

NBER WORKING PAPER SERIES

THE END OF OIL

Ryan Kellogg

Working Paper 33207

<http://www.nber.org/papers/w33207>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue

Cambridge, MA 02138

November 2024

I thank Christopher Costello, Derek Lemoine, Stephen Salant, and Rick van der Ploeg for comments on an earlier draft of this manuscript, and I thank a variety of seminar and conference participants for their comments and suggestions. I do not have any material financial relationships that are relevant to the research discussed in this manuscript. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2024 by Ryan Kellogg. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The End of Oil
Ryan Kellogg
NBER Working Paper No. 33207
November 2024
JEL No. L71, Q40, Q54

ABSTRACT

It is now plausible to envision scenarios in which global demand for crude oil falls to essentially zero by the end of this century, driven by improvements in clean energy technologies, adoption of stringent climate policies, or both. This paper asks what such a demand decline, when anticipated, might mean for global oil supply. One possibility is the well-known “green paradox”: because oil is an exhaustible resource, producers may accelerate near-term extraction in order to beat the demand decline. This reaction would increase near-term CO₂ emissions and could possibly even lead the total present value of climate damages to be greater than if demand had not declined at all. However, because oil extraction requires potentially long-lived investments in wells and other infrastructure, the opposite may occur: an anticipated demand decline reduces producers' investment rates, decreasing near-term oil production and CO₂ emissions. To evaluate whether this disinvestment effect outweighs the green paradox, or vice-versa, I develop a tractable model of global oil supply that incorporates both effects, while also capturing industry features such as heterogeneous producers, exercise of market power by low-cost OPEC producers, and marginal drilling costs that increase with the rate of drilling. I find that for model inputs with the strongest empirical support, the disinvestment effect outweighs the traditional green paradox. In order for anticipation effects on net to substantially increase cumulative global oil extraction, I find that industry investments must have short time horizons, and that producers must have discount rates that are comparable to U.S. treasury bill rates.

Ryan Kellogg
University of Chicago
Harris School of Public Policy
1307 East 60th Street
Chicago, IL 60637
and NBER
kellogg@uchicagoedu.mail.onmicrosoft.com

A code repository is available at https://github.com/kelloggrk/Public_EndOfOil

1 Introduction

For roughly a century, the vast majority of global transportation demand has been fueled by crude oil and petroleum products. Oil’s dominance as a transportation fuel is now, for the first time, 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 uncertainty about how quickly and deeply these developments will reduce crude oil demand, but it is now plausible to envision scenarios in which global oil demand falls 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 towards 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 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

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

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 speed with which demand decreases, the rates at which producers discount future profits, the remaining volumes of reserves, the magnitudes of production decline rates, and the extent of producers’ exercise of market power—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 Pifti (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 climate policy, finding that prices fall in most, but not all cases.⁶ This result is consistent

²A 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.

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

⁴Cairns 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.

⁵Lemoine (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.

⁶Norman and Schlenker (2024) finds that passage of renewable energy policies increases oil prices.

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 future drilling opportunities, and disinvestment arises because future demand reductions reduce the anticipated payoff from newly-drilled wells’ production streams. I nest this 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), and Balke, Jin and Yücel (2024) 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 shifts inward over 75 years until the quantity demanded is zero at any price. I also examine alternatives that involve faster, slower, delayed, 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.⁷ My focus on long-run decarbonization, motivated by increasing discussion of “zero 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.

Along a declining demand path, oil producers’ decisions will be affected both by the direct impact of lower current-period demand and by their anticipation of future demand reduc-

⁷My 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 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.

tions. I isolate the effect of anticipation—the net of the green paradox and disinvestment—and quantify its impact on cumulative extraction. To do so, I model producers’ beliefs in two ways. First, I simulate their drilling and extraction paths while assuming that they foresee the entire demand decline. Second, 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 these two simulations isolates how producers’ anticipation of the demand decline affects cumulative drilling and extraction. 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.

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 Synder (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.⁸ 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), and van der Meijden et al. (2023) 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.⁹ As discussed in section 3, I calibrate model parameters for each

⁸See, for instance, Herfindahl (1967), Asker, Collard-Wexler and De Loecker (2019), Heal and Schlenker (2019), and Asker, Collard-Wexler, De Canniere, 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.

⁹An 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

type based on academic literature and industry reports. Core OPEC, which includes Kuwait, Saudi Arabia, and the United Arab Emirates, acts strategically and is distinguished by a low annual real discount rate of 3% in the reference case model, reflecting these countries’ (and their national oil companies’) easy access to capital. I model all other types as price-takers, with a discount rate of 9%. Shale oil is distinguished by its wells’ high rate of production decline relative to the other types: 30% rather than 8% per year.

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 27.1%, relative to a baseline scenario in which demand does not fall. Oil is “left in the ground” for all four resource types. Relative to a scenario in which the demand decline is unanticipated, the anticipated decline reduces cumulative extraction by an additional 4.8%. Thus, the disinvestment effect is overall stronger than the green paradox. However, there is heterogeneity across resource types. Disinvestment is most important for non-core OPEC+ and conventional non-OPEC oil, since these resources have long-lived assets and small initial scarcity rents at baseline.¹⁰ Neither disinvestment nor the green paradox is important for shale oil, since it has a small initial scarcity rent and because its investments are short-lived. In contrast, the green paradox is important for core OPEC, since this resource features a large initial scarcity rent at baseline that stems from its low discount rate in the reference case model.

Alternative specifications, presented in section 4.2, explore how the above result is sensitive to model inputs. Consistent with intuition, increasing the rate at which wells’ production declines weakens the disinvestment effect, and increasing producers’ initial stock of reserves weakens the green paradox. When I configure the model so that investments pay off immediately, thereby shutting down disinvestment, I find that the green paradox is small in magnitude: anticipation increases cumulative extraction by 2.3%. This result arises because at baseline, the initial scarcity rents for all resources other than core OPEC are only a few dollars per barrel. To obtain an economically large green paradox, I find that I must both shut down disinvestment and decrease producers’ discount rates so that their baseline scarcity rents are large. For instance, in a specification in which investments pay off immedi-

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 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.

¹⁰Adelman (1990) and Hart and Spiro (2011) argue that the behavior of petroleum markets has been consistent with a lack of scarcity rents.

ately and all producers have a 3% discount rate, anticipation increases cumulative extraction by 11.0%.

Sections 4.3 and 4.4 discuss 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 scenarios in which the onset of the demand decline is delayed by 5 years, reflecting the possibility that newly announced climate policy may have a lagged implementation. These results are related to empirical event studies such as Norman and Schlenker (2024) and show that anticipation of a delayed demand decline can lead near-term oil prices to either increase or decrease, depending on model parameters, though even in the latter case the effect of anticipation on cumulative extraction can still be negative. I also consider a scenario in which anticipation of the demand decline causes the OPEC cartel to break down, à la Rotemberg and Saloner (1986), 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 slope 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 setting 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 from 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

The global oil price $P(t)$ at time 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. I treat time as continuous in the exposition of the model, though I will ultimately discretize time when I implement the model quantitatively in section 3.

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 towards zero; (3) the drilling of new wells is price-responsive; and (4) the marginal cost of drilling investments increases in the rate of drilling.¹¹ 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 towards zero—at essentially all times.¹²

Following Anderson et al. (2018), I model production from drilled wells as following an exponential decline, at a rate $\lambda > 0$. Letting $D(t)$ denote the rate of drilling investment in new production capacity, the rate of production $Q(t)$ evolves per equation (1):

$$\dot{Q}(t) = D(t) - \lambda Q(t). \tag{1}$$

Flow drilling costs are $C(D(t))$, 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.

Let $r > 0$ denote the discount rate, and let $R(t) \geq 0$ denote the stock of remaining undrilled capacity. 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

¹¹For 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).

¹²Anderson 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 circumstances do not arise in this paper, so I take it as given that production is always equal to wells' capacity constraint.

equilibrium for a representative firm, is given by:

$$\max_{D(t)} \int_0^{\infty} e^{-rt} \left[\int_0^{Q(t)} P(s, t) ds - C(D(t)) \right] dt \quad (2)$$

subject to

$$\dot{Q}(t) = D(t) - \lambda Q(t), \quad Q(0) \text{ given}, \quad (3)$$

$$\dot{R}(t) = -D(t), \quad R(0) \text{ given} \quad (4)$$

$$D(t) \geq 0, \quad R(t) \geq 0. \quad (5)$$

Note that this problem has two state variables: the production rate $Q(t)$ and the stock $R(t)$.

