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

Nicholas Vreugdenhil *

August 5, 2024

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. 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 and Lucas Davis, for helpful comments. Thank you also to Ryan Kellogg, Giulia Brancaccio, Stephen Ryan, and Adam Spencer, for discussing the paper.

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.[1] 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 particular 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 who 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.[2]

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

[1]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 analysed 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/`)

[2]For the theoretical literature see Baylis, Fullerton and Karney (2013). Leakage through a produce 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).

Figure 1: Capital reallocation channels

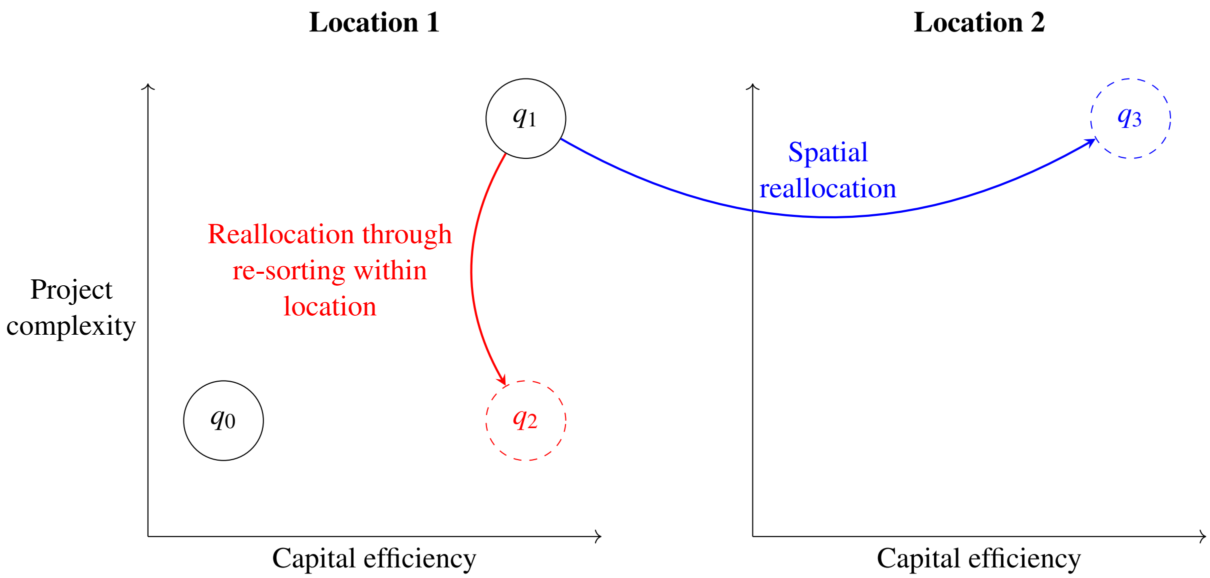

Note: This figure illustrates two channels for how capital reallocation generates leakage in the presence of two-sided vertical heterogeneity. Here, there are two pieces of physical capital (in this paper, these are drilling rigs): one high-efficiency and one low-efficiency. Initially there are two matches in location 1: the high-efficiency capital is allocated to a complex project q_1 and the low-efficiency capital is allocated to a simple project q_0. Consider the effects of a ban on drilling complex projects in location 1. The direct effect is to eliminate the match q_1. However, capital may reallocate to q_2, where it is less well-matched within location 1. Or, it may reallocate across space to a complex project in location 2.

do so I exploit an unusually detailed dataset of the universe of contracts and capital movements in the global market for deepwater drilling rigs. Such data are notable because key information (such as contracts between firms, prices, and allocations) on firm-to-firm markets are typically confidential, which has made studying capital reallocation difficult more generally.[3]

I apply the data to a new empirical framework to study capital reallocation within and across space. The framework 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 illustrate two main channels that result from the interaction of spatial reallocation and sorting in Figure 1. The framework also allows for extensive margin effects i.e. entry of projects and exit of rigs from the global market. The results illustrate the central role of capital reallocation in how proposed incomplete supply-side policies affect profits and carbon emissions, and the

[3]See e.g. Collard-Wexler and De Loecker (2015), Vreugdenhil (2023) who make similar points in alternative contexts.

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.[4] Furthermore, different locations contain different types of projects.

I begin with an analysis of the 2010 US offshore drilling moratorium, which was active in the months after the Deepwater Horizon/BP oil spill in the US Gulf of Mexico. The 2010 moratorium was a particularly stark example of incomplete regulation: while drilling was temporarily halted in the US market, it was allowed to continue in other locations around the world. The data show that rig owners responded to this difference in regulation by temporarily relocating out of the US market to other locations where they could continue to drill.

Motivated by this illustrative example, 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 2008-2016. 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 —

[4]Note that in the paper I use the terms 'capital' and 'rig' interchangeably. Similarly, I use the terms 'project' and 'well' interchangeably.

target the type of rig that best matches with their well type.[5]

In 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.

