

# **Supply-Side Climate Policies and Capital Reallocation: Evidence from the Offshore Oil and Gas Industry**

Nicholas Vreugdenhil \*

## **Abstract**

Supply-side climate policies — such as a drilling moratorium — aim to mitigate climate change by keeping fossil fuels ‘in the ground’. I examine how capital reallocation impedes the effectiveness of incomplete supply-side policies in the global offshore oil and gas industry. I develop a framework of a decentralized capital market which extends the location choice and dynamic matching literature to accommodate two-sided vertical heterogeneity. Applying the framework to a novel dataset of contracts and projects, I find that policy designs that do not account for capital reallocation are substantially less effective, and there are significant gains from a global agreement.

---

\*Department of Economics, Arizona State University. Email: [nvreugde@asu.edu](mailto:nvreugde@asu.edu). This is a heavily revised chapter from my dissertation. Thank you to my advisors Rob Porter, Mar Reguant, and Gaston Illanes, and seminar audiences at NBER Summer Institute EEE, Northwestern, Queens, Rice, U Arizona, UC Davis, and UT Austin, as well as Mike Abito, Lucas Davis, Matt Gentry, and Nicolai Kuminoff, for helpful comments. Thank you also to Ryan Kellogg, Giulia Brancaccio, Stephen Ryan, and Adam Spencer, for discussing the paper. Finally, thank you to the editor Alessandro Gavazza and two referees for their thoughtful comments.

# 1 Introduction

Supply-side climate policies — regulations that aim to keep fossil fuels ‘in the ground’ and are often targeted at the \$6.6 trillion global oil and gas industry — are a growing, and controversial, solution to climate change.<sup>1</sup> Advocates argue that these policies will aid the energy transition away from fossil fuels. However, these regulations are typically incomplete because they only cover a subset of the global market whereas drilling takes place in fields across the world. As a result, a concern is that this uncoordinated action may cause leakage when economic activity moves to locations or sectors with weaker regulation. Central to these questions, and of critical importance due to the sheer scale of active and potential supply-side policies, is understanding the exact channels by which leakage takes place.

In this paper I argue that incomplete supply-side policies in the oil and gas industry are particularly susceptible to leakage through a capital reallocation channel. Oil and gas wells are drilled by movable physical capital: drilling rigs, which are leased by oil companies that match with rigs in a decentralized global market. Here, local regulation decreases the profitability of local capital inputs, which causes capital to reallocate to unregulated markets; this decreases the price of inputs elsewhere and spurs production.

Despite its emphasis in theoretical work, the capital reallocation channel is fundamentally different to the leakage through a product market channel that is considered in the existing empirical literature.<sup>2</sup> As well as determining leakage in the oil and gas industry — which is critically important in its own right since it is responsible for around 40% of carbon emissions (International Energy Agency, 2023) — understanding the capital reallocation channel in this context is also potentially useful for predicting leakage in other markets. Specifically, given that the oil and gas industry is a market where physical capital is as mobile as can be, a result of less than full leakage may point to factors that potentially limit leakage in other settings.

Overall, this paper answers the questions: to what extent does capital reallocation reduce the efficacy of supply-side regulation, and would more complete regulation improve outcomes? To do so I develop an empirical

---

<sup>1</sup>For example, the Biden Administration has dramatically reduced the number of offshore drilling leases for sale in the US market because of the climate effects of the oil and gas which will be produced and consumed from the resulting wells (Friedman (2023)). In economics, supply-side climate policies have been analyzed theoretically in Harstad (2012), amongst others. Empirical work includes Covert and Kellogg (2023), Prest (2022), Prest and Stock (2023). Across the world there are numerous similar proposed policies; a useful summary is detailed in Ahlvik, Andersen, Hamang and Harding (2022). The valuation number counts both onshore and offshore production in 2022. (<https://www.ibisworld.com/global/market-size/global-oil-gas-exploration-production/>)

<sup>2</sup>For the theoretical literature see Baylis, Fullerton and Karney (2013). Leakage through a product market channel works because local regulation raises the relative price of tradeable goods, leading to increases in production and emissions elsewhere. Examples in industries such as electricity and cement include Abito, Knittel, Metaxoglou and Trindade (2022), Fowlie and Reguant (2018), Fowlie, Reguant and Ryan (2016).

framework that extends the location choice and dynamic matching literature in industrial organization to a setting with two-sided vertical heterogeneity in matches leading to sorting. I apply the framework to an unusually detailed dataset of the universe of contracts and capital movements in the global market for deepwater drilling rigs. Counterfactual results illustrate the central role of capital reallocation in how proposed incomplete supply-side policies affect profits and carbon emissions, and the gains to a global agreement.

The international offshore oil and gas industry is an archetypal global dirty industry with a movable form of capital: drilling rigs. Offshore oil rigs are ‘marine vessels’ that are explicitly designed to be easily transportable between locations. The industry is decentralized and oil companies such as BP and Chevron do not own the capital required to drill oil and gas projects. Instead, they contract out drilling to a rig owner.

The market is shaped by geographical space: oil field locations are situated across the world and rig owners must choose the most profitable location for their capital. It is also shaped by two-sided vertical heterogeneity in capital types and drilling projects: rigs can be ranked by their *efficiency* (their on-board drilling technology) and oil and gas projects can be ranked by their *complexity*. The match complementarities matter, with more efficient rigs sorting towards more complex projects.<sup>3</sup> Furthermore, different locations contain different types of projects.

I begin with a descriptive analysis of the industry, documenting four key facts. The first fact I document is empirical sorting patterns in the raw data, which show that more efficient rigs are matched towards more complex wells. This suggests complementarities in the matching function. Second, I show that more complex wells tend to produce significantly more oil. Third, I document that all types of rigs move location frequently, and that the stationary distribution of rigs is consistent with rigs moving towards locations where they are best matched. Finally, I show that the deepwater market is relatively stable within the sample period.

Next, I estimate a model of the global deepwater drilling market using data on the positions and status (including information about contracts) of all deepwater rigs worldwide between 2008-July 2015. In the model there are several spatial locations worldwide (oil fields). Locations differ by demand (potential projects), as well as costs relating to the operational expenditures of the rig owner (e.g. salaries and accommodation for the rig crew). Within each location oil companies first choose whether to enter. These potential projects then contact rigs and — given the types of available rigs and relative prices in the location — target the type of rig that best matches with their well type.<sup>4</sup>

---

<sup>3</sup>Note that in the paper I use the terms ‘capital’ and ‘rig’ interchangeably. Similarly, I use the terms ‘project’ and ‘well’ interchangeably.

<sup>4</sup>In equilibrium, the two-sided vertical heterogeneity in this market results in complex wells targeting high-efficiency rigs and simple wells

On the supply side of the model, capital owners are forward-looking. Within a location they may be contacted by an oil company to undertake a contract, in which case they will be unable to match for the duration of the contract. If the rig is not currently in use, it can move to a new location looking to match with a new project or stay in the current location. Key to this decision is the quality of potential matches for each particular capital type in each location. The model also allows for rigs to exit the global market and be scrapped. Overall, the model captures two main channels that result from the interaction of spatial reallocation and sorting. First, spatial reallocation, where capital moves towards unregulated markets with the greatest complementarities. Second, if policies disproportionately target certain types of projects, capital can reallocate by re-sorting toward unregulated project types within a location.

I use the model to test several counterfactual policies. These center around proposed ‘supply-side’ policies that ban complex wells. This counterfactual serves as a good proxy for proposed bans which correspond to sales of *new* offshore leases such as those implemented by the Biden Administration.<sup>5</sup>

I first test the effects of a US-only ban. Leakage through the capital reallocation channel is substantial and reduces the efficacy of regulation by -36.6 percent. Decomposing this into two leakage channels, for every ton of carbon dioxide saved through banning complex wells, within-location capital reallocation generates 0.11 more tons, while an additional 0.26 tons are generated through reallocation across space to unregulated locations.<sup>6</sup>

A global agreement would eliminate leakage across space (but not within each location). Overall, it would be substantially more effective than a US-only ban, with capital reallocation reducing the efficacy of the regulation by -10.8 percent. However, global coordinated regulation may not be politically feasible. Therefore, I also consider a coalition of richer countries incorporating the US, Europe, Australia, and South America. This grouping would also substantially improve the efficacy of the regulation: overall, capital reallocation limits the reduction in emissions by -17.0 percent.

---

targeting low-efficiency rigs.

<sup>5</sup>Over time drilling has broadened to geological formations which are more difficult to drill. Therefore, these new leases tend to correspond to more complex wells than existing leases. For a history of the industry and changes in leasing patterns see Gramling and Freudenberg (2012).

<sup>6</sup>Although wells produce both oil and natural gas, which both result in carbon emissions once consumed, this paper focuses only on the emissions from burning the oil content and not the natural gas content. The rationale is that deepwater wells are predominantly in fields where the hydrocarbon content is mainly oil. Furthermore, the chemistry of burning natural gas results in relatively less emissions per energy unit than oil: <https://www.eia.gov/energyexplained/natural-gas/natural-gas-and-the-environment.php>. Overall, this implies the results are conservative.

## 1.1 Contributions and related literature

Overall, this paper makes three main contributions. The first is a new framework of location-choice and matching in a decentralized capital market which allows for two-sided vertical heterogeneity. This is a key difference to previous work using location choice models in, for example, bulk shipping and taxis, where agents are relatively homogeneous (e.g. Brancaccio, Kalouptsi and Papageorgiou (2020), Buchholz (2022)).<sup>7</sup>

Furthermore, while there is a literature that emphasizes the role of two-sided vertical heterogeneity in capital markets (e.g. Vreugdenhil (2026)), combining this feature (amongst others) into a location choice model is challenging. Specifically, it requires computing a nested equilibrium, solving for both the equilibrium sorting patterns within a location given the distribution of rig types that locate there, and an across-location equilibrium where different types of capital optimally reallocate given the equilibrium within each location.

This results in rich equilibrium dynamics that result from the interaction of the within and across-location components. For example, regulation (like banning complex wells) may target heterogeneous projects. The direct effect is captured by the two previously discussed channels for capital reallocation in the model (spatial reallocation and re-sorting within regulated locations). The framework also captures that, in equilibrium, these individual decisions then reverberate throughout the global market. For instance, the entry of high-efficiency rigs in unregulated locations then causes a re-sorting of rigs and wells within that location; this may then spur additional exit of mid-efficiency rigs into other locations.

Moreover, accounting for two-sided vertical heterogeneity is essential to estimating leakage in this setting. For example, as I show in the counterfactuals section, a model with homogeneous agents would substantially over-predict efficiency losses from capital reallocation.<sup>8</sup>

The second contribution is the analysis of a detailed dataset of firm-to-firm contracts, movements, and projects, in a global capital market. These markets are typically difficult to study because micro-data on contracts and allocations are often confidential. At the same time, many types of physical capital are traded in decentralized markets (Gavazza (2016)). Moreover, movements of capital across space are thought to be a key method of capital reallocation more broadly (Ramey and Shapiro (2001)). The data provide a granular picture of the inner workings

---

<sup>7</sup>An exception is Yang (2024) which considers a location-choice model where trucks have horizontally differentiated preferences over locations.

<sup>8</sup>A second factor that may limit leakage is the relative abundance of rigs compared to projects. In this case, the movement of a rig to a new location may simply cannibalize existing demand in that location rather than spurring additional matches by, for example, relaxing capacity constraints.

of a real-world capital market.

The third contribution is a set of new findings about the efficacy of incomplete supply-side regulation with capital reallocation in the global offshore oil and gas industry. This is connected to several strands of literature. The first is the previously mentioned literature that investigates how incomplete environmental regulation operates through an alternative production channel for leakage.

A second strand of literature is in international trade where many papers investigate the relationship between environmental regulation and the patterns of trade (see Copeland and Taylor (2003) for a summary). The literature has detected effects in ‘footloose’ industries using more aggregated data, e.g., Ederington, Levinson and Minier (2005). My results are consistent with this high-level finding as the offshore oil and gas industry is extremely ‘footloose’ because capital is highly movable.<sup>9</sup> Davis and Kahn (2010) find that used vehicles that fail emissions testing in California are more likely to be exported to Mexico.

This contribution also builds on existing research into the oil and gas industry. For example, in addition to the papers already mentioned in Footnote 1, see also Kellogg (2011), Lewis (2019), and Agerton (2020). The papers Corts (2008) and Corts and Singh (2004) work with a more aggregated version of offshore rig data, and these data contain fewer covariates for the projects undertaken under each contract. Vreugdenhil (2026) uses contract data in the US Gulf of Mexico to study how booms and busts affect mismatch in the shallow water market; this paper uses similar contract data but in the global *deepwater* floater market, focusing on capital reallocation across oil fields in response to regulation.

## **2 Market description and data**

Offshore drilling is segmented into shallow water (< 500ft water depth) and deepwater drilling (> 500ft water depth). I follow industry practice and treat these two segments as separate markets due to the differences in capital types, geographical locations, and the scale of engineering required to drill a well. In this paper I focus solely on the deepwater drilling segment of the industry. Due to the extreme water depths, deepwater wells are drilled by ‘floater’ drilling rigs (called either Semi-submersibles or Drillships), which float on the ocean’s surface and are anchored at the well site. This is in contrast to the shallow water market detailed in Vreugdenhil (2026) which uses Jackup rigs which extend their legs to the seabed.

---

<sup>9</sup>The onshore oil and gas industry also uses drilling rigs which are designed to be movable across locations, and is similarly the target of supply-side policies.

Oil rigs are ships that move around the ocean drilling wells. Long-distance moves between fields (for instance, from the US Gulf of Mexico to the North Sea) are usually undertaken using a ‘dry tow’ where the rig is maneuvered onto a special ship and this ship then transports the rig. Figure 1 shows an example of a deepwater oil rig moving using a dry tow.