2.1.2 Solution

The current-value Hamiltonian for the maximization problem is:

$$H = \int_0^{Q(t)} P(s, t) ds - C(D(t)) + \theta(t)(D(t) - \lambda Q(t)) + \mu(t)(-D(t)), \quad (6)$$

where $\theta(t)$ and $\mu(t)$ are the co-state variables on the state variables $Q(t)$ and $R(t)$, expressed as current values. Necessary conditions are:

$$D(t) \geq 0, \quad \theta(t) - c(D(t)) - \mu(t) \leq 0, \quad D(t)(\theta(t) - c(D(t)) - \mu(t)) = 0 \quad (7)$$

$$\dot{\theta}(t) = -P(t) + (r + \lambda)\theta(t) \quad (8)$$

$$\dot{\mu}(t) = r\mu(t) \quad (9)$$

$$Q(t)\theta(t)e^{-rt} \rightarrow 0 \text{ and } R(t)\mu(t)e^{-rt} \rightarrow 0 \text{ as } t \rightarrow \infty. \quad (10)$$

The co-state variable $\theta(t)$ represents the shadow value of production capacity at time t , accounting for the fact that capacity that is newly invested today will produce oil at a declining rate into the future. Solving the differential equation of FOC (8) (with the endpoint given by the transversality condition (TVC) of FOC (10)) thus yields that $\theta(t)$ is equal to the present value of the future stream of oil prices, from t to infinity, discounted at $r + \lambda$:

$$\theta(t) = \int_t^{\infty} P(\tau) e^{-(r+\lambda)(\tau-t)} d\tau. \quad (11)$$

With the above understanding of $\theta(t)$, FOC (7) has the interpretation that whenever the rate of drilling investment is strictly positive, the difference between the marginal value

of capacity $\theta(t)$ and the marginal cost of drilling $c(D(t))$ must equal the shadow value of undrilled capacity $\mu(t)$. If the stock constraint on undrilled wells, $R(t) \geq 0$, is binding in the optimal program, then $\mu(t)$ will be strictly positive and per FOC (9) 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 in finite time (assuming demand has a finite choke price), with production continuing on a decline to zero beyond the time at which drilling ceases.


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 drilling $D(0)$, as illustrated in figure 1.¹³ First, anticipation of lower demand will reduce the initial scarcity value $\mu(0)$ of reserves. Per equation (7), $D(0)$ is determined by equating the sum of $\mu(0)$ and the marginal drilling cost $c(D(0))$ with $\theta(0)$, the marginal value of capacity. Because $c(D(0))$ is strictly increasing, the decrease in $\mu(0)$ 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 $\theta(0)$ of capacity that will still be producing after demand has fallen. This reduction in $\theta(0)$ 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 $\mu(0)$ more than it reduces $\theta(0)$, 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 clear. For instance, if λ is small 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

¹³I thank Christopher Costello for suggesting this figure.

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



Note: The figure shows the potential effects of an anticipated decrease in demand at some time $t > 0$. The initial scarcity rent $\mu(0)$ falls to $\mu'(0)$, and the initial marginal value of capacity $\theta(0)$ falls to $\theta'(0)$. The change in the initial rate of drilling from $D(0)$ to $D'(0)$ follows the first-order condition of equation (7), which sets $c(D(0)) + \mu_0 = \theta_0$. Depending on the relative magnitudes of the reductions in $\mu(0)$ and $\theta(0)$, the model can deliver a result that the green paradox outweighs disinvestment (panel (a), $D'(0) > D(0)$), or the opposite (panel (b), $D'(0) < D(0)$).

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.

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 sensitivity analyses in the quantitative exercise. Second, the model is fast to

compute because it only includes two state variables—current production and the stock of undrilled wells—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 below, 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.

2.2 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. The model is inspired by the original open-loop Nash-Cournot extraction game proposed in Salant (1976), which I extend by incorporating investment and allowing for multiple types of fringe producers.¹⁴ Each fringe type is indexed by $f \in \{1, \dots, N - 1\}$, and within each type the firms are modeled as a representative firm. Let m denote the dominant cartel, and let $i \in \{1, \dots, N - 1, m\}$ index the full set of producers, with $d_i(t)$ and $q_i(t)$ denoting the investment and production rates of type i at time t . Each type's discount rate, production decline rate, initial stock of undrilled wells, and drilling marginal cost function are denoted r_i , λ_i , $R_i(0)$, and $c_i(d_i(t))$. Let $D(t) = \sum_i d_i(t)$ and $Q(t) = \sum_i q_i(t)$ denote the total rates of drilling and extraction at t .

The dominant cartel + competitive fringe equilibrium is then defined as follows:

1. Each representative fringe firm f chooses a time path of drilling investment $\{d_f(t)\}$ to maximize

$$\int_0^\infty e^{-r_f t} [q_f(t)P(t) - C_f(d_f(t))] dt \quad (12)$$

taking as given the time path of prices $\{P(t)\}$, its initial production rate $q_f(0)$, and its initial stock $R_f(0)$. $q_f(t)$ obeys $\dot{q}_f(t) = d_f(t) - \lambda_f q_f(t)$, and $R_f(t)$ obeys $\dot{R}_f(t) = -d_f(t)$. The firms must satisfy $d_f(t) \geq 0$ and $R_f(t) \geq 0 \forall t$.

2. The dominant cartel m chooses a time path of drilling investment $\{d_m(t)\}$ to maximize

$$\int_0^\infty e^{-r_m t} [q_m(t)P(Q(t), t) - C_m(d_m(t))] dt \quad (13)$$

¹⁴I follow Salant (1976) in adopting the Cournot assumption that the firms choose strategies simultaneously rather than as Stackelberg leader-followers. In the Stackelberg model, the dominant firm will be tempted to revise its investment path downwards after $t = 0$, since doing so would be a best response to the fringe's investments. The dominant firm's Stackelberg strategy is therefore not credible absent a commitment mechanism. This problem does not arise in Nash-Cournot because all firms' open-loop strategies are dynamically consistent (Hanley, Shogren and White, 1997).

taking as given its initial production rate $q_m(0)$ and stock and $R_m(0)$, and the fringe firms' initial production rates $q_f(0)$ and time paths of investment $\{d_f(t)\}$, for $f \in \{1, \dots, N-1\}$. $q_m(t)$ obeys $\dot{q}_m(t) = d_m(t) - \lambda_m q_m(t)$, and $R_m(t)$ obeys $\dot{R}_m(t) = -d_m(t)$. The cartel must satisfy $d_m(t) \geq 0$ and $R_m(t) \geq 0 \forall t$.

3. The $P(t)$ equal inverse demand: $P(t) = P(Q(t), t) \forall t$.

The necessary conditions for each fringe firm are the same as those given in equations (7) through (10) for the representative firm version of the model. For the dominant cartel, equation (8) is different because it does not behave as a price-taker. Instead, the equation of motion for $\theta_m(t)$ is

$$\dot{\theta}_m(t) = - \left(P(t) + q_m(t) \frac{\partial P(Q(t), t)}{\partial Q(t)} \right) + (r_m + \lambda_m) \theta_m(t), \quad (14)$$

implying that $\theta_m(t)$ is the present value of future marginal revenue, discounted at $r_m + \lambda_m$:

$$\theta_m(t) = \int_t^\infty \left(P(\tau) + q_m(\tau) \frac{\partial P(Q(\tau), \tau)}{\partial Q(\tau)} \right) e^{-(r_m + \lambda_m)(\tau - t)} d\tau. \quad (15)$$

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.¹⁵ 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, Benckekroun 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.

The model is open-loop in that the game calls on the firms to choose full time series of drilling paths $\{d_i(t)\}$ at $t = 0$. This structure restricts firms from manipulating other firms' actions by changing their own state variables, as opposed to a game defined by a Markov perfect equilibrium, in which firms' actions at each time t are governed by Markov (also referred to as "feedback") strategies that are state-dependent policy functions. However, I show in appendix A that with one strategic player, the equilibrium investment and extraction out-

¹⁵Stiglitz (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.

comes from the open-loop game are the same as those from a Markov perfect equilibrium.¹⁶ The argument draws from Eswaran and Lewis (1985) and relies on the assumption that all but one player act as price-takers.¹⁷ Intuitively, because the fringe producers do not believe that their actions or state variables affect the equilibrium price path or the actions of any other producer, they 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 specification of the problem.¹⁸ This equivalence result is useful because the open-loop game is considerably easier to solve computationally than the game involving Markov strategies.¹⁹

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.²⁰ 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 towards 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.

For computation, I switch to discrete time because the time series of equilibrium prices

¹⁶This equivalence of outcomes with a single strategic player has been proven for a specific case in Benčekroun and Withagen (2012). I have been unable to find a general proof in the literature.

¹⁷Benčekroun 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 $N \rightarrow \infty$. Thus, price-taking needs to be assumed *a priori*.

¹⁸I thank Stephen Salant for thoughts that helped clarify this intuition.

¹⁹The 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.

²⁰Benčekroun 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.

$\{P(t)\}$ is a non-analytic function of time, making it difficult to evaluate the $\theta_i(t)$ in continuous time.²¹ I use a time interval between periods of just 6 months, so that any loss of precision is not substantial. I model a one-period delay between investment in new capacity and production from that capacity, reflecting the fact that oilfield investments typically require time-to-build.²² I also henceforth denote variables as having time subscripts rather than as functions of time; for example, q_{it} denotes the extraction rate of resource i in period t .

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 choice. The inner loop takes as given the firms’ shadow values μ_{it} and solves for the equilibrium series of d_{it} , q_{it} , P_t , and θ_{it} 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_{it} . That is, the equilibrium 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 discount rates, investment time-horizons, investment costs, and strategic versus price-taking behavior. I summarize the types below, where all production volumes are for 2023 and are compiled from country-level data in Energy Institute (2024) and US shale oil data in U.S. Energy Information Administration (2024c).²³

I. Core OPEC (15.4 mmbbl/d): This resource type includes three Middle East low-cost,

²¹For 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 time interval between periods goes to zero.

