The End of Oil

Ryan Kellogg*

September 6, 2026

Abstract

Even as global oil demand has increased in recent years, it is plausible to envision scenarios in which clean energy technologies or climate policies drive demand to essentially zero by the century's end. This paper asks what such a demand decline, when anticipated, might mean for global oil supply. One possibility is a "green paradox": producers accelerate extraction. However, because extraction requires durable capital investments, the opposite may occur: producers reduce their investment rate, decreasing extraction. To evaluate the relative strengths of these opposing mechanisms, I develop a model of global oil supply that incorporates both, among other industry features. For model inputs with the strongest empirical support, both effects are modest, with disinvestment typically outweighing the green paradox. In order for the green paradox to substantially increase cumulative global oil extraction, investments must have short time horizons, and producers' discount rates must be less than 4% real.

*Harris School of Public Policy, University of Chicago; and National Bureau of Economic Research; kelloggr@uchicago.edu. I thank Christopher Costello, Derek Lemoine, Robert Pindyck, Stephen Salant, Rick van der Ploeg, the Editor Bård Harstad, and four anonymous referees for comments on earlier drafts, and I thank a variety of seminar and conference participants for their comments and suggestions.

1 Introduction

For roughly a century, the vast majority of global transportation demand has been fueled by crude oil and petroleum products. But even as oil demand is projected to continue to grow, albeit modestly, through the remainder of this decade (International Energy Agency, 2024), its long-run dominance as a transportation fuel is now arguably under threat. Alternative technologies — most prominently, electric vehicles — have experienced large cost reductions and commensurate increases in consumer take-up, and climate policies have increased in global scope and intensity. There is considerable uncertainty about how quickly and deeply these developments will reduce crude oil demand, but it is now plausible to envision scenarios in which demand's trajectory shifts to a decline to zero, or at least near zero, by the end of the 21st century. The goal of this paper is to explore how such a decrease in oil demand might affect oil producers' behavior and global oil production during the transition.

My main question of interest concerns how producers' anticipation of a long-run decline in global oil demand might affect their behavior. One possibility is that producers increase their current rates of extraction, via a mechanism that dates to Sinclair (1992) and was coined the *green paradox* by Sinn (2008).[1] This concern is rooted in the idea that oil reserves are an exhaustible resource that producers will extract in a way that maximizes their long-run present discounted value, per Hotelling (1931). Hence, an anticipated future reduction in oil demand will induce producers to accelerate extraction toward the present. This intertemporal shift would at least partially offset any future emissions reductions and perhaps even lead to an increase in the total present value of climate damages.

It is also possible, however, that an anticipated demand decline will cause oil producers to reduce near-term extraction. A suite of papers in economics dating at least to Nystad (1987) and Adelman (1990) has documented that the oil industry is characterized by up-front investments in wells and other infrastructure that enable oil production at a low marginal cost and subject to a binding capacity constraint. If producers anticipate a fall in the future demand for oil, this belief will reduce their incentive to make long-lived investments, leading to lower extraction even in the near term. This potential *disinvestment effect* has drawn considerable industry commentary (Manley and Heller, 2021; Jain and Palacios, 2023; Ryan, 2023; Salzman, 2023), and a recent survey of oil and gas executives (Federal Reserve Bank of Dallas, 2023) found that a majority of respondents believed that an energy transition would increase rather than decrease the price of oil over a five-year horizon.

Disinvestment and the green paradox push in opposite directions, and the goal of this

[1]See Pittel, van der Ploeg and Withagen (2014), Jensen, Mohlin, Pittel and Sterner (2015), and van der Ploeg and Withagen (2015) for reviews of the subsequent literature on the green paradox.

research is to quantitatively assess, using a dynamic equilibrium model of the global oil market, the relative strengths of these anticipatory mechanisms when oil producers are faced with a long-term decline in global oil demand. Given the considerable uncertainties inherent to projecting oil production out to the end of the century, my primary aim is not to compute and defend a single point prediction for what is likely to transpire. Instead, I hope to shed light on what factors — such as the magnitudes of production decline rates, the remaining volumes of reserves, the rates at which producers discount future profits, and the delay before demand is anticipated to begin falling — are most relevant in governing the magnitudes of anticipation effects, as well as the overall reduction in oil extraction associated with a long-run decline in oil demand.[2]

Despite the discussion of disinvestment within the oil industry and trade press, this phenomenon has received, relative to the green paradox, little attention in the academic literature. Some work has modeled coal-fired electricity generator disinvestment in the face of future environmental policy (Lemoine, 2017; Bauer, McGlade, Hilaire and Ekins, 2018; Baldwin, Cai and Kuralbayeva, 2020; Gowrisankaran, Langer and Zhang, forthcoming),[3] and Bruno and Hagerty (2024) models groundwater extraction, finding that the green paradox and disinvestment effects are essentially zero in an open-access setting. Cairns and Smith (2019) and Kollenbach and Schopf (2022) are the only papers I am aware of in the green paradox literature that incorporate investment into a Hotelling model of fossil fuel extraction, though with research questions and approaches that are different from mine in a variety of ways.[4] Another set of papers complements my model-based approach by conducting empirical event studies of oil producers' reactions to changes in the likelihood of future climate policy.[5] Adolfsen, Heissel, Manu and Vinci (2024) and Bogmans, Pescatori and Prifti (2024) study producers' responses to the signing of the Paris Climate Agreement in 2015, drawing opposite conclusions about how the agreement affected investment. Norman and Schlenker (2024) studies a variety of news shocks about increased stringency of future

[2]A code package that allows users to simulate this paper's model with their own parameter assumptions is available at `https://github.com/kelloggrk/Public_EndOfOil`.

[3]McKinley (2024) argues that the incentive to rapidly depreciate generation assets through use is a countervailing force.

[4]Cairns and Smith (2019) considers a hypothetical small price-taking field that requires initial investment in development, whereas I model the global oil market equilibrium. Kollenbach and Schopf (2022) shows that the need for initial exploration investment can mitigate the green paradox, in a setting where the green paradox arises via "leakage" of emissions to an unregulated jurisdiction. Both papers restrict the green paradox by forcing the initial investment to occur at time zero rather than be spread continuously over time and determined endogenously in equilibrium, as is the case here (though Cairns and Smith (2019) allows for endogenously-timed enhanced oil recovery). My approaches to calibration, alternative specifications, and isolating the effects of anticipation are also considerably different from this prior work.

[5]Lemoine (2017) studies coal, finding that a decrease in the likelihood of U.S. climate legislation increased coal prices, but due primarily to a storage mechanism rather than the traditional green paradox.

climate policy, finding that prices fall in most, but not all cases. This result is consistent with a green paradox, though Lemoine (2017) points out that it is also consistent with disinvestment from oil-consuming assets.

This paper's workhorse is a model, presented in section 2, that combines the investment and production dynamics from my earlier work in Anderson, Kellogg and Salant (2018) with Salant's (1976) Nash-Cournot model of the global oil market. Anderson et al. (2018) shows that a model in which firms invest in wells, and then production from those wells declines exponentially with zero marginal extraction cost, can match a large number of industry features. The core logic from Hotelling (1931) remains, but it relates to the timing of drilling investment rather than extraction itself. In this model, the green paradox arises because an anticipated reduction in demand reduces the scarcity rent associated with undeveloped reserves, and disinvestment arises because future demand reductions reduce the anticipated payoff from newly-drilled wells' production streams. I derive comparative statics for how the net of these two effects — the overall effect of an anticipated future demand decline on the initial rate of drilling — is moderated by the production decline rate, the volume of undeveloped reserves, the discount rate, and the delay before the demand decline begins.

I nest the investment model within Salant's (1976) "open-loop" Nash-Cournot model, which features a dominant cartel and a competitive fringe. This model has been used extensively to study dynamic oil extraction games (see Benchekroun, van der Meijden and Withagen (2019), Benchekroun, van der Meijden and Withagen (2020), Balke, Jin and Yücel (2024), and De Cannière (2025) for recent applications). I build on Salant (1976) so that oil producers compete in investment rather than in extraction, and to allow for multiple heterogeneous types of fringe firms.

I use the model to quantify how producers respond to a long-run decrease in global oil demand. As a reference case, I model an affine demand curve that follows a "baseline" trajectory of modest demand growth for 4 years before shifting inward toward zero at a constant rate until the quantity demanded at a price of zero reaches zero after 75 years elapse. I also examine alternatives that involve faster, slower, or incomplete declines. This reduction in demand is meant to capture the effects of increasingly stringent climate policy and improvements in alternative technologies. Because my goal is to study oil supply, my analysis is agnostic to the specific driver of the demand shift, and I do not put any further structure on demand.[6] My focus on long-run decarbonization, motivated by increasing discussion of "zero

[6]My approach implicitly assumes that alternative technologies are not perfect substitutes for oil, since I do not model technology improvement as a decrease in a "backstop price" above which the quantity of oil demanded falls discontinuously to zero, as has been common in past work on oil extraction and climate (see, for instance, Fischer and Salant (2017), Heal and Schlenker (2019), Benchekroun et al. (2020), and van der Meijden, Benchekroun, van der Ploeg and Withagen (2023)). Casual observation of the market for electric

emissions" as a policy goal, is a departure from much of the green paradox literature, which has tended to consider marginal changes in climate policies or the cost of clean substitute technologies.

When I specify and calibrate the model, a guiding principle is that modeling choices and input parameters should be disciplined by established facts about oil production and should rationalize the current (2023) observed global oil market equilibrium. Three choices I make are especially important. First, I model the marginal cost of investing in new wells within a given resource type as an increasing function of the rate of investment. This feature is documented empirically in Kaiser and Snyder (2012), Kellogg (2014), Toews and Naumov (2015), Anderson et al. (2018), and Vreugdenhil (forthcoming), and rationalizes the fact that the industry exhibits simultaneous drilling and production from low-cost and high-cost fields. Simple Hotelling models of competitive producers with heterogeneous but constant marginal costs instead predict that drilling (or extraction in models without investment dynamics) proceeds in strict order from lowest to highest marginal cost.[7] Second, I allow a subset of producers to exert market power, a feature also considered in prior work on the green paradox (see Benchekroun et al. (2019), Curuk and Sen (2023), van der Meijden et al. (2023), and De Cannière (2025) for recent examples). Third, I enforce that in the baseline demand scenario — when producers believe that oil demand is not yet declining — the model reproduces the observed 2023 oil price and observed 2023 investment and production rates across different resources.

I categorize global oil producers as falling into four broad resource types, listed here in order from lowest to highest investment costs: core OPEC, non-core OPEC+, conventional non-OPEC, and shale oil.[8] As discussed in section 3, I calibrate model parameters for each

vehicles implies that EVs and conventional vehicles are far from perfectly substitutable, and Anderson (2012) finds that consumers do not even treat alternative liquid fuels as perfect substitutes for gasoline. Curuk and Sen (2023) also argues for imperfect substitution between oil and alternatives.

[7]See, for instance, Herfindahl (1967), Asker, Collard-Wexler and De Loecker (2019), Heal and Schlenker (2019), and Asker, Collard-Wexler, De Cannière, De Loecker and Knittel (2024). Models in which some producers are strategic players with market power can rationalize simultaneous drilling or production with constant marginal costs, but among the competitive fringe drilling still proceeds sequentially from lowest to highest cost in equilibrium.

[8]An alternative approach would be to model individual fields, as done in Asker et al. (2019), Heal and Schlenker (2019), and Bornstein, Krusell and Rebelo (2023), using cost estimates from consultancies such as Rystad or Wood Mackenzie. I eschew this approach for three reasons. First, the considerable uncertainty about factors such as production decline rates, discount rates, global reserves, and the decline of future oil demand swamps any gain in fidelity that would be achieved by modeling the industry all the way down to the field level. Second, Asker et al. (2019) has shown that field-level cost estimates cannot rationalize observed producers' behavior when applied directly in an equilibrium model: global oil extraction across fields does not occur in strict order of costs as reported by Rystad. The paper attributes this phenomenon to unobserved frictions or to differences between the Rystad data and true marginal costs (and to market power in the case of OPEC producers), though even then increasing costs are still needed to rationalize simultaneous production among competitive firms unless the frictions or differences perfectly equalize costs

type based on academic literature and industry reports. Core OPEC, which includes Kuwait, Saudi Arabia, and the United Arab Emirates, acts strategically, and I model all other types as price-takers. Shale oil is distinguished by its wells' high rate of production decline relative to the other types: 30% rather than 8% per year. The reference case specification assumes that all resource types discount the future at a rate of 9% per year, based on estimates from the literature on oil companies' discount rates. Alternative specifications consider both higher and lower rates.

I present simulation results from the reference case model in section 4.1. Overall, an anticipated 75-year decline in global oil demand to zero reduces cumulative extraction (from 2023 onward) by 25.0% relative to a baseline scenario in which demand grows slowly through 2030 and then stabilizes (and reserves are completely exhausted). Oil is "left in the ground" for all four resource types. During the initial 4-year period before demand begins to decline, I find that drilling investment is 4.3% lower than at baseline, indicating that the disinvestment effect outweighs the green paradox, though it would be fair to characterize neither effect as terribly large. Disinvestment occurs for the three conventional resource types but not for shale oil, since shale wells' production declines quickly and thus investments are short-lived. For all resources, the green paradox is weak in magnitude because the baseline initial scarcity rents are small: all lower than \$4/bbl. These small scarcity rents are consistent with a literature dating to Adelman (1990) that argues that crude oil resources are abundant (Rogner, Aguilera, Archer, Bertani, Bhattacharya, Dusseault, Gagnon, Haberl, Hoogwijk, Johnson, Rogner, Wagner, Yakushev, Arent, Bryden, Krausmann, Odell, Schillings and Shafiei, 2012; Sorrell, Speirs, Bentley, Miller and Thompson, 2012) and thus that exhaustion constraints and scarcity rents have not been quantitatively important for crude oil supply (Hart and Spiro, 2011; Prest, 2022; Prest, Fell, Gordon and Conway, 2024; Lemoine, 2025), and with a literature that documents the poor empirical evidence for scarcity rents in oil markets (Krautkraemer, 1998; Slade and Thille, 2009).

I also quantify the effect of anticipation — the net of the green paradox and disinvestment — on long-run cumulative extraction. Along a declining demand path, oil producers' decisions are affected both by the direct impact of lower current-period demand and by their anticipation of future demand reductions. To isolate the effect of anticipation, I construct drilling and extraction paths where, in each period, producers believe that future demand will not decline from its current state, so that all future demand decreases are unanticipated. The difference in outcomes between this simulation and one in which the demand decline is fully anticipated isolates the effects of producers' anticipation of future decreases in demand.

across fields. Finally, the use of proprietary field-level data would erect barriers — in terms of both data purchase costs and computational time — to future researchers who might want to use this paper's model.

I find that anticipation reduces cumulative extraction by 6.2% relative to the unanticipated scenario. Thus, the effect of disinvestment on cumulative extraction outweighs that of the green paradox in the reference case model. This approach to quantifying the effects of anticipation is an innovation on the green paradox literature, which has heretofore focused on how anticipation affects producers' initial rate of extraction rather than cumulative extraction.

Alternative specifications, presented in section 4.2, explore how the above results are sensitive to model inputs, thereby quantifying the comparative static results from the analytical model. Increasing the rate at which wells' production declines or increasing the delay before the demand decline begins weakens the disinvestment effect, and decreasing producers' initial stock of reserves strengthens the green paradox, but disinvestment remains the stronger force within the range of reasonable parameter inputs. The discount rate turns out to be the most important parameter. Within the range of discount rates that oil producers say, in surveys, that they use when evaluating investments, even large per-barrel marginal profits at the time of resource exhaustion are discounted down to single-digit dollar-per-barrel scarcity rents in present day. To obtain large scarcity rents and hence a green paradox that outweighs the disinvestment effect, I find that I must assume a real discount rate below 4%, which roughly corresponds to the bottom of the range of estimates of the industry's weighted average cost of capital. In particular, if I endow all resource types with a rate of 3%, I find that the initial scarcity rents exceed $25/bbl across all resource types, and that anticipation increases cumulative extraction by 4.8%. Going further, if I also configure the model so that investments pay off immediately, thereby shutting down disinvestment, I find a large green paradox: anticipation increases cumulative extraction by 12.6%. This case is not a realistic fit to the oil industry, but it serves to illustrate how an economically meaningful green paradox can be generated when disinvestment is shut down and producers discount the future at a rate below 4%.