I use the model to test several counterfactual policies. These center around proposed 'supply-side' policies that ban complex wells. This counterfactual corresponds directly to drilling bans in the industry like the 2010 US moratorium, which target deepwater wells that are usually the most complex to drill but also tend produce the most hydrocarbons. It also serves as a good proxy for proposed bans which correspond to sales of *new* offshore leases such as those implemented by the Biden Administration.[6]

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 -34.8 percent. Decomposing this into two leakage channels, for every ton of carbon dioxide saved through banning complex wells, within-location capital reallocation generates 0.16 more tons, while an additional 0.19 tons are generated through reallocation across space to unregulated locations.[7]

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

---

[5]In equilibrium, the two-sided vertical heterogeneity in this market results in complex wells targeting high-efficiency rigs and simple wells targeting low-efficiency rigs.

[6]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 Freudenburg (2012).

[7]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 predominately 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.

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 improve substantially the efficacy of the regulation: overall, capital reallocation limits the reduction in emissions by -15.5 percent.

Finally, all these options would still be costly due to capital misallocation: incomplete regulation generates misallocation across space, but complete regulation generates misallocation primarily through its effect on sorting within locations. For example, high-specification rigs who were previously matched to complex wells re-target and match to wells with less complementarities.

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. Frechette, Lizzeri and Salz (2019), Buchholz (2022), Brancaccio, Kalouptsidi and Papageorgiou (2020)).

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 arrive at a loss of efficiency from capital reallocation at -70.5 percent in the US-only counterfactual, around double the "true" value. Theoretically, as I detail further in the counterfactuals, removing two-sided vertical heterogeneity could cause the model to under- or over-predict leakage. However, a dominant effect is that a "homogeneous agents" model would eliminate match quality considerations. As a result rigs and wells would be more substitutable within and across locations, which would then incorrectly amplify leakage.

The second contribution is the analysis of a detailed dataset of firm-to-firm contracts, movements, and projects, in a global capital market. As previously mentioned, 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 is thought to be a key method of capital reallocation more broadly (Ramey and Shapiro (2021)). The data provide a granular picture of the inner-workings 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. Most notably, these papers seek to test the 'pollution haven hypothesis' which is that stringent regulation in developed countries like the US has caused industries to relocate to less regulated developing countries (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.[8] Davis and Kahn (2010) finds 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, as well as the papers already mentioned, Kellogg (2011), Asker, Collard-Wexler and De Loecker (2019), Lewis (2019). 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 (2023) 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

[8]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.

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 (2023) which uses Jackup rigs which extend their legs to the seabed.

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 manoeuvred onto a special ship and this ship then transports the rig. Figure 2 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.[9] 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

[9] 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.

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 comes from a proprietary dataset from Rigzone (an industry data provider). The full dataset consists of the status of marketed drilling rigs worldwide 2000-2016. I cut the data to only deepwater rigs (defined as those with a maximum drilling 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 who owns the well. Contracts are almost always fixed price per day for a given duration and rarely contain performance incentives.

The data sample covers the years 2008-2016 and Table 1 provides summary statistics. I choose this period because dayrates are stable over the sample period (as documented in Online Appendix Figure A-2). Similarly, the number of contracts drilled each year are relatively stable.[10] Motivated by this fact I model the market as in a steady-state equilibrium for this sample period.

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

[10]For example, towards the start of sample in the oil price bust in 2009 the number of contracts was 161. When the oil price returned to a boom in 2011, the total number of deepwater drilling contracts was 186.

Figure 2: 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	1241	0.35	0.15
Contract Duration	Days	1241	171	158
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	0.3	0.34	0.43
Average Match: Well Complexity	1.0	2.4	3.6

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 B.

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', 'medium' 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 depth and bottomhole pressure, and ranks wells on a one-dimensional index of drilling complexity. I detail steps taken build this metric in Appendix B.3.3.

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

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.

3 Descriptive analysis of the deepwater rig market

To motivate the model I make two observations of the raw data.

Figure 3: Effects of the 2010 moratorium on the market for deepwater drilling rigs