²²Newell 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.

²³I use Energy Institute’s (2024) “crude oil and condensate” production data, which exclude natural gas liquids and refinery processing gain.

high-output OPEC members: Kuwait, Saudi Arabia, and the United Arab Emirates.²⁴ 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”.²⁵ 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. These producers are also distinguished by arguably having a relatively low discount rate.

- II. Non-core OPEC+ (30.4 mmbbl/d): This group includes all other OPEC+ members.²⁶ Relative to core OPEC, I model these producers as having higher discount rates (equal to those of types III and IV below) and 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.

²⁴Balke et al. (2024) also includes Qatar in its core OPEC group, but Qatar is no longer an OPEC member (Energy Institute, 2024).

²⁵Russia would be another country to potentially include 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 and low discount rates. Still, a natural sensitivity that I later examine is the inclusion of Russia in the core OPEC producer type.

²⁶OPEC+ 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).

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).²⁷ 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, Fell, Gordon and Conway (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 (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 towards 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. 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). The reference case demand decline scenario then shifts the affine demand curve inward towards zero at a constant rate, relative to the baseline demand path, for 75 years.²⁸ 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) estimates an annual decline rate for onshore Texas conventional wells of 8%. Thompson (2001) finds

²⁷I use the 2023 average Brent crude oil spot price.

²⁸More precisely, I start with the 2023 demand curve and compute a time path that decreases the demand intercept by the same amount per period, such that demand reaches zero after 75 years elapse. I then increase demand by 0.44% per year through 2030 (demand still decreases from 2023–2030 on net because the rate of decline exceeds the baseline growth rate).

Table 1: Calibrated parameters for the reference case model

Parameter		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	λ_i	per year	0.08	0.08	0.08	0.30
Discount rate	r_i	per year	0.03	0.09	0.09	0.09
Reserves	x_{i0}	billion bbl	394	542	563	125
Initial 2023 drilling rate*	d_{i0}	mmbbl/d	0.66	1.31	1.22	1.39
Drilling cost intercept	a_i	\$/bbl	5.00	10.00	20.00	30.00
Drilling cost intercept*	α_i	\$bn per mmbbl/d	16.39	21.52	43.04	26.42
Drilling cost slope	g_i	\$/bbl per mmbbl/d	30.39	53.89	50.09	36.26
Drilling cost slope*	γ_i	\$bn per (mmbbl/d) ²	99.62	115.97	107.79	31.94
Drilling cost elasticity*			0.80	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. $\alpha_i = \Omega_i a_i$ and $\gamma_i = \Omega_i g_i$, where $\Omega_i = Y/(1000(\hat{r}_i + \hat{\lambda}_i))$, Y is the number of days per period, and \hat{r}_i and $\hat{\lambda}_i$ are per-period rates. The drilling cost elasticity is evaluated at the 2023 investment rate for each type. “mmbbl/d” denotes millions of barrels per day.

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%. Taking these estimates together, my reference case uses an annual decline rate of $\lambda_i = 8\%$ 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 $\lambda_{IV} = 30\%$.

3.4 Discount rates

I consider discount rates separately for private oil companies (resource types III and IV), core OPEC (type I), and non-core OPEC+ (type II). The discount rate is especially important for governing the magnitude of both the green paradox and disinvestment, so I ultimately consider a range of alternative values around the reference case calibration.

Manley and Heller (2021) surveys industry estimates of annual discount rates for private oil companies, finding a range of 10–14.8% nominal. Damodaron (2024) computes a real weighted average cost of capital of 7.34% via CAPM. Cairns and Smith (2019) uses a real annual discount rate of 8%, “typical of what an [integrated oil company] would use to evaluate an investment”, and Anderson et al. (2018) uses 10%, citing Society of Petroleum Evaluation Engineers (1995). I adopt an annual real discount rate of 9%, roughly in the center of these estimates, for the reference case model.

It is not possible to observe the discount rate that core OPEC countries make when considering oilfield investments. One possibility is that they invest similarly to private oil companies. But it is also possible that their discount rate is comparable to a riskless discount rate, given that major decisions are controlled by national oil companies that can leverage their countries’ strong bond ratings and easy access to capital markets. For instance, yields on Saudi Arabian government bonds have been about 1 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). My reference case therefore assumes that core OPEC countries use a real annual discount rate of $r_I = 3\%$. I also investigate alternative specifications in which core OPEC is assumed to use the same discount rate as private oil companies.

Non-core OPEC+ consists primarily of countries whose production is controlled by national oil companies. Unlike core OPEC, the other countries in OPEC+ 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. While it is thus clear that non-core OPEC+ producers are unlikely to discount the future at anything like a riskless rate, it is less clear what is the appropriate rate to use. My reference case assumes that they use the same discount rate as private oil companies, 9%, though a higher rate is plausible.

3.5 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).²⁹ 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).

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.³⁰ 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.

Because the model’s scarce resource is the stock of remaining undrilled capacity rather than the volume of oil reserves, I convert the above reserves in barrels, which I denote x_{i0} , to reserves measured in available capacity, R_{i0} . Letting $\eta_i = 1000\hat{\lambda}_i/Y$, where $\hat{\lambda}_i$ is the per-period decline rate and Y is the number of days per period,³¹ the available reserves of undrilled capacity for resource i , in mmbbl/d, are given by $R_{i0} = \eta_i x_{i0} - q_{i0}$.

3.6 Marginal investment costs

There exist a variety of industry estimates 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 don’t use these breakevens 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

²⁹Rystad 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.

³⁰I 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.

³¹ $\hat{\lambda}_i = 1 - (1 - \lambda_i)^{1/T}$, where $T = 2$ is the number of simulated periods per year.

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: $c_i(d_{it}) = \alpha_i + \gamma_i d_{it}$, with α_i and $\gamma_i > 0$. I set the α_i based on industry estimates, and I then calibrate the γ_i so that the model’s simulated equilibrium under the baseline demand scenario matches each type’s observed 2023 investment and production. I execute this latter 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 rates or discount rates), I re-calibrate the γ_i so that the baseline simulation reproduces 2023 observed production and investment.

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 α_i 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 α_i based on break-evens that are at or near the bottom of the reported ranges within each resource type. I use break-evens of \$5/bbl for core OPEC, \$10/bbl for non-core OPEC+, \$20/bbl for non-OPEC conventional, and \$30/bbl for shale. Denote these cost intercepts, expressed in break-even form, as a_i .

The model’s marginal cost intercepts α_i are the marginal cost of an investment in capacity, expressed in units of \$ billion per mmbbl/d of investment. To convert the a_i break-even prices in \$/bbl to the α_i , I compute the investment cost such that a capacity investment would break even, in the sense of having a present discounted profit of zero, were the price a_i paid for all subsequent production. Define $\Omega_i = Y/(1000(\hat{r}_i + \hat{\lambda}_i))$, where Y is the number of days per period, and \hat{r}_i and $\hat{\lambda}_i$ are per-period rates.³² Then $\alpha_i = \Omega_i a_i$. Note that in alternative specifications in which I vary the λ_i or r_i , I hold the a_i fixed at their calibrated values and let the α_i change with the changing Ω_i . This approach holds fixed the drilling cost intercepts in terms of the cost per amortized barrel of production.³³

I calibrate the marginal cost slopes γ_i by matching the model’s initial 2023 equilibrium

³²The formula for Ω_i accounts for the fact that in the model, production starts the period after capacity is invested. $\hat{r}_i = (1 + r_i)^{1/T} - 1$, where $T = 2$ is the number of simulated periods per year.

³³If I were to instead hold the α_i fixed while I changed the λ_i or r_i , I would (perhaps substantially) change the economic cost of drilling. For instance, in the alternative specifications with larger values of λ_i , holding the α_i fixed would lead drilling to be expensive relative to wells’ lifetime revenue (which falls with λ_i), hence causing simulated drilling at baseline to fall below the level that matches observed 2023 investment.

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 capacity addition 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.³⁴ I then find the set of γ_i for which the model’s simulated initial investment rates match these d_{i0} .

The resulting reference case estimates of the γ_i , given in table 1, are difficult to interpret because they are in units of \$ billion per (mmbbl/d)². I also provide a more interpretable $g_i = \gamma_i/\Omega_i$, which can be interpreted as the change in the \$/bbl break-even price for a 1 mmbbl/d increase in capacity investment. The estimated g_i are 30.39, 53.89, 50.09, and 36.26 \$/bbl per mmbbl/d for resource types I, II, III, and IV, respectively. Yet another way to interpret these values is to compute the implied elasticities of drilling costs to investment at the observed initial investment rate. These elasticities are 0.80, 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 Synder, 2012; Toews and Naumov, 2015; Anderson et al., 2018; Vreugdenhil, forthcoming).³⁵

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 drilled. Reserves of undrilled shale resources are exhausted first due to their

³⁴Per the discrete time analog of equation (3), the necessary initial investment rate for each resource type is given by $d_{i0} = q_{i0}((1.0044 + \lambda_i)^{1/T} - 1)$. The initial investment rates are 0.66, 1.31, 1.22, and 1.39 mmbbl/d of new capacity addition for types I, II, III, and IV, respectively.