The process of drilling a deepwater well and procuring a rig is as follows. Oil companies like BP and Chevron lease areas of the seabed from national governments, which provide them the option to drill a well. Using geological surveys and (if available) information about other existing wells in nearby leases, these oil companies decide whether to drill a well and determine the potential well design. Since oil companies do not own the oil rigs they use to drill with, they need to match with an appropriate drilling rig. Oil rigs are rented under simple dayrate contracts for the time it takes to drill a well.<sup>10</sup> After the well is completed (around 6 months) it is typically connected to an undersea pipe for continuing extraction. The rig then moves on to its next job.

Drilling responsibilities are precisely delineated in this industry. While rig owners are responsible for furnishing the rig in good working order, and paying expenses for the salaries and accommodation of the crew onboard the rig, they do not pay for any of the drilling costs of materials. Instead, the oil company owns the well, is the beneficiary of selling the produced hydrocarbons, and bears responsibility for drilling expenses like materials. The oil company has a representative (called the ‘company man’) who lives on the rig and represents the oil company’s interests.

Finally, the deepwater drilling industry is highly fragmented. Both the rig owner side and the oil company side of the global market are unconcentrated. Therefore, I do not allow for either side of the industry to exert market power in the model.

## 2.1 Data

The contract and status data come from a proprietary dataset from Rigzone (an industry data provider). The full dataset consists of the status of marketed drilling rigs worldwide from 2000 to July 2015. I cut the data to only deepwater rigs (defined as those with a maximum water depth of >500 feet). I observe the country and region that each drilling rig is currently in at each point in time, and whether a rig is idle or under contract. If a rig is under contract, then I observe key covariates for the contract including price, duration, and the oil company that owns the well. Contracts are almost always fixed price per day for a given duration and rarely contain performance

---

<sup>10</sup>An alternative contracting form is sometimes used in the industry: a turnkey contract where a rig is hired to drill a set number of wells rather than for a period of time (Corts and Singh, 2004). I have additional data from IHS on whether a contract is a turnkey or dayrate contract for the US market. In the period of time studied, for the deepwater market, all of the contracts are dayrate contracts.

**Figure 1:** A deepwater drilling rig moving between locations



*Note:* This picture shows a deepwater drilling rig (called the ‘Deepwater Nautilus’) undergoing a dry tow between locations. Source: <https://2b1stconsulting.com/wp-content/uploads/2012/04/nautilus-dry-tow.jpg>.

**Table 1:** Summary Statistics

Variable	Units	N	Mean	Std. Dev.
Daily Rig Activities	Millions	1.75		
Status Updates	Unique status changes	4564		
Contract Price	Millions USD/day	1211	0.36	0.14
Contract Duration	Days	1211	173	159
Prob. of Relocation	Events	805	0.43	0.25

**Table 2:** Summary Statistics: Heterogeneity

	Capital Type (Efficiency)		
	Low	Mid	High
Prob. of Relocation	0.29	0.35	0.68
Utilization	0.86	0.86	0.89
Dayrate	0.31	0.36	0.44
Average Match: Well Complexity	1.0	2.4	3.6

incentives.

The data sample covers the years 2008-July 2015 and Table 1 provides summary statistics. Although most rigs operate under relatively short-run contracts (around 6 months) and are rented over time by many different oil companies, there are a small number of rigs that operate continually under very long-term contracts. As a result, I delete rigs that operate under contracts of duration greater than two years. In total I have 4564 ‘status updates’ for deepwater rigs, which amount to 1.75 million daily rig activities. I provide more detail about these status updates and the data cleaning steps in Appendix A-2.

**Rig heterogeneity** As is the convention in the industry, rigs can be ranked by their maximum drilling depth, which is a proxy for capital efficiency since it is highly correlated with onboard technology, age, and other factors. I aggregate capital heterogeneity into three types by maximum drilling depth and call these types ‘low’, ‘mid’, and ‘high’ specification rigs.

Table 2 describes some ways that these rig differences matter. High-specification rigs fetch higher prices than other rig types and also tend to relocate more frequently. However, all capital types have relatively similar levels of utilization.

**Well heterogeneity** Wells can be ranked in terms of how complex they are to drill using an engineering model called the ‘Mechanical Risk Index’. This index takes well covariates such as drilling depth of the well, water depth, and bottomhole pressure, and ranks wells on a one-dimensional index of drilling complexity. The Mechanical Risk Index was developed by Conoco engineers in the 1980s (Kaiser (2007)). Well complexity is directly related to the cost of drilling a well: these wells run an increased risk of technical issues which may require new materials or result in blowouts. I follow Vreugdenhil (2026) to construct this index, which draws directly from Kaiser (2007).

As is apparent in Table 2, more efficient rigs sort towards more complex wells. In addition, more complex wells tend to produce more oil (and therefore more emissions when this oil is consumed). I describe these patterns in more detail in Section 3.

**Location heterogeneity** I aggregate capital locations into eight large regions across the world; within these regions the main oil fields are relatively close to each other.

## 2.2 Quantifying oil production and emissions

I provide an overview here of how matches in the model are mapped into changes in global oil production and emissions, with the details presented in Appendix A-1. Note that I discuss long-run Hotelling dynamics considerations in Section 7.3.

Wells produce both oil and natural gas in different quantities, and both result in carbon emissions once burned. This paper focuses only on emissions from burning the oil content. The reason is that deepwater fields predominantly produce oil, and furthermore burning oil produces a far greater magnitude of emissions than burning the equivalent energy unit of natural gas.<sup>11</sup>

The complexity of an individual project is mapped into a production volume of oil using the empirical relationship that more complex projects tend to produce more oil (which I document further in the next section). Then, given the equilibrium number and types of matches predicted by the model in each location, the model predicts a total volume of oil produced in the deepwater market.

I then convert changes in oil production in the deepwater market into a change in carbon emissions globally in two steps. As I explain further in Appendix A-1, I first convert the decrease in supply from the deepwater market to an equilibrium global change in oil produced and consumed through changes in global oil prices. This incorporates demand responses (since removing deepwater supply would shift out the residual demand curve facing other regions) as well as supply responses from non-deepwater fields. Second, I convert this global change in output to carbon emissions by scaling by the EPA's Greenhouse Gas Equivalencies Calculator.

## 3 Descriptive analysis of the deepwater rig market

I now document patterns in the data which motivate the setup of the model.

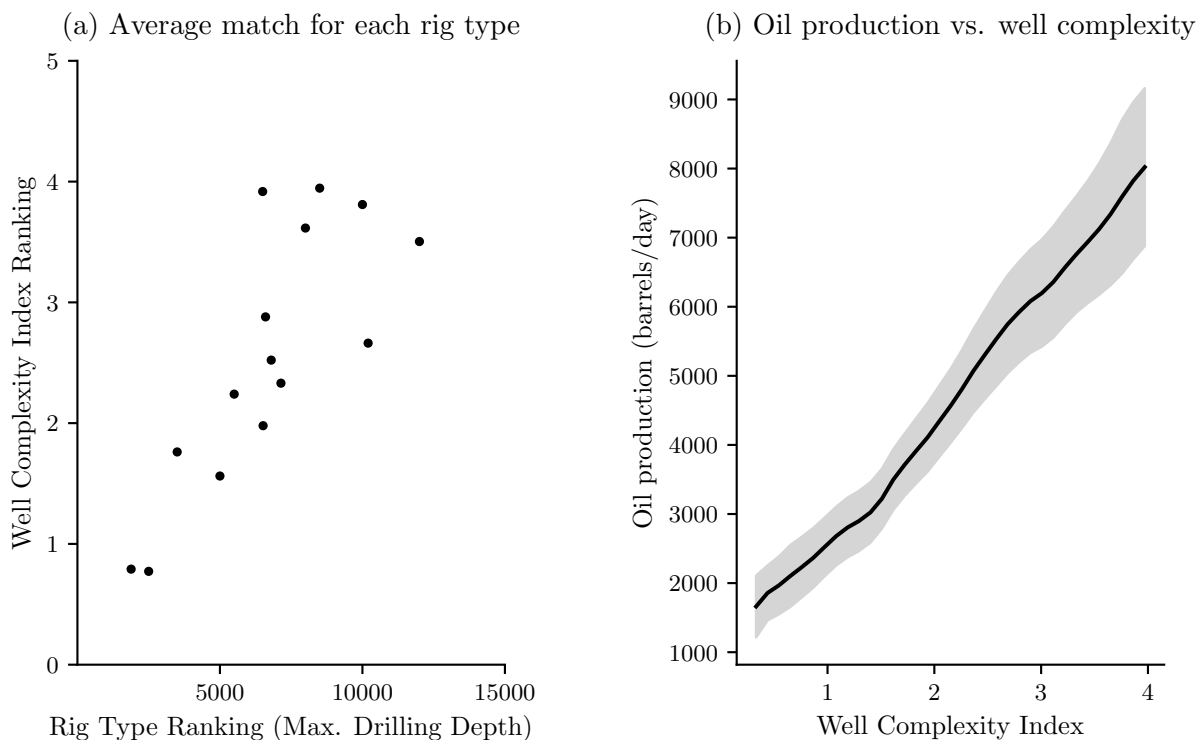
### **Observation 1: Positive sorting patterns suggest that match complementarities matter.**

First, I show patterns related to two-sided vertical heterogeneity, which is linked to the 're-sorting within location' channel. Figure 2(a) illustrates the sorting patterns between capital (rigs) and projects (wells) in the US market. Recall that wells can be ranked vertically by their complexity using an engineering model called the 'Mechanical Risk Index', and that rigs can also be vertically ranked by their efficiency (proxied by their maximum drilling depth). Figure 2(a) illustrates that more complex projects tend to match with more efficient rigs, both on average

---

<sup>11</sup>See, e.g. <https://www.eia.gov/energyexplained/natural-gas/natural-gas-and-the-environment.php>.

**Figure 2:** Within location choices: sorting patterns and oil production



*Note:* Figure (a) presents positive sorting patterns in terms of the average match for each rig type. The x-axis is the rig efficiency ranking (where rig efficiency is proxied for by the maximum drilling depth) and the y-axis is the project complexity ranking (the ‘Mechanical Risk Index’ which is an engineering model used in the industry that maps well covariates into a one-dimensional index for how difficult the well is to drill). Each point on the x-axis corresponds to a particular maximum drilling depth. These maximum drilling depths are typically given as round-number increments (e.g. maximum drilling depth of 6000 feet) and so each point on the graph corresponds to all the projects undertaken by the many rigs which share a particular drilling depth. Note that I drop one rig with a very small number (2) of observations here. Figure 2(b) displays a local polynomial regression of the relationship between well complexity and oil production. I present this relationship in more detail in Appendix A-1.

and over the entire distribution of project types.<sup>12</sup> These pictures suggest that match complementarities matter.


Where do these match complementarities come from in this industry? Broadly, more efficient rigs — through their better on-board technology — generate cost efficiencies once allocated to complex wells. For example, a complex well may involve drilling around a difficult geological formation, which involves a greater probability of risks like a “stuck pipe”. The better technology of efficient rigs allows them to drill these difficult formations, and more readily deal with unexpected events as they occur. I further discuss complementarities in Section 6.

**Observation 2: More complex wells produce more oil.**

Figure 2(b) displays a local polynomial regression of the relationship between well complexity and oil production. I present this relationship in more detail in Appendix A-1; in this section I also show that this relationship is robust

<sup>12</sup>These sorting patterns are also apparent for shallow-water rigs (Vreugdenhil, 2026).

**Figure 3: Distribution of Contracts by Region and Rig Efficiency**



to rig-type controls. The figure illustrates that more complex wells also produce more oil. Therefore, for example, a ban on complex wells that causes rigs to reallocate towards simpler wells within a location could result in less oil being produced even if the total number of wells drilled remains the same.

**Observation 3: Locations are heterogeneous and attract different distributions of capital types.**

Next, I show patterns related to spatial reallocation. As documented in the summary statistics in Table 2, rigs clearly move around the world. Figure 3 plots the resulting empirical distribution of which kinds of rigs drill in each location.

Figure 3 illustrates substantial heterogeneity across regions in terms of which types of rigs are most active. Although I do not directly observe the distribution of well complexity for regions outside the United States, these patterns are indicative of the kinds of wells that are available for rigs in each location. For example, North America, Africa, and Central America are part of what is called the “golden triangle” of ultra-deepwater drilling, and this is reflected in the kinds of rigs that locate there.<sup>13</sup> On the other hand, the European market (i.e. the North Sea) is known to have much shallower and simpler wells to drill, and so this region attracts relatively low-specification rigs.


<sup>13</sup>See e.g. <https://www.offshore-energy.biz/us-firm-scores-two-long-term-gigs-with-key-drilling-players-in-golden-triangle-area/> for this terminology.

These aggregate patterns are also useful to understand qualitatively which coalitions of countries may be the most effective in reducing leakage through the spatial capital reallocation channel. For example, the most politically feasible coalition — which I later consider in the model counterfactuals — involves the US, Europe, Australia, and South America. Figure 3 suggests that the coalition may not be so successful in eliminating spatial leakage. Concretely, while it covers complex wells in North and South America, the African market also contains such wells but is left unregulated, and so efficient rigs may relocate here. Similarly, Europe contains mainly simple wells, but these wells are also present in Asia which is unregulated.

Finally, I present direct evidence in Appendix Figure A-1 that the entry of a rig into a location stimulates the drilling of additional wells.

**Observation 4: The deepwater market is relatively stable within the sample period.**

**Figure 4:** Prices are relatively stable over the sample period



*Note:* The figure shows a local polynomial regression of dayrates (the price paid by the oil company to the rig) over time. The shaded area is the 95% confidence interval. Overall, the figure shows that dayrates are relatively stable over time which is consistent with the market being relatively stable for the period in which the model is estimated.

One reason that I choose the sample period 2008-July 2015 is that dayrates are stable over the sample period, as documented in Figure 4. Similarly, the number of contracts drilled each year is relatively stable.<sup>14</sup> Practitioners in the industry have also noted that the deepwater market, due to the kinds of projects that are undertaken, is “...less sensitive to short-term fluctuations in oil prices than onshore development of shale gas and tight oil resources”

<sup>14</sup>For example, towards the start of the sample in the oil price bust in 2009 the number of contracts was 157. When the oil price returned to a boom in 2011, the total number of deepwater drilling contracts was 181.