Section 4.3 discusses a broader set of alternative specifications and parameterizations. I show that the speed of the demand decline is important: faster declines reduce cumulative extraction and strengthen the disinvestment effect. I also consider a scenario in which anticipation of the demand decline weakens the OPEC cartel's ability to collude, a la Rotemberg and Saloner (1986), Haltiwanger and Harrington (1991), and De Cannière (2025), providing an alternative mechanism for a green paradox. Anticipation still reduces cumulative extraction on net in this scenario. Finally, I find that adjusting the oil demand elasticity or the shape of producers' marginal drilling cost functions does not substantially affect the results. I conclude in section 5 by discussing features outside the model — such as uncertainty and exploration — that are strong candidates for future research.

2 Model

My goals in specifying this paper's model are to: (1) incorporate a tension between the green paradox and disinvestment; (2) reflect important real-world features of the global oil industry; (3) maintain clear intuition; and (4) allow for fast computation. Some of these goals are in conflict, and I will discuss how I have attempted to balance competing objectives.

I begin in section 2.1 by discussing a simplified version of the model in which oil is supplied by a single representative firm. This model incorporates investment dynamics from Anderson et al. (2018) and is sufficient to set up the paper's core tension between the green paradox and disinvestment. In section 2.2 I then expand the model to feature a dominant cartel and a competitive fringe consisting of a finite number of heterogeneous firm types, building on Salant (1976). This expanded model does not fundamentally alter the intuition behind the green paradox or disinvestment effects, but instead adds value by helping the model capture salient cross-producer heterogeneity.

2.1 Representative agent model

2.1.1 The model

Time is discrete. The global oil price P_t in period t is given by the inverse demand function $P(Q_t, t)$, where Q_t is the rate of global oil consumption (equal to production) at t. $P(Q_t, t)$ is strictly decreasing in Q_t. I include t as a direct input to $P(Q_t, t)$ to allow for the possibility that demand shifts over time.

The supply side of the model draws heavily from Anderson et al. (2018), which in turn is based on evidence that: (1) the industry is highly capital-intensive; (2) production from drilled wells does not respond to oil price shocks but rather declines asymptotically toward zero; (3) the drilling of new wells is price-responsive; and (4) the marginal cost of drilling investments increases in the rate of drilling.[9] Anderson et al. (2018) rationalizes these facts with a model in which drilling investment is the primary choice variable for the firm, and it is optimal to set the rate of oil production from drilled wells at the wells' capacity constraint — which declines toward zero — at essentially all times.[10]

[9]For evidence on the importance of investment rather than production decisions in the oil industry, also see Nystad (1987), Adelman (1990), Thompson (2001), Kellogg (2014), and Newell and Prest (2019). Anderson et al. (2018) further emphasizes that its model reproduces observed dynamics of how drilling, production, and oil prices respond to demand shocks (see figure 8 in that paper). I have verified that the richer model that I calibrate in section 3 reproduces these dynamics as well.

[10]Anderson et al. (2018) presents a model in which both new investment and production from existing wells are choice variables, and goes on to show that the optimal program will almost always involve setting production equal to the wells' capacity constraint. The exception involves situations in which the oil price is anticipated to increase more quickly than the discount rate for an extended interval of time. Such circum-

Let $R_t \geq 0$ denote the stock of undeveloped oil reserves as of period t. Let D_t denote the volume of reserves that becomes developed in period t via investment in the drilling of new wells, so that undeveloped reserves evolve via $R_{t+1} = R_t - D_t$.[11] Wells' production follows an exponential decline at a per-period rate $\lambda \in (0, 1)$. Thus, when a volume D_t of reserves is developed, these reserves produce at a rate λD_t in their first period of production, then $(1 - \lambda)\lambda D_t$ the next period, $(1 - \lambda)^2 \lambda D_t$ the period after that, and so on, so that cumulative production is D_t.

The rate of production Q_t therefore evolves per equation (1) below. Note that this law of motion incorporates a one-period delay between drilling investment and initial production from the newly drilled wells, reflecting the fact that oilfield investments typically require time-to-build.[12]

$$Q_t = \lambda D_{t-1} + (1 - \lambda)Q_{t-1}. \tag{1}$$

The period t drilling cost is $C(D_t)$, with a strictly increasing marginal cost $c(D_t) = dC(D_t)/dD_t$. Per Anderson et al. (2018), I assume that the marginal cost of production is zero. At $t = 0$, producers inherit an initial production rate $Q_0 \geq 0$ (and thus an initial volume of developed reserves Q_0/λ) from wells drilled in the past.

Let $r > 0$ denote the per-period discount rate, with a corresponding discount factor $\beta = 1/(1 + r)$. The planner's problem for maximization of total surplus, which per the first welfare theorem has a solution that leads to the same outcomes as the competitive equilibrium for a representative firm, is then given by:

$$\max_{\{D_t\}} \sum_{t=0}^{\infty} \beta^t \left(\int_0^{Q_t} P(s, t)ds - C(D_t) \right), \tag{2}$$

stances do not arise in this paper, so I take it as given that production is always equal to wells' capacity constraint.

[11] Anderson et al. (2018) defined reserves R_t as the stock of undrilled production capacity rather than the stock of oil remaining to be produced from undrilled wells, and it defined D_t as the amount of capacity investment, rather than the amount of reserve development. I have re-defined these variables here — essentially by dividing them both by the decline rate λ — because this paper is primarily concerned with the quantity of oil developed and produced (and hence burned to form CO_2) rather than capacity installed. In particular, defining R_t in terms of barrels of oil remaining, rather than undrilled production capacity, has the property that the total quantity of oil (and thus CO_2) remaining is invariant to λ, holding R_t constant. This property is useful because I later consider comparative statics with respect to λ, and while doing so I wish to hold constant the total amount of oil available to be extracted.

[12] Newell and Prest (2019) finds an average of 2.7 months elapse between the commencement of drilling and first production for U.S. onshore conventional wells, and 4.5 months for unconventional wells. Time-to-build may be considerably longer for large installations such as offshore platforms. In the computational implementation of the model, each period has a duration of 6 months.

subject to

$$Q_{t+1} = \lambda D_t + (1-\lambda)Q_t,\ Q_0 \geq 0 \text{ given}, \tag{3}$$

$$R_{t+1} = R_t - D_t,\ R_0 > 0 \text{ given}, \tag{4}$$

$$D_t \geq 0,\ R_t \geq 0. \tag{5}$$

2.1.2 Solution

Let θ_t and μ_t denote the current value co-state variables associated with the two state variables: developed reserves Q_{t+1}/λ, and undeveloped reserves R_t. First-order necessary conditions are:

$$D_t \geq 0,\ \theta_t - c(D_t) - \mu_t \leq 0,\ D_t(\theta_t - c(D_t) - \mu_t) = 0 \tag{6}$$

$$(1-\lambda)\theta_t = (1+r)\theta_{t-1} - \lambda P_t \tag{7}$$

$$\mu_{t+1} = \mu_t(1+r) \tag{8}$$

$$Q_t\theta_t\beta^t \to 0 \text{ and } R_t\mu_t\beta^t \to 0 \text{ as } t \to \infty. \tag{9}$$

θ_t represents the time t shadow value of developed reserves at $t+1$ (reflecting the one-period lag between investment and production). Solving the difference equation of FOC (7) (with the endpoint given by the transversality condition of FOC (9)) thus yields that θ_t equals the present value of the future stream of oil prices, starting at $t+1$ and accounting for the exponential production decline at rate λ:

$$\theta_t = \frac{\lambda}{1-\lambda}\sum_{s=t+1}^{\infty} P_s \left(\frac{1-\lambda}{1+r}\right)^{s-t}. \tag{10}$$

With the above understanding of θ_t, FOC (6) has the interpretation that whenever the rate of drilling investment is strictly positive, the difference between the marginal value of newly developed reserves θ_t and the marginal cost of drilling $c(D_t)$ must equal the shadow value of undeveloped reserves μ_t. If the stock constraint $R_t \geq 0$ is binding in the optimal program, then μ_t will be strictly positive and per FOC (8) will increase over time at the discount rate r, following the original insight from Hotelling (1931). Thus, the marginal profit from drilling, $\theta_t - c(D_t)$, must increase at r whenever $D_t > 0$. It is in this sense that this model recasts oil supply as a Hotelling investment timing problem rather than as a Hotelling extraction timing problem. Conversely, if it is optimal to incompletely exhaust the stock (as may happen with declining oil demand), then $\mu_t = 0\ \forall\ t$, and the firm sets $\theta_t = c(D_t)$ at all times when $D_t > 0$. In either case, the optimal path will typically be characterized by a rate of drilling that decreases to zero by some finite time T (assuming demand has a finite choke price), with production continuing on a decline to zero beyond the time at which drilling ceases.[13]

The model is of course an abstraction. In particular, while I think it does a reasonable job of capturing the dynamics of investment in new oil wells and their subsequent production, it is not well-tailored to investments in lumpy, large-scale infrastructure such as deepwater oil platforms or other processing facilities. Nor does it incorporate exploration for new reserves. These investments typically involve long time lags between the initial capital expenditure and the date of first production, suggesting that they are especially prone to disinvestment effects in the face of declining oil demand. Nonetheless, I see two significant advantages from deriving the investment model from Anderson et al. (2018) rather than further enriching the investment dynamics. The first is conceptual: the "shelf life" of new investments, which moderates the disinvestment effect, is captured by a single parameter, λ, clarifying intuition and facilitating both the comparative statics presented in section 2.1.3 below and the sensitivity analyses in the quantitative exercise. Second, the model is fast to compute because it only includes two state variables — developed and undeveloped reserves — with no need to account for the full historical time series of investment. This speed, which is obtained despite incorporating market power and heterogeneous types into the model as discussed in section 2.2, allows me to quantify numerous alternative specifications. It also allows for simulation of outcomes under an unanticipated demand decline, which requires re-solving the model at each time step.

2.1.3 The green paradox and disinvestment

Before turning to the full model in section 2.2, it is useful to pause here to discuss how the representative firm model can capture both the green paradox and disinvestment. Consider a situation in which future oil demand (at some time $t > 0$) is anticipated to be permanently lower than current ($t = 0$) demand. This belief will generate two opposing effects on the initial rate of investment D_0, as illustrated in figure 1.[14] First, anticipation of lower demand will reduce the initial scarcity value μ_0 of reserves. Per equation (6), D_0 is determined by equating the sum of μ_0 and the marginal drilling cost $c(D_0)$ with θ_0, the marginal value of developed reserves. Because $c(D_0)$ is strictly increasing, the decrease in μ_0 will increase D_0, which then increases subsequent extraction and is the mechanism behind the green paradox. Second, the anticipated demand reduction reduces the value θ_0 of developed reserves that

[13] See Anderson, Kellogg and Salant (2014) for a full characterization of the equilibrium in this representative agent model when demand is stationary.

[14] I thank Christopher Costello for suggesting this figure.

Figure 1: Illustration of the green paradox and disinvestment in the representative firm model

(a) Green paradox outweighs disinvestment

(b) Disinvestment outweighs the green paradox

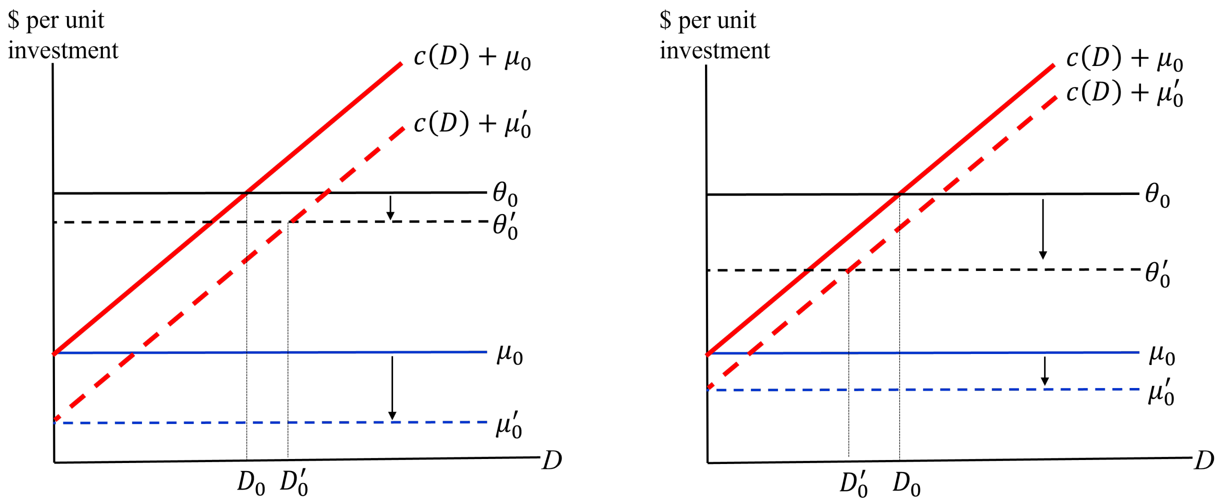

Note: The figure shows the potential effects of an anticipated decrease in demand at some time $t > 0$. The initial scarcity rent μ_0 falls to μ_0', and the initial marginal value of developed reserves θ_0 falls to θ_0'. The change in the initial rate of investment from D_0 to D_0' follows the first-order condition of equation (6), which sets $c(D_0) + \mu_0 = \theta_0$. Depending on the relative magnitudes of the reductions in μ_0 and θ_0, the model can deliver a result that the green paradox outweighs disinvestment (panel (a), $D_0' > D_0$), or the opposite (panel (b), $D_0' < D_0$).

will still be producing after demand has fallen. This reduction in θ_0 will decrease D_0, and hence subsequent extraction, and is the mechanism behind disinvestment.

In panel (a) of figure 1, the anticipated demand decline reduces μ_0 more than it reduces θ_0, so that the green paradox dominates. Panel (b) shows the reverse case. The model can generate either of these outcomes, and the goal of this paper's quantification exercise is to assess conditions under which the green paradox is likely to outweigh disinvestment, or vice versa.

Some of the model's comparative statics are intuitive. If λ is large so that investments are short-lived, then the disinvestment effect will be weak. On the other hand, if reserves R_0 are large so that the initial scarcity rent is small, then the green paradox will be weak. Other comparative statics are less clear. For example, a higher discount rate r will weaken both the green paradox and disinvestment, and which of these effects is weakened more will depend on other parameters. For instance, if λ is large and R_0 is small, then the disinvestment effect is not quantitatively important so that increasing r would primarily impact the green paradox. On the other hand, if λ is small, then assets are long-lived so that changes in r

will strongly affect the extent of disinvestment.

This intuition can be formalized in a parameterized version of the model. Per the computational model that I construct in section 3, let the marginal investment cost function $c(D)$ be affine, with the form $c(D) = \frac{\lambda}{r+\lambda}(c_0 + c_1 D)$, where c_0 and c_1 are both strictly positive scalars. The multiplication by $\lambda/(r+\lambda)$ is a normalization that allows $c_0 + c_1 D_t$ to be interpreted as an amortized per-barrel marginal extraction cost of reserves developed.[15] This normalization ensures that the amortized per-barrel marginal cost of extracting oil, given an investment rate D_t, is invariant to r or λ.[16]

For demand, in the baseline scenario I assume that producers face an exogenous, constant price $P_0 > c_0$. This simplified model can be interpreted as a model of a small, price-taking producer.[17] I then model an anticipated decline in demand as a sequence of price decreases $\delta_t \geq 0$ each period, starting in some period $\tau \in (0, T)$, where T is the last period of strictly positive investment in the optimal program for the baseline scenario. Specifically, $P_t = P_0 - \sum_{s=0}^{t} \delta_s$, where $\delta_s = 0\ \forall s < \tau$, and $\delta_\tau > 0$.

The model then yields the following comparative statics on how λ, R_0, r, and τ moderate how the anticipated demand decline affects the initial rate of drilling investment:

Theorem 1

Let D_0 denote the initial $t = 0$ rate of drilling at baseline, let D_0' denote the initial rate of drilling in the scenario with an anticipated future decline in demand, and define $\Delta D_0 \equiv D_0' - D_0$ (so that $\Delta D_0 > 0$ indicates that the green paradox outweighs the disinvestment effect, and vice versa). Suppose that in the demand decline scenario, the shifters δ_t are large

[15] An up-front payment of $x\lambda/(r+\lambda)$ for developing an additional barrel of reserves is equivalent in present value to paying x per barrel eventually produced from those reserves at the time of extraction.