³⁵Kaiser and Synder (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

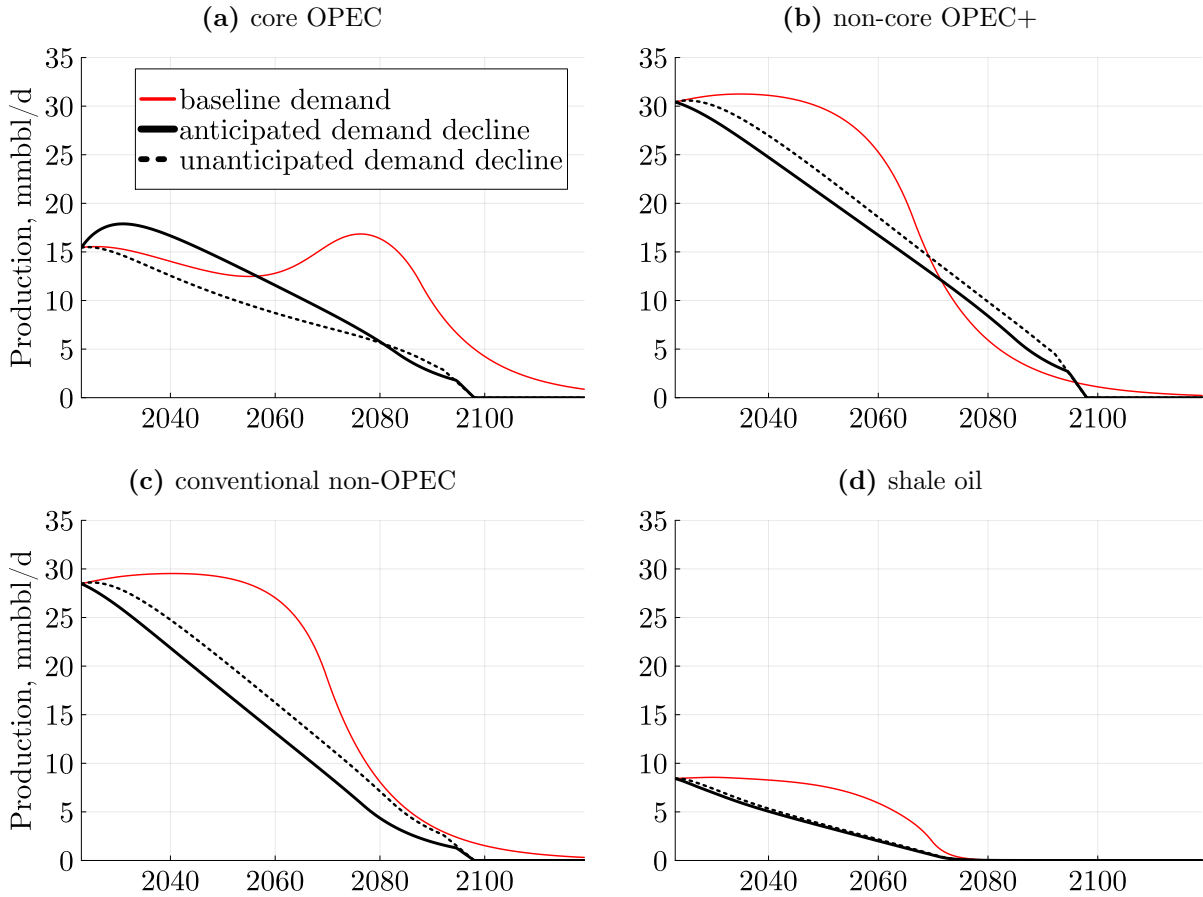


Note: “mmbbl/d” denotes millions of barrels per day. See section 3 for a discussion of the reference case calibration. In the demand decline counterfactuals, the demand curve shifts inward to zero over 75 years.

relatively high cost and short investment cycle. All undrilled resources outside of core OPEC are exhausted in 2070, and core OPEC drilling ceases in 2087. Production then declines asymptotically to zero, fully exhausting global oil reserves (1624 billion bbl) in the limit, with the price of oil rising to the demand choke price.

Resources are exhausted most slowly by core OPEC in the baseline simulation for two reasons: it acts strategically in the world oil market, and its relatively low discount rate

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 to zero over 75 years.

gives it a large initial Hotelling scarcity rent of \$30.88/bbl in 2023.³⁶ Initial scarcity rents for the other three resource types are all less than \$4/bbl.

Under an anticipated 75-year decline in global oil demand, cumulative global drilling investment falls by 35.2%, leading to a reduction in cumulative global oil extraction of 27.1% (from 1624 to 1185 billion bbl).³⁷ The initial scarcity rent for all four resource types falls to

³⁶The raw scarcity rents from the model are in units of dollars per mmbbl/d of capacity. I convert these scarcity rents to more interpretable \$/bbl by dividing them by the Ω_i .

³⁷As 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 figure 2 as the kink in the production path a few years before production reaches zero.

zero because all four types leave oil in the ground. The production decrease, in percentage terms, is greatest for shale oil (41.5%) and smallest for non-core OPEC+ (17.1%).

An important feature of this simulation is that aggregate investment in the initial period is slightly lower than at baseline: 4.50 rather than 4.59 mmbbl/d (initial production is identical to q_{i0} for each resource type by construction). Demand in the initial period is unchanged; thus, this result stems from producers' anticipation of the coming decline in demand. The fact that anticipation results in reduced rather than increased investment indicates that the disinvestment effect is outweighing the green paradox.

I next quantify the effects of anticipation on cumulative drilling and extraction, not just first-period drilling, by evaluating the full time path of drilling and extraction when the demand decline is unanticipated. To do so, in each simulated period I model producers as believing that demand will not decline going forward.³⁸ I then isolate the net anticipation effect by subtracting the drilling and production outcomes of this simulation from those obtained when I simulate producers as anticipating the demand decline. I find that cumulative drilling and extraction are 7.3% and 4.8% lower, respectively, when the demand decline is anticipated rather than unanticipated. Thus, in the reference case model, aggregating across resource types, the disinvestment effect outweighs the green paradox.

Figure 3 shows how the difference between oil extraction paths under the anticipated and unanticipated demand declines varies across resource types. Disinvestment, and hence reduced output, is greatest for non-core OPEC+ and non-OPEC conventional producers, which have long investment time horizons and low initial scarcity rents. Shale producers exhibit nearly zero effect of anticipation on net. For these producers, the disinvestment effect is weak because production from new shale wells declines rapidly, and the green paradox is weak because they have small initial scarcity rents at baseline. In contrast, anticipatory effects result in increased production from core OPEC. For these producers, the disinvestment effect is outweighed by the green paradox and by their reaction to the decreases in drilling and production by the other producers.

The reference case results for cumulative global extraction 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 oil extraction, where I apply a discount rate of 3% (the same as that assumed for core OPEC

³⁸In 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.

in the reference case model). I find that the anticipated demand decline reduces present discounted emissions by 18.9% relative to baseline and by 2.7% relative to the unanticipated demand decline.

4.2 Alternative decline rates, reserves, and discount rates

This section presents results from alternative specifications that moderate the relative strengths of the disinvestment effect and green paradox by varying the model’s assumed production decline rates, reserve volumes, and discount rates. These results are summarized in rows 2–11 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 γ_i per the discussion in section 3.6, so that the initial investment and production rates in the baseline simulations always match observed values in 2023.³⁹ Table C.1 in the appendix presents the calibrated drilling cost elasticities for each specification.

Per the intuition from section 2.1, the disinvestment effect is strengthened the longer-lived are investments. Accordingly, row 2 of table 2 shows that when I re-calibrate 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 cumulative extraction by more than in the reference case: 5.3% rather than 4.8%. Further decreasing each region’s production decline rate to 6% causes the net anticipation effect to further increase in magnitude, so that cumulative extraction falls by 5.5%. On the other hand, assigning shale’s 30% decline rate to all resources reduces the disinvestment effect so that it is nearly balanced with the green paradox on net: anticipation reduces cumulative extraction by just 1.4% (row 4 of table 2). Row 5 shows that when I take this intuition to the limit by removing all investment dynamics, anticipation increases cumulative extraction by 2.3%. This result is consistent with the intuition that in this specification, the green paradox is the only remaining force governing anticipation effects. Its magnitude is modest, however, reflecting the fact that at baseline, the initial scarcity rents are small for all but core OPEC producers.

Row 6 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 3.5 for a discussion of reserve estimates). Following the intuition from section 2.1, this increase in reserves lowers producers’ initial scarcity rents at baseline, weakening the green paradox. Accordingly, anticipation reduces cumulative extraction by 7.9% in this model. Conversely, assuming that reserves are just 1283 billion bbl leads to increased scarcity rents at baseline and a

³⁹For example, in the specification with lower initial reserves (row 7 of table 2), firms’ simulated initial drilling would be lower than observed 2023 drilling were I to hold the γ_i fixed. The γ_i are therefore smaller in this specification than in the reference case, as shown in appendix table C.1.