(EIA, 2016).<sup>15</sup> Motivated by this fact I model the market as in a steady-state equilibrium for this sample period.

## 4 Model

### 4.1 Setup

There are locations  $l \in L$  across the world, each of which corresponds to an oil field. Agents are projects (wells)  $x$  and capital (rigs)  $y$ . Capital is differentiated by efficiency  $y \in Y = \{low, mid, high\}$  and projects are differentiated by their complexity  $x$ . The model is dynamic with one period equal to one month. Agents have the discount factor  $\beta$ .

To keep notation simple I suppress time, capital, and project-specific subscripts. Instead, I write the model components as just a function of the types of each capital and project (e.g.  $x$  instead of  $x(i)$  for project  $i$ ,  $y$  instead of  $y(j)$  for rig  $j$ ) and explicitly discuss any cases that deviate from this convention.

In order to drill a project, a project owner needs to match with capital. Denote the number of type- $y$  rigs in location  $l$  by  $n_{l,y}$ . Each rig has a queue (a ‘backlog’) of projects and if the queue is sufficiently short — specifically, if the number of contracted months in the backlog is below a critical value  $t_{backlog}$  — then the rig is ‘available to match’.<sup>16</sup> The timing in each period within each location  $l$  is as follows:

1. Project entry in each location The number of new potential projects in each period is given by a draw from a Poisson distribution with a location-specific mean  $\lambda_l$ . The type of each of these potential projects (pre-entry) is characterized by an independent draw from a distribution  $x \sim f_{l,x}$  of project complexity. If a potential project chooses to enter it pays an entry cost  $c_{entry}$ .
2. Targeting Each potential project that enters knows its type  $x$  and chooses which kind of rig  $y$  to match with (‘target’).
3. Matching Within a period, potential projects match sequentially in the (random) order in which they are drawn with the capital type that they choose to target. If there are no more available rigs, then unmatched potential projects immediately exit.<sup>17</sup> Otherwise, a match is formed.

---

<sup>15</sup>Note that it could still be possible that large long-term changes in oil prices could eventually affect the industry. For example, EIA (2016), which is published after my data ends, notes that the sustained decline in oil prices around that time may affect decisions. However, I see no such effects in my data in the earlier sample period I use for estimation.

<sup>16</sup>The constraint that projects will refuse to match if the backlog is too long arises mainly from oil company preferences; rig owners tend to prefer longer backlogs since they reduce the risk of a rig not being utilized. In fact, it is common for rig owners to actively advertise their deepwater rig backlogs to shareholders in annual reports as a positive signal about their firm’s financial health.

<sup>17</sup>The immediate exit of unmatched potential projects is not an assumption but rather optimal behavior given the setup of the problem. Specifically, if a potential project is unmatched it implies that there is not enough available capital (i.e. matching with any capital would

4. Production If a potential project successfully matches with capital, the  $\tau$  periods of the contract are added to the capital's backlog. The total per-period payoff is given by  $m_{x,y} - c_{l,y}$  for each of the  $\tau$  periods of the contract, where the function  $m_{x,y}$  is the match value and  $c_{l,y}$  is a location-specific and capital type-specific cost. As I explain further in the Estimation section, the matching function also incorporates that some matches cannot occur due to engineering constraints. Prices are determined by Nash bargaining.
5. Relocations Capital not currently under contract can either stay in the current location  $l$ , or move to a new location  $l'$ . Moving to a different location incurs a cost that is dependent on distance between locations  $d_{l,l'}$  but not capital type.<sup>18</sup>

As well, I include an extensive margin response for the total number of rigs in the global market. Overall, I model the market as a dynamic steady-state in which the primitives of demand and supply in each location, for each rig type, are stationary, but there are idiosyncratic shocks that induce rigs to move around the globe.

**Discussion of key assumptions/properties** I now discuss four aspects of the model setup. First, I assume that agents make their decisions based on long-run averages in the market. Specifically, for potential projects, they use the long-run average probability that a rig type is at the capacity constraint to determine rig selection and prices, and rigs use long-run averages for the probability of matching and prices in their location choice. One justification for the above assumption — as previously mentioned in Section 2.1 — is that the deepwater market is relatively stable in the sample period. I also show in Appendix Table A-2 that the distribution of rigs across space is stable across time, splitting the sample into the first vs last half of the data. However, this assumption is still in contrast to an alternative set-up where agents can condition their behavior on a more transient state of the market (such as exactly how many other potential projects entered in the same period, or the exact number of months in the backlog of every rig). The benefit of this assumption is computational; allowing agents to condition their behavior on a more transient state of the market would add substantial complexity to the decisions of potential projects and capital, and generate a large state space for capital's dynamic decisions, which would result in a curse of dimensionality. In addition, this assumption is arguably realistic for this market. For example, according to my data provider, contracts are eventually fully reported but there is often a delay, so the data are somewhat 'stale' and the current state of the marketplace is unknown at any point in time.

---

produce a backlog longer than  $t_{backlog}$ ). But then this implies that waiting an additional period for capital to become available would require waiting longer than  $t_{backlog}$  to drill a well.

<sup>18</sup>Typically long-distance moves of the rigs are accomplished using tow boats, not the rigs' internal engines, and so I do not allow moving costs to depend on the capital type  $y$ .

The second notable assumption is that agents can target their best match, and I do not allow for search frictions. Instead, I micro-found the matching process in the model through a queuing simulation that is tailored to the institutional details of the industry. If there is an available rig and it is being targeted by a potential project, then these agents will meet. Capital unemployment is generated solely due to the Poisson draws in demand: several successive low draws may result in rig unemployment. The ‘no search frictions’ assumption is different to previous work in other markets like taxis, which involve matching with much larger numbers of agents searching on both sides of the market, which leads to coordination frictions. Unlike these markets, the scale of the deepwater rig market is much smaller and so matching arguably involves fewer opportunities for frictions.

Third, I assume that capital receives a match only after entering a location. I experimented with an alternative assumption where rigs start the period matched, but I found numerically that it made little difference to rigs’ location choice decision. This is because matches are relatively short-term compared to the overall time that a rig spends within each location.

Fourth, a property of the model is that agents do not reject matches. This is *not* an assumption. Rather, it follows without loss of generality from the setup that (i) potential projects make an entry decision and (ii) potential projects can direct their search towards their best match. Therefore, potential projects will only enter if the eventual match will be accepted. This property is shared with the broader literature on directed search, e.g., Moen (1997).

## 4.2 Demand: How projects match with capital

I first discuss rig choices and entry choices for potential projects, which correspond to the two choices that these potential projects make in the model.

After entering, the ex-ante payoff to targeting capital of type  $y$  is:

$$\Pi_{l,x,y}^{project} = q_{l,y}^{project} \underbrace{\left( \sum_{s=0}^{\tau-1} \beta^s (m_{x,y} - p_{l,x,y}) + \varepsilon_y \right)}_{\text{Match value with type } y \text{ capital}} \quad (1)$$

The term  $q_{l,y}^{project}$  is the long-run probability of matching capital type- $y$  in location  $l$  (and  $1 - q_{l,y}^{project}$  is the probability that capital is at its capacity constraint),  $m_{x,y}$  is the value of a match between project type  $x$  and capital type  $y$ , and  $\varepsilon_y$  is an idiosyncratic error for each capital type  $y$  distributed i.i.d. extreme value. Note that I am suppressing individual project subscripts, but the  $\varepsilon_y$  is drawn independently for each searching project (as well as for each rig type  $y$ , and for each time period). A potential project contacts the capital type that offers the highest

expected value:  $\text{argmax}_y \{ \Pi_{l,x,y}^{project} \}$ .<sup>19</sup>

Potential projects will enter if the value of entering is greater than the entry cost  $c_{entry}$  (where  $c_{entry}$  embeds the costs of, for example, drawing up a detailed well plan):  $\max_k \{ \Pi_{l,x,k}^{project} \} \geq c_{entry}$ .

Integrating over the distribution of project complexity  $f_{l,x}$ , the share of potential wells that target capital type  $y$  is:

$$s_{l,y} = \int 1[y = \text{argmax}_k \{ \Pi_{l,x,k}^{project} \}] 1[\max_k \{ \Pi_{l,x,k}^{project} \} \geq c_{entry}] f_{l,x} dx \quad (2)$$

Since agents can target their best match and choose whether to enter, this implies that no matches are rejected (otherwise the project would have a negative payoff from entering the market).

I compute the probabilities of matching for projects and capital,  $q_{l,y}^{project}$  and  $q_{l,y}^{capital}$ , that result from the above targeting decision using a matching simulation. I briefly discuss this simulation here and leave a more detailed description to Appendix A-3.2. Overall, I simulate a queue, for each rig type  $y$ . If the match is at the front of the queue, then it takes the duration of the contract  $\tau$  to complete the match. For each rig type queue, there are  $n_{l,y}$  rigs that can complete projects, so there are  $n_{l,y} t_{backlog} / \tau$  places in the queue. The ‘queuing discipline’ is first-in-first-out. Amongst other things, the matching simulation implies that if there is more capital of a certain type in a location, then other capital of that type will match with a lower probability and wells targeting that type of capital will match with a higher probability.

### 4.3 Supply: Location decision

The location decision of an unemployed piece of capital of type  $y$  is to either stay in the same location  $l$ , or to choose to move to a different oil field  $l'$ . The values of these options are:

$$\text{Value of moving to } l' : \quad -c_d d_{l,l'} + \beta V_{l',y} + \sigma_\varepsilon \varepsilon_{l'} \quad (3)$$

$$\text{Value to staying in location } l : \quad b_{stay} + \beta V_{l,y} + \sigma_\varepsilon \varepsilon_l \quad (4)$$

Here  $c_d$  is the per-mile transport cost,  $d_{l,l'}$  is the distance,  $\varepsilon_{l'}$  is the idiosyncratic logit error, and  $\sigma_\varepsilon$  is the scale parameter of the errors. Although I am suppressing rig-specific and time-specific subscripts, the logit shocks are

---

<sup>19</sup>Note that this rig selection choice implicitly assumes that projects will immediately match with rigs if the rig is not at a capacity constraint. I experimented with a more complicated model for rig selection and prices where potential projects discount rig types based on the average delay due to backlog. The results were relatively unchanged, so I chose not to incorporate this more complicated feature.

drawn independently for each individual rig, each time period, and each location. The second term is the value of staying put in location  $l$ . In this term,  $b_{stay}$  is a parameter that reflects unobserved benefits of remaining unmatched in the same location such as labor savings. These values deliver multinomial logit conditional choice probabilities for moving location which I later use for estimation; I provide more details about the exact form of these equations in Appendix A-3.1. Using the location decision in Equations (3) and (4) I can write the ex-ante value function for unemployed capital (that is, the value function before the  $\varepsilon_l$  shocks are drawn):

$$U_{l,y} = \sigma_\varepsilon \log \left( \sum_{l' \neq l} \exp \left( \frac{-c_d d_{l,l'} + \beta V_{l',y}}{\sigma_\varepsilon} \right) + \exp \left( \frac{b_{stay} + \beta V_{l,y}}{\sigma_\varepsilon} \right) \right) + \sigma_\varepsilon \gamma^{euler} \quad (5)$$

where  $\gamma^{euler}$  is Euler's constant.

The value function  $V_{l,y}$  (the value of being in location  $l$  before matching takes place) is given by:

$$V_{l,y} = q_{l,y}^{capital} \underbrace{\left( \sum_{s=0}^{\tau-1} \beta^s \delta_{l,y} + \sigma_\varepsilon \gamma^{euler} + \beta^\tau V_{l,y} \right)}_{\text{Expected value to matching for the rig}} + (1 - q_{l,y}^{capital}) U_{l,y} \quad (6)$$

Here,  $q_{l,y}^{capital}$  is the long-run average probability that capital of type  $y$  matches with a well in location  $l$ . The expected value to the rig of being in a contract in each period is  $\delta_{l,y} = \bar{p}_{l,y} - c_{l,y}$ .<sup>20</sup> Note that I do not allow for unobserved demand shocks. This is mainly due to industry-specific reasons which suggest that such shocks are less important: for instance, as previously mentioned, the total number of contracts per year over the sample period is relatively stable.

#### 4.4 Supply: extensive margin

Unemployed capital will *not* exit if the value of remaining unemployed in a location  $U_{l,y}$  is greater than the value of exit (scrapping the rig):

$$U_{l,y} \geq b_{scrap} \quad (7)$$

---

<sup>20</sup>Note that although the costs are assumed to be fixed over the period of the sample, the model could be extended to accommodate time-varying cost shocks by separately estimating the model year-by-year. But this would introduce an additional computational burden. Furthermore, if there were large time-varying cost shocks, they would bias upwards the estimate of  $\sigma_\varepsilon$ , but as I later document this value is relatively small.

While the total number of rigs is relatively stable over the sample period, this channel is most important for the counterfactuals, where demand is reduced which may result in capital exit.

Note that, while it would be possible to also incorporate a capital entry margin (at the cost of more computational complexity), I do not do so for two reasons. First, time to build, which is several years for a deepwater rig (Kaiser and Snyder (2013)), makes it difficult for the market to respond (at least in the medium term) to policy changes. Second, the counterfactuals considered tend to reduce the lifetime present value of an active rig, and so it is potential rig exit that is the key margin for drilling bans.

## 4.5 Prices

Since prices are determined by Nash bargaining, the price  $p_{l,x,y}$  of an  $(x,y)$  match in location  $l$  is determined by:

$$\operatorname{argmax}_p \left[ \sum_{s=0}^{\tau-1} \beta^s [m_{x,y} - p] - \beta(1 - P_{exit})W_{l,x,y} \right]^{1-\eta} \left[ \sum_{s=0}^{\tau-1} \beta^s (p - c_{l,y}) + \beta^\tau V_{l,y} - U_{l,y} \right]^\eta \quad (8)$$

Here  $\tau$  is the length of a contract in months,  $V_{l,y}$  is the ex-ante value of available capital,  $U_{l,y}$  is the value of unemployed capital,  $\eta$  is the Nash bargaining parameter, and  $P_{exit}$  is an exit shock for unmatched potential projects. So,  $1 - P_{exit}$  is the probability that the unmatched potential project does not exit and continues to search if the match is rejected. Note that this event occurs off the equilibrium path since all matches are accepted.