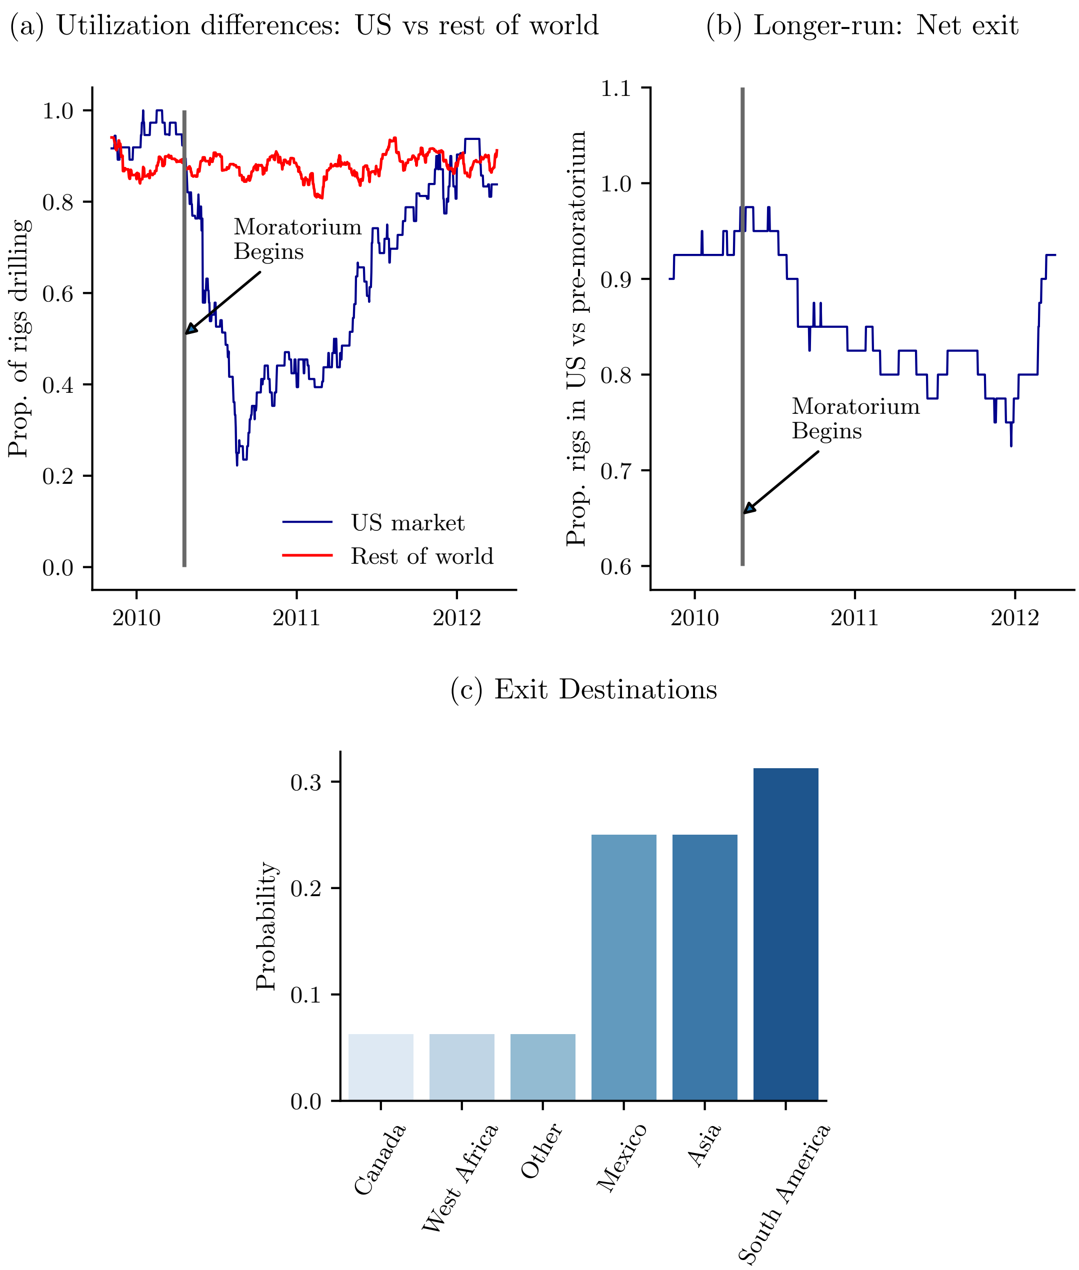

Note: For part (a) and (b) I indicate when the moratorium started. However, I do not plot where the moratorium ended since the exact date is hard to determine: although the moratorium officially ended in October 2010, a 'defacto' moratorium persisted where no permits were awarded for new wells until February 2011. The permit approval slowdown ended around mid-2011. Part (c) shows the destinations of the rigs that exited after the 2010 moratorium.

Observation 1: Rigs respond to differences in regulation by changing location.

To make this observation, I analyze the effects of the 2010 offshore drilling moratorium as a case study.[11] On April 20 the Macondo prospect that the Deepwater Horizon oil rig was drilling blew out, discharging oil into the Gulf of Mexico in the largest oil spill in US history.[12] On April 29 the Obama Administration announced it would issue no new drilling permits until an investigation was completed and I date the start of the moratorium from this date. Later, a continuation of this moratorium was introduced in May 30 2010. Although the moratorium officially ended in October 2010 a 'defacto moratorium' persisted until at least February 2011 with no new drilling permits awarded (Broder and Krauss (2011)).

Figure 3(a) plots the short-run effects of the moratorium focusing on rig utilization (the proportion of rigs that are actively drilling). The moratorium had a dramatic effect, causing utilization to fall from around 95% to 20% *only in the US market* (it is difficult to safely stop all drilling and so some rigs continued to drill). However, as is documented by the red line, utilization in other locations remained relatively high, providing strong incentives for capital to relocate. Over time, through rig exit (as well as a slow return of permitting), utilization climbed to its pre-moratorium level.