Table 2: Simulated cumulative oil extraction and present discounted extraction for the reference case and alternative specifications, aggregated across resource types

Model	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Q	PVQ	Q	decline	Q	decline	(3)-(1)	(4)-(2)	(3)-(5)	(4)-(6)
1. Reference case	1624	798	1185	648	1244	665	-27.1%	-18.9%	-4.8%	-2.7%
Alternative declines, reserves, and discount rates										
2. All regions decline at 8%	1624	798	1191	651	1258	672	-26.7%	-18.4%	-5.3%	-3.1%
3. All regions decline at 6%	1624	792	1221	661	1292	684	-24.8%	-16.6%	-5.5%	-3.4%
4. All regions decline at 30%	1624	802	1134	630	1151	632	-30.2%	-21.4%	-1.4%	-0.3%
5. No investment	1624	803	1117	627	1092	608	-31.2%	-21.9%	+2.3%	+3.0%
6. High reserves (2602 billion bbl)	2602	949	1160	635	1260	672	-55.4%	-33.1%	-7.9%	-5.5%
7. Low reserves (1283 billion bbl)	1283	704	1212	663	1158	640	-5.5%	-5.8%	+4.6%	+3.6%
8. All resource types discount at 9%	1624	807	1163	636	1248	668	-28.4%	-21.2%	-6.8%	-4.8%
9. All resource types discount at 3%	1624	783	1264	692	1221	655	-22.2%	-11.6%	+3.5%	+5.6%
10. No investment + low reserves	1283	712	1141	639	1072	601	-11.1%	-10.2%	+6.4%	+6.5%
11. No investment + discounting at 3%	1624	782	1214	682	1093	608	-25.2%	-12.8%	+11.0%	+12.2%
Alternative demand declines										
12. 50-year demand decline	1624	798	803	521	878	557	-50.5%	-34.7%	-8.5%	-6.4%
13. 100-year demand decline	1624	798	1530	730	1529	727	-5.8%	-8.6%	+0.1%	+0.4%
14. Decline delayed 5 years	1624	798	1266	687	1311	697	-22.1%	-13.9%	-3.5%	-1.5%
15. Decline delayed 5 years; 9% discounting	1624	807	1242	674	1320	702	-23.5%	-16.4%	-5.9%	-3.9%
16. 15% of demand remains after decline	1624	798	1586	660	1602	677	-2.4%	-17.3%	-1.0%	-2.5%

Note: “Q” denotes cumulative oil extraction in billions of barrels, and “PVQ” denotes present (2023) discounted cumulative extraction (also in billions of bbl), discounted at 3%. See section 3 for a discussion of the reference case calibration, and sections 4.2 and 4.3 for discussions of the alternative specifications. For each alternative specification, a figure analogous to figure 2 is presented in appendix C. For each specification, the re-calibrated drilling cost elasticities are presented in appendix table C.1.

stronger green paradox. For instance, the baseline scarcity rent for conventional non-OPEC in this specification is \$9.70/bbl, rather than just \$2.77/bbl in the reference case. This scarcity causes the green paradox to outweigh the disinvestment effect, so that anticipation increases cumulative extraction by 4.6% (row 7 of table 2).


An increase in the discount rate attenuates both the disinvestment effect and the green paradox. Rows 8 and 9 of table 2 show that the impact of the discount rate on the green paradox is relatively more important, according with intuition that this effect operates on a longer time horizon than does the disinvestment effect. When I increase core OPEC's discount rate from 3% to 9% (equal to the discount rate of all other resource types in the reference case), the green paradox weakens so that anticipation reduces cumulative extraction by 6.8% (versus 4.8% in the reference case). Alternatively, in row 9 I set the discount rate of all resource types to 3%, leading them to have substantial scarcity rents at baseline. For example, the initial scarcity rent for conventional non-OPEC becomes \$33.70/bbl, rather than \$2.77/bbl in the reference case. Now, anticipation increases cumulative production by 3.5%. This result is smaller in magnitude than net anticipation under the “low reserves” scenario, despite the large initial scarcity rents, because the low discount rate in this simulation is also increasing the magnitude of the disinvestment effect.

Figure 4 elaborates on these results by showing how the effect of anticipation on cumulative extraction varies with the discount rate that is assumed to hold for all resource types. Starting from a discount rate of 3%, for which anticipation causes an increase in cumulative extraction, increasing the discount rate attenuates the green paradox more than it attenuates disinvestment. The net anticipation effect becomes negative for a discount rate of 4% and continues to become more negative for higher discount rates, up to a discount rate of 10%. At this point, the initial baseline scarcity rents are quite small, so that further increases in the discount rate primarily serve to decrease the magnitude of the disinvestment effect. Thus, the net anticipation effect decreases in magnitude as the discount rate increases beyond 10%, as shown in figure 4.

Finally, I study two specifications that both remove all investment dynamics (like the model from table 2, row 5) and employ parameters that lead to high initial scarcity rents at baseline. Row 10 of table 2 combines row 5 with the “low reserves” model from row 7, finding that the green paradox increases cumulative extraction by 6.4%. Row 11 combines row 5 with the 3% discounting model from row 9, finding that the green paradox increases cumulative extraction by 11.0%. This last model leads to a significant green paradox because the low discount rates lead the model to infer substantial scarcity rents in the baseline simulation: \$21.06, \$25.23, \$23.05, and \$24.43/bbl for resource types I, II, III, and IV, respectively.⁴⁰

⁴⁰In the model with no investment and low reserves, initial baseline scarcity rents are \$29.99, \$3.32, \$3.67,

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 γ_i , 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 cumulative global extraction when the demand decline is anticipated relative to when it is unanticipated.

This paper’s model is therefore able to deliver an economically large green paradox, but only when the disinvestment effect is shut down and discount rates are set to rates comparable to those for riskless securities, leading to initial scarcity rents that are large.

4.3 Alternative models for the decline in oil demand

This section considers alternative paths for the decline in demand, relative to the reference case that assumes that demand falls at a constant rate, reaching zero in 75 years. I first consider an accelerated 50-year and a lengthened 100-year decline. The results from these simulations are shown in rows 12 and 13 of table 2. The speed of the demand decline, not surprisingly, substantially affects cumulative extraction, which I find to be 803 billion barrels under an anticipated 50-year decline and 1530 billion barrels under an anticipated 100-year

and \$5.35/bbl for resource types I, II, III, and IV, respectively. Changing assumptions about reserves is therefore less powerful at affecting the magnitudes of the initial scarcity rents than is changing assumptions about discount rates.

decline (out of 1624 billion bbl of total reserves).⁴¹ When I focus on anticipation effects, I find that relative to the reference case, they become more negative with a 50-year decline (a 8.5% decrease in cumulative extraction) and essentially zero with a 100-year decline (a 0.1% increase 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 under the anticipated demand decline are zero in both the accelerated decline and reference case models.

I also study an alternative shape for the demand decline path. Holding fixed the time required for demand to fall all the way to zero (75 years in the reference case), I model a scenario in which the onset of the decline is delayed: demand follows the baseline path for 5 years before declining smoothly to zero (over the remaining 70 years). This scenario captures the possibility that once a new climate policy or technology is announced, it may be several years before oil demand is actually impacted.

Under a delayed demand decline, the disinvestment effect is weakened during the first few years of the simulation, since producers anticipate that their $t = 0$ investments will substantially depreciate away by the time the demand decline begins. Consequently, I find that under the delayed decline path, the initial investment rate is primarily affected by the green paradox and is thus larger than that in the baseline simulation (see appendix figure C.14). 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 evidence from event studies of climate policy announcements (Adolfson et al., 2024; Norman and Schlenker, 2024) that find increases in oil company investments and a decrease in the oil price. However, as time passes in this simulation and the onset of the demand decline draws near and then occurs, the disinvestment effect strengthens. As a result, I find that anticipation still reduces cumulative extraction on net, though by less than in the reference case: 3.5% rather than 4.8% (see row 14 of table 2).

I also study a specification in which the demand decline is delayed and all resource types discount at 9% (table 2, row 15). In this simulation, the green paradox is sufficiently weak that investment and extraction during the initial 5-year delay period is now lower than that in the baseline simulation (consistent with Bogmans et al. (2024)), and oil prices are higher. For instance, at the start of year 5, the simulated oil price at baseline is \$83.18/bbl, while that under the anticipated demand decline is \$84.50/bbl. Thus, while the price increase is modest in magnitude, the model is capable of producing the result predicted by oil executives surveyed in Federal Reserve Bank of Dallas (2023) that an anticipated decline in demand

⁴¹In the 100-year anticipated decline model, the non-OPEC resources leave oil in the ground so that their initial scarcity rents fall to zero, while the OPEC resources are fully exhausted.

can increase rather than decrease the near-term price of oil.

Finally, I 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. Row 16 of table 2 therefore presents results from a model in which the demand decline follows the reference case until 15% of the initial demand remains, after which point the demand curve is stationary. When this partial demand decline is anticipated, producers eventually achieve near-complete extraction, since only shale oil resources are incompletely extracted.⁴² However, because long-run demand is so low, drilling and extraction proceed over hundreds of years, creating a large divergence between cumulative extraction (1586 billion bbl) and present discounted extraction (660 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 1.0% and 2.5%, respectively.