The value  $\beta(1 - P_{exit})W_{l,x,y}$  is the project's outside option. For simplicity I assume that if a project rejects a match then it will target the same type of capital and so  $W_{l,x,y}$  has a capital  $y$  subscript as well as the location  $l$ . The value  $W_{l,x,y}$  is given as:

$$W_{l,x,y} = q_{l,y}^{project} \sum_{s=0}^{\tau-1} \beta^s (m_{x,y} - p_{l,x,y}) \quad (9)$$

which is the probability that the capital type is not at its capacity constraint, multiplied by the payoff to the well of matching. Note that if the capital is at its capacity constraint (which happens with probability  $1 - q_{l,y}^{project}$ ) then the project exits immediately since the backlog is too long, receiving a payoff of 0, and so this term disappears.

## 4.6 Equilibrium

I formally define the equilibrium here. Note that rigs and project owners only internalize the private benefits and costs of their location choices and not the emissions produced once the hydrocarbons are consumed.

Equilibrium is defined as a set of prices  $p_{l,x,y}$ , matching probabilities  $q_{l,y}^{project}$  and  $q_{l,y}^{capital}$ , and a spatial capital distribution  $\{n_{l,y}\}_{l \in L, y \in \{low, mid, high\}}$ , that satisfies:

1. Demand Side Equilibrium Optimal entry and targeting decisions by potential projects — given equilibrium prices and the equilibrium total number of rigs  $n_{l,y}$  in each location — that satisfy Equations (1) - (2) and the queuing model.
2. Supply Side Equilibrium Optimal location decision by rigs subject to the equilibrium average prices  $\bar{p}_{l,y}$  in each location and the equilibrium probability of matching  $q_{l,y}^{capital}$  in each location, resulting in a spatial distribution of capital satisfying Equations (3) - (6).
3. Extensive margin for the supply side governed by Equation (7).
4. Prices  $p_{l,x,y}$  determined by Nash bargaining, defined in Equations (8) and (9).
5. Expectations of agents consistent with the long-run equilibrium.

## 5 Estimating the model

### 5.1 Overview

I provide an overview of the parametric assumptions used, and whether the parameters are estimated or calibrated, in Table 3.

**Justification for the calibrated values** The discount factor is not identified, as is typically the case in dynamic discrete choice models: Magnac and Thesmar (2002). So, I set the monthly discount factor  $\beta = 0.99$ . I calibrate the contract length  $\tau = 6$ , which is approximately the mean contract length in the data.

I calibrate the *maximum* backlog in a queue to  $t_{backlog} = 12$  months, which is around the 75th percentile of backlog in the deepwater US market.<sup>21</sup> What this means is that there is at most one current project that the rig is working on, which takes six months, and another project sitting in the queue, which also takes six months. (In the queuing literature the backlog is often written as the number of spots for matches that are not currently being processed i.e.  $t_{backlog} = 6$ ; I write the backlog incorporating the existing match here for expositional clarity to a broader economics audience.) Note that I do not fix the backlog itself and it can change in the equilibrium: for example, if there is excess capital (due to a moratorium), then the matching simulation will result in a lower backlog for other capital in that location.

---

<sup>21</sup>In my dataset the US market is the only market where backlog data are systematically available.

**Table 3:** Overview of how the parameters are computed

Object	Notation	Parameterization	Param.	Method
Calibrated params.	$\beta, \tau, t_{backlog}, \eta$ $P_{exit}, c_d, c_{entry}, b_{scrap}$			Calibrated
Costs	$c_{l,y}$			Estimate step 1
Remain in loc.	$b_{stay}$		$b_{stay}$	Estimate step 1
Preference shock	$\sigma_\varepsilon$	Logit	$\sigma_\varepsilon$	Estimate step 1
Demand distribution	$f_{l,x}$	Log-normal	$\mu_l, \sigma_l$	Estimate step 2
Demand draws	$D_l$	Poisson dist.	$\lambda_l$	Estimate step 2
Match value*	$m_{x,y}$	$m_{0,y} + m_{1,y}x$	$m_{0,y}, m_{1,y}$	Estimate step 2

Note:\* The match value  $m_{x,y}$  is also constrained so that low-specification and mid-specification rigs can only match with the 99% empirical quantile of the well matches of that rig type. This captures engineering constraints that may make some lower-efficiency rig and high-complexity well matches infeasible in counterfactuals.

I also calibrate the moving cost parameter  $c_d$ . Long-range capital movements are usually accomplished by a ‘dry tow’, which means that the capital is loaded onto a ship and moved to the new location. The speed of a dry tow is typically 14 knots (16.11 miles per hour) (Golson (2014)). Since rigs are moved by similar tow boats, and the cost of towing is proportional to the distance, I divide the distance between fields by the tow speed and calibrate the per-day cost of towing as  $c_d = \$0.25$  million.<sup>22</sup>

I need to calibrate the bargaining parameter and I assume that the parties split the match surplus equally and set this to  $\eta = 0.5$ .<sup>23</sup> In addition, I need to calibrate the exogenous exit rate in the well’s outside option. This is difficult: as previously discussed, it is optimal for all matches to be accepted and so ‘taking the outside option’ occurs off the equilibrium path. I choose a value of  $P_{exit} = 0.5$ .<sup>24</sup>

I calibrate the potential project entry cost  $c_{entry}$  to \$13.15 million USD using Hossain (2015).<sup>25</sup> I calibrate the scrap value  $b_{scrap} = \$5$  million USD using the figure from Kaiser and Snyder (2013). (Moreover, when discussing this market, this survey also mentions “Very few drilling contractors scrap rigs” which would make it difficult to estimate this value directly rather than calibrating it).

<sup>22</sup>I choose this value based on the assumed dayrate for a heavy lift marine transport ship undertaking a ‘wet tow’ suggested by industry practitioners in Terpstra, Hellinga and Leerdam (2013). While a ‘dry tow’ may be more expensive than a ‘wet tow’ since it is faster, industry practitioners suggest that there are also substantial other cost savings to using a dry tow (Dockwise (2012)) and so I assume that overall these values are comparable.

<sup>23</sup>This is somewhat close to the value of 0.37 used for the bargaining parameter in the shallow water analysis in Vreugdenhil (2026).

<sup>24</sup>Brancaccio, Kalouptsi and Papageorgiou (2020) also need to calibrate a similar value for their ‘exporter survival rate’.

<sup>25</sup>Hossain (2015) puts pre-spud drilling costs at around 18% of total expenses. Using this number, and setting other expenses to the mean total payment to a rig, which is \$59.9 million, I calibrate the entry cost as  $(0.18/0.82) \times 59.9 = 13.15$  million dollars. The paper Vreugdenhil (2026) follows a similar procedure using numbers in the shallow water market.

**Discussion of parametric assumptions** I assume that the distribution of complexity for new wells  $f_{l,x}$  is given by a log-normal distribution with mean  $\mu_l$  and standard deviation  $\sigma_l$ .

I assume that the match value is given by the functional form  $m_{x,y} = m_{0,y} + m_{1,y}x$ , where  $m_{0,y}$  and  $m_{1,y}$  are parameters that depend on the type of rig  $y$ . Importantly, the parameter  $m_{1,y}$  indexes the complementarities between applying a type- $y$  rig to a type- $x$  well. I further discuss these complementarities in the estimation results (Section 6). In addition to this affine functional form, I also impose the engineering constraint that low-specification and mid-specification rigs cannot match with wells of complexity above the 99% empirical quantile of that rig type. This is to capture engineering constraints that may inhibit some well and rig matches. The overall effect is that this restriction makes the model more conservative about leakage predictions, since it prevents particular wells from reallocating to particular rigs.

I now discuss the two estimation steps in more detail. I only use data in estimation in the period outside the US 2010 drilling moratorium.

## 5.2 Step 1: Computing supply-side parameters

This section is similar in spirit to the estimation strategy in Brancaccio, Kalouptsi and Papageorgiou (2020). I recover  $\delta_{l,y}$ ,  $\sigma_\varepsilon$ , and  $b_{stay}$  using the observed choice probabilities of moving between locations for each rig type using maximum likelihood. I can then back out  $c_{l,y} = \bar{p}_{l,y} - \delta_{l,y}$ . I provide more information about how I compute the value functions and the exact algorithm for estimation in Appendix A-3.1.

Formally, these parameters are identified by the following argument. Suppose that there are  $n \geq 3$  locations and focus on identifying the parameters for a particular rig type  $y$ . There are  $n + 2$  distinct parameters to identify: one markup parameter for each of the  $n$  locations as well as  $b_{stay}$  and  $\sigma_\varepsilon$ . In each location there are  $n$  choice probabilities (the rig can choose to stay or move to one of the  $n - 1$  different locations), and  $n - 1$  degrees of freedom since these probabilities must sum to 1. This leads to  $n(n - 1)$  degrees of freedom in total. So long as the number of locations  $n \geq 3$ , then  $n(n - 1) \geq (n + 2)$  and the parameters are identified.

Overall, the identification intuition is that  $b_{stay}$  is identified by the probability of an available rig remaining in the same location. The  $\delta_{l,y}$  parameters are identified by matching the probability of a single location choice per location (e.g. the choice probability of a move from Asia to Africa for high-specification rigs would identify  $\delta_{Asia,high}$ ). In one location (the US market) I have information on deepwater rig operational expenses from Kaiser and Snyder (2013), and so I incorporate this information into estimation by calibrating  $\delta_{US,y}$  based on this.

Finally, given the other parameters, there are many remaining location choice probabilities to pin down  $\sigma_\varepsilon$ . Intuitively, the model matches these choices ‘on average’, with higher  $\sigma_\varepsilon$  corresponding to choice probabilities that generate a more ‘spread out’ stationary distribution of rigs in each location. Lower values of  $\sigma_\varepsilon$  yield location choices where rigs predominantly choose locations with the highest markups (as well as the highest probability of matching).

### 5.3 Step 2: Computing the match value and the distribution of potential wells

A key challenge is that, although I observe contracts (price, duration, and the parties) in each location, I only have matched contract-project data where I see the exact well type drilled in the US market. For non-US markets, I therefore employ a strategy of estimating the distribution of potential wells  $f_{l,x}$  from the price/contract data alone. As well, I also estimate the Poisson parameter  $\lambda_l$  that governs the number of draws  $D_l$  from this distribution in each location. Intuitively, this strategy requires knowing the mapping between prices and well types so the distribution of prices identifies the underlying distribution of well types. Therefore, I split Step 2 into two substeps: I first retrieve the parameters that underlie the match value function (as well as  $f_{l,x}$  and  $\lambda_l$ ) in the US market using simulated method of moments. Then, I use the estimated match value parameters and data on prices and utilization to estimate the equilibrium distribution of potential projects, as well as the number of draws  $\lambda_l$ , in the other markets using simulated method of moments.

**Moments and identification** The six match value parameters ( $\{m_{0,y}, m_{1,y}\}_{y \in \{low, mid, high\}}$ ) are determined by two sets of moments constructed using data from the US market. First, I include moments that match coefficients from the following auxiliary regression of prices on project complexity and capital type for each contract in the US market where project characteristics are observed:

$$price_i = \beta_{0,y} + \beta_1 complexity_i + \beta_2 \cdot (\max \text{drilling depth})_i \cdot complexity_i + \varepsilon_i \quad (10)$$

where  $\beta_{0,y}$  is a capital-specific fixed effect. Equation (10) captures the relationship between prices and contract characteristics through the match value in the Nash bargaining solution. Intuitively — for a given match and after adjusting for the outside options of the parties — a higher price corresponds to a higher match value. I match three coefficients from this regression:  $\beta_{0,low}, \beta_1, \beta_2$ . Note that in this regression ‘max drilling depth’ is measured in feet. Second, I include the average price for each capital type (3 moments). Intuitively, the average price moments identify  $m_{0,y}$ . Fitting the remaining three moments from the auxiliary regression identifies the  $m_{1,y}$  parameters that govern complementarities between capital type and project type.

Once the match value parameters are pinned down, the parameters that characterize the distribution of potential projects  $(\mu_l, \sigma_l)$ , as well as the number of draws of potential wells  $\lambda_l$ , are identified from observed matches and rig utilization rates. To understand the intuition behind identification here, consider the case of the US market, and for expositional purposes first consider a simplified version of the model with a single rig type. Recall from the discussion in Section 4.1 that, since wells can direct their search towards their best match and there is an entry condition, no matches are rejected in equilibrium. However, potential wells may still fail to match with a rig if the rig queue is at its maximum backlog (where this maximum backlog is calibrated). Since the entry process into the queue is at random, the probability that a well fails to match is constant across all well types  $x$ . Therefore — in the special case of a single rig type — the distribution of wells that enter is exactly equal to the distribution of observed matches. With the entry condition known, the underlying distribution of potential wells  $f_{l,x}$  (which may enter or not enter) can then be recovered. Furthermore, the number of draws of potential wells  $\lambda_l$  can be mapped through the queuing simulation to a unique utilization rate of rigs (so long as the utilization rate is  $< 1$ ).

This basic intuition then extends to the case with multiple rig types. Here, I need to account for the fact that entered wells will target their best-matched rig, and there will be different queues for each rig type. Therefore, the distribution of observed matches will no longer be equal to the distribution of entered wells. However, since the targeting condition in Equation 1 can be pinned down using the value of a match (which is identified from the price moments), and the equilibrium  $q_{l,y}^{project}$  is determined from the matching simulation, there is a unique mapping between  $f_{l,x}$  and observed matches through (i) the entry condition, (ii) the targeting condition, (iii) the queuing simulation. This mapping can then be inverted to recover the  $f_{l,x}$ . Similarly, there is a unique mapping from  $\lambda_l$  through (i), (ii), (iii), to a utilization rate in each location.