[16] Absent this normalization, an increase in λ would mechanically cause the amortized per-barrel marginal extraction cost to fall, since extraction occurs more quickly, increasing the optimal investment rate even in the baseline demand (constant price) scenario and complicating the interpretation of the comparative statics. With the normalization, in the baseline demand scenario the initial rate of drilling is invariant to λ, and it is affected by r only through the Hotelling scarcity rent (see appendix A). A related attractive property of this normalization is that, as $\lambda \to 1$ and as the length of each discrete period t goes to zero (so that the per-period discount rate r also goes to zero), the model converges to a standard Hotelling model of extraction, with a marginal per-barrel extraction cost given by $c_0 + c_1 Q_t$.

[17] See appendix A.5 for a discussion of the more general model in which a representative firm faces downward-sloping demand. This model is not sufficiently tractable to derive the comparative static results given in theorem 1, since the Jacobian matrix associated with the FOCs for drilling (equation (6)) is full rather than diagonal, as drilling in any period k affects the value of the FOC for all periods t, not just $t = k$ (this feature is not present in models of extraction rather than drilling investment). However, the constant, exogenous price model (for which the Jacobian is diagonal) is the limit of the downward-sloping demand model as the slope of inverse demand goes to zero, and appendix A.5 shows that for the reference case numerical model that I ultimately implement, the diagonal terms of the Jacobian are much larger than the off-diagonals.

enough that, under the optimal program, the reserves are not fully extracted.[18] *Then:*

(a) $\frac{d\Delta D_0}{d\lambda} \geq 0$, *with the inequality strict if* $\delta_t > 0$ *for any* $t > 1$.[19] *As investments become short-lived, the disinvestment effect is weakened. Also,* $\frac{d\Delta D_0}{d\lambda} \to 0$ *as* $\lambda \to 1$, *since for large* λ *the production decline is so rapid that* $t = 0$ *investments are nearly completely depreciated away by* $t = 2$, *so that further increases in* λ *become immaterial.*

(b) $\frac{d\Delta D_0}{dR_0} < 0$. *As reserves become less scarce, the green paradox is weakened.*

(c) $\frac{d\Delta D_0}{dr}$ *can be positive or negative. For* $\lambda \to 1$, $\frac{d\Delta D_0}{dr} < 0$, *and for* $R_0 \to \infty$, $\frac{d\Delta D_0}{dr} > 0$. *The effect of an increase in the discount rate primarily weakens the green paradox when* λ *is large, and primarily weakens the disinvestment effect when* R_0 *is large. In addition, for the case studied in the quantitative model — a gradual demand decline to zero with a non-trivial production decline rate* λ *— the implicit function* $\Delta D_0(r)$ *will be U-shaped. i.e.,* $\frac{d\Delta D_0}{dr} < 0$ *for small* r, *and* $\frac{d\Delta D_0}{dr} > 0$ *for large* r.

(d) *Suppose that the entire demand decline path is shifted later in time by a small amount* α, *so that the demand shifter for each period* t *is given by* $\alpha\delta_{t-1} + (1-\alpha)\delta_t$.[20] *Then* $\frac{d\Delta D_0}{d\alpha} > 0$. *Postponing the demand decline weakens the disinvestment effect.*

All proofs are in appendix A.

2.2 Full model with a dominant cartel and a competitive fringe

In the full model, there are two classes of producers: a single dominant cartel that acts strategically in the world oil market, and a set of competitive fringe producers who act as price takers. Each fringe type is indexed by $f \in \{1, ..., N-1\}$, and within each type the firms are modeled as a single representative firm. Let m denote the dominant cartel, and let $i \in \{1, ..., N-1, m\}$ index the full set of producers, with d_{it} and q_{it} denoting the investment and production rates of type i at time t. Each type's discount factor, production decline rate, initial undeveloped reserves, and drilling cost function are denoted β_i, λ_i, R_{i0}, and $C_i(d_{it})$. Let $D(t) = \sum_i d_{it}$ and $Q(t) = \sum_i q_{it}$ denote the total rates of drilling and extraction at t, and let q_t and R_t denote the sets of state variables for the fringe firms at t (i.e., $q_t = \{q_{1t}, q_{2t}, ..., q_{N-1,t}\}$).

[18]A necessary condition is that $\sum_{t=0}^{\infty} \delta_t > P_0 - c_0$.

[19]If there is only a one-time drop in demand at $t = \tau = 1$, then a change in λ does not affect incentives since the production decline does not begin until after $t = 1$.

[20]I consider a small amount α rather than a full one-period shift in τ to ensure that the final period of extraction is unchanged in the counterfactual, facilitating the comparative statics.

I assume that the firms' strategies form a Markov perfect equilibrium (MPE), per definition 1 below. This MPE is similar to that proposed in Benchekroun and Withagen (2012), but extended to incorporate capital investment and heterogeneity among the fringe firms.

Definition 1
In the MPE of the dominant firm – competitive fringe investment game:

1. *Each representative fringe firm f chooses a strategy $d_{ft} = \phi_f(t, q_{ft}, R_{ft})$ to maximize*

$$\sum_{s=t}^{\infty} \beta_f^s \left[q_{fs} P_s - C_f(\phi_f(s, q_{fs}, R_{fs}))\right] \tag{11}$$

   at each t, taking as given the time path of prices $\{P_s\}$, its initial production rate q_{ft}, and its initial stock R_{ft}.

2. *The dominant cartel m chooses a strategy $d_{mt} = \phi_m(t, q_{mt}, R_{mt}, q_t, R_t)$ to maximize*

$$\sum_{s=t}^{\infty} \beta_m^s \left[q_{ms} P(Q_s, s) - C_m(\phi_m(s, q_{ms}, R_{ms}, q_s, R_s))\right] \tag{12}$$

   at each t, taking as given the fringe firms' strategies $\phi_f(t, q_{ft}, R_{ft})$ for $f \in \{1, ..., N-1\}$, the initial production rates q_{mt} and q_t, and the initial stocks R_{mt} and R_t.

3. *For all types, q_{is} obeys $q_{i,s+1} = \lambda_i d_{is} + (1 - \lambda_i) q_{is}$, and R_{is} obeys $R_{i,s+1} = R_{is} - d_{is}$. $d_{is} \geq 0$ and $R_{is} \geq 0$ $\forall s$.*

4. *The P_s equal inverse demand: $P_s = P(Q_s, s)$ $\forall s$.*

This definition assumes that fringe firms take the entire price path $\{P_t\}$ as given, rather than just the current price P_t or the rule by which prices are determined by future production quantities. Consequently, fringe firms believe that neither their current decisions nor their future state variables influence the price path. This assumption follows Benchekroun and Withagen (2012) and, as argued therein, is sensible because the alternatives would imply that fringe firms do not account for their choices' impact (small as it may be) on the current price but do account for how their stocks may influence future prices.

An alternative solution concept for the game is an "open-loop" equilibrium in which, at $t = 0$, each firm chooses a time path of investment $\{d_{it}\}$. This version of the game is given by definition 2 below and follows from Salant's (1976) open-loop Nash-Cournot extraction game, extended to incorporate investment and multiple types of fringe producers.

Definition 2

The open-loop equilibrium of the dominant firm - competitive fringe investment game includes parts 3 and 4 of the MPE from definition 1, and it replaces parts 1 and 2 from definition 1 with the below:

1. *Each representative fringe firm f chooses a time path of drilling investment $\{d_{ft}\}$ to maximize*
$$\sum_{t=0}^{\infty} \beta_f^t \left[q_{ft}P_t - C_f(d_{ft})\right] \tag{13}$$
taking as given the time path of prices $\{P_t\}$, its initial production rate q_{f0}, and its initial stock R_{f0}.

2. *The dominant cartel m chooses a time path of drilling investment $\{d_{mt}\}$ to maximize*
$$\sum_{t=0}^{\infty} \beta_m^t \left[q_{mt}P(Q_t, t) - C_m(d_{mt})\right] \tag{14}$$
taking as given its initial production rate q_{m0} and stock R_{m0}, and the fringe firms' initial production rates q_0 and time paths of investment $\{d_{ft}\}$, for $f \in \{1, ..., N-1\}$.

While the open-loop game is restrictive compared to the MPE, the outcomes from these two definitions of the game coincide:[21]

Theorem 2

The outcomes from the MPE (definition 1) coincide with those from the open-loop equilibrium (definition 2).

Proof: In the MPE, because each fringe firm believes that it does not influence the price path $\{P(t)\}$, its strategy $\phi_f(t, q_{ft}, R_{ft})$ devolves to a choice of drilling rate as a function of time $\phi_f(t)$, given $\{P_t\}$, q_{f0}, and R_{f0}. Thus, the fringe firms' Markov strategies can be expressed as open-loop strategies $\{d_{ft}\}$, with $d_{ft} = \phi_f(t)$.

Then from the cartel's perspective, once it conditions on the fringe firms' drilling paths and initial quantities q_0, additionally conditioning its strategy on the fringe's state variables R_t and q_t adds no value. Thus, the cartel's problem can be simplified to a choice of strategy $\phi_m(t, q_{mt}, R_{mt})$, taking the fringe firms' drilling paths $\{d_{ft}\}$ and q_0 as given. But this problem

[21]This equivalence of outcomes with a single strategic player has been proven for a specific case in Benchekroun and Withagen (2012). I have been unable to find a general proof in the literature, though I draw from arguments presented in Eswaran and Lewis (1985).

is then the same as choosing a drilling rate as a function of time $\phi_m(t)$, given $\{d_{ft}\}$, q_{m0}, R_{m0}, and q_0, which is an open-loop strategy.

We therefore have that for both the cartel and fringe firms, the equilibrium Markov strategies can be expressed as open-loop strategies, implying that the outcomes of the Markov perfect and open-loop equilibria will coincide.

The keys to theorem 2 are the presence of only one strategic player and the assumed myopic behavior of the fringe firms.[22] Absent these features, strategic interactions would cause equilibrium outcomes to differ across the two game definitions, as shown in simulations in Eswaran and Lewis (1985) and Benchekroun and Withagen (2012) (though the quantitative differences tend to be minor). Here, however, the fringe firms are effectively solving a single-agent dynamic problem, the solution to which does not depend on whether their strategies take the form of open-loop investment paths or a Markov policy function. Then, from the dominant firm's perspective, it is simply optimizing against a sequence of residual demand curves, so that its optimal investment path is also the same in either characterization of the game.[23]

The equivalence result from theorem 2 is useful because the open-loop game is considerably easier to solve computationally than the game involving Markov strategies.[24] I therefore adopt the open-loop characterization of the game through the remainder of the paper.

In the open-loop equilibrium, the necessary conditions for each fringe firm are the same as those given in equations (6) through (9) for the representative firm version of the model. For the dominant cartel, equation (7) is different because it does not behave as a price-taker. Instead, the equation of motion for θ_{mt} is

$$(1-\lambda_m)\theta_{mt} = (1+r_m)\theta_{m,t-1} - \lambda_m\left(P_t + q_{mt}\frac{\partial P(Q_t,t)}{\partial Q_t}\right), \tag{15}$$

implying that θ_{mt} is the present value of future marginal revenue, accounting for both the discount rate r_m and the production decline λ_m:

$$\theta_{mt} = \frac{\lambda_m}{1-\lambda_m}\sum_{s=t+1}^{\infty}\left(P_s + q_{ms}\frac{\partial P(Q_s,s)}{\partial Q_s}\right)\left(\frac{1-\lambda_m}{1+r_m}\right)^{s-t}. \tag{16}$$

[22]Benchekroun and Withagen (2012) shows that the open-loop equilibrium outcomes are not the limit of the Markov perfect equilibrium outcomes of a game with N strategic firms as $N \to \infty$. Thus, price-taking needs to be assumed *a priori*.

[23]I thank Stephen Salant for thoughts that helped clarify this intuition.

[24]The computational model involves nine state variables for the cartel (the demand state and two states per resource type) and three states for each fringe type, so that solving for the policy functions and Markov perfect equilibrium would be computationally expensive. The open-loop equilibrium, in contrast, simply requires solving for each firm's initial scarcity rent.

Per logic dating to Stiglitz (1976), market power will tend to induce the dominant firm to extract its resource more slowly than a competitive firm would, all else equal.[25] And work on the dominant firm + competitive fringe extraction game has shown that if the resource whose owner exercises market power has a relatively low extraction cost, then that resource may be produced simultaneously with or even after production from higher-cost fringe resources, violating the Herfindahl (1967) principle of sequential extraction in the first-best (see, for instance, Benchekroun et al. (2020)). Relative to this literature, my model: (1) studies competition in investment rather than competition in extraction; (2) includes heterogeneous types of firms within the fringe rather than a single homogeneous type; and (3) includes strictly upward-sloping rather than constant marginal investment costs, so that different firm types will invest simultaneously in equilibrium even in the absence of market power.

A full characterization of the model's equilibrium is beyond the scope of this paper. In general, different marginal cost functions, reserves, decline rates, and discount rates will yield different patterns of extraction, with transitions between phases in which different subsets of resource types are simultaneously drilling.[26] Given the observation that the oil industry is characterized by simultaneous production and drilling across a wide range of heterogeneous fields, the model's implementation in this paper will naturally be steered toward specifications in which all producer types initially exhibit strictly positive drilling in equilibrium. In practice, this outcome is achieved in the model by allowing for marginal drilling costs that increase sufficiently steeply with the rate of drilling that low-cost resources with large reserves do not initially "crowd out" higher-cost resources on the equilibrium path.

2.3 Quantitative implementation of the model

I close this section by summarizing how I quantitatively implement the model on the computer. A more detailed discussion of the computational algorithm is given in appendix B.

I define each period t as having a duration of 6 months.[27] Given a set of input parameters and a time profile of demand, I solve for the equilibrium using two nested loops, similar to the procedure proposed in Salant (1982) for a model of production rather than investment

[25]Stiglitz (1976) shows that if oil demand is log-concave, or if demand is constant elasticity and extraction costs are non-zero, a monopolist will initially extract more slowly than a competitive firm.

[26]Benchekroun et al. (2019), for instance, discusses a variety of equilibrium sequences in a simpler oligopoly-fringe game with constant marginal extraction costs. Depending on costs and reserves, the equilibrium may start with simultaneous supply or with the oligopolists being the sole producers, with subsequent transitions to simultaneous and then oligopoly-only or fringe-only supply.

[27]For the alternative specifications in which I turn off all investment dynamics, I use continuous time. I have verified that the results from the discrete time model approach those of the no-investment continuous time model as the per-period production decline rates λ_i approach one and the duration of each period goes to zero.

choice. The inner loop takes as given the firms' initial shadow values μ_{i0} and solves for the equilibrium series of d_{it}, q_{it}, P_t, and θ_{it} using the FOCs of the open-loop equilibrium, while ignoring the firms' resource stock constraints. This process involves iterating through the d_{it}, q_{it}, P_t, and θ_{it} series until the equilibrium prices converge in the sup norm. The outer loop solves the mixed complementarity problem for the μ_{i0}. That is, the equilibrium initial shadow values, for each firm, either equate cumulative investment to initial reserves or leave the resource constraint slack, with the shadow value equal to zero.

3 Calibration

3.1 Overview

I next quantitatively apply the model developed in section 2 to the global oil market. I categorize producers into four types — core OPEC, non-core OPEC+, conventional non-OPEC, and shale oil — to capture what I view as important heterogeneity in investment time horizons, investment costs, strategic behavior, and potentially discount rates. I summarize the types below, where all production volumes are for 2023 and are compiled from country-level data in Energy Institute (2024) and U.S. shale oil data in U.S. Energy Information Administration (2024c).[28]

I. Core OPEC (15.4 mmbbl/d): This resource type includes three Middle East low-cost, high-output OPEC members: Kuwait, Saudi Arabia, and the United Arab Emirates.[29] I model core OPEC as the only set of producers that acts strategically in the global oil market, following arguments from Balke et al. (2024) that these countries account for most of OPEC's production and have lower costs than other OPEC producers, which "act more like competitive fringe producers", and following findings from De Cannière (2025) that OPEC as a whole colludes only imperfectly.[30] The fact that reserve-to-production ratios are higher for these countries than for other OPEC members (Rystad Energy, 2023) is also consistent with this treatment.

[28] I use Energy Institute's (2024) "crude oil and condensate" production data, which exclude natural gas liquids and refinery processing gain.

[29] Balke et al. (2024) also includes Qatar in its core OPEC group, but Qatar is no longer an OPEC member (Energy Institute, 2024). In 2026, the United Arab Emirates left OPEC (Nereim and Elliott, 2026). The impact on the simulation results of the UAE leaving core OPEC would be the reverse of the impact of adding Russia to core OPEC (which I consider in section 4.3), albeit smaller in magnitude.