4.4 Additional alternative specifications

A final suite of alternative specifications explores how the simulation results depend on market power, the oil demand function, and the slope of the marginal drilling cost function. Simulated outcomes for cumulative production are given in table 3. Overall, these outcomes are only modestly different from those in the reference case.

I first examine how the results differ when I increase core OPEC’s market power by including Russia’s initial production and reserves volume in this resource group.⁴³ The results of this simulation are presented in row 17 of table 3. To isolate the effects of market power when running this simulation, rather than also capture the effect of increasing the size of the resource base to which I assigned a low discount rate in the reference case, I assign all resource types (including core OPEC) a discount rate of 9%. Thus, these results should be compared to those from row 8 of table 2 rather than to the reference case.

I find that including Russia in core OPEC modestly decreases total cumulative extraction under an anticipated decline (from 1163 to 1157 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. On net for this specification, anticipation causes cumulative extraction to be 6.5% less under an anticipated demand decline than under an unanticipated decline. This anticipation effect is slightly

⁴²Shale oil is the only resource whose marginal cost curve intercept a_i (in \$/bbl terms) sits above the long-run demand choke price.

⁴³The “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).

Table 3: Simulated cumulative oil extraction and present discounted extraction for alternative specifications discussed in section 4.4, aggregated across resource types

Model	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Q	PVQ	Q	Anticipated decline	Q	Unanticipated decline	(3)-(1)	(4)-(2)	(3)-(5)	(4)-(6)
	Q	PVQ	Q	PVQ	Q	PVQ	(1)	(2)	(5)	(6)
1. Reference case	1624	798	1185	648	1244	665	-27.1%	-18.9%	-4.8%	-2.7%
Alternative specifications										
17. Russia in core OPEC, 9% discounting	1624	808	1157	633	1237	663	-28.8%	-21.7%	-6.5%	-4.5%
18. No market power	1624	800	1190	649	1255	669	-26.7%	-18.9%	-5.2%	-3.1%
19. No market power in anticipated decline	1624	798	1215	670	1244	665	-25.2%	-16.0%	-2.3%	+0.7%
20. High demand elasticity (-0.6)	1624	799	1182	647	1251	668	-27.2%	-19.1%	-5.5%	-3.3%
21. Low demand elasticity (-0.4)	1624	797	1188	649	1233	661	-26.9%	-18.5%	-3.7%	-1.7%
22. No demand growth	1624	788	1150	630	1213	650	-29.2%	-20.0%	-5.2%	-3.0%
23. Higher demand growth (1.5% per year)	1624	822	1271	690	1315	702	-21.8%	-16.0%	-3.4%	-1.7%
24. High cost function intercepts	1624	797	1111	626	1171	643	-31.6%	-21.5%	-5.2%	-2.6%

Note: “Q” denotes cumulative oil extraction in billions of barrels, and “PVQ” denotes present (2023) discounted cumulative extraction (also in billions of bbl), discounted at 3%. See section 3 for a discussion of the reference case calibration, and section 4.4 for discussions of the alternative specifications. The “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 constant. The “No market power” specification, in contrast, models core OPEC as behaving competitively at baseline and for both demand declines. Demand growth in the baseline simulations is through 2030. The last row uses values of \$20, \$25, \$35, and \$45/bbl for the cost function intercepts a_i for $i = I, II, III,$ and $IV,$ respectively. For each specification, the re-calibrated drilling cost elasticities are presented in appendix table C.1.

smaller in magnitude than the 6.8% result reported in row 8 of table 2, in which Russia was not included in core OPEC.

Row 18 of table 3 reports outcomes when I eliminate OPEC’s market power altogether by modeling the core OPEC resource type as behaving competitively (discount rates for row 18 and all later rows are as in the reference case). Mirroring the results from row 17 in which 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 5.2% on net, slightly more than in the reference case.

My next specification examines an alternative mechanism for a green paradox: what if an anticipated demand decline causes the OPEC cartel to lose its discipline and begin behaving competitively, and hence increase its collective extraction rate? The intuition behind such a potential outcome comes from Rotemberg and Saloner (1986): because cartel cooperation is motivated by the potential for future profits, an anticipated decrease in demand will tighten and potentially break the cartel members’ incentive compatibility constraint, causing a breakdown 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. The results of this exercise are presented in table 3, row 19, where extraction at baseline and under an unanticipated demand decline are the same as in the reference case, by construction. The cartel’s breakdown increases cumulative extraction under an anticipated demand decline to 1215 billion bbl (versus 1185 billion bbl in the reference case), though this value remains less than (by 2.3%) cumulative extraction under an unanticipated decline.

Rows 20 through 23 of table 3 present results from alternative specifications for 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 22), anticipation of the demand decline reduces cumulative extraction by 5.2%, more than the 4.8% 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 effect of anticipation on cumulative output is still negative on net (row 23).

Finally, in table 3, row 24 I consider an alternative specification in which I make producers’ marginal drilling cost functions “flatter” by increasing their cost intercepts α_i . Instead of using the (\$ per barrel equivalent) reference case intercepts of \$5/bbl for core OPEC,

\$10/bbl for non-core OPEC+, \$20/bbl for non-OPEC conventional, and \$30/bbl for shale, I use values of \$20, \$25, \$35, and \$45/bbl, respectively. Consequently, the resource types' drilling cost elasticities fall to 0.27, 0.69, 0.57, and 0.44 (versus 0.80, 0.88, 0.75, and 0.63 in the reference case specification).

Flatter cost functions increase firms' responsiveness to anticipated future demand changes, strengthening 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 modest strengthening of the disinvestment effect over the green paradox.

5 Concluding discussion

This paper develops a model to assess how global oil producers might respond to an anticipated, long-run decline in oil demand towards zero. My emphasis in constructing the model is on capturing two important, 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, disinvestment effects outweigh the green paradox. In the reference case model, an anticipated demand decline reduces cumulative oil production by 4.8% more than what would be reduced from an unanticipated decline. Disinvestment occurs primarily for conventional oil resources outside of core OPEC, since these resources do not have large scarcity rents, and their extraction involves investments with long time-horizons. Core OPEC resources have large initial scarcity rents and thus exhibit a green paradox on net. For shale oil, investments are short-cycle and scarcity rents are small, so overall anticipation effects are small.

In order for this paper's model to deliver an economically large green paradox on net, I find that two conditions need to hold. First, oilfield investments need to be very short-cycle, at least as short-cycle as shale, thereby neutralizing the disinvestment effect. Second, initial scarcity rents under the baseline (non-declining) demand forecast must be large throughout the globe, not just for core OPEC. Satisfying this condition requires an assumption that all producers, including private firms and national oil companies of countries with poor access to international credit markets, have discount rates that are similar to U.S. treasury bill rates. These two conditions seem unlikely to hold individually, let alone jointly, suggesting

that the green paradox is unlikely to 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 and reserves 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 enhanced oil recovery. Accounting for exploration and surface facilities, which can be long-term investments that also involve substantial time-to-build, would enhance the disinvestment effect, while the impact of enhanced oil recovery on this paper’s results is not clear *a priori*. My model also assumes a perfectly inelastic fixed supply of undrilled reserves rather than allow for investment costs that increase as the stock is depleted, which would potentially lead to incomplete exhaustion even with stationary demand. Results in van der Ploeg (2016) suggest that the fixed stock assumption strengthens the green paradox. 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. 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 CO₂ 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.

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

- Hamilton, James D.**, “Understanding Crude Oil Prices,” *Energy Journal*, 2009, 30 (2), 179–206.
- Hanley, Nick, Jason F. Shogren, and Ben White**, *Environmental Economics in Theory and Practice*, Red Globe Press, 1997.
- 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.
- 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. Synder**, “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.

- 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 Economic Statistics*, 2017, 99 (3), 499–513.
- 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.
- McKinley, Andrew**, “Short Circuited: Hotelling Meets EPA Rulemaking,” working paper, 2024.
- 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., 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.
- 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.
- Society of Petroleum Evaluation Engineers**, “Survey of Economic Parameters Used in Property Evaluations,” 1995.
- Stiglitz, Joseph E.**, “Monopoly and the Rate of Extraction of Exhaustible Resources,” *American Economic Review*, 1976, *66* (4), 665–661.
- 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/petroleum-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, “Second-Best Carbon Taxation in the Global Economy: The Green Paradox and Carbon Leakage Revisited,” *Journal of Environmental Economics and Management*, 2016, 78, 85–105.

- **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.

Online appendix for “The End of Oil”

A Equivalent outcomes of the open-loop Nash-Cournot and Markov perfect equilibria

Here I show that the open-loop Nash-Cournot investment equilibrium defined in section 2.2 of the main text yields the same outcomes as a Markov perfect equilibrium (MPE). The proof draws from arguments in Eswaran and Lewis (1985) regarding dynamic oligopoly; I apply those arguments here to the case of a dominant cartel and competitive fringe, in the case of an investment game.