Figure 3(b) plots the cumulative change in the number of drilling rigs in the Gulf of Mexico. After the moratorium is implemented rigs quickly exit for other oil fields not under a moratorium where they will be more fully utilized, and I document the exact locations in Figure 3(c) . When the moratorium is lifted, rigs reenter the region. Overall, this shows that rigs are responsive to differences in regulation across markets.[13]

---

[11] Note that my main analysis incorporates information of *all* movements of rigs worldwide over a longer period of time to test the effects of regulation.

[12] Deepwater Horizon was owned by Transocean and was drilling a well for BP.

[13] Market analysts also predicted rig relocation as a consequence of the moratorium. For example, in May 2010: *"[The rigs] cost 500,000 to 1 million a day to lease, says Michael King of FMC Technologies in Houston. He presumes many of their owners will break their contracts and ship them to places with ongoing demand. "There are oil fields off West Africa, off Brazil and in the North Sea," he said. "That might be the most efficient use of a rig over the next six months."* (Ludden (2010)).

Figure 4: Sorting patterns for deepwater rig efficiency vs project complexity

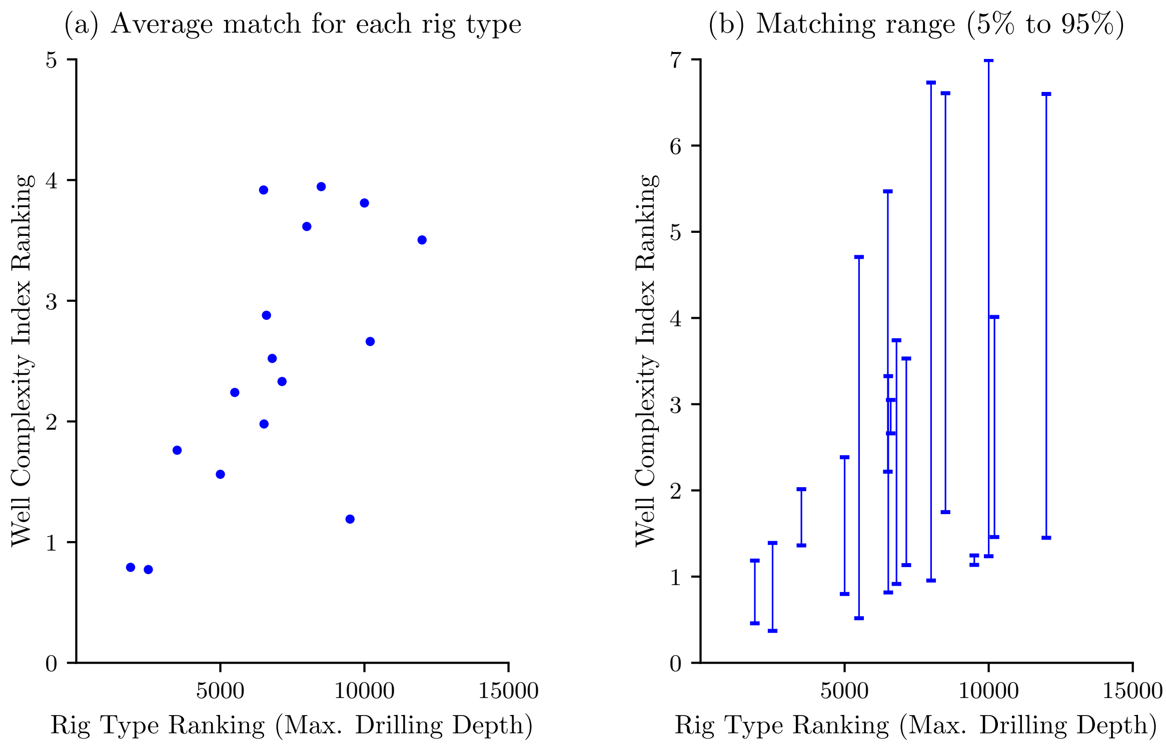

Note: For both figures, the x-axis is the rig efficiency ranking (where rig efficiency is proxied for by the maximum drilling depth) and the y-axis is the project complexity ranking (the 'mechanical risk index' which is an engineering model used in the industry that maps well covariates into a one-dimensional 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. Figure (a) presents positive sorting patterns in terms of the average match for each rig type. Figure (b) presents positive sorting patterns in terms of the entire distribution of projects that rigs match to.

Observation 2: Positive sorting patterns suggest that match complementarities matter.

Figure 4 illustrates the sorting patterns between capital (rigs) and projects (wells) in the US market. Recall that we can rank wells vertically by their complexity using an engineering model called the 'mechanical risk index' and we can also vertically rank rigs by their efficiency (proxied by their maximum drilling depth). Figure 4 illustrates that more complex projects tend to match with more efficient rigs, both on average and over the entire distribution of project types.[14] 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.

Different markets have different distributions of well complexity, which make them suitable for relocating different types of rigs. For example, the European market (the North Sea) is known to have relatively simple projects. Therefore, in the empirical distribution of rigs across the world, this region has proportionally more lower-efficiency rigs.

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 β.

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

[14] These sorting patterns are also apparent for shallow-water rigs Vreugdenhil (2023).

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'.[15] 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 λ_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.[16] Otherwise, a match is formed.