In practice, I use the following moments to identify  $(\mu_l, \sigma_l)$ . For the US market, these parameters are identified by moments related to the average well-complexity match for each capital type (3 moments). These moments also ensure that the model matches the sorting patterns between capital and projects.

For the non-US markets these parameters are identified by matching the average prices for low and high capital types (2 moments per market). Finally, the Poisson parameter for new project entry  $\lambda_l$  is determined by the average capital utilization in each market (1 moment per market); higher values of  $\lambda_l$  correspond to higher capital utilization.

**Computation** I provide information about how I compute the demand-side equilibrium in Appendix A-3.2. Using this algorithm, I first compute the equilibrium in the US which returns  $f_{l,x}$  and  $\lambda_l$  in the US market and the

match-value function. Using this match-value function, I next estimate  $f_{l,x}$  and  $\lambda_l$  in each remaining location. To fit the parameters I use the standard GMM criterion function with the weight matrix as the identity matrix, except for the average well-complexity match moments which I weight by 0.1 to ensure they are of the same scale as the other moments. I also provide a Monte Carlo simulation in Appendix Table A-1 to illustrate that the estimator performs well at recovering  $f_{l,x}$  and  $\lambda_l$ .

## 6 Results

**Supply side parameters** Table 4 presents the estimated parameters from both the supply side and the demand side. Values for costs, match value parameters, and the other parameters are given in millions of dollars per day. The values for the preference shock  $\sigma_\varepsilon$  and the stay put benefit  $b_{stay}$  in Table 4(c) are both relatively low. For example, scaling up the  $\sigma_\varepsilon$  to a per-month value (i.e. the value per period in the model) is \$3.3 million; the total price paid to a rig on average per match (with a six-month contract) is around \$60 million.

In Table 4(a) I report the average cost over all rig types within a location; I report costs also broken out by rig type and location in Appendix Table A-4. Overall, the estimates reveal heterogeneity in rig operational costs across regions. For example, Europe has some of the highest drilling costs globally, consistent with this region having higher employment standards and salary requirements for workers, which are a major component of rig owner operational expenses.

**Demand side and match value parameters** The fit of the model to the targeted moments is detailed in Appendix Table A-3; since the model is exactly identified, the model closely fits the empirical moments. I also perform a model validation exercise centered around predicting the average prices of mid-specification rigs in each location. These moments are not used in estimation (with the exception of the US market). I plot the fit to these untargeted moments in Appendix Figure A-2. The model also closely fits the untargeted moments with a median difference of only 3.4 percent.

The demand parameter results are in the second, third, and fourth columns of Table 4(a). The estimates reveal substantial differences in project complexity distributions across the world. For example, the mean project complexity terms  $\mu_l$  are consistent with the utilization of different types of capital in different fields. For instance, Africa is a primary market for high-specification rigs and contains complex projects that involve drilling deep and high-pressure wells.

Next, consider the match value results for  $m_{0,y}$  and  $m_{1,y}$  in Table 4(b). Theoretically, these estimates must

**Table 4:** Estimation results

(a) Location-specific Parameters				
	Costs (av. over y)	# Entry	Mean	Std. dev
	$c_{l,y}$	$\lambda_l$	$\mu_l$	$\sigma_l$
Africa	0.16 <sup>†</sup>	6.93	0.97	1.28
	(0.0042)	(0.44)	(0.24)	(0.21)
Asia	0.11 <sup>†</sup>	5.90	0.63	0.84
	(0.0040)	(0.45)	(0.28)	(0.22)
Australia	0.17 <sup>†</sup>	3.17	0.58	0.86
	(0.0056)	(0.47)	(0.24)	(0.21)
Central Am.	0.18 <sup>†</sup>	3.59	0.81	0.85
	(0.0054)	(0.46)	(0.17)	(0.22)
Europe	0.20 <sup>†</sup>	14.13	0.68	1.19
	(0.0093)	(1.95)	(0.23)	(0.27)
Mid East	0.14 <sup>†</sup>	2.81	0.60	0.75
	(0.0045)	(0.47)	(0.49)	(0.31)
South Am.	0.11 <sup>†</sup>	13.07	0.56	0.60
	(0.0055)	(0.88)	(0.15)	(0.19)
US	0.13 <sup>†</sup>	6.07	0.63	0.89
	(0.0000)	(0.54)	(0.09)	(0.07)

(b) Match value parameters			(c) Other parameters	
	$m_{0,y}$	$m_{1,y}$		
Low-spec	0.697	-0.382	Preference shock ( $\sigma_\varepsilon$ )	0.11
	(0.051)	(0.046)		(0.02)
Mid-spec	0.497	-0.0319	Stay put benefit ( $b_{stay}$ )	0.10
	(0.059)	(0.023)		(0.03)
High-spec	0.347	0.016		
	(0.063)	(0.020)		

*Note:* Standard errors in brackets computed using 200 bootstrap replications. The † symbol on the cost estimates indicates that reported here is the average of the costs over the rig types in each location. The full cost matrix — which is used in the counterfactuals — is reported in the Appendix.

satisfy increasing differences to generate the empirical positive sorting patterns between capital efficiency and project complexity. The empirical estimates satisfy this increasing differences requirement. Beyond increasing differences, however, the exact sign and ordering of the coefficients is theoretically ambiguous since the match value represents both costs and benefits of drilling different well complexities. For example, for the match value slope parameter  $m_{1,y}$ , this may be negative for some rig types (e.g. complex projects incur more costs to the well owners — such as drilling delays or the need to replace a damaged part of the well — and this may differ with rig type), or it may be positive (more complex projects tend to produce more oil).<sup>26</sup> Overall, the match value estimates indicate that low-specification rigs have a comparative advantage in drilling simple projects. Conversely, high-specification rigs have a comparative advantage in drilling more complex wells.

## 7 Counterfactuals

I investigate counterfactual scenarios that center around a moratorium on drilling complex wells. This policy corresponds to real-world potential regulations in the industry. For example, it reflects the practical effects of a ban on new drilling permits or new leases, as has been proposed but not fully implemented by the Biden Administration: Friedman (2023). Specifically, over time the industry has expanded into deeper waters (which are more complex to drill, involving higher pressure formations and greater depths), and so new permits and lease sales tend towards these kinds of wells.<sup>27</sup>

I evaluate the efficacy of this policy if it were implemented using incomplete versus more complete regulation. I begin with US-only regulation, motivated by the fact that many proposed policies for the domestic US industry are not developed cooperatively with other regions. I also consider a global agreement, as well as coordinated regulation through a coalition of ‘richer countries’ (incorporating the US, Europe, Australia, and South America). These parties to a regional agreement align approximately with what is known as the ‘regulated areas’ of the global oil and gas industry (see e.g. Holand (2017) who uses this terminology).

As mentioned, the method to compute emissions in Appendix A-1 accounts for the effects of changes in global oil and gas prices due to changes in deepwater production.

---

<sup>26</sup>Given that the relationship between oil production and well complexity is robust to rig-type controls, as highlighted in Appendix Section A-1, in this setting it is likely that complementarities are due to differences in costs borne by the well owner — like the cost of extra materials if there is a drilling incident — when wells of the same complexity are drilled by different rigs. This is apparent in estimates for  $m_{1,y}$ , which indexes how the match value changes with well complexity. Specifically, these estimates are either small and not significantly different from zero (in the case of high-specification rigs), or negative.

<sup>27</sup>Recall that in order to drill a well, a tract on the ocean floor needs to be leased from the government and then a permit to drill needs to be granted. This system is typically used throughout the world. A lease and permit grant an oil and gas company the option to drill a well. Gramling and Freudenburg (2012) provide a summary of the evolution of the industry in the US.

**Computation** Unlike in estimation, where I was able to leverage empirical objects like the probability of matching and prices to simplify the computation, in the counterfactuals I need to re-solve for the entire equilibrium. I provide the algorithm in Appendix A-3.3.

## 7.1 Discussion: Ban on complex wells

I implement the ban on drilling complex wells by eliminating wells with a complexity index greater than 4.0, which is around the upper third of well complexity globally. The results are in Figure 5.

**US-only ban** I first consider the effect of a ban on complex wells only in the US market. The results in Figure 5(a) show that the regulation decreases carbon dioxide emissions by reducing oil production in the US by -21.2% but also reduces profits by -14.3%.

If a policymaker looked at the effects of the regulation on the US market in isolation (as is typically the case when doing cost-benefit analyses in this industry for the offshore oil and gas leasing program and other regulation, e.g., BOEM (2016)) and did not allow for the possibility of capital reallocation, they might conclude that the regulation is effective in reducing pollution, albeit expensive. However, looking at the total effect reveals that the regulation is -36.6 percent less effective — as measured by the reduction in total emissions — due to capital reallocation.

The leakage analysis in Figure 5(a) shows where this inefficiency is coming from. Re-sorting of rigs to other matches within the US implies that for every unit of carbon dioxide saved due to the regulation, 0.11 units are generated through increased drilling of other well types. Leakage across space is also important: for every unit of carbon dioxide saved by decreasing production in the US, 0.26 units are produced elsewhere.

Moreover, because the regulation spurs movement of rigs to other locations to which they are worse matched, or causes rigs to re-sort to wells where there is less complementarity, the regulation generates capital misallocation. The total effect is illustrated in the change in profit numbers in Figure 5(a).

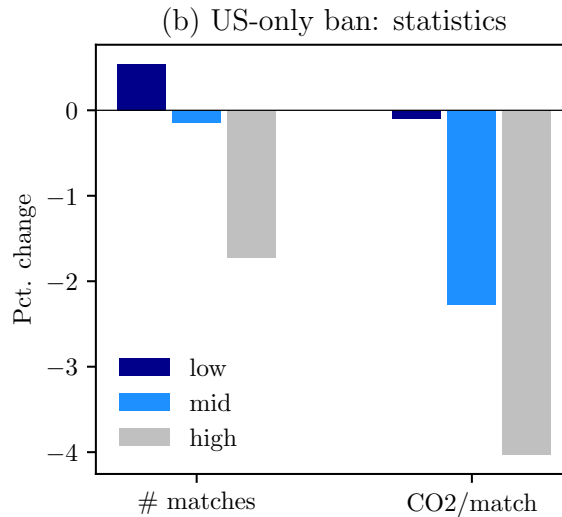
Figure 5(b) illustrates some key statistics that underlie the results. While high-specification rigs are matched with a lower probability due to the ban, because many of these rigs relocate out of the US, the remaining low-specification rigs actually have a higher number of matches due to within-location reallocation.

This reallocation affects the sorting patterns, which affects which kinds of wells are drilled, which then affects carbon emissions. As illustrated by the results in the carbon dioxide per match in Figure 5(b), a ban on complex wells causes high and mid-specification rigs to reallocate towards simpler wells. Although the rigs are less well-

**Figure 5: Counterfactual: complex well ban**

(a) Overall results

Counterfactual	Leakage (per unit)		Total change (percent)		Total change from reallocation (percent)	
	Re-sorting within regulated locations	Spatial	Regulated locations	Global	Full model	No het.
<b>US-only ban</b>						
CO <sub>2</sub>	0.11	0.26	-21.2	-1.7	-36.6	-73.6
Profits	0.38	0.08	-14.3	-1.5	-46.1	-71.6
<b>Coalition ban</b>						
CO <sub>2</sub>	0.08	0.09	-21.5	-10.7	-17.0	-30.6
Profits	0.20	0.02	-13.6	-8.2	-21.5	-30.1
<b>Global ban</b>						
CO <sub>2</sub>	0.11	0.00	-24.5	-24.5	-10.8	-37.8
Profits	0.23	0.00	-16.4	-16.4	-23.3	-37.4



*Note:* Part (a): Leakage is defined as the increase in emissions (or profits) generated through unregulated economic activity, for a one unit decrease in emissions (or profits) in the regulated activity. I decompose this into: (i) leakage due to capital reallocation from re-sorting within regulated locations, where I define the “regulated market” as complex wells in the US market, and the unregulated market as all other wells in the US market; (ii) leakage from the movements of capital to other locations, where I define the “regulated market” as the US market and the “unregulated market” as all other locations. The column ‘Total change from reallocation’ is the total decrease in effectiveness compared to a ‘no reallocation benchmark’ where sorting patterns and rig locations are fixed. Part (b): Summarizes heterogeneous effects of two key components that underlie the results: (i) the change in the number of matches for each rig type; (ii) for each match the average carbon dioxide emitted when the resulting oil and gas is burned (which can change as rigs sort to other wells). Appendix A-3 has detail on this figure for the other counterfactuals.

matched here, there is an environmental benefit: these wells tend to produce less oil and therefore fewer emissions once the oil is burned.

Note that in this counterfactual no rigs choose to exit from the global market. In part, this is due to the fact that capital is able to reallocate to other wells and regions not under regulation, which moderates the effect on the change in  $U_{l,y}$  for each rig type, and this is robust to an order of magnitude change in  $b_{scrap}$ , as I further discuss in Section 7.4. Indeed, I find that no rigs would exit in any of the counterfactuals. Nevertheless, it is still important to include this extensive margin in the model, as well as the extensive margin for wells. This is because if rigs did exit, but the channel was not included, then the model would tend to overpredict leakage.

**Global ban** I also consider how a global agreement would affect the market. Since regulation is now uniform, there are no unregulated locations. However, there are still well types which are not banned within each location. As a consequence, although there is no leakage across space, shutting down this channel exacerbates within-location leakage to 0.11 tons of carbon dioxide produced elsewhere for every ton saved directly from the ban.

Globally, the change in emissions is -24.5 percent. This is mainly coming from more locations under regulation. Reallocation still lowers the efficacy of the regulation by -10.8 percent.

**Coalition ban** Although a global ban does reduce leakage, it may not be politically feasible. Therefore, I also consider a more pragmatic agreement involving a coalition of rich countries. A coalition ban would be substantially more effective than a US-only ban, with capital reallocation undercutting the efficacy of the regulation in terms of total global emissions by -17.0 percent. Spatial leakage is still relatively high at 0.09: although more regions are under regulation, the coalition does not encompass all locations with complex wells.