[30] De Cannière (2025) estimates an OPEC conduct parameter of 0.17 (where a value of 1.0 signifies perfect collusion). My model implicitly assumes a conduct parameter of 1.0 for core OPEC; translating De Cannière's (2025) parameter to just core OPEC suggests a conduct parameter of 0.36. When I simulate the model with this conduct parameter, I find that the results move toward an alternative specification in which core OPEC behaves competitively (conduct parameter = 0), shown in row 17 of table 3.

II. Non-core OPEC+ (30.4 mmbbl/d): This group includes all other OPEC+ members.[31] Relative to core OPEC, I model these producers as behaving competitively and having somewhat higher investment costs.

III. Conventional non-OPEC, including deepwater (28.5 mmbbl/d): These producers have still higher investment costs than non-core OPEC+.

IV. Shale oil producers (8.4 mmbbl/d): This resource type has the highest investment costs and is distinguished by production from drilled wells having rapid decline rates.

To calibrate the model, my approach is to first set the demand curve and each type's decline rate, discount rate, and reserve volume based on estimates from the literature and industry sources. Then, I calibrate each type's investment cost function so that the baseline simulation of the model — in which global oil demand does not decline — reproduces observed 2023 production and investment for each type.

I focus the discussion below on what I will refer to as the "reference case" calibration. The calibrated parameters for the reference case are summarized in table 1. Section 4 will later present results from an extensive set of alternative specifications.

3.2 Global oil demand

I model global oil demand using an affine demand curve that passes through the 2023 market outcome: production and consumption of 82.8 mmbbl/d, with a price of $82.49/bbl (U.S. Energy Information Administration, 2024b).[32] I choose affine demand, rather than constant elasticity, to reflect intuition that alternative liquid fuels should be able to fully substitute for oil at a high but finite price, and that quantity demanded will not go to infinity as the price falls to zero.

To calibrate the slope of the demand curve, I refer to the literature on the long-run elasticity of demand for oil. A recent survey of this literature is provided in Prest et al. (2024). Most of the recent studies cited therein study only short-run oil price fluctuations, but Krupnick, Morgenstern, Balke, Brown, Herrera and Mohan (2017) and Balke and Brown

[31]OPEC+ members (other than members of OPEC itself) include Azerbaijan, Bahrain, Brunei, Kazakhstan, Malaysia, Mexico, Oman, Russia, South Sudan, and Sudan (U.S. Energy Information Administration, 2023). Energy Institute (2024) includes Angola in its OPEC category, but I drop it because Angola left OPEC at the end of 2023 (Reed, 2023). Russia could arguably instead be included in the core OPEC group, as it is the only OPEC+ member whose production rate is comparable to that of Saudi Arabia. However, Russia's invasion of Ukraine in 2022 has sharply curtailed its access to capital and technology, implying that it does not share core OPEC's low investment costs. Still, a natural sensitivity that I later examine is the inclusion of Russia in the core OPEC producer type.

[32]I use the 2023 average Brent crude oil spot price.

(2018) study the long-run elasticity, finding estimates of -0.53 and -0.51, respectively. An earlier survey (Hamilton, 2009) cites long-run elasticity estimates ranging from -0.21 to -0.86. In the reference case model, I set the slope of the demand curve so that its elasticity is -0.5 at the 2023 price and quantity demanded. In alternative specifications I use elasticities with lower and higher values at this point.

I model modest demand growth in the baseline demand scenario, and a gradual fall toward zero in the demand decline scenario. At baseline, I assume that demand increases (i.e., shifts right) at 0.44% per year through 2030 and remains stationary thereafter (and that producers anticipate this trajectory), per the International Energy Agency's (2024) recent forecast and a similar projection from BP's (2024) "current trajectory" model. Modest growth is also consistent with the median forecast from Federal Reserve Bank of Dallas (2023). I also evaluate alternative specifications that involve no demand growth or a higher rate of growth (1.5% per year through 2030) per Organization of the Petroleum Exporting Countries (2023).

For the demand decline scenario, I specify that global oil demand follows the baseline trajectory for 4 years before commencing to shift inward toward zero at a constant rate, relative to the baseline demand path, holding the slope constant.[33] Oil demand reaches zero quantity demanded at any price after 75 years (including the initial 4-year delay period). The 4-year delay corresponds to the policy lag of the U.S. "Waxman-Markey" climate legislation studied in Lemoine (2017) and Norman and Schlenker (2024), and the 75-year overall decline leaves zero oil demand by roughly the end of the century. Alternative specifications consider other demand decline paths.

3.3 Production decline rates

When calibrating wells' production decline rates λ_i, I distinguish between conventional wells (resource types I, II, and III) and shale wells (type IV). Anderson et al. (2018) directly estimates an annual decline rate for onshore Texas conventional wells of 8%. Thompson (2001) finds that annual production from U.S. wells tends to consistently be about 11% of reserves. For Saudi Arabia, I am not aware of estimated decline rates in the academic literature, though Lynch (2019) states that the decline rate for its largest field, Ghawar, is thought to be 8%, and Gnana (2022) cites a statement by Saudi Aramco's CEO that the global natural decline rate is 7%. International Energy Agency (2025) estimates a global natural decline rate of 8.5%. Taking these estimates together, my reference case uses an

[33] More precisely, I start with the 2023 demand curve and compute a time path that decreases the demand intercept by the same amount per period, starting in 2027, such that demand reaches zero after 75 years elapse. I then increase demand by 0.44% per year through 2030 (demand still decreases from 2027–2030 on net because the rate of decline exceeds the baseline growth rate).

Table 1: Calibrated parameters for the reference case model

Parameter	Notation	Units	Value by type I	II	III	IV
Demand parameters						
Initial 2023 oil price	P_0	\$/bbl	82.49			
Initial 2023 consumption	Q_0	mmbbl/d	82.8			
Elasticity at (P_0, Q_0)			-0.5			
2023 choke price*		\$/bbl	247.47			
Growth rate through 2030		per year	0.0044			
Supply parameters						
Initial 2023 production	q_{i0}	mmbbl/d	15.4	30.4	28.5	8.4
Production decline rate	$1-(1-\lambda_i)^2$	per year	0.08	0.08	0.08	0.30
Discount rate	$(1+r_i)^2-1$	per year	0.09	0.09	0.09	0.09
Total reserves	see note	bn bbl	394	542	563	125
Initial 2023 investment*	d_{i0}	bn bbl/year	5.95	11.71	10.96	3.12
Drilling cost intercept	c_{0i}	\$/bbl	5.00	10.00	20.00	30.00
Drilling cost slope	c_{1i}	see note	7.68	6.00	5.57	16.27
Drilling cost elasticity*			0.90	0.88	0.75	0.63

Note: * indicates that the parameter is derived from the others. Type I is core OPEC, type II is non-core OPEC+, type III is conventional non-OPEC, and type IV is shale. The demand curve and marginal drilling cost curves are affine. λ_i and r_i are defined as per-period rates, and there are 2 periods per year in the model. Undeveloped reserves R_{i0} are equal to the shown total reserves minus initial developed reserves $\frac{Y}{1000}\frac{q_{i0}}{\lambda_i}$, where Y denotes the number of days per period. The drilling cost elasticity is evaluated at the 2023 investment rate for each type. "mmbbl/d" denotes millions of barrels per day, and "bn bbl" denotes billions of barrels. The units for the drilling cost slopes are \$/bbl per billion bbl developed per year.

annual decline rate of 8% for resource types $i =$ I, II, III.

Shale wells are well-known to exhibit production declines that are initially steeper than for conventional wells. Jacobs (2020) finds that shale wells exhibit a decline rate of at least 50% in their first year of production and then 30% in year two. The annual decline rate then falls to 17% by year 5 and below 10% after 8–10 years. My reference case model approximates this nonlinear decline with an exponential decline curve with an annual decline rate of 30%.

3.4 Reserves

Oil reserves are notoriously difficult to estimate, as doing so requires a forecast of future resource discoveries and the economic viability of those discoveries at future oil prices. I therefore explore using a range of reserve volumes in various specifications of the model.

I employ country-level crude oil reserve estimates from Rystad Energy (2023). The reference case uses Rystad's "2PCX" reserve numbers, which include reserves from existing fields, contingent resources in discovered but undeveloped fields, and "prospective resources in yet undiscovered fields". Thus, these values capture the spirit of what the resource stocks represent in the model: the ultimate exhaustible volume of reserves. Rystad reports 2PCX reserves of 394, 542, 563, and 125 billion bbl for resource types I, II, III, and IV, respectively, with total global reserves of 1624 billion bbl (implying a reserves to current production ratio of 54 years).[34]

As an alternative specification with lower reserves, I use Rystad's "2PC" estimate of 1283 billion bbl, which excludes prospective resources in undiscovered fields. That said, Rystad's reserve estimates are overall conservative relative to those published by U.S. Energy Information Administration (2024a) and International Energy Agency (2023), which both estimate global proven reserves that exceed even Rystad's 2PCX estimate. As an alternative specification with higher reserves, I use International Energy Agency's (2023) technically recoverable resource estimate of 2602 billion bbl.[35] The IEA does not break out its reserve estimates by country, so I allocate the total IEA reserves to the four resource types in proportion to the reference case Rystad estimates.

Finally, because the model's undeveloped reserves R_{i0} do not include reserves that are already developed at time $t = 0$, I compute the R_{i0} by subtracting off $\frac{Y}{1000}\frac{q_{i0}}{\lambda_i}$ from the above reserve numbers, where Y denotes the number of days per period.

3.5 Discount rates

Two broad approaches exist for determining an appropriate discount rate by which oil producers evaluate investments and the value of reserves. First, the weighted average cost of capital (WACC) approach uses the CAPM to compute a firm's opportunity cost of capital based on its costs of debt and equity financing. Damodaran (2024) computes an average real WACC of 5.75% for global oil and gas producers based on data through 2023, and Texas Comptroller of Public Accounts (2026) computes an average nominal WACC of 11.04% for large producers operating in Texas.[36] Second, there exist a variety of surveys that ask pro-

[34]Rystad Energy (2023) does not separately break out reserves for Azerbaijan, Bahrain, Brunei, Malaysia, Oman, South Sudan, and Sudan (all in my "non-core OPEC+" type), leaving them as part of "other non Opec". I infer these countries' reserves by assuming proportionality to their production share relative to production from the full set of countries in Rystad's "other non Opec" set.

[35]I exclude International Energy Agency's (2023) estimated technically recoverable resources of natural gas liquids, bitumen, and kerogen, which would add another 3541 billion bbl to the total.

[36]Damodaran (2024) computes a nominal WACC of 7.34% and assumes an expected inflation rate of 1.50%. The Damodaran WACCs vary over time and range from 3.66% to 8.56% (real) over 2022–2026. I have not been able to locate a 2023 (or other historical) version of Texas Comptroller of Public Accounts (2026). That report provides a pre-tax average WACC of 13.98%; I convert this value to after-tax using the report's oil and gas tax rate of 26.6%. The difference between sources is primarily due to Texas Comptroller

ducers to report the discount rates they use to evaluate assets and potential investments. Society of Petroleum Evaluation Engineers (2023) surveyed industry evaluators and found a mean "standard" discount rate of 10% nominal, but also found that many evaluators add risk adjustments that push applied discount rates to 16% on average (and to more than 20% for undeveloped reserves).[37] Bureau of Ocean Energy Management, Economics Division (2020) and Manley and Heller (2021) discuss a variety of other surveys that report discount rates, with values ranging from 10% to 18% nominal. Texas Comptroller of Public Accounts (2026) also cites survey evidence and concludes that "a discount rate range of 13.25 to 23.52 percent [nominal] is generally suitable for the appraisal of oil and gas properties". Across methods, there is thus a wide range of discount rates that could be plausibly applied. For the reference case model, I use an annual real discount rate of 9%, roughly in the center of the above estimates and comparable to my past work in Kellogg (2014) and Anderson et al. (2018), for resource types III and IV.[38] I ultimately conduct a wide range of sensitivity analyses.

It is not possible to observe or compute the discount rates used by national oil companies, which control production from all of core OPEC and most of non-core OPEC+. A natural default is that they apply the same rate as private oil companies, given that they are producing an identical commodity. Consequently, the reference case also assumes an annual real discount rate of 9% for resource types I and II. Arguably, core OPEC producers could leverage their countries' strong bond ratings and easy access to capital markets. For instance, yields on Saudi Arabian government bonds have been about one percentage point greater than those on U.S. treasuries: 5.1% (nominal) for 10-year bonds and 5.9% for 30-year bonds (Omar, 2024). Access to cheap capital does not imply that these nations should apply a low discount rate to reserve development decisions; however, their leaders have made public pronouncements that investment decisions are made with long-run, intergenerational considerations in mind (Hamilton, 2009).[39] I will address this possibility with an alternative specification in which core OPEC adopts a low real discount rate of $r_I = 3\%$.

of Public Accounts's (2026) use of a larger equity premium and oil industry CAPM beta.

[37]The gap between stated discount rates and WACCs for oil companies parallels a similar finding reported by Gormsen and Huber (2025) for a broad set of industries during 2002–2021. Gormsen and Huber (2025) finds that firm-level variation in discount rates is associated with variation in investments, conditional on the reported WACC.

[38]Kellogg (2014) found that a nominal annual discount rate of 12.5% fit Texas producers' responses to oil price volatility shocks better than rates of 10.5% or 14.5%, consistent with the use here of a 9% real rate, though the improvement in fit was not statistically significant.

[39]The non-core countries in OPEC+, on the other hand, tend to have poorer debt ratings and less access to capital. Adelman (1986) argues that national oil companies should, if anything, have greater discount rates than private oil companies, since it is difficult for the sovereign to diversify given oil's large share in national wealth. The reference case assumes that they use the same discount rate as private oil companies, 9%, though a higher rate is plausible.

3.6 Marginal investment costs

There exist a variety of industry assessments of development costs for new reserves. These estimates typically take the form of "break-even" prices: the price of oil at which an investment in new production would break even, accounting for development costs and discounting. I do not use these break-evens directly as estimates of the model's marginal investment cost functions $c_i(d_i)$, since they are given as constant values rather than functions of investment rates, and since they do not on their own rationalize observed production (Asker et al., 2019). Instead, I calibrate the cost functions by fitting the model to observed 2023 production data, using the industry estimates to inform the values of the cost intercepts $c_i(0)$.

Specifically, for each resource type I model the marginal investment cost as an affine function of the rate of investment, following section 2.1.3: $c_i(d_{it}) = \frac{\lambda_i}{r_i+\lambda_i}(c_{0i} + c_{1i}d_{it})$, with c_{0i} and $c_{1i} > 0$. I set the c_{0i} based on industry estimates, and I then calibrate the c_{1i} so that the model's simulated equilibrium under the baseline demand scenario matches each type's observed 2023 investment and production.

I execute the calibration step for every model specification I employ, not just the reference case. That is, for every alternative specification of the model (e.g., alternative production decline rate or discount rate parameters), I re-calibrate the c_{1i} so that the baseline simulation reproduces 2023 observed production and investment. While this approach has the drawback of muddying the link between these specifications and the analytic comparative statics from theorem 1, it confers the benefit of anchoring each specification to the variables that we quantitatively know most precisely: actual 2023 production and investment levels for each resource. Still, for comparison, I will also present results for the main alternative specifications that do not re-calibrate the c_{1i}.

Estimated break-even prices for new developments, broken out coarsely by geography (e.g., "Middle East" vs. "shale" vs. "deepwater"), are reported by Wood Mackenzie and S&P Global (Dukes, 2019; Brady, 2021). These prices are provided as ranges of break-evens drawn from field-level break-evens within each geography. Because the model's cost function intercepts represent the marginal investment cost at a rate of zero investment — i.e., significantly less investment than that implied by production data — I set the c_{0i} equal to break-evens that are at or near the bottom of the reported ranges within each resource type. I use break-evens of c_{0i} =\$5/bbl for core OPEC, \$10/bbl for non-core OPEC+, \$20/bbl for non-OPEC conventional, and \$30/bbl for shale.

I calibrate the marginal cost slopes c_{1i} by matching the model's initial 2023 equilibrium investment rates in the baseline simulation to those implied by observed 2023 production data and the assumed initial rate of production growth, which I take as equal to assumed baseline demand growth. Specifically, I solve for the implied initial investment rates d_{i0} as the rates of reserve development that would grow production from observed 2023 levels (i.e., those given in section 3.1, which I take as the model's q_{i0}) at a 0.44% annual rate for the first simulated period, for each resource type.[40] I then find the set of c_{1i} for which the model's simulated initial investment rates match these d_{i0}.