4. Production If a potential project successfully matches with capital the τ 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 τ 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

[15]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 it reduces 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.

[16]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 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.

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 to 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.[17]

As well, I include an extensive margin response for the total number of rigs in the global market.

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. 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.

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

[17] 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.

queueing 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:

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

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 ε_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 ε_y is drawn independently for each searching project (as well as for each rig type y). A potential project contacts the capital type that offers it the highest expected value: $\max_y \left\{\Pi_{l,x,y}^{project}\right\}$.[18]

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

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

$$s_{l,y} = \int 1\left[y = \underset{k}{\text{argmax}} \left\{\Pi_{l,x,k}^{project}\right\}\right] 1[\underset{k}{\text{argmax}} \left\{\Pi_{l,x,k}^{project}\right\} \geq c_{entry}] f_{l,x} dx \tag{2}$$

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 probability of matching for projects and capital $q_{l,y}^{project}$, $q_{l,y}^{capital}$ that results from the above targeting decision using a matching simulation. I briefly discuss this simulation here and leave a more detailed description to Appendix D.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 contract duration τ to complete the match. For each rig type queue, there are $n_{l,y}$ rigs that projects can be completed at, so there are $n_{l,y} t_{backlog}/\tau$ places in the queue. The 'queuing discipline' is first-in-first-out.

[18]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 and so I choose to not incorporate this more complicated feature.

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'. Mathematically, this choice is:

$$U_{l,y} = \max\left\{ \max_{l' \neq l} \{ \underbrace{-c_d d_{l,l'} + \beta V_{l',y} + \sigma_\varepsilon \varepsilon_{l'}}_{\text{Value of moving to } l'} \}, \ \underbrace{b_{stay} + \beta V_{l,y} + \sigma_\varepsilon \varepsilon_l}_{\text{Value to staying in location } l} \right\} \tag{3}$$

Here the first term is the value of moving from location l to l', where c_d is the per-mile transport cost, $d_{l,l'}$ is the distance, $\varepsilon_{l'}$ is the idiosyncratic logit error, and σ_ε is the scale parameter of the errors. Although I am suppressing rig-specific subscripts, the logit draws are drawn independently for each individual rig as well as 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. Equation (3) delivers 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 D.1.

Using the location decision in Equation (3) I can write the ex-ante value function for unemployed capital (that is, the value function before the ε_l shocks are drawn):

$$U_{l,y} = \sigma_\varepsilon \log\left(\sum_{l' \neq l} \exp\left(\frac{-c_d d_{l,l'} + \beta V_{l',y}}{\sigma_\varepsilon} \right) + \exp\left(\frac{b_{stay} + \beta V_{l,y}}{\sigma_\varepsilon} \right) \right) + \sigma_\varepsilon \gamma^{euler} \tag{4}$$

where γ^{euler} is Euler's constant.

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

$$V_{l,y} = q^{capital}_{l,y} \underbrace{\left(\sum_{s=0}^{\tau-1} \beta^s \delta_{l,y} + \sigma_\varepsilon \gamma^{euler} + \beta^\tau V_{l,y} \right)}_{\text{Expected value to matching for the rig}} + (1 - q^{capital}_{l,y}) U_{l,y} \tag{5}$$

Here, $q^{capital}_{l,y}$ 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 $\delta_{l,y} =$

$\bar{p}_{l,y} - c_{l,y} + \xi_{l,y}$, where $\xi_{l,y}$ accounts for unobserved cost shocks.[19] 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 are 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):

$$U_{l,y} \geq b_{scrap} \tag{6}$$

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:

$$\underset{p}{\operatorname{argmax}} \left[\sum_{s=0}^{\tau-1} \beta^s [m_{x,y} - p] - \beta(1 - P_{exit}) W_{l,x,y} \right]^{1-\eta} \left[\sum_{s=0}^{\tau-1} \beta^s \delta_{l,y} + \beta^\tau V_{l,y} - U_{l,y} \right]^{\eta} \tag{7}$$

[19] This formulation is useful later in estimation when I parameterize $c_{l,y}$. 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 σ_ε, but as I later document this value is relatively small.

Recall that prices are embedded in $\delta_{l,y} = p - c_{l,y} + \xi_{l,y}$. Here τ 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, η 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:

$$W_{l,x,y} = q_{l,y}^{project} \sum_{s=0}^{\tau-1} \beta^s (m_{x,y} - p_{l,x,y}) \tag{8}$$

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 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. 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 predominately produce oil, and furthermore burning oil produces a far greater magnitude of emissions than burning the equivalent energy unit of natural gas.[20]

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. Then, given the equilibrium number and types of matches predicted by the model in each location, the model

[20]See, e.g. https://www.eia.gov/energyexplained/natural-gas/natural-gas-and-the-environment.php.

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, I first convert the change in supply from the deepwater market to an equilibrium global change in oil produced and consumed, incorporating demand responses 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 Gases Equivalencies Calculator.

4.7 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. In Appendix F I provide an example of the role of two-sided vertical heterogeneity in determining the within-location equilibrium response to across-location entry and exit of capital.

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 satisfies Equations (1) - (2) and the queuing model detailed in Appendix D.2.