## 7.2 Discussion: Role of two-sided vertical heterogeneity

In theory, incorporating two-sided vertical heterogeneity could amplify or diminish the effects of leakage. It could amplify the effects if the regulation causes high-specification rigs to disproportionately exit to other locations. This is because the entry of these rigs into other locations would reduce capacity constraints on the most complex wells, which also tend to produce more oil and gas. On the other hand, incorporating two-sided vertical heterogeneity makes rigs and wells less substitutable, which would tend to reduce leakage from reallocation.

To quantify the empirical relevance of these channels, in the final column of Figure 5(a) I present the results of the model if two-sided vertical heterogeneity were eliminated and all rigs and all wells were the same. In this exercise, for simplicity, I also set location-specific costs to their average, so locations are differentiated only by the number

of potential well draws  $\lambda_j$ .

Comparing the results in the final column of Figure 5(a) to the full model, the ‘homogeneous agents’ model would vastly over-predict leakage in the counterfactuals. For example, in the US-only counterfactual, reallocation would be predicted to undercut the change in emissions by -73.6 percent, around double the value in the full model.

### 7.3 Discussion: Long-run Hotelling dynamics

One may ask about the implications of long-run Hotelling dynamics. Specifically, if the decision of oil and gas companies is just *when* to extract oil from their reserves, the movement of rigs to other locations would simply bring forward in time when the oil is extracted, but would not affect the long-term cumulative amount of oil produced in unregulated regions.

Consistent with other papers that study supply-side policies in the oil and gas industry, including Prest (2022), Prest and Stock (2023), and Covert and Kellogg (2023), I do not include these Hotelling dynamics in the model. The reason is that — similarly to these papers which study the onshore oil and gas industry — deepwater oil and gas reserves do not seem to be declining. This is due to continued discovery of new reserves and productivity improvements that mirror those documented in the onshore oil and gas industry (e.g. Agerton (2020)). Consequently, and similar to the arguments made in these papers, increases in production in unregulated regions come from reserves that would otherwise be kept ‘in the ground’ for an extremely long time. Arguably, this time-span is potentially long enough that improvements in clean technologies or a global climate agreement may then significantly reduce oil demand.

### 7.4 Robustness

I run several robustness tests, focusing on how the counterfactuals are affected by different values of the calibrated parameters. Across all the robustness tests, the main finding — that leakage under US-only regulation is relatively high but a coalition or global regulation would improve outcomes — is qualitatively unchanged.<sup>28</sup> The leakage measure is a ratio whose numerator and denominator both shift with parameter changes, so the response of the ratio is not generally obvious from theory. Furthermore, as discussed, the equilibrium is determined by many complex effects (e.g., two-sided sorting interacting with location choices across heterogeneous locations) and so any single parameter change can echo through the model in non-obvious ways. However, for the two parameters with the largest quantitative effects ( $c_d$  and  $\eta$ ), I sketch the mechanism behind the response in their respective

---

<sup>28</sup>A key reason behind this is that although the counterfactual effect changes after a parameter change, so does the baseline effect. This moderates the overall difference in the results.

discussions below.

I first consider robustness with respect to the scrap value  $b_{scrap}$ . I find that the value of the parameter where firms would actually scrap rigs — defined as  $\min_l \{U_{l,y}\}$  — is on the order of hundreds of millions of dollars across all the counterfactuals. Therefore, the kinds of values of  $b_{scrap}$  that would cause rig exit are implausibly high compared to the calibrated value (which is based on industry studies). Specifically, an order of magnitude change to the assumed scrap value (\$50 million) or even a 50-fold increase (\$250 million) would not induce any rig exit in any of the counterfactuals.<sup>29</sup>

Next, I consider if reducing the calibrated value of  $P_{exit}$  by 20% would change the results in Appendix Table A-5a. This reduction of  $P_{exit}$  is economically meaningful. For example, consider the probability that a rig is at its capacity constraint — and so a project will be unable to match and exit —  $1 - q_{l,y}^{project}$ . A reduction of  $P_{exit}$  by 20% is comparable to moving from the median value of  $q_{l,y}^{project}$  for a high-specification rig across locations to the maximum value across locations. Appendix Table A-5a illustrates that the leakage results are robust and qualitatively similar under this test, in the sense that there are substantial gains to reducing leakage moving from unilateral US-only action to a coalition or a global agreement.

I also investigate robustness to the calibrated value of  $c_{entry}$  in Appendix Table A-5b. Here I increase and decrease  $c_{entry}$  by 20%. These changes are also economically meaningful. For example, compare these changes to the estimated location-specific drilling costs  $c_{l,y}$  in Table 4, scaling the daily  $c_{l,y}$  up over a typical 6-month contract. Moving from the lower bound to the upper bound of  $c_{entry}$  considered in this test is a comparable cost change to moving drilling from the United States — which has a relatively low  $c_{l,y}$  — to the African market, which has a more moderate  $c_{l,y}$ . Overall, I find that the results are robust to changing this value.

As well, I test robustness to increasing and decreasing the value of the Nash bargaining parameter  $\eta$  by 20% in Appendix Table A-5d. I choose this change in  $\eta$  because the lower bound approximates the estimated value of the Nash bargaining parameter in the shallow water market in Vreugdenhil (2026). Intuitively, a lower value of  $\eta$  gives more rents to drilling projects. These rents change incentives for where rigs spatially locate in the baseline and counterfactuals. I find that lowering  $\eta$  raises the leakage estimate under US-only regulation. This is primarily due to the across-location channel: the equilibrium baseline shifts more rigs into the US, so the US ban displaces

---

<sup>29</sup>Note that the values for  $U_{l,y}$  seem reasonable, as they should approximate the value of a second-hand floater rig. In 2008-2010 (which overlaps with the sample period 2008-2015), this value was on the order of hundreds of millions of dollars (Kaiser and Snyder, 2013). Furthermore, as mentioned, the difference between the scrap value and  $U_{l,y}$  in the model baseline links with the empirical observation that “very few drilling contractors scrap rigs” (Kaiser and Snyder, 2013).

a larger pool of US rigs into non-US locations. Again, the results are qualitatively robust to changing this value, with US-only regulation producing substantially more leakage than coalition or global agreement counterfactuals.

Finally, I consider robustness to changing the calibrated distance cost  $c_d$  by 50% in Appendix Table A-5c. Such changes to the parameters are quite extreme, and so these tests arguably represent loose bounds on the true quantitative magnitudes of leakage. Despite these large changes to  $c_d$ , the qualitative ordering of the counterfactuals, as well as the main finding that leakage under US-only regulation is relatively high but could be reduced significantly through a coalition or global agreement, persists.

Quantitatively, leakage is still relatively large under the robustness tests involving  $c_d$ . Increasing the value of  $c_d$  appears to amplify the costs of US-only regulation compared to coordinated action. Intuitively, increasing  $c_d$  changes where rigs leave after US-only regulation is imposed. Specifically, rigs (and particularly high-specification rigs) tend to relocate to South America, which is relatively closer to the US market; this changes the amount of leakage generated elsewhere after a ban. However, even in the case most favorable to US-only regulation — reducing  $c_d$  by 50% — leakage is still around one-quarter in the US-only counterfactual. Moreover, a coordinated agreement of rich countries reduces leakage in terms of pollution by over one-third, and leakage in terms of profits by around one-quarter. Furthermore, note that the calibrated value of  $c_d$  is arguably conservative since it relies on the value of a dry tow but does not include other (smaller) moving costs such as insurance or rig preparation. This therefore implies that the counterfactual results which highlight the benefits of coordinated action are likely to also be conservative.

## 8 Conclusion

Supply-side climate policies in the global oil and gas industry are increasingly being proposed and implemented throughout the world. Since these policies are usually incomplete, a key question is whether leakage might undermine their efficacy. In this paper I quantify the role of leakage via capital reallocation and the potential gains to a global agreement. To do so I develop a framework that extends the literature on spatial matching models in industrial organization to incorporate two-sided vertical heterogeneity of firms leading to sorting. I apply the framework to a previously unexplored dataset of contracts and relocation decisions in the market for offshore deepwater drilling rigs.

I find that supply-side climate policies, when implemented through US-only incomplete regulation, induce substantial responses through capital reallocation. This reallocation undercuts the environmental benefits of regulation, causing oil to be produced elsewhere in the world while inducing spatial misallocation. A global ban

would be more effective, as would a more politically feasible coalition of rich countries implementing coordinated regulation. The results illustrate that capital reallocation is an important channel for leakage and should be a central consideration in the design of supply-side policies in the oil and gas industry. More generally, the results have implications for quantifying leakage in other settings where capital may be less mobile. Concretely, even in the setting of this paper where the physical capital is as mobile as can be, there is still less than full leakage. This highlights a key factor that limits leakage in settings where agents need to match to produce output and there is two-sided vertical heterogeneity: the imperfect substitutability between capital and projects within and across space.

## References

- Abito, Jose Miguel, Christopher R. Knittel, Konstantinos Metaxoglou, and André Trindade**, “The Role of Output Reallocation and Investment in Coordinating Environmental Markets,” *International Journal of Industrial Organization*, July 2022, 83.
- Agerton, Mark**, “Learning Where to Drill: Drilling Decisions and Geological Quality in the Haynesville Shale,” *Working Paper*, April 2020.
- Ahlvik, Lassi, Jørgen Juel Andersen, Jonas Hveding Hamang, and Torfinn Harding**, “Quantifying Supply-side Climate Policies,” *CAMP Working Paper Series*, 2022, (1).
- Baylis, Kathy, Don Fullerton, and Daniel H. Karney**, “Leakage, Welfare, and Cost-Effectiveness of Carbon Policy,” *American Economic Review*, May 2013, 103 (3), 332–37.
- BOEM**, “Economic Analysis Methodology for the 2017-2022 Outer Continental Shelf Oil and Gas Leasing Program,” *Bureau of Ocean Energy Management, Technical Report*, November 2016.
- Brancaccio, Giulia, Myrto Kalouptsi, and Theodore Papageorgiou**, “Geography, Transportation, and Endogenous Trade Costs,” *Econometrica*, March 2020, 88, 657–691.
- Buchholz, Nicholas**, “Spatial Equilibrium, Search Frictions, and Dynamic Efficiency in the Taxi Industry,” *The Review of Economic Studies*, March 2022, 89, 556–591.
- Copeland, Brian R. and M. Scott Taylor**, “Trade and the Environment: Theory and Evidence,” *Princeton Series in International Economics*, 2003.
- Corts, Kenneth S.**, “Stacking the Deck: Idling and Reactivation of Capacity in Offshore Drilling,” *Journal of Economics & Management Strategy*, 2008, 17, 271–294.
- and **Jasjit Singh**, “The Effect of Repeated Interaction on Contract Choice: Evidence from Offshore Drilling,” *Journal of Law, Economics, and Organization*, April 2004, 20, 230–260.
- Covert, Thomas and Ryan Kellogg**, “Environmental Consequences of Hydrocarbon Infrastructure Policy,” *Becker Friedman Institute for Economics Working Paper*, October 2023, (2023-138).
- Davis, Lucas W. and Matthew E. Kahn**, “International Trade in Used Vehicles: The Environmental Consequences of NAFTA,” *American Economic Journal: Economic Policy*, November 2010, 2 (4), 58–82.
- Dockwise**, “Dockwise Vanguard: New Generation of Heavy Marine Transport,” *NOIA Meeting*, October 2012.
- Ederington, Josh, Arik Levinson, and Jenny Minier**, “Footloose and Pollution-free,” *The Review of Economics and Statistics*, February 2005, 87 (1), 92–99.
- EIA**, “Trends in U.S. Oil and Natural Gas Upstream Costs,” *Report*, March 2016.
- Fowlie, Meredith and Mar Reguant**, “Challenges in the Measurement of Leakage Risk,” *AEA Papers and Proceedings*, May 2018, 108, 124–29.
- , —, and **Stephen P. Ryan**, “Market-Based Emissions Regulation and Industry Dynamics,” *Journal of Political Economy*, 2016, 124 (1), 249–302.
- Friedman, Lisa**, “Biden Said He Would Stop Drilling. Then Reality Hit.” *New York Times*, September 2023.
- Gavazza, Alessandro**, “An Empirical Equilibrium Model of a Decentralized Asset Market,” *Econometrica*, 2016, 84 (5), 1755–1798.
- Golson, Jordan**, “The Enormous Ship that Submerges Itself to Carry Entire Oil Rigs,” *Wired Magazine*, August 7, 2014.
- Gramling, Robert and William R. Freudenburg**, “A Century of Macondo: United States Energy Policy and the BP Blowout Catastrophe,” *American Behavioral Scientist*, 2012, 56 (1), 48–75.
- Harstad, Bård**, “Buy Coal! A Case for Supply-Side Environmental Policy,” *Journal of Political Economy*, February 2012, 120 (1), 77–115.
- Holand, Per**, “Loss of Well Control Occurrence and Size Estimators, Phase I and II,” *Bureau of Safety and Environmental Enforcement, Technical Report*, May 2017.
- Hossain, M. Enamul**, “Drilling Costs Estimation for Hydrocarbon Wells,” *Journal of Sustainable Energy Engineering*, 2015, 3 (1).
- International Energy Agency**, “Emissions from Oil and Gas Operations in Net Zero Transitions,” May 2023.
- Kaiser, Mark**, “A Survey of Drilling Cost and Complexity Estimation Models,” *International Journal of Petroleum Science and Technology*, 2007, 1 (1), 1–22.
- and **Brian Snyder**, *The Offshore Drilling Industry and Rig Construction in the Gulf of Mexico*, Springer-Verlag, London, 2013.
- Kellogg, Ryan**, “Learning by Drilling: Interfirm Learning and Relationship Persistence in the Texas Oilpatch,” *The Quarterly Journal of Economics*, 2011, 126, 1961–2004.
- Lewis, Eric**, “Patchwork Policies, Spillovers, and the Search for Oil and Gas,” *American Economic Journal: Economic Policy*, 2019, 11 (1), 380–405.
- Magnac, Thierry and David Thesmar**, “Identifying Dynamic Discrete Decision Processes,” *Econometrica*, 2002, 70, 801–

- Moen, Espen R.**, “Competitive Search Equilibrium,” *The Journal of Political Economy*, April 1997, 105 (2), 385–411.
- Prest, Brian C.**, “Supply-Side Reforms to Oil and Gas Production on Federal Lands: Modeling the Implications for CO<sub>2</sub> Emissions, Federal Revenues, and Leakage,” *Journal of the Association of Environmental and Resource Economists*, July 2022, 9 (4).
- **and James H. Stock**, “Climate Royalty Surcharges,” *Journal of Environmental Economics and Management*, July 2023, 120, 102844.
- **, Harrison Fell, Deborah Gordon, and TJ Conway**, “Estimating the Emissions Reductions from Supply-side Fossil Fuel Interventions,” *Resources for the Future Working Paper*, 2023, (23-11).
- Ramey, Valerie A. and Matthew D. Shapiro**, “Displaced Capital: A Study of Aerospace Plant Closings,” *Journal of Political Economy*, 2001, 109 (5), 958–992.
- Terpstra, T., E. Hellinga, and H. Leerdam**, “Offshore Dry-docking of FPSOs; A Response to Industry Needs,” *Offshore Technology Conference*, 2013.
- Vreugdenhil, Nicholas**, “Booms, Busts, and Mismatch in Capital Markets: Evidence from the Offshore Oil and Gas Industry,” *Journal of Political Economy*, 2026, 134 (5), 1468–1505.
- Yang, Ron**, “(Don’t) Take Me Home: Home Preference and the Interstate Trucking Market,” *Working Paper*, August 2024.