The resulting reference case estimates of the c_{1i} are 7.68, 6.0, 5.57, and 16.27 $/bbl per billion bbl per year developed for resource types I, II, III, and IV, respectively. One way to interpret these values is to compute the implied elasticities of marginal drilling costs to investment at the observed initial investment rates. These elasticities are 0.9, 0.88, 0.75, and 0.63 for types I, II, III, and IV, respectively. These values fall within the range of drilling cost elasticity estimates reported in the literature (Kaiser and Snyder, 2012; Toews and Naumov, 2015; Anderson et al., 2018; Vreugdenhil, forthcoming).[41] That said, the literature's elasticities are all short- to medium-run estimates, and the long-run elasticities that are most relevant to my simulations may be smaller. I will therefore consider alternative specifications that use lower drilling cost elasticities than the reference case.

4 Results

4.1 Reference case model

I begin by discussing simulation results from the reference case model, starting with outcomes under the baseline scenario for demand (modest growth followed by stationarity), then outcomes under an anticipated demand decline, and then finally outcomes under an unanticipated demand decline. Subsequent subsections discuss alternative specifications.

Figure 2 presents time series of total global drilling, total global oil production, and oil prices under the reference case model. Figure 3 shows production broken out across the four resource types. At baseline, without a long-run demand decline, all available resources are eventually developed and produced. All undeveloped resources outside of core OPEC are developed by 2071, and core OPEC drilling ceases in 2078. Production then declines asymptotically to zero, fully exhausting global oil reserves (1624 billion bbl) in the limit,

[40] Per equation (3), the necessary initial investment rate, in billions of bbl per year, for each resource type is given by $d_{i0} = 2q_{i0}(1.0044^{1/2} - 1 + \lambda_i)Y/1000\lambda_i$, where Y denotes the number of days per period. The initial reserve development rates are 5.9, 11.7, 11.0, and 3.1 billion bbl per year for types I, II, III, and IV, respectively.

[41] Kaiser and Snyder (2012) estimates cost elasticities of 0.87 and 1.4 for offshore jackup and floating rigs. Toews and Naumov (2015) estimates an elasticity of about 0.3 using global data from Wood Mackenzie. Anderson et al. (2018) estimates that Texas onshore drilling and rig dayrates have elasticities with respect to the oil price of 0.7 and 0.77, respectively, implying a cost elasticity slightly greater than 1. Vreugdenhil (forthcoming) suggests a cost elasticity near 1. Bornstein et al. (2023) also estimates cost elasticities, but with respect to the extraction rate rather than investment rate, and thus significantly larger in magnitude.

Figure 2: Simulation results from the reference case model, aggregated across resource types

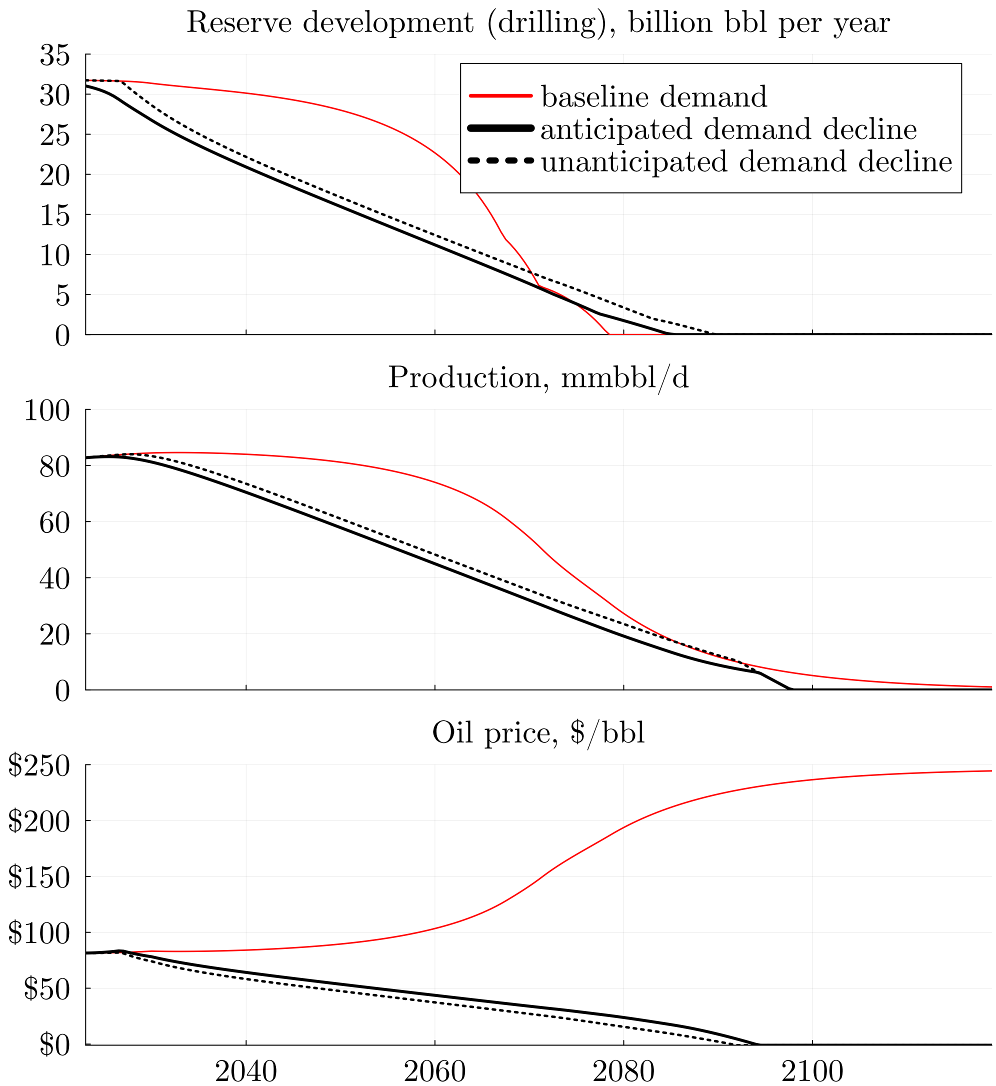

Note: "mmbbl/d" denotes millions of barrels per day. See section 3 for a discussion of the reference case calibration. In the demand decline counterfactuals, the demand curve shifts inward toward zero starting in 2027, reaching zero in 2098.

with the price of oil rising to the demand choke price.

At $t = 0$ (2023), the initial scarcity rents for the four resource types are \$1.56/bbl, \$3.28/bbl, \$2.53/bbl, and \$2.10/bbl for core OPEC, non-core OPEC+, conventional non-OPEC, and shale oil, respectively.[42,43] These scarcity rents are small because at 9% dis-

[42]The raw scarcity rents μ_{i0} from the model are in \$/bbl of undeveloped reserves. The numbers quoted in the text are amortized per-bbl scarcity rents, making them comparable to the "break-even" costs $c_0 + c_1 d_{it}$. I convert the model's μ_{i0} to amortized per-bbl scarcity rents by dividing them by $\frac{\lambda_i}{r_i+\lambda_i}$.

[43]Oil suppliers in the reference case model's baseline scenario still enjoy substantial profits despite these

Figure 3: Simulation results from the reference case model: time paths of extraction for each resource type

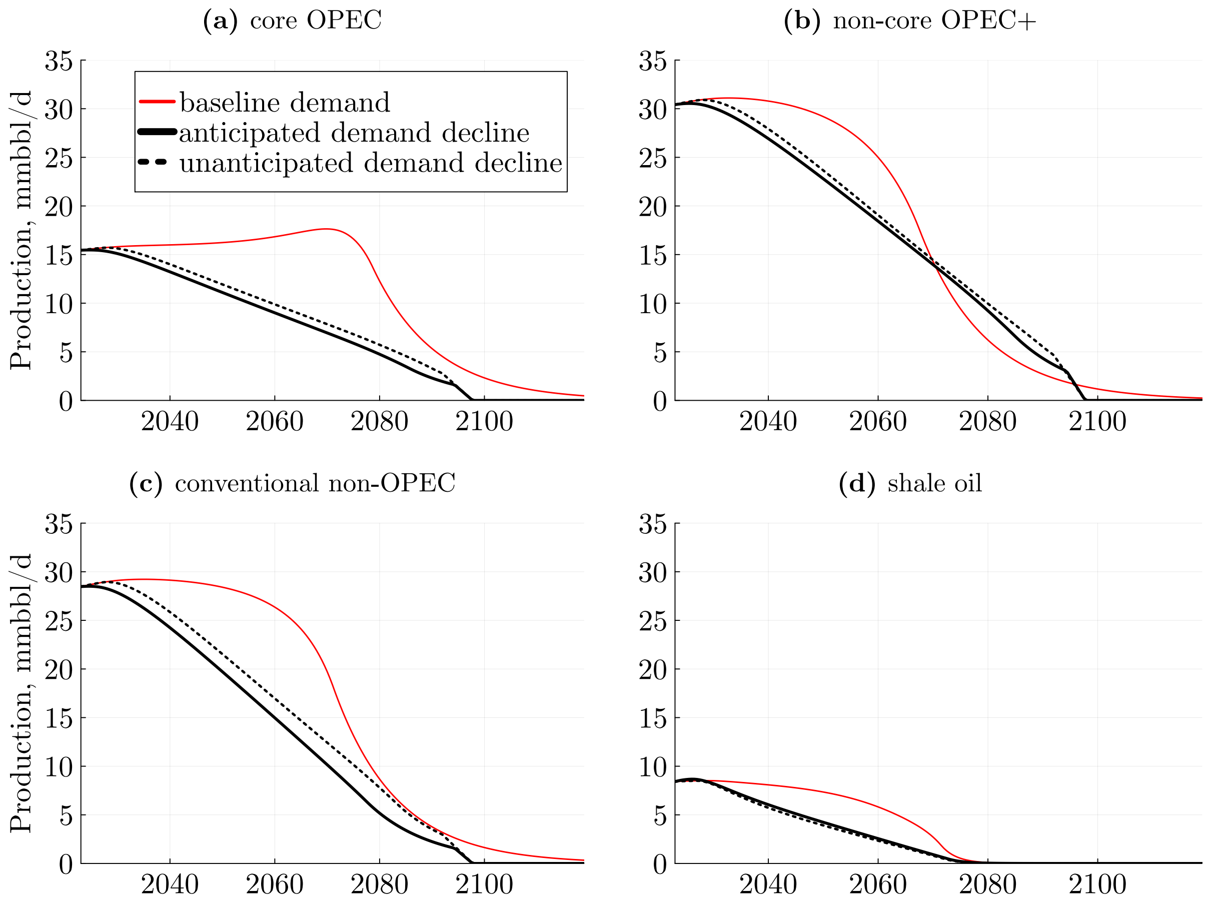

Note: "mmbbl/d" denotes millions of barrels per day. See section 3 for a discussion of the reference case calibration. In the demand decline counterfactuals, the demand curve shifts inward toward zero starting in 2027, reaching zero in 2098.

counting, even the large per-bbl investment profits at the time of exhaustion (recall that the demand choke price is \$247.47/bbl) amount to only a few dollars per barrel in present value. The scarcity rent is lowest for core OPEC because its strategic withholding of production from the market leads it to have a low marginal valuation of undeveloped reserves. The wedge between the $t = 0$ oil price P_0 (\$82.49/bbl) and core OPEC's marginal $t = 0$ investment cost of $c_{01} + c_{11}d_{10} = \$50.69$/bbl is primarily accounted for by its strategic markup,

small scarcity rents. Core OPEC producers naturally profit from their exercise of market power, but even competitive producers earn substantial profits from inframarginal rents, since the marginal investment cost function is strictly upward-sloping. The annualized first-period (2023) flow profits are \$386, \$662, \$591, and \$118 billion for resource types I–IV, respectively.

not its scarcity rent.[44] In contrast, the $t = 0$ marginal investment costs for all other resource types are on par with P_0, exceeding \$80/bbl.[45]

In the demand decline scenario, the decrease in demand is sufficiently large that, when producers anticipate it, the initial scarcity rents for all four resource types fall to zero, and all four types leave oil in the ground.[46] Relative to the baseline scenario, cumulative global drilling investment falls by 29.8%, leading to a reduction in cumulative global oil extraction of 25.0%, from 1624 to 1218 billion bbl.

At $t = 0$, I find that when producers anticipate the demand decline, they reduce their initial investment rate by 2.2%, from 31.7 to 31.0 billion bbl/year (their $t = 0$ production rates q_{i0} are unchanged by construction). Demand in $t = 0$ is unchanged; thus, this result stems from producers' anticipation of the coming decline in demand. The fact that anticipation results in decreased rather than increased investment indicates that the disinvestment effect is outweighing the green paradox at $t = 0$. This finding, moreover, extends throughout the 4-year delay period before the demand decline begins: total drilling investment during this time is 4.3% lower than at baseline. The reduced investment leads to lower extraction during the delay period, which in turn leads to higher oil prices. At the end of year 4, just before the demand decline begins, the simulated oil price is \$84.43/bbl when the demand decline is anticipated, relative to \$82.91/bbl at baseline. These results are consistent with the empirical result in Bogmans et al. (2024) that oil supply investments fell in anticipation of strengthened climate policy, and with the prediction of the oil executives surveyed in Federal Reserve Bank of Dallas (2023) that an anticipated decline in demand would increase rather than decrease the near-term price of oil.

To quantify the effects of anticipation on long-run, not just initial, cumulative drilling and extraction, I evaluate the full time path of drilling and extraction when the demand decline is unanticipated: in each simulated period I model producers as believing that demand will not decline going forward.[47] I find that cumulative drilling and extraction are 9.3% and 6.2%

[44] In the case discussed in section 4.2 below in which core OPEC uses a discount rate of 3%, I find that its marginal $t = 0$ investment cost is \$25.21/bbl.

[45] The $t = 0$ values of $(r_i + \lambda_i)\theta_i/\lambda_i$ for resource types $i =$ I, II, III, and IV are \$52.20, \$83.50, \$83.50, and \$82.74/bbl, respectively.

[46] As the demand curve reaches zero, a point is reached in the simulations in which the resource owners no longer drill, but production from previously drilled wells continues, with the oil price at zero. Because the demand curve is shifting inward more quickly than wells' production naturally declines, I curtail wells' production so that it does not exceed the quantity demanded at $P = 0$. This curtailment is visible in figures 2 and 3 as the kink in the production path a few years before production reaches zero.

[47] In periods before 2030, I model producers as believing that demand will increase at 0.44% per year through 2030, starting at its current level. And after 2030, I model producers as believing that future demand will be identical to current demand. Simulating equilibrium outcomes under the unanticipated demand decline requires re-solving the model in every period, based on the new demand belief and remaining reserves in each period.

lower, respectively, when the demand decline is anticipated rather than unanticipated. Thus, in the reference case model, aggregating across resource types, the effect of disinvestment on cumulative extraction outweighs that of the green paradox. Given the small baseline initial scarcity rents, this result should be characterized as driven primarily by the green paradox being weak, rather than by the disinvestment effect being overwhelmingly large.

Figure 3 shows how the difference between oil extraction paths under the anticipated and unanticipated demand declines varies across resource types. The three non-shale resource types all exhibit disinvestment, both in the short run (when the demand decline is anticipated but has not yet occurred) and in the long run (extraction is lower throughout the demand decline when future demand decreases are anticipated). Shale producers, in contrast, exhibit more drilling and extraction when the demand decline is anticipated rather than unanticipated, though the magnitude of this effect is small. This outcome is consistent with shale's high production decline rate (30% per year) and the intuition from section 2.1.3 that the disinvestment effect will be weak for producers with short investment time horizons.

The reference case results are summarized in row 1 of table 2. From a climate perspective, it can be useful to account for the fact that emissions in the near future generate a greater present value of damages than emissions in the distant future. I therefore also evaluate the simulation results in terms of present discounted cumulative oil extraction, where I apply a discount rate of 3%. I find that the anticipated demand decline reduces present discounted emissions by 17.7% relative to baseline and by 4.2% relative to the unanticipated demand decline. Results from all specifications in terms of present discounted cumulative extraction are presented in appendix tables C.2 and C.3.

4.2 Main alternative specifications: decline rates, reserves, discount rates, and demand decline timing

This section presents results from alternative specifications of the model that vary key parameters governing the relative strengths of the disinvestment effect and green paradox, as highlighted in theorem 1: the production decline rates λ_i, reserve volumes R_{i0}, discount rates r_i, and delay τ before the demand decline begins. The results are summarized in rows 2–15 of table 2, and figures showing simulated drilling, production, and oil price time series are provided in appendix C.