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) - (5).

3. Extensive margin for the supply side governed by Equation (6).

4. Prices $p_{l,x,y}$ determined by Nash bargaining, defined in Equations (7) and (8).

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.[21] 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 queueing 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.)

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 the similar tow boats, and the cost of towing is proportional to the distance, I convert the distance between fields by the tow speed and calibrate the per-day cost of towing as $c_d = \$0.25$ million.[22]

I need to calibrate the bargaining parameter and I assume that the parties split the match surplus

[21] In my dataset the US market is the only market where backlog data are systematically available.

[22] 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, Hellingaand 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.

Table 3: Overview of how the parameters are computed

Object	Notation	Parameterization	Param.	Method
Calibrated params.	$\beta, \tau, t_{max}, \eta$			Calibrated
	$P_{exit}, c_d, c_{entry}, b_{scrap}$			
Costs	$c_{l,y}$	$c_{l,y} = \gamma \bar{p}_{l,y}$	γ	Estimate step 1
Remain in loc.	b_{stay}		b_{stay}	Estimate step 1
Preference shock	σ_ε	Logit	σ_ε	Estimate step 1
Demand distribution	$f_{l,x}$	Log-normal	μ_l, σ_l	Estimate step 2
Demand draws	D_l	Poisson dist.	λ_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.

equally and set this to $\eta = 0.5$.[23] 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 $P_{exit} = 0.5$.[24]

I calibrate the potential project entry cost c_{entry} to \$13.15 million USD using Hossain (2015).[25] 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).

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 μ_l and standard deviation σ_l.

[23]This is somewhat close to the $\delta = 0.37$ used in the shallow water analysis in Vreugdenhil (2023).

[24]Brancaccio, Kalouptsidi and Papageorgiou (2020) also need to calibrate a similar value for their 'exporter survival rate'.

[25]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 (2023) follows a similar procedure using numbers in the shallow water market.

Estimating a separate cost for each capital type in each location would require estimating 18 different costs. However, given there are several markets where the number of relocations for a given capital type are small, I choose to parameterize costs in the following way: $c_{l,y} = \gamma \bar{p}_{l,y}$. Here, the term $\gamma \in [0,1]$ is a scale parameter that relates costs to the average price in each location. Note that $1-\gamma$ is the capital's markup and so a low γ corresponds to a high markup and a high γ corresponds to a low markup. One objection to this parameterization is that it may be inconsistent with Nash bargaining, because prices also incorporate outside options which may vary with demand in each location or by rig type. To address this concern I resimulate the entire equilibrium of the baseline model at the estimated demand and supply parameters. I then regress model-predicted average prices by rig-location on a constant and the cost parameters. The resulting R^2 =0.92 which illustrates that most of the variation in prices is driven by differing costs, and not outside options.[26]

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 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.

[26]This may be unsurprising when considering that rigs are incentivized to move to locations with high markups. This then tends to make outside options across locations more equal in equilibrium. For example, if demand was very high in one location such that outside options were also high, rigs would enter this location, endogenously reducing the outside option.

5.2 Step 1: Computing supply side parameters

This section is similar in spirit to the estimation strategy in Brancaccio, Kalouptsidi and Papageorgiou (2020) except that I depart from it to allow for unobserved cost shocks. To do so I split estimation into two sub-steps. In Substep 1.1 I recover $\delta_{l,y}$, σ_ε, and b_{stay} using the observed choice probabilities of moving between locations. In Substep 1.2 I use an instrumental variables strategy to compute the markup γ — which can be used to back out $c_{l,y}$ — in the presence of unobserved cost shocks which may generate price endogeneity.

Substep 1.1: I estimate the parameters in this substep by fitting the empirical location choice probabilities for each rig type using maximum likelihood. I provide more information about how I compute the value functions and the exact algorithm for estimation in Appendix D.1. I provide a more formal proof of identification in Appendix E.1.

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 the 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 σ_ε. Intuitively, the model matches these choices 'on average', with higher σ_ε corresponding to choice probabilities that generate a more 'spread out' stationary distribution of rigs in each location. Lower values of σ_ε yield location choices where rigs predominately choose locations with the highest markups (as well as the highest probability of matching).

Substep 1.2: Using the parameterization of costs in terms of a markup over the average price $c_{l,y} = \gamma \bar{p}_{l,y}$, I then use the values of $\delta_{l,y}$ from Substep 1.1 to estimate γ using the equation $\delta_{l,y} = (1-\gamma)\bar{p}_{l,y} + \xi_{l,y}$. Since prices may be endogeneous and a function of the unobserved cost shocks $\xi_{l,y}$ I use an instrumental variables strategy.

Specifically, I instrument prices with a demand shifter: unexploited oil and gas reserves by location and water depth.[27] The intuition for why this instrument is independent of the rig cost shocks ξ_{jt} stems from industry norms around the contractual division of responsibility between oil companies and rig owners detailed in Section 2.[28]

As a result, even though the volume of hydrocarbons may be correlated with underlying geological conditions that make the well more costly to drill, this is *not* a cost that would affect ξ_{jt}. Rather, the cost shocks in ξ_{jt} may come from, for example, local labor market conditions that make workers more costly to hire; these are unlikely to be correlated with the geology of deepwater oil and gas fields. For similar reasons, this instrument also satisfies the exclusion restriction (i.e. does not directly enter the utility function of the rig): rig owners do not directly benefit from selling the hydrocarbons the well produces.