## Appendix

### A-1 Model for emissions

I require a model that maps changes in the number and type (i.e. the complexity) of the wells in each location into changes in global carbon dioxide emissions. I do this in three steps.

Note that, as I set out in more detail below, Step 2 and Step 3 reduce down to multiplying changes in oil and gas production in the deepwater market by a scale factor which is calibrated based on the best-available empirical estimates. I evaluate the counterfactuals in the paper mainly by either the *percent changes* in emissions, or leakage statistics that are the change in emissions generated in the unregulated markets divided by the change in emissions generated in the regulated market. These metrics are scale-free measures, and so typically do not rely on specific assumptions for the scale factors applied in Steps 2 and 3. The only exception is the set of numbers under the ‘Total change (percent)’ column of Figure 5(a) for the CO<sub>2</sub> rows, which would differ based on the factor chosen in Step 2.

#### 1. Map matches into oil and gas production in the deepwater market

For the first step, I run the regression:

$$\text{oil production}_i = \beta_0 + \beta_1 \text{well complexity}_i + \varepsilon_i \quad (\text{A-1})$$

which exploits the strong relationship between complex wells and increased oil production. In order to measure oil production, I take the average number of barrels produced per day for the first 10 years of production. The results from this regression are  $\beta_0 = 1558.2(1123.4)$  and  $\beta_1 = 1340.8(405.2)$  (standard errors in parentheses).

These results are robust to including controls for rig type. For example, including an additive control for the (centered) maximum water depth of the rig — which defines rig type — results in  $\beta_0 = 1849.1(1259.1)$  and  $\beta_1 = 1225.2(463.5)$  (standard errors in parentheses), and the control is not significant at the 10% level. These results for  $\beta_0$  and  $\beta_1$  are not statistically different from the baseline results at the 10% level, and the point estimates are relatively close. Similarly, including an interaction of complexity with the centered maximum depth of the rig results in  $\beta_0 = 1035.9(1256.2)$  and  $\beta_1 = 1596.5(489.5)$ , and the interaction control is not significant at the 10% level. Again, these results are not statistically different from the baseline results at the 10% level, and the point estimates are relatively close.

## 2. Map changes in oil and gas production in the deepwater market into global changes

In this step I follow Prest et al. (2023). I assume that the supply and demand curves in the world market are well-approximated (locally) by constant elasticity supply and demand curves. Then, a partial equilibrium change of an extra barrel produced in the deepwater market will result in a total equilibrium change to the equilibrium barrels produced and consumed worldwide of  $1 - L$ , where  $L = e_S / (e_S - e_D)$  defines ‘leakage’ due to changes in global oil and gas prices, and  $e_D$  and  $e_S$  denote global demand and supply elasticities. That is, I multiply the partial equilibrium change in production in the deepwater market by  $1 - L$  (e.g. if the regulation results in a change of -100 barrels the global change is  $-100(1 - L)$ ).

Prest et al. (2023) provides a comprehensive analysis of various  $e_D$  and  $e_S$  used in the literature and computes a leakage of approximately  $L = 0.56$  on average. Therefore, I use a multiplicative factor of 0.44.

## 3. Map changes in global oil and gas consumption into changes in global carbon dioxide emissions

In this step I convert this global change in output to carbon emissions by scaling by the EPA’s Greenhouse Gas Equivalencies Calculator. This factor is 0.43 metric tons carbon dioxide/barrel oil.<sup>30</sup>

## A-2 Data construction

I use three main datasets for the analysis: 1. Rigzone data of rig status updates and contracts; 2. Permit, borehole, and production data from the BSEE (Bureau of Safety and Environmental Enforcement); 3. CERDI dataset of bilateral distances between locations.

**Cleaning the Rigzone data** Although most rigs operate under relatively short-run contracts, a small number of rigs operate under extremely long-run contracts for a single oil company (e.g. a 10-year contract). I delete rigs that operate under contracts that are longer than 2 years, treating these very long-run contracts as essentially a different type of market than those deepwater rigs which perform short-run work. Specifically, I delete 227 contracts that are longer than 2 years, comprising 15.8% of the total number of contracts. I also delete 30 contracts (less than 2.5% of the total dataset) in the South American location which have an unusually low dayrate. This is potentially due to data reporting issues in this location. Concretely, I use a cutoff of  $\$ < 0.1$  million/day but many of these contracts have a reported dayrate of  $\$0$  per day. The 2-year cutoff reduces the total number of contracts from 1438 to 1211.

**Merging the Rigzone data with the BSEE data** Critical to my estimation strategy are data that link rig and well characteristics to contract prices (and other contract details). In the US market I have data on both well characteristics (from the BSEE) and contract details (from Rigzone). In this section I describe how I merge these two datasets.

I begin with a dataset of 218 contracts for the US market, and 607 wells. I successfully match 218 wells with contracts, and collapse these to the contract level (many contracts contain multiple wells and for these cases I use the average complexity of the matched wells). Why are some wells and contracts not matched? Sometimes the rig name is recorded differently between the well dataset and the contract dataset, for example the rig’s name might change and the current name may be used in the contract dataset, whereas the original name is used in the well dataset. I try to connect as many rig and well names as possible by including previous rig names and accounting for simple typographical differences between the names in the datasets.

I use these contracts for the auxiliary regression in the model estimation (removing contracts during the period

---

<sup>30</sup>The calculator is here: <https://www.epa.gov/energy/greenhouse-gas-equivalencies-calculator-calculations-and-references>

of the moratorium). Note that for the graphs in Figure 2 I use the ‘uncollapsed’ sample of 218 wells. Also, note that for many metrics such as average prices, utilization, etc., I do not require the matched contracts and so where possible I use the full sample of contracts.

## A-2.1 Constructing metrics

**Rig relocations** I construct rig relocations simply by looking at changes in rig statuses (for example, a rig is in one region in one status and then in the next status it is in a different region). Typically, these relocations are that the rig is drilling in one region and then moves to drill in a different region. However, there are some exceptions to this. Notably, I include several movements from a shipyard (e.g. the rig is ‘Under Construction’ in Asia) to a different region (e.g. the rig is then ‘Drilling’ in USA) as a relocation. My justification for including this as a relocation choice is that these rigs could have alternatively just remained and drilled in the field in the region in which they were constructed (my definition of a region is large enough that it is always a possibility that this could occur) or potentially could have chosen to move after construction to a different location. Therefore, I assume that these relocations contain useful information about the relative costs and benefits of drilling in different locations and so include them in the relocation data.

**Rig utilization** Rig utilization is defined as the average proportion of time that a rig is being used in a given location. I begin by converting the statuses of every rig (which in the raw data look like, for example, a rig is “Drilling” between a block of two dates at a given dayrate for a particular oil and gas company in a particular location) into a daily dataset of activity at the individual rig level, for every day between 2008-July 2015. I then classify these daily statuses into whether a rig is utilized or not. I consider a rig utilized if it is in the statuses of “Drilling”, “Workover”, “Production”, “Modification”, “Inspection”. I consider the total number of rigs in a location (the denominator for utilization) as all these statuses except if a rig is “Under Construction” or “Cold Stacked” (mothballed). The main status of rigs not utilized in a location is that they are “Ready Stacked” (staffed and ready to drill with little delay).

Note that to compute the moments in estimation, I delete the period of time when the 2010 US drilling moratorium was active for the entire world market.

## A-3 Algorithms

### A-3.1 Algorithm for computing the supply-side equilibrium

I compute the parameters  $\delta_{l,y}$ ,  $\sigma_\varepsilon$ ,  $b_{stay}$  using maximum likelihood. I compute the likelihood function for each guess of the parameters as follows. Denote the  $k$ -th iteration of the value functions for a searching rig and an unemployed rig, respectively, by  $V_{l,y}^k$ ,  $U_{l,y}^k$ . Then, I compute the likelihood using predicted choice probabilities using the following algorithm:

1. Guess initial value functions at iteration  $k = 0$ :  $V_{l,y}^0$ ,  $U_{l,y}^0$ .
2. Using Equation (6), update the value of searching in each location  $l$ , for each type of rig  $y$ ,  $V_{l,y}^{k+1}$ . Use the empirical probabilities of matching in each location as the  $q_{l,y}^{capital}$ . Also, using Equation (5), update the value of unemployment in each location  $l$ , for each type of rig  $y$ ,  $U_{l,y}^{k+1}$ .
3. Repeat from Step 2. until the value functions converge.
4. Compute predicted choice probabilities for moving from location  $l$  to  $l'$  for a rig of type  $y$ , using  $V_{l,y}$  and

Equations (A-2) and (A-3). These choice probabilities are:

$$P_{l,l',y} = \frac{\exp\left(\left(-c_d d_{l,l'} + \beta V_{l',y}\right)/\sigma_\varepsilon\right)}{\exp\left(\left(b_{stay} + \beta V_{l,y}\right)/\sigma_\varepsilon\right) + \sum_{l' \neq l} \exp\left(\left(-c_d d_{l,l'} + \beta V_{l',y}\right)/\sigma_\varepsilon\right)} \quad (\text{A-2})$$

$$P_{l,l,y} = \frac{\exp\left(\left(b_{stay} + \beta V_{l,y}\right)/\sigma_\varepsilon\right)}{\exp\left(\left(b_{stay} + \beta V_{l,y}\right)/\sigma_\varepsilon\right) + \sum_{l' \neq l} \exp\left(\left(-c_d d_{l,l'} + \beta V_{l',y}\right)/\sigma_\varepsilon\right)} \quad (\text{A-3})$$

5. Compute the likelihood:  $L = \sum_y \sum_l \sum_{l'} \log P_{l,l',y}^{n_{l,l',y}}$ . Here, the value  $n_{l,l',y}$  is the number of observations (i.e. months) that I observe a type- $y$  rig move from location  $l$  to location  $l'$ .

### A-3.2 Algorithm for computing the demand-side equilibrium

Here I set out the algorithm to compute the demand-side equilibrium that I recompute at every iteration of the objective function in the second step of estimation. Note that rig costs are known from the first step of the estimation. I also know the empirical total number of rigs in each location  $n_{l,y}$ . Given these objects, and candidate match value parameters and the demand parameters, I compute an equilibrium as a fixed point in the probability that a project matches with each capital type  $\{q_{l,y}^{project}\}_{y \in \{low, mid, high\}}$ .

Note that the demand-side equilibrium computation is separable over each location (since I fix the equilibrium number of rigs  $n_{l,y}$  in each location). Therefore, the following algorithm centers on how to compute the equilibrium in one location.

1. Guess the matching probability  $\{q_{l,y}^{project,k}\}_{y \in \{low, mid, high\}}$ , where  $k$  denotes the iteration and  $k = 0$  denotes the initial guess.
2. Since rig costs are known, and for a candidate vector of demand-side parameters and the  $q_{l,y}^{project,k}$ , I can compute prices using the Nash bargaining solution. Therefore, I can compute the value of a project of type  $x$  targeting a rig of type  $y$  using Equation (1):  $\Pi_{l,x,y}^{project}$ .
3. Update the probability of a project matching in the type- $y$  rig submarket to iteration  $k + 1$ ,  $\{q_{l,y}^{project,k+1}\}_{y \in \{low, mid, high\}}$ , using a *matching simulation* (detailed below).
4. Iterate from Step 2 until convergence.

#### A-3.2.1 Matching simulation

The matching process outlined in the paper fits into a queuing framework for each rig  $y$  submarket. The queue has a ‘service time’ at the contract duration  $\tau$ ,  $n_{l,y}$  ‘servers’,  $n_{l,y} t_{backlog} / \tau$  places in the queue, and a queuing discipline of first-in-first-out.

Denote the iteration of the queue by  $h$ , where an ‘iteration’ can be thought of as a snapshot of the queue in a particular time period (a month) and an update to  $h + 1$  of the queue can be thought of as the transition of the queue to the next time period. I simulate the queue for each rig type  $y$  submarket, in each location  $l$ , in the following way:

1. Initialize the backlog of each of the  $n_{l,y}$  rigs at 0. Therefore, the state of the queue is a  $n_{l,y}$ -length vector where each element is the backlog of a particular rig. Denote each element (for the  $i$ -th rig) by  $b_{l,y}^{i,h}$ .