Recall that for each specification, I re-calibrate the marginal cost slope parameters c_{1i} per the discussion in section 3.6, so that the initial investment and production rates in the baseline simulations always match observed values in 2023. Table C.1 in the appendix presents the calibrated drilling cost elasticities for each specification. With the exception of the alter-

Table 2: Simulated initial drilling rates and cumulative extraction for the reference case and main alternative specifications, aggregated across resource types

	Model	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
		Baseline			Anticipated decline			Unant. decline	$\frac{(4)-(1)}{(1)}$	$\frac{(5)-(2)}{(2)}$	$\frac{(6)-(3)}{(3)}$	$\frac{(6)-(7)}{(7)}$
		$\mathbf{D_0}$	$\mathbf{D_L}$	**Q**	$\mathbf{D_0}$	$\mathbf{D_L}$	**Q**	**Q**				
1.	**Reference case**	31.7	127	1624	31.0	121	1218	1299	-2.2%	-4.3%	-25.0%	-6.2%
Alternative production decline rates λ (annualized)												
2.	All resources decline at 8%	31.9	127	1624	31.1	121	1225	1314	-2.5%	-4.8%	-24.6%	-6.8%
3.	Conventional resources decline at 6%	32.2	129	1624	31.2	121	1246	1327	-3.3%	-5.6%	-23.3%	-6.1%
4.	Conventional resources decline at 12%	31.2	125	1624	31.0	122	1192	1262	-0.8%	-2.6%	-26.6%	-5.6%
5.	All resources decline at 30%	30.6	123	1624	30.9	123	1166	1204	+0.8%	-0.2%	-28.2%	-3.1%
6.	No investment ($\lambda \to 1$)	30.2	121	1624	30.4	122	1159	1151	+0.4%	+0.5%	-28.7%	+0.7%
Alternative initial reserves $\mathbf{R_0}$												
7.	High reserves (2602 billion bbl)	31.7	127	2602	30.5	120	1211	1313	-3.8%	-5.8%	-53.4%	-7.7%
8.	Low reserves (1283 billion bbl)	31.7	126	1283	31.9	124	1164	1198	+0.7%	-1.5%	-9.3%	-2.8%
Alternative discount rates r (annualized)												
9.	All resources discount at 6%	31.7	126	1624	32.1	125	1230	1290	+1.3%	-1.2%	-24.3%	-4.7%
10.	All resources discount at 3%	31.7	125	1624	40.6	151	1319	1259	+27.9%	+20.5%	-18.8%	+4.8%
11.	Core OPEC discounts at 3%	31.7	126	1624	34.0	130	1241	1292	+7.1%	+3.1%	-23.6%	-3.9%
12.	3% discounting and no investment	30.2	121	1624	33.7	136	1287	1143	+11.4%	+12.1%	-20.7%	+12.6%
Alternative delays τ before decline starts												
13.	One-year delay	31.7	32	1624	29.3	29	1126	1213	-7.5%	-8.0%	-30.6%	-7.1%
14.	Seven-year delay	31.7	221	1624	31.7	216	1312	1381	-0.1%	-2.4%	-19.2%	-5.0%
15.	Fifteen-year delay	31.7	469	1624	32.2	469	1529	1554	+1.4%	+0.1%	-5.9%	-1.6%

Notes: $\mathbf{D_0}$ denotes the initial $t = 0$ rate of reserve development ("drilling"), in billions of bbl per year.

$\mathbf{D_L}$ denotes cumulative reserves developed, in billions of bbl, during the periods prior to the start of the demand decline.

Q denotes cumulative oil extraction in billions of barrels, over all time periods.

In the unanticipated decline, $\mathbf{D_0}$ and $\mathbf{D_L}$ are the same as at baseline, by construction. In rows 6 and 12, the values in the D_0 and D_L columns are the initial rates of extraction, Q_0 and Q_L in bn bbl/year. The cost function slopes c_{1i} are re-calibrated in each specification 2–12; the re-calibrated drilling cost elasticities are presented in appendix table C.1. See section 3 for a discussion of the reference case calibration, and section 4.2 for discussions of the alternative specifications. For each alternative specification, a figure analogous to figure 2 is presented in appendix C.

native discount rate specifications, the re-calibrated elasticities change little relative to the reference case.[48] Table C.4 presents results without re-calibration of the c_{1i} for comparison.

Production decline rates λ_i

Part (a) of theorem 1 shows, using a restricted version of the model, that the disinvestment effect is strengthened when λ_i is smaller, so that investments are longer-lived. Accordingly, row 2 of table 2 shows that when I set the decline rate of shale oil wells to be as slow as that of conventional wells — 8% rather than 30% — I find that anticipation reduces drilling and extraction by more than in the reference case. Relative to the situation in which the demand decline is unanticipated, initial investment D_0 falls by 2.5%, cumulative drilling prior to the demand decline D_L falls by 4.8%, and cumulative extraction Q falls by 6.8%. The corresponding values from the reference case are 2.2%, 4.3%, and 6.2%, respectively.[49]

Row 3 of table 2 shows results when the decline rate for conventional resources (types I, II, and III) is assumed to be 6%, at the bottom of its empirical support discussed in section 3.3 (shale is assumed to have a 30% decline, as in the reference case). As expected, the disinvestment effect strengthens.[50] Increasing the decline rates has the opposite effect. In row 4 of table 2, I set $\lambda = 12\%$ for the conventional resources, at the top of the empirical support. Anticipation now decreases D_0, D_L, and Q by 0.8%, 2.6%, and 5.6%, respectively, all smaller in magnitude than in the reference case. Going further, in row 5 I set the decline rates of all resource types to that of shale oil, 30%. This value is not plausible for conventional oil production, but it serves to illustrate the intuition that when all investments are short-lived, the disinvestment effect is weak. Here, anticipation actually increases the initial investment rate D_0 by 0.8%, indicating that the green paradox is the stronger effect on net. However, D_L and total extraction Q still both decrease, by 0.2% and 3.1%, relative to the case when the demand decline is unanticipated.

Finally, row 6 of table 2 takes this intuition to the limit by taking $\lambda \to 1$ and removing the one-period lag between expenditure and production. This model therefore removes all investment dynamics and is akin to a traditional Hotelling-style model of extraction. Accordingly, anticipation of the demand decline now increases the initial extraction rate Q_0,

[48] For example, in the specification in which all resource types discount at 3% (row 10 of table 2), firms' simulated initial drilling would be lower than observed 2023 drilling were I to hold the c_{1i} fixed, since the lower discount rate increases the initial shadow costs μ_{i0}. In the specifications that vary τ, the baseline simulation is unaffected, so the values of the c_{1i} are the same as in the reference case.

[49] The pattern of results that Q is more negative than D_L, which in turn is more negative than D_0, is consistent with part (d) of theorem 1, which states that disinvestment becomes relatively stronger when the start of the demand decline is nearer in the future.

[50] The percent decrease in cumulative extraction in row 3 is smaller than in the reference case, but this result simply reflects the fact that the initial developed reserve base, equal to q_{i0}/λ_i for each resource, has increased. Total cumulative investment falls by more in this specification than in the reference case.

cumulative extraction prior to the demand decline Q_L, and total cumulative extraction Q by 0.4%, 0.5%, and 0.7%, respectively.[51] This result is consistent with the fact that the green paradox is now the only remaining force governing anticipation effects. The magnitudes are modest, however, reflecting the fact that at baseline, the producers' initial scarcity rents are small.

Undeveloped reserves R_{i0}

Row 7 of table 2 considers a model in which global oil reserves are 2602 billion bbl, rather than just 1624 billion bbl as in the reference case (recall section 3.4 for a discussion of reserve estimates). Following the intuition underpinning part (b) of theorem 1, this increase in reserves lowers producers' initial scarcity rents at baseline, weakening the green paradox. Accordingly, anticipation reduces D_0, D_L, and Q by 3.8%, 5.8%, and 7.7%, respectively, in this model. Conversely, assuming that reserves are just 1283 billion bbl leads to increased scarcity rents at baseline and a stronger green paradox. For instance, the baseline scarcity rent for conventional non-OPEC in this specification is \$9.10/bbl, rather than just \$2.53/bbl in the reference case. This scarcity strengthens the green paradox relative to the disinvestment effect, so that anticipation increases D_0 by 0.7%, though D_L and Q remain decreased, by 1.5% and 2.8%, respectively (row 8 of table 2).

Discount rates r_i

The discount rate r_i used by each resource type turns out to play a key role in determining the relative strengths of the disinvestment effect and green paradox. Figure 4 shows how the effects of anticipation on D_0, D_L, and Q vary with the discount rate that is assumed to hold for all resource types (with a rate of 9% corresponding to the reference case). For a discount rate below 11%, increasing the discount rate attenuates the green paradox more than it attenuates disinvestment, so that all three lines in the graph are downward-sloping. Above this rate, the lines are upward-sloping, as further increases in the discount rate attenuate disinvestment more than the green paradox. This U-shape in figure 4 is consistent with part (c) of theorem 1: once the baseline initial scarcity rents are driven to essentially zero, further increases in the discount rate only serve to decrease the magnitude of the disinvestment effect.

[51] The 0.4% increase in Q_0 in the non-investment model ($\lambda \to 1$) is smaller than the 0.8% increase in D_0 when the λ_i all equal 30%. The direction of this change is the opposite of that predicted by theorem 1. The discrepancy arises because theorem 1 is derived under the assumption of perfectly elastic demand, and because even under the conditions of theorem 1, the predicted magnitude of $\frac{d\Delta D_0}{d\lambda}$ goes to zero as $\lambda \to 1$ (so that the implications of the perfectly elastic demand assumption become quantitatively important for large λ). I find that if I instead assume that demand is nearly perfectly elastic (an elasticity of -50), then the percentage increase in Q_0 in the non-investment model (0.6%) is greater than the percentage increase in

Figure 4: Sensitivity of the net anticipation effect to the discount rate that is applied to all resource types

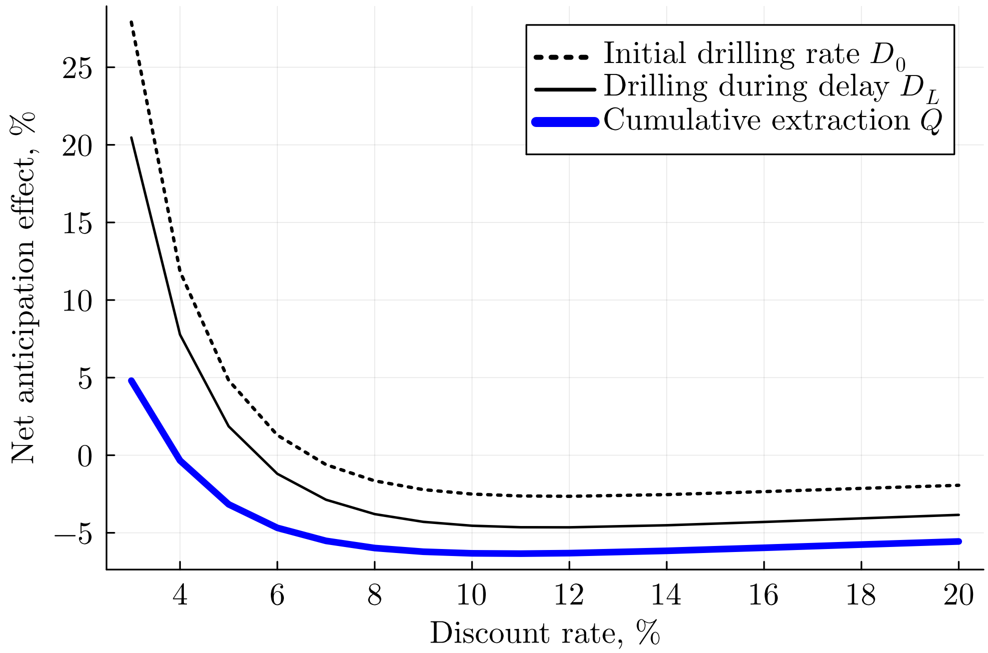

Note: Each point on the figure shows the result from a simulation in which all producers have the real discount rate given by the horizontal axis. All other model inputs are per the reference case (save the drilling cost slopes c_{1i}, which vary with the discount rate to match 2023 production and investment). The net anticipation effect on the vertical axis is the percentage change in initial drilling investment (D_0), cumulative investment (D_L) during the delay before the demand decline begins, or total cumulative global extraction (Q) when the demand decline is anticipated relative to when it is unanticipated.

Figure 4 also shows that for sufficiently low discount rates, the green paradox can outweigh disinvestment, such that initial investment and cumulative extraction increase rather than decrease in the anticipated demand decline scenario. For initial investment D_0, this result occurs for discount rates below 6.7%, and for D_L and Q the required discount rates are below 5.6% and 3.9%, respectively. Per row 9 of table 2, at a discount rate of 6%, anticipation increases D_0 by 1.3% rather than decreases it, and the reductions in D_L and Q are 1.2% and 4.7%, respectively, smaller in magnitude than in the reference case. And if all producers were to discount at 3%, I find that anticipation of the demand decline causes D_0, D_L, and Q to increase by 27.9%, 20.5%, and 4.8%, respectively (row 10). This level of patience is below the range of discount rates discussed in section 3.5, but it serves to illustrate the quantitatively large role that discount rates can play in determining the relative strengths

D_0 when the λ_i all equal 30% (+0.0%), consistent with the theory.

of the disinvestment effect and green paradox. The mechanism follows the intuition that underpins part (c) of theorem 1: as the discount rate is lowered, the initial baseline scarcity rents can eventually reach high levels, leading to a green paradox on net. For instance, the initial scarcity rent for conventional non-OPEC is \$33.70/bbl at a 3% discount rate, versus just \$2.53/bbl in the reference case.

Row 11 of table 2 presents results from a plausible case in which core OPEC discounts at 3%, while all other producers remain at the reference case rate of 9%. In this specification, anticipation increases D_0 and D_L by 7.1% and 3.1%, but decreases cumulative extraction Q by 3.9%. Thus, when core OPEC is patient, the green paradox outweighs disinvestment in the near term, since early investments will have substantially depreciated away by the time the demand decline begins. This increase in investment leads the near-term extraction rate to rise and the near-term oil price to fall relative to their baseline paths. These results are consistent with the empirical results in Adolfsen et al. (2024) and Norman and Schlenker (2024) that oil supply investments increased and oil prices decreased following climate policy announcements. However, as time passes in this simulation and the onset of the demand decline draws near and then occurs, the disinvestment effect strengthens, resulting in an overall reduction in extraction over the full time horizon.

Finally, row 12 of table 2 evaluates a specification that sets the discount rate to 3% for all resource types and removes all investment dynamics by taking $\lambda \to 1$ and eliminating the lag between expenditure and production (as in row 6). In this specification, anticipation causes initial extraction Q_0 to increase by 11.4%, cumulative extraction prior to the demand decline Q_L to increase by 12.1%, and total cumulative extraction Q to increase by 12.6%. This specification is not realistic, but it demonstrates that this paper's model is able to deliver an economically large green paradox for total cumulative extraction when both the disinvestment effect is shut down and the assumed discount rate is small.

Delay τ before the demand decline begins

Table 2, rows 13–15, presents results in which I vary the delay τ before the demand decline begins. In each specification, the year in which the demand decline ends varies one-for-one with the change in τ. All other parameters are per the reference case.

These results accord with part (d) of theorem 1, which shows that as the delay τ decreases, disinvestment becomes relatively stronger. When the delay is just 1 year, anticipation decreases D_0, D_L, and Q by -7.5%, -8.0%, and -7.1%, respectively (row 13). These decreases are all larger in magnitude than in the reference case, in which the delay is 4 years. Row 14 shows that for a 7-year delay, the magnitudes are instead reduced. Row 15 considers a delay of 15 years, finding that anticipation now increases D_0 and D_L by 1.4% and 0.1%, since in this case the delay is sufficiently long that the disinvestment effect is quite weak near $t = 0$. However, cumulative extraction Q still decreases by 1.6%, indicating that disinvestment is still stronger over the full time horizon.

4.3 Additional alternative specifications

Table 3 presents results from additional alternative specifications of the model. In row 16, I examine how the results differ when I increase core OPEC's market power by including Russia's initial production and reserve volume in this resource group.[52] I find that doing so modestly decreases total cumulative extraction under an anticipated decline (from 1218 to 1212 billion bbl), consistent with the notion that an increase in market power will cause core OPEC to extract its reserves more slowly. This intuition also applies to the case of an unanticipated demand decline, so that on net, cumulative extraction is 6.1% less under an anticipated demand decline than under an unanticipated decline. This effect is slightly smaller in magnitude than the 6.2% result from the reference case.