5.3 Step 2: Computing the match value and demand

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 demand from the price/contract data alone (recall that 'demand' is the underlying distribution of wells). 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 spilt Step 2 into two substeps: I first retrieve the parameters that underlie the match value function (as well as demand) 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 in the other markets using simulated method of moments.

[27] I use data from the 'Global Oil and Gas Extraction Tracker' from the Global Energy Monitor: https://globalenergymonitor.org/projects/global-oil-gas-extraction-tracker/. Within each location, I convert gas reserves to 'barrels of oil equivalent'. I then split up the total reserves into three quantiles of water depth and compute the share of unexploited reserves below the maximum of each water depth quantile.

[28] Specifically, rig owners are responsible for rig operating expenses like wages for the crew; oil companies pay for well-related drilling expenses but own the hydrocarbons the well produces.

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:

$$\text{price}_i = \beta_{0,y} + \beta_1 \text{complexity}_i + \beta_2 \cdot (\text{max drilling depth})_i \cdot \text{complexity}_i + \varepsilon_i \tag{9}$$

where $\beta_{y,0}$ is a capital-specific fixed effect. Equation (9) 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: $\beta_{low,0}, \beta_1, \beta_2$. 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 (μ_l, σ_l) are identified from observed matches. Intuitively, the entry condition and the matching simulation generate a mapping between the distribution of potential projects in a location and the distribution of observed matches. In practice, I use the following moments to identify (μ_l, σ_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 market 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 λ_l is determined by the average capital utilization in each market (1 moment per market); higher values of λ_l correspond to higher capital utilization.

Computation I provide information about how I compute the demand-side equilibrium in Appendix D.2. Using this algorithm, I first compute the equilibrium in the US which returns demand in the US market and the match-value function. Using this match-value function, I next estimate demand 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.

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 σ_{ε} and the stay put benefit b_{stay} in Table 4(c) are both relatively low. For example, scaling up the σ_{ε} 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 \$59.9 million.

The 'unexploited oil reserves' instrument used in the second substep to estimate γ is a strong instrument, with a first-stage F-statistic of 94. 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 Online Appendix Table A-2. 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 is 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-1; 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-4. The model also closely fits the untargeted moments with a median

Table 4: Estimation results

(a) Location-specific Parameters

	Costs (av. over y) $c_{l,y} = \gamma \bar{p}_{l,y}$	# Entry λ_l	Mean μ_l	Std. dev σ_l
Africa	0.17†	6.58	0.59	1.07
	(0.0015)	(0.40)	(0.27)	(0.22)
Asia	0.14†	5.75	0.52	0.83
	(0.0013)	(0.41)	(0.29)	(0.21)
Australia	0.16†	3.31	0.59	0.92
	(0.0015)	(0.44)	(0.27	(0.24)
Central Am.	0.17†	3.44	0.76	0.96
	(0.0016)	(0.46)	(0.17)	(0.23)
Europe	0.18†	14.16	0.56	1.02
	(0.0017)	(1.62)	(0.22)	(0.26)
Mid East	0.15†	2.83	0.55	0.82
	(0.0014)	(0.59)	(0.50)	(0.32)
South Am.	0.12†	11.81	0.44	0.77
	(0.0011)	(0.95)	(0.15)	(0.15)
US	0.16†	6.29	0.66	0.87
	(0.0014)	(0.57)	(0.09)	(0.07)

(b) Match value parameters

	$m_{0,y}$	$m_{1,y}$
Low-spec	0.660	-0.377
	(0.057)	(0.050)
Mid-spec	0.457	-0.0287
	(0.058)	(0.024)
High-spec	0.376	0.018
	(0.063)	(0.020)

(c) Other parameters

Scale parameter (γ)	0.570
	(0.004)
Preference shock (σ_{ε})	0.11
	(0.02)
Stay put benefit (b_{stay})	0.10
	(0.03)

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 Online Appendix.

difference of only 2.5 percent.

The demand parameter results are in the second, third, and fourth columns of Table 4(a). The estimates reveal substantial differences in demand across the world. For example, the mean project complexity terms μ_l are consistent with the utilization of different types of capital in different fields. For instance, Europe is a primary markets for low-specification rigs, and contains simpler projects. By contrast, the US 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 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 in 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). Overall, the match value estimates indicate that low-specification rigs have an advantage in drilling simple projects. Conversely, high-specification rigs have an advantage in drilling more complex wells.

7 Counterfactuals

I investigate counterfactual scenarios that centre 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.[29] As well, drilling bans in the industry like the 2010 US moratorium primarily targeted complex wells (e.g. while a ban on drilling shallow water/simpler wells was lifted after a few days, the broader moratorium lasted for months) .