I first define the MPE. I then show that under the assumed behavior of the fringe firms, the MPE outcomes are the same as those from the open-loop game.

The MPE that I adopt is similar to that proposed in Benckroun and Withagen (2012), adapted here to incorporate capital investment and heterogeneity among the fringe firms. Let $q(t)$ and $R(t)$ denote the sets of state variables for the fringe firms at t (e.g., $q(t) = \{q_1(t), q_2(t), \dots, q_{N-1}(t)\}$). $Q(t) = \sum_i q_i(t)$ as in the main text. Define the equilibrium as follows:

1. Each representative fringe firm f chooses a strategy $d_f(t) = \phi_f(t, q_f(t), R_f(t))$ to maximize

$$\int_t^\infty e^{-r_f \tau} [q_f(\tau)P(\tau) - C_f(\phi_f(\tau, q_f(\tau), R_f(\tau)))] d\tau \quad (\text{A.1})$$

taking as given the time path of prices $\{P(\tau)\}$, its initial production rate $q_f(t)$, and its initial stock $R_f(t)$. $q_f(\tau)$ obeys $\dot{q}_f(\tau) = d_f(\tau) - \lambda_f q_f(\tau)$, and $R_f(\tau)$ obeys $\dot{R}_f(\tau) = -d_f(\tau)$. The firms must satisfy $d_f(\tau) \geq 0$ and $R_f(\tau) \geq 0 \forall \tau$.

2. The dominant cartel m chooses a strategy $d_m(t) = \phi_m(t, q_m(t), R_m(t), q(t), R(t))$ to maximize

$$\int_t^\infty e^{-r_m \tau} [q_m(\tau)P(Q(\tau), \tau) - C_m(\phi_m(\tau, q_m(\tau), R_m(\tau), q(\tau), R(\tau)))] d\tau \quad (\text{A.2})$$

taking as given the fringe firms' strategies $\phi_f(t, q_f(t), R_f(t))$ for $f \in \{1, \dots, N-1\}$, the initial production rates $q_m(t)$ and $q(t)$, and the initial stocks $R_m(t)$ and $R(t)$. $q_m(\tau)$ obeys $\dot{q}_m(\tau) = d_m(\tau) - \lambda_m q_m(\tau)$, and $R_m(\tau)$ obeys $\dot{R}_m(\tau) = -d_m(\tau)$. The cartel must satisfy $d_m(\tau) \geq 0$ and $R_m(\tau) \geq 0 \forall \tau$.

3. The $P(t)$ equal inverse demand: $P(t) = P(Q(t), t) \forall t$.

I highlight that 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 Benckroun and Withagen (2012) and, as argued therein, is sensible because the alternatives would imply that fringe firms are unaware of their choices' impact (small as it may be) on the current price but are aware of how their stocks may influence future prices.

It then follows that each fringe firm's strategy $\phi_f(t, q_f(t), R_f(t))$ devolves to a choice of drilling rate as a function of time $\phi_f(t)$, given $\{P(t)\}$, $q_f(0)$, and $R_f(0)$. Thus, the fringe firms' Markov strategies can be represented as open-loop strategies $\{d_f(t)\}$, with $d_f(t) = \phi_f(t)$.

Then for the cartel, conditioning its strategy on the fringe's state variables $q(t)$ and $R(t)$ adds no value because given $\{P(t)\}$ and the initial states, the fringe firms' drilling time paths are fully determined. Thus, the cartel's problem can be simplified to a choice of strategy $\phi_m(t, q_m(t), R_m(t))$, taking the fringe firms' drilling paths $\{d_f(t)\}$ as given. But this problem is then the same as choosing a drilling rate as a function of time $\phi_m(t)$, given $\{d_f(t)\}$, $q_m(0)$, and $R_m(0)$, which is an open-loop strategy. We therefore have that for both the cartel and fringe firms, the equilibrium Markov strategies are just functions of time, implying that the outcomes of the Markov perfect and open-loop equilibria will coincide.

I emphasize that the key to the above argument is the assumed myopic behavior of the fringe firms. Absent that, firms would have an incentive to influence each others' future actions via their choices regarding their own future state variables. These strategic interactions would cause equilibrium outcomes to differ across the two game definitions, as shown in simulations in Eswaran and Lewis (1985) and Benckroun and Withagen (2012) (though the quantitative differences tend to be minor). Here, however, the fact that there is only a single strategic player implies that the game in Markov strategies does not involve these strategic interactions, so that its equilibrium strategies collapse to open-loop strategies.

B Computational algorithm

I implement the model in discrete time, with a time interval of 6 months between periods. I simulate a total of 200 years, which is sufficient for simulated production to decline to essentially zero in all specifications.¹

The solution algorithm uses two nested loops, similar to the algorithm for production games proposed in Salant (1982). It proceeds as follows:

1. Make an initial guess of the initial shadow values μ_{i0} for each resource type.
2. Execute the inner loop to solve for equilibrium drilling, production, and prices given the provisional μ_{i0} :
 - (a) Make an initial guess of the equilibrium price and cartel production time series $\{P_t^0\}$ and $\{q_{mt}^0\}$.
 - (b) Compute the θ_{it} from the $\{P_t^0\}$, $\{q_{mt}^0\}$, and each resource's decline and discount rate, using equation (15) for m and equation (11) for each fringe type f .
 - (c) In each period t , use the first-order condition (7) to solve for the d_{it} , given the θ_{it} and the current shadow values μ_{it} (in turn based on the provisional initial shadow values μ_{i0}).
 - (d) Compute the production paths $\{q_{it}\}$ using the discrete time analogs of the equations of motion $\dot{q}_{it} = d_{it} - \lambda_i q_{it}$, the q_{i0} , and the $\{d_{it}\}$ from the previous step. Sum across resource types to obtain the aggregate production time series $\{Q_t\}$.
 - (e) Compute a new price series $\{P_t\}$ using the demand curve $P(Q_t, t)$ and the $\{Q_t\}$ from the previous step.
 - (f) For each t , compute $P'_t = (1 - \phi)P_t^0 + \phi P_t$ and $q'_{mt} = (1 - \phi)q_{mt}^0 + \phi q_{mt}$, where $\phi \in (0, 1]$. I use $\phi = 0.2$. Larger values of ϕ converge more quickly but can lead to numerical instability.
 - (g) If $\{P'_t\}$ is sufficiently close to $\{P_t\}$ in the sup norm, stop. Otherwise, return to step (a), replacing the $\{P_t^0\}$ and $\{q_{mt}^0\}$ with the $\{P'_t\}$ and $\{q'_{mt}\}$. I use a tolerance of 1e-05 \$/bbl.
3. Solve the mixed complementarity problem for the μ_{i0} . Specifically, given a set of provisional shadow values μ_{i0} , let Z_i denote the difference between initial reserves

¹The one exception is the specification in which demand declines only to 15% of its initial level (results shown in row 16 of table 2). I run this specification for 300 years because the persistent low level of demand in the decline scenarios leads to a long period of low but non-trivial production.

R_{i0} for each resource and cumulative drilling computed from the inner loop of the routine. The equilibrium shadow values are then those that simultaneously satisfy: (1) $\mu_{i0} \geq 0$; (2) $Z_i \geq 0$; and (3) $\mu_{i0}Z_i = 0$, for all i . I solve these transversality conditions simultaneously using a tolerance of 1e-04 billion bbl on the Z_i . In practice I search on the $\log \mu_{i0}$ rather than the μ_{i0} ; I treat the condition $\mu_{i0} = 0$ as being satisfied for any resource type whenever $\mu_{i0} < 1e-04$ \$ per mmbbl/d.


C Additional tables and figures

Table C.1: Calibrated drilling cost elasticities

Parameter	Value by type			
	I	II	III	IV
1. Reference case	0.80	0.88	0.75	0.63
2. All regions decline at 8%	0.80	0.88	0.75	0.62
3. All regions decline at 6%	0.79	0.87	0.75	0.62
4. All regions decline at 30%	0.83	0.88	0.76	0.63
5. No investment	0.84	0.88	0.75	0.63
6. High reserves (2602 billion bbl)	0.88	0.88	0.76	0.64
7. Low reserves (1283 billion bbl)	0.68	0.87	0.74	0.60
8. All resource types discount at 9%	0.90	0.88	0.75	0.63
9. All resource types discount at 3%	0.82	0.80	0.64	0.46
10. No investment + low reserves	0.77	0.87	0.75	0.61
11. No investment + discounting at 3%	0.84	0.83	0.66	0.48
17. Russia in core OPEC, 9% discounting	0.83	0.88	0.75	0.63
18. No market power	0.91	0.88	0.75	0.63
20. High demand elasticity (-0.6)	0.85	0.88	0.75	0.63
21. Low demand elasticity (-0.4)	0.64	0.88	0.75	0.63
22. No demand growth	0.81	0.88	0.75	0.63
23. Higher demand growth (1.5% per year)	0.77	0.88	0.75	0.63
24. High cost function intercepts	0.27	0.69	0.57	0.44


Note: Type I is core OPEC, type II is non-core OPEC+, type III is conventional non-OPEC, and type IV is shale. The drilling cost elasticities are evaluated at the 2023 investment rate for each type. The row numbering corresponds to the specifications presented in tables 2 and 3 in the main text. Rows not included here are specifications that use the same calibrated values for the γ_i as the reference case, by construction. Demand growth in the baseline simulations is through 2030. The last row uses values of \$20, \$25, \$35, and \$45/bbl for the cost function intercepts a_i for $i = \text{I, II, III, and IV}$, respectively.

