J. D. (2009). Understanding crude oil prices. *The Energy Journal*, 30(2):179–206.

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

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

Kellogg, R. (2024). The end of oil. NBER Working Paper 33207, National Bureau of Economic Research.

Krupnick, A., Morgenstern, R., Balke, N., Brown, S. P., Herrera, A. M., and Mohan, S. (2017). Oil supply shocks, us gross domestic product, and the oil security premium. *Report. Resources for the Future*.

Long, N. V. and Sinn, H.-W. (1985). Surprise price shifts, tax changes and the supply behaviour of resource extracting firms. *Australian Economic Papers*, 24:278–289.

McGlade, C. and Ekins, P. (2015). The geographical distribution of fossil fuels unused when limiting global warming to 2 c. *Nature*, 517(7533):187–190.

Pedersen, L. H. (2025). Can sustainable finance save the planet? *Journal of Finance, forthcoming*.

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

Trout, K., Muttitt, G., Lafleur, D., Van de Graff, T., Mendelevitch, R., Mei, L., and Meinshausen, M. (2022). Existing fossil fuel extraction would harm the world beyond 1.5 °c. *Environmental Research Letters*, 17.

van der Ploeg, F. and Withagen, C. (2012). Is there really a green paradox? *Journal of Environmental Economics and Management*, 64(3):342–363.

Vatter, M. H., Van Vactor, S. A., and Coburn, T. C. (2022). Price responsiveness of shale oil: a bakken case study. *Natural Resources Research*, 31(1):713–734.

Welsby, D., Price, J., Pye, S., and Ekins, P. (2021). Unextractable fossil fuels in a 1.5 degrees celsius world. *Nature*, 597:230–234.