2. Denote the realization of the number of wells in iteration  $h$  that enter by  $d_l^h$ . Compute this  $d_l^h$  by taking a draw from  $Poisson(\lambda_l)$ . Then take  $d_l^h$  draws from the project complexity distribution  $f_{l,x}$ . Denote the type of each of these draws by  $x_l^{j,h}$ .
3. For each complexity draw  $x_l^{j,h}$ , compute the payoff to matching with each particular type of rig using Equation (1). (Note here that the  $\varepsilon_y$  draws in Equation (1) are project  $j$  and rig type  $y$  specific.)
4. For each  $j$ , find the rig type  $y$  where the payoff to matching is maximized. Determine for each  $j$  if this payoff satisfies the entry condition (i.e. the payoff is greater than  $c_{entry}$ ).
5. Denote the total number of wells that enter and target rig type  $y$  by  $d_{l,y}^h$ . Then, compute matches:
  - If the number of wells  $d_{l,y}^h$  is less than or equal to the total number of available rigs, then all wells will match with a rig. Therefore, add the contract length  $\tau$  to the backlogs of  $d_{l,y}^h$  available rigs. Note that in this iteration of the queue the probability of a well matching will be equal to 1.
  - If the number of wells  $d_{l,y}^h$  is greater than the number of available rigs, then not all of the wells will match. In this case, allocate wells to available rigs in the order of entry until the backlogs are completely full (i.e. adding one more match with contract length  $\tau$  would cause the backlog of a rig to be greater than the critical value  $t_{backlog}$ ). Note that in this iteration of the queue the probability of a well matching will be less than 1.
6. Update the queue to the next period  $h + 1$  by removing 1 month of each rig  $i$  with a backlog  $> 1$  (backlogs cannot be negative so if the current backlog is 0, there will be no change in  $b_{l,y}^{i,h}$ ).
7. Repeat from Step 2.

In practice, I start by ‘burning-in’ the matching simulation to remove dependence on the initial guess. I then iterate over  $h$  in the above algorithm many times. From this, I generate the long-run probability of a project matching in the type- $y$  capital submarket ( $q_{l,y}^{project}$ ), which is then used in the demand-side algorithm. As a side-product, the queuing algorithm also delivers  $q_{l,y}^{capital}$ , rig utilization (the proportion of rigs with a positive backlog), and the average match for each rig type, amongst other things. These objects are useful for computing counterfactuals and moments in the demand estimation.

### A-3.3 Algorithm for computing the counterfactuals

Overall, the algorithm for the counterfactuals involves recomputing the entire equilibrium of the global market for deepwater rigs. Unlike in estimation, I can no longer leverage empirical objects in the data like empirical probabilities of matching, because these will change in the counterfactuals.

1. Initialize the algorithm at iteration  $k = 0$  with a guess of the number of rigs of type  $y$  in each location  $l$ , denoted  $n_{l,y}^{k=0}$ . Also guess the probability of a project matching in each rig type  $y$  submarket in each location  $q_{l,y}^{project,k=0}$ .
2. For each location  $l$ , compute the demand-side equilibrium at  $n_{l,y}^{k=0}$  and also for ‘perturbations’ around  $n_{l,y}^{k=0}$ :
  - To compute the perturbations, I use the demand-side algorithm detailed above in Section A-3.2.
  - These perturbations correspond to varying the number of rigs of a particular type, holding the number of other rigs fixed *and* the  $q_{l,y}^{project,k}$  fixed when projects make their targeting choices. Since  $q_{l,y}^{project,k}$  is fixed, this requires computing steps 2. to 5. in the demand-side algorithm only once, and therefore computing the matching simulation only once.
  - For each perturbation, and for each rig type  $y$ , record the equilibrium price and the probability of a rig matching.
3. Given these perturbations, update the supply side for each rig type  $y$  *across* locations in the following way.

Note that the following is an inner loop which iterates over the distribution of rigs in each location, and I denote the iteration of this inner loop by  $h$ .

- (a) Initialize the inner loop at  $h = 0$  using the current distribution of rigs across space  $n_{l,y}^{h=0}$ .
  - (b) Using linear interpolation over the perturbations computed in Step 2., get the expected price in each location, as well as the probability of matching for rigs.
  - (c) Using these prices and the probabilities of matching, get the rig value functions  $U_{l,y}^h$ ,  $V_{l,y}^h$  and the corresponding conditional choice probabilities using the supply-side algorithm in Section A-3.1.
  - (d) Using the conditional choice probabilities, construct a transition probability matrix (i.e. a matrix where each row is the probability that a type  $y$  rig in location  $l$  will move to an alternative location  $l'$ ).
  - (e) Update the number of rigs in each location *once* to produce a new distribution of rigs across locations  $n_{l,y}^{h+1}$ .
  - (f) Repeat from sub-step 3(b) until convergence of the distribution of rigs of type  $y$  across locations.
4. Given the new supply-side equilibrium, update the demand side equilibrium using the demand-side algorithm detailed above in Section A-3.2.
  5. Repeat from step 2. until convergence in the distribution of rigs across space in the outer loop.

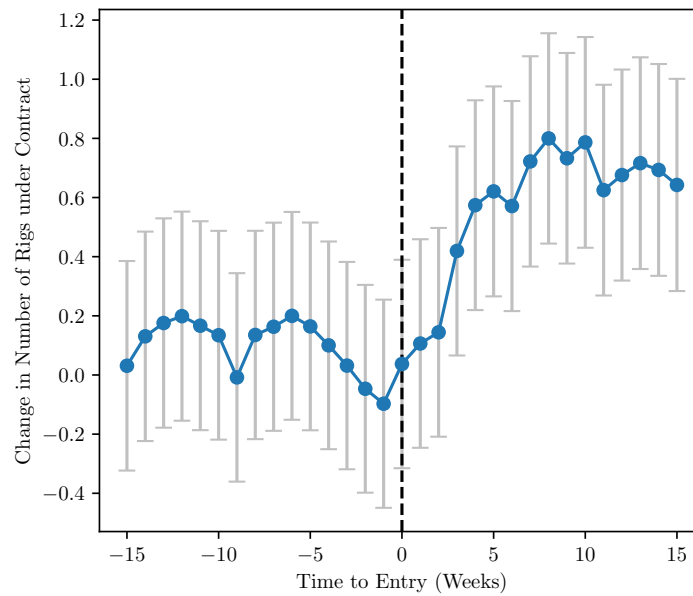
#### A-4 Additional tables and figures

**Table A-1:** Monte Carlo simulation for  $\lambda_l, f_{l,x}$

Variable	True value	Estimates		
		Mean	Std. dev	Coef. of Var.
$\mu_l$	0.6295	0.6286	0.0077	0.0122
$\sigma_l$	0.8901	0.8924	0.0142	0.0159
$\lambda_l$	6.0651	6.0561	0.2728	0.0450

*Note:* This table presents a Monte Carlo simulation to illustrate that the parameters underlying  $\lambda_l$  and  $f_{l,x}$  (which are  $\mu_l, \sigma_l, \lambda_l$ ) are identified. As explained further in the main paper, these parameters are identified by simulating demand in each location, so for simplicity this Monte Carlo simulation focuses on the US market. I start by initializing the three estimated parameter values in the US market as the “true values” of the parameters. Then I simulate 500 “datasets” at these parameters. Each of these “datasets” has 120 months of simulated data for the US location, which is of a similar time span and frequency to the empirical dataset. I compute moments from each of these simulated datasets. For each Monte Carlo simulation, I then optimize over  $\mu_l, \sigma_l, \lambda_l$  in order to match these simulated moments using the estimation procedure in the main paper for the demand-side of the market in the US, keeping other parameters at their empirical baselines (since the objective of this exercise is to highlight the identification of  $\mu_l, \sigma_l, \lambda_l$ ). As is apparent from the table, the estimates for these simulated values closely match the “true” values. Furthermore, the magnitudes of the standard deviation and coefficient of variation across the replications are relatively low. I provide further intuition about identification of the parameters in the main paper.

**Figure A-1: Effect of Rig Entry on Change in Number of Wells Being Drilled in a Location**



*Note:* The figure shows how the number of matches within a location changes after a rig enters. Each blue dot is from a regression coefficient of the number of matches in a location given the number of weeks before/after the rig enters. The regression also includes controls for a cubic time trend and region fixed effects. Overall, after a rig enters a location (at time = 0), there is a period of transition where the rig searches for a match, and then the change in the number of additional matches in a location flattens to around 0.7.

**Table A-2: Distribution of rigs across space in the first vs last half of data**

Region	First half (% total)	Second half (% total)
Africa	15.2	15.9
Asia	14.0	13.2
Australia	5.2	3.7
Central America	4.0	4.3
Europe	18.9	16.7
Mid East	4.2	4.2
S. America	22.5	23.8
US	16.1	18.2


*Note:* I break up the sample by time into two halves. Each column shows the distribution of rigs across space within a time period, with each number corresponding to the percentage of rigs that are in the location, on average. Overall, the table shows that the distribution of rigs across space is very stable across time: a regression of the share in the first half on the share in the second half returns an  $R^2$  of 0.97.

**Table A-3:** Fit of the moments: simulated vs data

	Simulated	Data		Simulated	Data
<b>US</b>			<b>Australia</b>		
Av. Price: Low	0.304	0.309	Av. Price: Low	0.346	0.345
Av. Price: High	0.377	0.401	Av. Price: High	0.429	0.429
Utilization	0.802	0.805	Utilization	0.915	0.917
Av. complexity: Low	1.002	1.031	<b>Central America</b>		
Av. complexity: Mid	2.379	2.363	Av. Price: Low	0.270	0.268
Av. complexity: High	3.588	3.583	Av. Price: High	0.465	0.497
$\beta_{0,low}$	0.313	0.282	Utilization	0.894	0.894
$\beta_1$	-0.018	-0.001	<b>Europe</b>		
$\beta_2$	0.019	0.018	Av. Price: Low	0.349	0.340
<b>Africa</b>			Av. Price: High	0.473	0.475
Av. Price: Low	0.315	0.302	Utilization	0.904	0.903
Av. Price: High	0.460	0.462	<b>Mid East</b>		
Utilization	0.829	0.834	Av. Price: Low	0.311	0.314
<b>Asia</b>			Av. Price: High	0.400	0.398
Av. Price: Low	0.243	0.242	Utilization	0.870	0.874
Av. Price: High	0.421	0.428	<b>South America</b>		
Utilization	0.820	0.819	Av. Price: Low	0.255	0.246
			Av. Price: High	0.415	0.452
			Utilization	0.921	0.922

Note: This table shows the fit of the simulated moments vs the empirical moments in the data.

**Figure A-2:** Fit to untargeted moments




Note: This figure shows the fit of the simulated untargeted moments vs the empirical untargeted moments in the data. Each dot represents the average price for a mid-specification rig in a location (with the exception of the US, which is a targeted moment).

**Table A-4:** Estimation results:  $c_{l,y}$  detail

	Drilling cost $c_{l,y}$		
	Low-type	Mid-type	High-type
	$c_{l,low}$	$c_{l,mid}$	$c_{l,high}$
Africa	0.099 (0.0885, 0.1102)	0.154 (0.1466, 0.1613)	0.216 (0.2110, 0.2202)
Asia	0.036 (0.0275, 0.0438)	0.094 (0.0858, 0.1041)	0.196 (0.1912, 0.2038)
Australia	0.161 (0.1506, 0.1700)	0.142 (0.1273, 0.1529)	0.204 (0.1968, 0.2145)
Central Am.	0.079 (0.0688, 0.0893)	0.197 (0.1860, 0.2101)	0.262 (0.2553, 0.2748)
Europe	0.147 (0.1309, 0.1573)	0.210 (0.1981, 0.2184)	0.241 (0.2327, 0.2592)
Mid East	0.127 (0.1153, 0.1370)	0.115 (0.1076, 0.1245)	0.167 (0.1609, 0.1738)
South Am.	0.063 (0.0497, 0.0752)	0.056 (0.0381, 0.0672)	0.209 (0.2028, 0.2135)
US	0.113 (0.1130, 0.1130)	0.137 (0.1370, 0.1370)	0.147 (0.1466, 0.1466)

Note: Confidence intervals at 95% using 200 bootstrap replications in brackets.

**Figure A-3:** Heterogeneous effects of regulation: detail

Note: Gray bar: change for high-specification rig. Light-blue bar: change for mid-specification rig. Dark-blue bar: change for low-specification rig. Overall, this figure shows detail about the change in the carbon emissions per match, and also the number of matches, for all the counterfactuals.

## A-5 Robustness tests

**Table A-5: Robustness checks**

<b>(a) Reducing <math>P_{exit}</math></b>			
	US-only	Coalition	Global
Leakage of pollution (%)			
Baseline	-36.6	-17.0	-10.8
$P_{exit}$ -20%	-38.6	-15.5	-5.1
Leakage of profit (%)			
Baseline	-46.1	-21.5	-23.3
$P_{exit}$ -20%	-43.7	-16.5	-11.3
<b>(b) Changing <math>c_{entry}</math></b>			
	US-only	Coalition	Global
Leakage of pollution (%)			
Baseline	-36.6	-17.0	-10.8
$c_{entry}$ +20%	-36.4	-17.2	-10.8
$c_{entry}$ -20%	-33.0	-16.7	-10.7
Leakage of profit (%)			
Baseline	-46.1	-21.5	-23.3
$c_{entry}$ +20%	-41.5	-21.2	-23.0
$c_{entry}$ -20%	-39.0	-20.1	-22.6
<b>(c) Changing <math>c_d</math></b>			
	US-only	Coalition	Global
Leakage of pollution (%)			
Baseline	-36.6	-17.0	-10.8
$c_d$ +50%	-45.1	-16.8	-11.0
$c_d$ -50%	-24.4	-14.1	-9.1
Leakage of profit (%)			
Baseline	-46.1	-21.5	-23.3
$c_d$ +50%	-65.7	-15.6	-23.2
$c_d$ -50%	-24.3	-18.3	-19.7
<b>(d) Changing <math>\eta</math></b>			
	US-only	Coalition	Global
Leakage of pollution (%)			
Baseline	-36.6	-17.0	-10.8
$\eta$ +20%	-34.2	-16.0	-8.9
$\eta$ -20%	-51.6	-15.2	-11.6
Leakage of profit (%)			
Baseline	-46.1	-21.5	-23.3
$\eta$ +20%	-29.1	-10.5	-17.2
$\eta$ -20%	-65.3	-21.7	-25.1