Row 17 of table 3 reports outcomes when I eliminate OPEC's market power altogether by modeling the core OPEC resource type as behaving competitively. Mirroring the results from row 16 in which OPEC's market power was strengthened, I find that eliminating market power modestly increases cumulative extraction, relative to the reference case, under both an anticipated and unanticipated decline. Overall, anticipation decreases cumulative extraction by 6.4% on net, slightly more than in the reference case.

The next specification, in row 18, examines an alternative mechanism for a green paradox: what if an anticipated demand decline weakens the OPEC cartel's ability to collude, leading it to increase its extraction rate? The intuition behind such a potential outcome comes from Rotemberg and Saloner (1986), Haltiwanger and Harrington (1991), and De Cannière (2025): because cartel cooperation is motivated by the potential for future profits, an anticipated decrease in demand will tighten the cartel members' incentive compatibility constraints, forcing either collusion at a lower price level or an outright return to competitive behavior. To model this possibility, I start with the reference case model and then, when modeling the anticipated demand decline, switch core OPEC producers' behavior to be competitive rather than strategic.[53] Doing so causes the initial investment rate D_0 to increase by 5.6%.

[52] The "Russia in core OPEC" specification moves Russia's 10.6 mmbbl/d of initial production and 143 billion bbl of reserves from non-core OPEC+ (type II) to core OPEC (type I).

[53] This "no market power in anticipated decline" specification uses the reference case model for the baseline and unanticipated demand decline simulations, and for the anticipated decline models core OPEC as behaving competitively, holding all other reference case parameters (including the cost slopes c_1) constant. The "No market power" specification (row 17), in contrast, models core OPEC as behaving competitively at baseline and for both demand declines, with the cost slopes recalibrated so that baseline investment matches actual investment.

Table 3: Simulated initial drilling rates and cumulative extraction for the reference case and additional alternative specifications, aggregated across resource types

		(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
		Baseline			Anticipated decline			Unant. decline	$\frac{(4)-(1)}{(1)}$	$\frac{(5)-(2)}{(2)}$	$\frac{(6)-(3)}{(3)}$	$\frac{(6)-(7)}{(7)}$
	Model	$\mathbf{D_0}$	$\mathbf{D_L}$	**Q**	$\mathbf{D_0}$	$\mathbf{D_L}$	**Q**	**Q**				
1.	**Reference case**	31.7	127	1624	31.0	121	1218	1299	-2.2%	-4.3%	-25.0%	-6.2%
Alternative specifications for market power												
16.	Russia in core OPEC	31.7	127	1624	31.2	122	1212	1290	-1.7%	-4.0%	-25.4%	-6.1%
17.	No market power	31.7	127	1624	31.0	121	1223	1307	-2.3%	-4.4%	-24.7%	-6.4%
18.	No market power in anticipated decline	31.7	127	1624	33.5	130	1259	1299	+5.6%	+2.4%	-22.4%	-3.1%
Alternative specifications for demand												
19.	High demand elasticity (-0.6)	31.7	127	1624	31.0	121	1218	1304	-2.4%	-4.5%	-25.0%	-6.6%
20.	Low demand elasticity (-0.4)	31.7	127	1624	31.1	122	1219	1293	-1.9%	-4.0%	-24.9%	-5.7%
21.	No demand growth	30.2	121	1624	29.5	115	1183	1266	-2.5%	-4.6%	-27.1%	-6.5%
22.	Higher demand growth (1.5 per year)	35.3	142	1624	34.8	137	1305	1375	-1.5%	-3.5%	-19.6%	-5.1%
23.	Demand decline ends in 50 years	31.7	127	1624	30.4	117	844	931	-4.3%	-7.5%	-48.0%	-9.3%
24.	Demand decline ends in 100 years	31.7	127	1624	31.3	123	1550	1578	-1.3%	-2.8%	-4.6%	-1.8%
25.	15% of demand remains after decline	31.7	127	1624	31.0	121	1606	1615	-2.2%	-4.3%	-1.1%	-0.6%
26.	Demand doubles rather than declines	31.7	127	1624	31.6	129	1624	1624	-0.5%	+1.6%	+0.0%	+0.0%
Alternative specifications for investment costs												
27.	High cost function intercepts	31.7	127	1624	31.1	121	1150	1235	-2.1%	-4.4%	-29.2%	-6.9%
28.	Cost dependence on aggregate drlg	31.7	127	1624	30.9	121	1220	1300	-2.6%	-4.7%	-24.8%	-6.1%
29.	Quadratic marginal cost function	31.7	127	1624	31.1	122	1301	1364	-2.0%	-3.5%	-19.9%	-4.6%
30.	Costs decline 2% per year for 50 years	31.7	129	1624	31.5	126	1370	1409	-0.9%	-3.0%	-15.7%	-2.8%
31.	Cost dependence on reserves	31.7	126	1624	31.1	121	1144	1222	-1.9%	-4.0%	-29.5%	-6.3%
32.	Seven regions	31.7	127	1624	31.2	122	1173	1249	-1.7%	-3.8%	-27.8%	-6.1%

Notes: $\mathbf{D_0}$ denotes the initial $t = 0$ rate of reserve development ("drilling"), in billions of bbl per year. $\mathbf{D_L}$ denotes cumulative reserves developed, in billions of bbl, during the periods prior to the start of the demand decline. **Q** denotes cumulative oil extraction in billions of barrels, over all time periods.

In the unanticipated decline, $\mathbf{D_0}$ and $\mathbf{D_L}$ are the same as at baseline, by construction. The cost function slopes c_{1i} are re-calibrated in each specification 16–22 and 27–32; the re-calibrated drilling cost elasticities are presented in appendix table C.1. For specification 26, columns (4) to (11) capture effects of a demand increase rather than a decline. See section 3 for a discussion of the reference case calibration, and section 4.3 for discussions of the alternative specifications. For each alternative specification, a figure analogous to figure 2 is presented in appendix C.

Overall, cumulative extraction under an anticipated demand decline increases relative to the reference case, from 1218 to 1259 billion bbl, though this value remains less than (by 3.1%) cumulative extraction under an unanticipated decline.

Rows 19 through 22 of table 3 present results from alternative specifications for baseline demand. Assuming a higher (in magnitude) demand elasticity of -0.6 causes the disinvestment effect to modestly strengthen relative to the green paradox, while assuming a lower elasticity of -0.4 does the opposite. Setting demand growth to zero at baseline, rather than the 0.44% annual growth rate through 2030 assumed in the reference case, diminishes the green paradox because the resources' scarcity values are lower at baseline. Thus, in this specification (row 21), anticipation of the demand decline reduces cumulative extraction by 6.5%, more than the 6.2% in the reference case. Conversely, increasing the demand growth rate to 1.5%, per Organization of the Petroleum Exporting Countries (2023), strengthens the green paradox, though the effects of anticipation on D_0, D_L, and Q are still negative on net (row 22).

I next evaluate the implications of accelerating or lengthening the demand decline, holding constant the 4-year delay before the decline begins. Row 23 of table 3 presents results from a model in which demand falls to zero over 50 years (including the delay period), and row 24 presents results when demand falls to zero over 100 years (versus 75 years in the reference case). The speed of the demand decline, not surprisingly, substantially affects cumulative extraction, which I find to be 844 billion barrels under an anticipated 50-year decline and 1550 billion barrels under an anticipated 100-year decline (out of 1624 billion bbl of total reserves).[54] When I focus on anticipation effects, I find that relative to the reference case, they become more negative with a 50-year decline (a 9.3% decrease in cumulative extraction) and smaller in magnitude with a 100-year decline (a 1.8% decrease in cumulative extraction). Intuitively, accelerating the demand decline enhances the disinvestment effect while not affecting the strength of the green paradox, since the initial scarcity rents in the anticipated decline scenario are zero for both the reference case and accelerated decline models.

I also consider the possibility that the demand decline may be incomplete, leaving some remaining demand for oil in perpetuity. For instance, it may not be possible to find cost-competitive substitutes for some petrochemicals or aviation fuels. Table 3, row 25, presents results from a model in which the demand decline follows the reference case until 15% of the initial demand remains, after which point demand is stationary. When this partial demand decline is anticipated, producers eventually achieve near-complete extraction, since only shale oil resources are incompletely extracted.[55] However, because long-run demand is so low,

[54] In the 100-year anticipated decline model, non-core OPEC+'s resources are fully exhausted.

[55] Shale oil is the only resource whose marginal cost curve intercept c_{0i} sits above the long-run demand

drilling and extraction proceed over hundreds of years, creating a large divergence between cumulative extraction (1606 billion bbl) and present discounted extraction (676 billion bbl). I find that the disinvestment effect still modestly outweighs the green paradox in this scenario, with anticipation effects reducing cumulative and present discounted extraction by 0.6% and 3.9%, respectively.

The last demand scenario I consider is the opposite of a demand decline: what if global oil demand instead increases substantially? Such a scenario is plausible: in Federal Reserve Bank of Dallas (2023), 25% of industry survey respondents indicated that they expect global oil consumption to "significantly increase" relative to its 2023 level. In row 26 of table 3, demand follows the baseline scenario for 4 years and then expands outward at a constant rate, so that its intercept doubles after 75 years elapse. When this demand increase is anticipated, there are two offsetting effects. First, what was formerly the disinvestment effect now becomes an *investment effect*, since forward-looking producers will increase their durable investments in advance of the shift. And the green paradox will work in reverse: anticipation of the demand increase will increase the initial scarcity rent, thereby reducing initial investment, as producers will hold back production in order to take advantage of higher future prices.[56] On net, the investment effect outweighs the reverse green paradox: anticipation of the demand increase leads to a 1.6% increase in investment D_L during the 4-year period before demand begins increasing. Comparing the anticipated and unanticipated demand increase scenarios, extraction is overall accelerated in the former relative to the latter, with both exhibiting complete exhaustion of the undeveloped reserves.

The final six rows of table 3 show results from alternative specifications for the marginal investment cost functions $c_i(d_{it})$. In row 27, I consider an alternative specification in which I increase the marginal investment cost intercepts c_{0i} to \$20/bbl for core OPEC, \$25/bbl for non-core OPEC+, \$35/bbl for non-OPEC conventional, and \$45/bbl for shale, versus values of \$5, \$10, \$20, and \$30/bbl in the reference case. These changes flatten the marginal investment cost functions when the model is calibrated because the initial marginal investment costs $c_i(d_{i0})$ must remain similar to those in the reference case, since the initial scarcity rents are not substantially changed.[57] The resulting marginal drilling cost elasticities are 0.61, 0.69, 0.57, and 0.44 for resource types I–IV, so that this specification speaks to the possibility that long-run drilling cost elasticities are smaller than the reference case elasticities (0.9, 0.88, 0.75, and 0.63) that matched short-run elasticities from the empirical literature.[58]

choke price.

[56]In the simulation corresponding to row 26 of table 3, the scarcity rents for the anticipated demand increase scenario are \$6.97, \$10.69, \$9.45, and \$7.45/bbl for resource types I–IV, respectively.

[57]The initial baseline scarcity rents in the "high cost function intercepts" model are \$1.54/bbl, \$2.96/bbl, \$2.33/bbl, and \$1.83/bbl for the four resource types.

[58]In appendix D, I further consider a specification in which marginal drilling costs depend on the current

Flatter cost functions increase firms' responsiveness to anticipated future demand changes, which strengthens both the disinvestment effect and the green paradox. They also reduce producers' initial scarcity rents at baseline, since their future marginal drilling costs will not decrease as much over time as in the reference case. The net effect of this change in specification is a strengthening of the disinvestment effect over the green paradox. Anticipation reduces cumulative extraction by 6.9% in this model, more than the 6.2% from the reference case.

In row 28, I consider a specification in which each resource type's marginal investment cost is an affine function of both its own investment and aggregate investment across types. I am not aware of empirical evidence with which to calibrate this additional parameter, so I adopt a value of \$1.00/bbl per billion bbl that is sufficient to force the own-investment drilling cost elasticities to meaningfully decrease (to 0.39, 0.63, 0.50, and 0.28 for the four types) when the model is calibrated. As with row 27, there is a modest strengthening of disinvestment. Row 29 presents results from another alternative functional form that models marginal costs as a function of squared investment.[59] In this case, the initial drilling cost elasticities are relatively large (1.80, 1.75, 1.50, and 1.25 for the four types), attenuating disinvestment.

Row 30 presents a specification in which technological progress causes marginal investment costs to decrease at an annual rate of 2% for 50 years, for all resources. Like specification 22, which increased the rate of demand growth, this specification causes the initial baseline scarcity rents to increase (in this case, by 61% on average across the four resources). This increase strengthens the green paradox relative to disinvestment, though anticipation still reduces initial investment and cumulative extraction on net.

Row 31 of table 3 presents results from a specification in which each resource type's marginal cost intercept increases linearly as undeveloped reserves are depleted, starting from its initial reference case value of c_{0i} and reaching c_{0i} + \$20 upon exhaustion. This "stock effect" causes the shadow values μ_{it} to increase at a rate strictly less than the discount rate,[60] leading the initial scarcity rents to roughly triple relative to the reference case.[61] This increase in baseline scarcity rents does not, however, meaningfully strengthen the green paradox because under the anticipated demand decline, the initial scarcity rents also increase.[62]

rate of drilling and the lagged rate of drilling, creating a difference between short-run and long-run responses.

[59]The functional form is $c_i(d_{it}) = c_{0i} + c_{2i}d_{it}^2$, where I calibrate the c_{2i} parameters so that simulated $t = 0$ baseline investment matches actual investment.

[60]With stock effects, the marginal investment cost function is given by $c_i(d_{it}) = c_{0i} + c_{1i}d_{it} + c_{Ri}R_{it}$, with $c_{Ri} < 0$. The equation of motion for the scarcity rent, in current value, becomes $\mu_{i,t+1} = \mu_{it}(1+r_i) + c_{Ri}d_{it}$, so that μ_{it} increases at a rate slower than r_i whenever investment is strictly positive.

[61]With stock effects, the initial baseline scarcity rents become \$5.43/bbl, \$8.68/bbl, \$7.44/bbl, and \$7.31/bbl for the four resource types.

[62]Even though depletion is incomplete under the anticipated demand decline, the initial scarcity rents are

Finally, in row 32 of table 3 I show that segmentation of the competitive resource types II–IV into sub-types with heterogeneous investment costs does not meaningfully change the quantitative results. This specification divides the reserves of each resource II–IV into two halves, so that there are seven total types. For one half, the cost intercept c_0 is reduced by \$5/bbl, and the initial production rate is 2/3 that of the original resource type. For the other half, the cost intercept is increased by \$10/bbl, and the initial production rate is 1/3 that of the original resource (thus, the production-weighted average cost for each resource is unchanged from the reference case). The results are similar to those from specification 31, where marginal investment costs rose continuously with depletion.

5 Concluding discussion

Even as most forecasts indicate that global oil demand will continue to grow modestly or at least hold steady, an alternative scenario in which a combination of climate policies and clean energy technologies substantially drives down the future demand for oil is now also plausible. This paper develops a model to assess how global oil producers might respond to an anticipated, long-run decline in oil demand toward zero. My emphasis in constructing the model is on capturing two opposing forces that govern the direction and magnitude of anticipation effects: the green paradox and disinvestment. The model also captures salient industry features such as market power and cross-resource heterogeneity. At the same time, the model is sufficiently simple that its mechanisms are intuitive and its numerical version computes quickly, enabling simulations in which the demand decrease is unanticipated, and allowing for the exploration of a large number of alternative specifications.

Using this model, I find that for parameter inputs with the strongest empirical support, the green paradox is small, such that it tends to be outweighed by the disinvestment effect. In the reference case model, an anticipated future demand decline reduces the initial investment rate by 2.2% and cumulative oil production by 6.2% more than would occur under an unanticipated decline. Disinvestment occurs for all conventional resources but not meaningfully for shale oil, since its investments are short-cycle.

The green paradox is quantitatively small in this model because initial scarcity rents in the baseline (non-declining) demand scenario are small, reflecting producers' real 9% discount rate and the large size of global oil reserves relative to oil demand. Obtaining a green paradox that substantially increases cumulative extraction requires assuming that

non-zero in the specification with stock effects: \$3.31/bbl, \$5.23/bbl, \$4.38/bbl, and \$4.88/bbl for the four resource types. In the equilibrium, the scarcity rents decline over time, reaching zero in the last period in which investment is strictly positive.

producers are patient, with a real discount rate below 4%, and that investments are short-lived. Overall, these results suggest that the green paradox may not be a significant problem for climate policies or green technologies that will reduce future oil demand.