Figure C.2: Simulation results, aggregated across resource types. Model in which production from all resources declines at 8% per year




Note: Figure corresponds to row 2 of table 2 in the main text. “mmbbl/d” denotes millions of barrels per day. In the demand decline counterfactuals, the demand curve shifts inward to zero over 75 years.

Figure C.3: Simulation results, aggregated across resource types. Model in which production from all resources declines at 6% per year




Note: Figure corresponds to row 3 of table 2 in the main text. “mmbbl/d” denotes millions of barrels per day. In the demand decline counterfactuals, the demand curve shifts inward to zero over 75 years.

Figure C.4: Simulation results, aggregated across resource types. Model in which production from all resources declines at 30% per year




Note: Figure corresponds to row 4 of table 2 in the main text. “mmbbl/d” denotes millions of barrels per day. In the demand decline counterfactuals, the demand curve shifts inward to zero over 75 years.

Figure C.5: Simulation results, aggregated across resource types.
 Model in which there are no investment dynamics




Note: Figure corresponds to row 5 of table 2 in the main text. “mmbbl/d” denotes millions of barrels per day. In the demand decline counterfactuals, the demand curve shifts inward to zero over 75 years.

Figure C.6: Simulation results, aggregated across resource types.
 Model in which initial reserves are high




Note: Figure corresponds to row 6 of table 2 in the main text. “mmbbl/d” denotes millions of barrels per day. In the demand decline counterfactuals, the demand curve shifts inward to zero over 75 years.

Figure C.7: Simulation results, aggregated across resource types.
 Model in which initial reserves are low




Note: Figure corresponds to row 7 of table 2 in the main text. “mmbbl/d” denotes millions of barrels per day. In the demand decline counterfactuals, the demand curve shifts inward to zero over 75 years.

Figure C.8: Simulation results, aggregated across resource types.
 Model in which all resource types discount at 9%



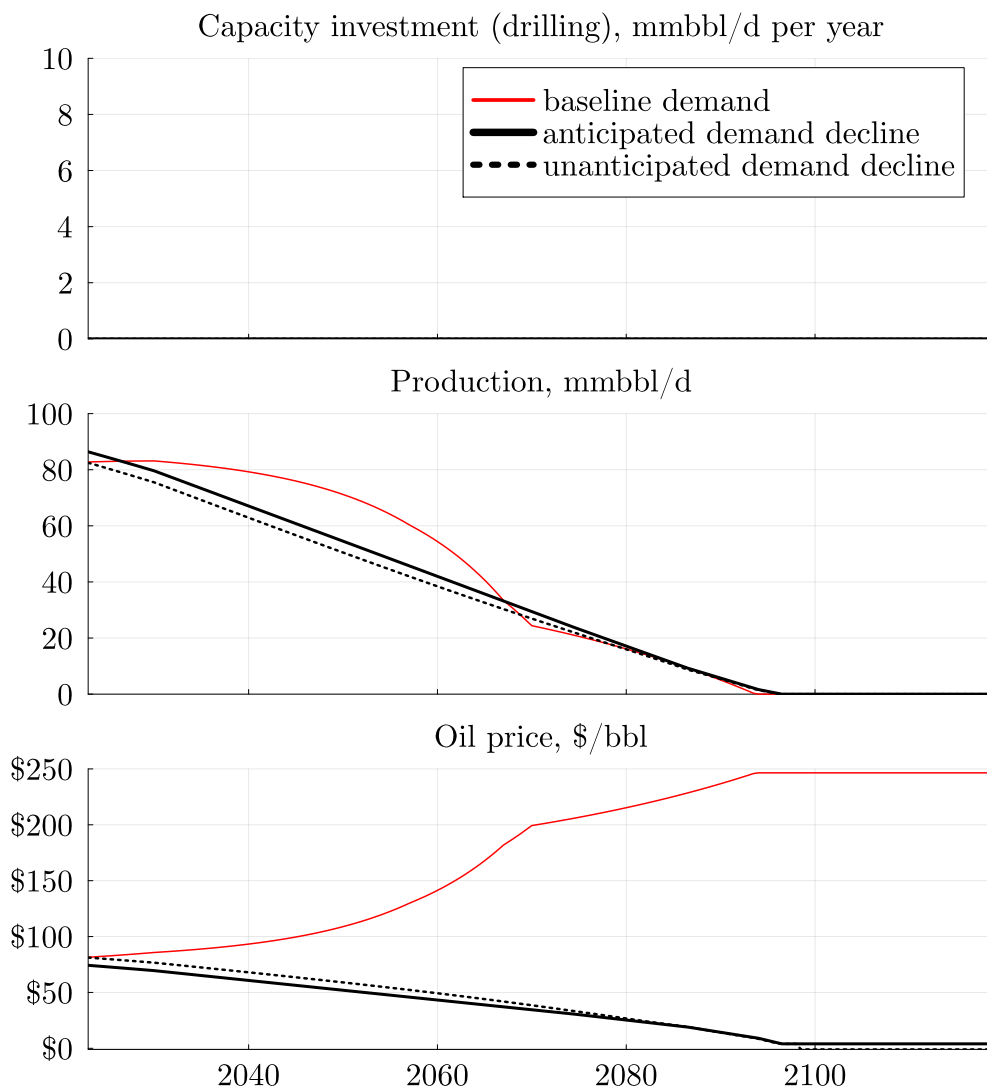
Note: Figure corresponds to row 8 of table 2 in the main text. “mmbbl/d” denotes millions of barrels per day. In the demand decline counterfactuals, the demand curve shifts inward to zero over 75 years.

Figure C.9: Simulation results, aggregated across resource types.
 Model in which all resource types discount at 3%




Note: Figure corresponds to row 9 of table 2 in the main text. “mmbbl/d” denotes millions of barrels per day. In the demand decline counterfactuals, the demand curve shifts inward to zero over 75 years.

Figure C.10: Simulation results, aggregated across resource types. Model in which there are no investment dynamics and initial reserves are low




Note: Figure corresponds to row 10 of table 2 in the main text. “mmbbl/d” denotes millions of barrels per day. In the demand decline counterfactuals, the demand curve shifts inward to zero over 75 years.

Figure C.11: Simulation results, aggregated across resource types.
 Model in which there are no investment dynamics and all resource types discount at 3%




Note: Figure corresponds to row 11 of table 2 in the main text. “mmbbl/d” denotes millions of barrels per day. In the demand decline counterfactuals, the demand curve shifts inward to zero over 75 years.

Figure C.12: Simulation results, aggregated across resource types.
 Model in which demand declines over 50 years




Note: Figure corresponds to row 12 of table 2 in the main text. “mmbbl/d” denotes millions of barrels per day. In the demand decline counterfactuals, the demand curve shifts inward to zero over 50 years.

Figure C.13: Simulation results, aggregated across resource types.
 Model in which demand declines over 100 years




Note: Figure corresponds to row 13 of table 2 in the main text. “mmbbl/d” denotes millions of barrels per day. In the demand decline counterfactuals, the demand curve shifts inward to zero over 100 years.

Figure C.14: Simulation results, aggregated across resource types. Model in which the start of the demand decline is delayed 5 years




Note: Figure corresponds to row 14 of table 2 in the main text. “mmbbl/d” denotes millions of barrels per day. In the demand decline counterfactuals, demand follows its baseline trajectory over 5 years and then declines inward, reaching zero 75 years after the start of the simulation.

Figure C.15: Simulation results, aggregated across resource types.
 Model in which the start of the demand decline is delayed 5 years and all resource types
 discount at 9%



Note: Figure corresponds to row 15 of table 2 in the main text. “mmbbl/d” denotes millions of barrels per day. In the demand decline counterfactuals, demand follows its baseline trajectory over 5 years and then declines inward, reaching zero 75 years after the start of the simulation.

Figure C.16: Simulation results, aggregated across resource types.
 Model in which demand declines incompletely, leaving 15% of initial demand in perpetuity



Note: Figure corresponds to row 16 of table 2 in the main text. “mmbbl/d” denotes millions of barrels per day. In the demand decline counterfactuals, the demand curve shifts inward to zero on the same trajectory as in the reference case (an 75-year decline), but the decline stops when demand is at 15% of its initial level.

References

- Benchekroun, Hassan and Cees Withagen**, “On Price-Taking Behavior in a Nonrenewable Resource Cartel-Fringe Game,” *Games and Economic Behavior*, 2012, *76*, 355–374.
- Eswaran, Mukesh and Tracy Lewis**, “Exhaustible Resources and Alternative Equilibrium Concepts,” *Canadian Journal of Economics*, 1985, *18* (3), 459–473.
- Salant, Stephen W.**, “Imperfect Competition in the International Energy Market: A Computerized Nash-Cournot Model,” *Operations Research*, 1982, *30* (2), 252–280.