I see two main limitations to my analysis, which future work can hopefully address. First, the Anderson et al. (2018) investment model that I adopt is simple, treating all oilfield investments as akin to drilling, wherein production declines exponentially from the invested capacity. In practice, firms make other forms of investment such as exploration, construction of surface facilities (including deepwater platforms) and rigs, development of human capital, and enhanced oil recovery. Accounting for these activities, which are often long-term investments that also involve substantial time-to-build, would enhance the disinvestment effect. Second, the model is deterministic. Extending this paper to a stochastic environment would be valuable given the importance of both demand and supply shocks in the world oil market, and it would relieve the tension inherent in coupling a deterministic model with risk-adjusted discount rates. Such an extension would, however, be a substantial undertaking, since the model would require solving for an equilibrium in Markov strategies, which is considerably more computationally intensive than solving for the open-loop equilibrium of this paper's deterministic model.

Finally, it would be valuable to evaluate green paradox versus disinvestment effects for resources other than oil. This paper's model would naturally extend to natural gas, which involves a similar production technology as oil. Studying coal seems especially important given its high CO_2 emissions intensity. Doing so would require a model distinct from that presented here, since coal extraction is not known to follow production decline curves analogous to oil well production declines.

Data and code availability

The data and code underlying this research are available on Zenodo at `https://dx.doi.org/10.5281/zenodo.22089220`.

References

Adelman, M. A., "Oil Producing Countries' Discount Rates," *Resources and Energy*, 1986, *8*, 309–329.

_ , "Mineral Depletion, with Special Reference to Petroleum," *The Review of Economics and Statistics*, 1990, *72* (1), 1–10.

Adolfsen, Jakob Feveile, Malte Heissel, Ana-Simona Manu, and Francesca Vinci, "Burn Now or Never? Climate Change Exposure and Investment of Fossil Fuel Firms," 2024. European Central Bank working paper #2945.

Anderson, Soren T., "The Demand for Ethanol as a Gasoline Substitute," *Journal of Environmental Economics and Management*, 2012, *63*, 151–168.

_ , Ryan Kellogg, and Stephen W. Salant, "Hotelling Under Pressure," 2014. NBER working paper 20280.

_ , _ , and _ , "Hotelling Under Pressure," *Journal of Political Economy*, 2018, *126*, 984–1026.

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

_ , _ , Charlotte De Cannière, Jan De Loecker, and Christopher R. Knittel, "Two Wrongs Can Sometimes Make a Right: The Environmental Benefits of Market Power in Oil," November 2024. NBER working paper 33115.

Baldwin, Elizabeth, Yongyang Cai, and Karlygash Kuralbayeva, "To Build or Not To Build? Capital Stocks and Climate Policy," *Journal of Environmental Economics and Management*, 2020, *100*, 102235.

Balke, Nathan S. and Stephen P.A. Brown, "Oil Supply Shocks and the U.S. Economy: An Estimated DSGE Model," *Energy Policy*, 2018, *116*, 357–372.

_ , Xin Jin, and Mine Yücel, "The Shale Revolution and the Dynamics of the Oil Market," *Economic Journal*, 2024, pp. 1–38.

Bauer, Nico, Christophe McGlade, Jérôme Hilaire, and Paul Ekins, "Disinvestment Prevails Over the Green Paradox When Anticipating Strong Future Climate Policies," *Nature Climate Change*, 2018, *8*, 130–134.

Benchekroun, Hassan and Cees Withagen, "On Price-Taking Behavior in a Nonrenewable Resource Cartel-Fringe Game," *Games and Economic Behavior*, 2012, *76*, 355–374.

_ , Gerard van der Meijden, and Cees Withagen, "An Oligopoly-Fringe Non-Renewable Resource Game in the Presence of a Renewable Substitute," *Journal of Economic Dynamics and Control*, 2019, *105*, 1–20.

_ , _ , and _ , “OPEC, Unconventional Oil and Climate Change — On the Importance of the Order of Extraction,” *Journal of Environmental Economics and Management*, 2020, *104*, 102384.

Bogmans, Christian, Andrea Pescatori, and Ervin Prifti, “The Impact of Climate Policy on Oil and Gas Investment: Evidence From Firm-Level Data,” *European Economic Review*, 2024, *165*, 104750.

Bornstein, Gideon, Per Krusell, and Sergio Rebelo, “A World Equilibrium Model of the Oil Market,” *Review of Economic Studies*, 2023, *90*, 132–164.

BP, “bp Energy Outlook: 2024 edition,” 2024.

Brady, Aaron, “Global Crude Oil Cost Curve Shows 90% of Projects Through 2040 Breaking Even Below $50/bbl,” *S&P Global*, September 2021.

Bruno, Ellen M. and Nick Hagerty, “Anticipatory Effects of Regulating the Commons,” 2024. working paper.

Bureau of Ocean Energy Management, Economics Division, “Recommended Special Case Royalty Relief Discount Rates for Deepwater Oil and Gas Projects Using Subsea Tiebacks Requiring Enhanced Flow Assurance Technologies,” 2020.

Cairns, Robert D. and James L. Smith, “The Green Paradox, A Hotelling Cul de Sac,” *Economics of Energy and Environmental Policy*, 2019, *8* (2), 1–20.

Curuk, Malik and Suphi Sen, “Climate Policy and Resource Extraction with Variable Markups and Imperfect Substitutes,” *Journal of the Association of Environmental and Resource Economists*, 2023, *10* (4), 1091–1120.

Damodaran, Aswath, “Costs of Capital by Industry Sector,” Online; `https://pages.stern.nyu.edu/~adamodar/New_Home_Page/datafile/wacc.html`; accessed 12 July, 2024.

De Cannière, Charlotte, “Pump it? Market Power and the Energy Transition in the Global Oil Market,” January 2025. working paper.

Dukes, R.T., “Energy Outlook, TXCPA Energy Conference,” *Wood Mackenzie*, August 2019.

Energy Institute, “Statistical Review of World Energy,” 2024.

Eswaran, Mukesh and Tracy Lewis, "Exhaustible Resources and Alternative Equilibrium Concepts," *Canadian Journal of Economics*, 1985, *18* (3), 459–473.

Federal Reserve Bank of Dallas, "Dallas Fed Energy Survey," 2023. https://www.dallasfed.org/research/surveys/des/2023/2303#tab-report.

Fischer, Carolyn and Stephen W. Salant, "Balancing the Carbon Budget for Oil: The Distributive Effects of Alternative Policies," *European Economic Review*, 2017, *99*, 191–215.

Gnana, Jennifer, "Saudi Aramco Raises 2022 Capex Guidance to $40-$50 Bil," *S&P Global*, March 2022.

Gormsen, Niels Joachim and Kilian Huber, "Corporate Discount Rates," *American Economic Review*, June 2025, *115* (6), 2001–49.

Gowrisankaran, Gautam, Ashley Langer, and Wendan Zhang, "Policy Uncertainty in the Market for Coal Electricity: The Case of Air Toxics Standards," *Journal of Political Economy*, forthcoming.

Haltiwanger, John and Joseph E. Harrington, "The Impact of Cyclical Demand Movements on Collusive Behavior," *RAND Journal of Economics*, 1991, *22* (1), 89–106.

Hamilton, James D., "Understanding Crude Oil Prices," *Energy Journal*, 2009, *30* (2), 179–206.

Hart, Rob and Daniel Spiro, "The Elephant in Hotelling's Room," *Energy Policy*, 2011, *39*, 7834–7838.

Heal, Geoffrey and Wolfram Schlenker, "Coase, Hotelling, and Pigou: The Incidence of a Carbon Tax and CO2 Emissions," NBER working paper #26086, 2019.

Herfindahl, Orris C., "Depletion and Economic Theory," in Mason Gaffney, ed., *Extractive Resources and Taxation*, University of Wisconsin Press, 1967.

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

International Energy Agency, "World Energy Outlook 2023," 2023.

_ , "Oil 2024: Analysis and Forecast to 2030," 2024.

_ , "The Implications of Oil and Gas Field Decline Rates," 2025.

Jacobs, Trent, "Life After 5: How Tight-Oil Wells Grow Old," *Journal of Petroleum Technology*, January 2020.

Jain, Gautam and Luisa Palacios, "Investing in Oil and Gas Transition Assets En Route to Net Zero," March 2023. Columbia SIPA Center on Global Energy Policy commentary.

Jensen, Svenn, Kristina Mohlin, Karen Pittel, and Thomas Sterner, "An Introduction to the Green Paradox: The Unintended Consequences of Climate Policies," *Review of Environmental Economics and Policy*, 2015, *9*, 246–265.

Kaiser, Mark J. and Brian F. Snyder, "Empirical Relations Characterize Rig Day Rates," *Offshore Magazine*, October 2012.

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

Kollenbach, Gilbert and Mark Schopf, "Unilaterally Optimal Climate Policy and the Green Paradox," *Journal of Environmental Economics and Management*, 2022, *113*, 102649.

Krautkraemer, Jeffrey A., "Nonrenewable Resource Scarcity," *Journal of Economic Literature*, 1998, *36* (4), 2065–2107.

Krupnick, Alan, Richard Morgenstern, Nathan Balke, Stephen P.A. Brown, Ana María Herrera, and Shashank Mohan, "Oil Supply Shocks, US Gross Domestic Product, and the Oil Security Premium," Resources for the Future report 2017.

Lemoine, Derek, "Green Expectations: Current Effects of Anticipated Carbon Pricing," *Review of Economics and Statistics*, 2017, *99* (3), 499–513.

_ , "Climbing the Energy Ladder: How Energy Resources Hinder, Facilitate, and Fuel Economic Growth," May 2025. NBER working paper 33864.

Lynch, Michael, "Declining Production At Saudi Arabia's Largest Oil Field Is Not Cause For Concern," *Forbes*, April 2019.

Manley, David and Patrick R.P. Heller, "Risky Bet: National Oil Companies in the Energy Transition," 2021. Natural Resource Governance Institute paper.

McKinley, Andrew, "Short Circuited: Hotelling Meets EPA Rulemaking," working paper, 2024.

Nereim, Vivian and Rebecca F. Elliott, "United Arab Emirates Says It Will Leave OPEC in Blow to Oil Cartel," *The New York Times*, April 2026.

Newell, Richard G. and Brian C. Prest, "The Unconventional Supply Boom: Aggregate Price Response from Microdata," *Energy Journal*, 2019, *40* (3), 1–30.

Norman, Maya A. and Wolfram Schlenker, "Empirical Tests of the Green Paradox for Climate Legislation," 2024. NBER working paper 32405.

Nystad, Arild N., "Rate Sensitivity and the Optimal Choice of Production Capacity of Petroleum Reservoirs," *Energy Economics*, 1987, *9* (1), 37–45.

Omar, Abeer Abu, "Saudi Arabia Issues $12 Billion Bond as EM Nations Load Up," *Bloomberg*, January 2024.

Organization of the Petroleum Exporting Countries, "2023 World Oil Outlook," 2023.

Pittel, Karen, Rick van der Ploeg, and Cees Withagen, *Climate Policy and Nonrenewable Resources: The Green Paradox and Beyond*, Cambridge, MA: MIT Press, 2014.

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*, 2022, *9*, 681–720.

_ **, Harrison Fell, Deborah Gordon, and TJ Conway**, "Estimating the Emissions Reductions from Supply-Side Fossil Fuel Interventions," *Energy Economics*, 2024, *136*, 107720.

Reed, Stanley, "Angola to Quit OPEC in Dispute Over Oil Production Levels," *The New York Times*, December 2023.

Rogner, Hans-Holger, Roberto F. Aguilera, Cristina L. Archer, Ruggero Bertani, S.C. Bhattacharya, Maurice B. Dusseault, Luc Gagnon, Helmut Haberl, Monique Hoogwijk, Arthur Johnson, Mathis L. Rogner, Horst Wagner, Vladimir Yakushev, Doug J. Arent, Ian Bryden, Fridolin Krausmann, Peter Odell, Christoph Schillings, and Ali Shafiei, "Energy Resources and Potentials," in "Global Energy Assessment - Toward a Sustainable Future," Cambridge, UK and New York, NY, USA: Cambridge University Press, 2012, pp. 425–512.

Rotemberg, Julio J. and Garth Saloner, "A Supergame-Theoretic Model of Price Wars During Booms," *American Economic Review*, 1986, *76* (3), 390–407.

Ryan, Carol, "Shell's New Strategy Avoids the Toughest Questions," *Wall Street Journal*, 14 June 2023.

Rystad Energy, "Recoverable Oil Reserves Top 1,600 Billion Barrels, Capable of Warming the Planet an Extra 0.2°C by 2100," June 2023. Online; `https://www.rystadenergy.com/news/recoverable-oil-reserves-billion-barrels-warming-planet`; accessed 9 July, 2024.

Salant, Stephen W., "Exhaustible Resources and Industrial Structure: A Nash-Cournot Approach to the World Oil Market," *Journal of Political Economy*, 1976, *84* (5), 1079–1093.

_ , "Imperfect Competition in the International Energy Market: A Computerized Nash-Cournot Model," *Operations Research*, 1982, *30* (2), 252–280.

Salzman, Avi, "COP28 Deal Points to Oil's Eventual Decline and Investment Risks," *Barron's*, 13 December 2023.

Sinclair, Peter, "High Does Nothing and Rising is Worse: Carbon Taxes Should Keep Declining To Cut Harmful Emissions," *The Manchester School*, 1992, *60* (1), 41–52.

Sinn, Hans-Werner, "Public Policies Against Global Warming: A Supply Side Approach," *International Tax and Public Finance*, 2008, *15* (4), 360–394.

Slade, Margaret E. and Henry Thille, "Whither Hotelling: Tests of the Theory of Exhaustible Resources," *Annual Review of Resource Economics*, August 2009, *1* (1), 239–259.

Society of Petroleum Evaluation Engineers, "Survey of Economic Parameters Used in Property Evaluation," 2023.

Sorrell, Steve, Jamie Speirs, Roger Bentley, Richard Miller, and Erica Thompson, "Shaping the global oil peak: A review of the evidence on field sizes, reserve growth, decline rates and depletion rates," *Energy*, 2012, *37* (1), 709–724. 7th Biennial International Workshop "Advances in Energy Studies".

Stiglitz, Joseph E., "Monopoly and the Rate of Extraction of Exhaustible Resources," *American Economic Review*, 1976, *66* (4), 655–661.

Texas Comptroller of Public Accounts, "2026 School District Property Value Study: Discount Rate Range for Oil and Gas Properties," 2026.

Thompson, Andrew C., "The Hotelling Principle, Backwardation of Futures Prices and the Values of Developed Petroleum Reserves — the Production Constraint Hypothesis," *Resource and Energy Economics*, 2001, *23*, 133–156.

Toews, Gerhard and Alexander Naumov, "The Relationship Between Oil Price and Costs in the Oil Industry," *Energy Journal*, 2015, *36* (1), 237–254.

U.S. Energy Information Administration, "What is OPEC+ and How Is It Different From OPEC?," 2023. [Online; `https://www.eia.gov/todayinenergy/detail.php?id=56420`; accessed 9 July, 2024].

— , "International Petroleum and Other Liquids," https://www.eia.gov/international/data/world/petroleu and-other-liquids/annual-crude-and-lease-condensate-reserves?; accessed 10 July, 2024.

— , "Petroleum & Other Liquids Spot Prices," https://www.eia.gov/dnav/pet/pet_pri_spt_s1_a.htm; accessed 11 July, 2024.

— , "Tight Oil Production Estimates By Play," https://www.eia.gov/outlooks/steo/xls/Fig42.xlsx; accessed 9 July, 2024.

van der Meijden, Gerard, Hassan Benchekroun, Frederick van der Ploeg, and Cees Withagen, "Do Strong Oligopolies Reverse Green Paradox Effects?," *European Journal of Political Economy*, 2023, *79*, 102434.

van der Ploeg, Frederick and Cees Withagen, "Global Warming and the Green Paradox: A Review of Adverse Effects of Climate Policies," *Review of Environmental Economics and Policy*, 2015, *9* (2), 285–303.

Vreugdenhil, Nicholas, "Booms, Busts, and Mismatch in Capital Markets: Evidence from the Offshore Oil and Gas Industry," *Journal of Political Economy*, forthcoming.