I evaluate the efficiency 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 coordination 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).

The counterfactual results are reported in terms of percentage changes. The benefit (discussed further in Appendix A) is that converting from deepwater market production to emissions essentially involves scaling by a multiplicative factor. Percentage changes are scale-free and therefore robust to the choice of this multiplicative factor.

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 D.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 regulations decrease carbon dioxide emissions by reducing oil production in the US by -36.7% but also reduce profits by -17.2%.

[29]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.

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.
US-only ban						
CO_2	0.16	0.19	-36.7	-3.6	-34.8	-70.5
Profits	0.37	0.08	-17.2	-1.7	-44.5	-70.5
Coalition ban						
CO_2	-0.02	0.18	-36.2	-18.3	-15.5	-14.5
Profits	0.03	0.17	-16.0	-8.7	-20.3	-14.5
Global ban						
CO_2	0.11	0.00	-33.8	-33.8	-10.8	-30.6
Profits	0.23	0.00	-13.6	-13.6	-23.5	-30.6

(b) US-only ban: statistics

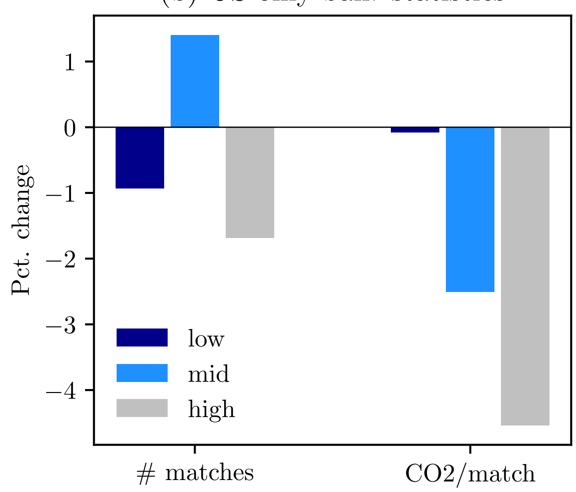

Note: Part (a): Leakage is defined as the increase in emissions (or profits) generated through unregulated economic activity, for a one unit decrease in emissions (or profits) in the regulated activity. I decompose this into: (i) leakage due to capital reallocation from re-sorting within regulated locations, 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 'Δ Total 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 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-5 has detail on this figure for the other counterfactuals.

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 -34.8 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, 0.16 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, 0.19 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 are less complementarities, 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 exit, the remaining mid-specification rigs actually have a higher number of matches due to within-location reallocation. Low-specification rigs occasionally drill more complex wells and the net effect here is that they have a lower probability of matching.

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-matched here, there is an environmental benefit: these wells tend to produce less oil and therefore less 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. Indeed, I find that no rigs would exit in any of the counterfactuals, (and this would persist even if the scrap value parameter was doubled). 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 leakage within-location to 0.11 tons of carbon dioxide produced elsewhere for every ton saved directly from the ban.

Globally, the change in emissions is -33.8 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 -15.5 percent (which is similar to the number in the global agreement counterfactual). Spatial leakage is still relatively high at 0.19: although more regions are under regulation, the coalition does not encompass all locations with complex wells.

Within-location leakage in terms of emission is slightly negative — implying that the indirect effect of the ban is to reduce emissions further — and this results from the interaction of a number of mechanisms. Concretely, when high-specification rigs leave the regulated area (which happens disproportionally with a complex-well ban), they generate a lower probability of matching for complex wells just below the cutoff of the ban. This then causes these wells to enter with a lower probability. As a result, the remaining rigs in the location match with wells that produce less oil and resulting emissions.

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 disproportionally exit for 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 was 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 D_l. (One side-effect of this simpler version of the model is that the percentage changes in CO_2 and profits are the same because the total effects are driven solely by the number of homogeneous matches in each location.)

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 US-only and global ban counterfactuals. For example, in the US-only counterfactual, reallocation would be predicted to undercut the change in emissions by -70.5 percent, around double the value in the full model. These results are consistent with the simpler model failing to capture the imperfect rig and well substitutability (and the resulting effects on rig location choice) of the full model. However, capturing heterogeneous rig relocations — which can amplify leakage — is important too. This is apparent in the coalition ban counterfactual, where the two theoretical effects approximately cancel each other out and leakage would be slightly under-predicted by the 'homogeneous agents' model.

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 efficiency. 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 be a more politically feasible coalition of rich countries implementing coordinated regulation. Overall, 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.

References

Abito, Jose Miguel, Christopher R. Knittel, Konstantinos Metaxoglou, and André Trindade, "The Role of Output Reallocation and Investment in Coordinating Externality Markets," *International Journal of Industrial Organization*, July 2022, *83.*

Ahlvik, Lassi, Jørgen Juel Andersen, Jonas Hvending Hamang, and Torfinn Harding, "Quantifying Supply-side Climate Policies," *CAMP Working Paper Series*, 2022, (1).

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

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 Kalouptsidi, and Theodore Papageorgiou, "Geography, Transportation, and Endogenous Trade Costs," *Econometrica*, March 2020, *88*, 657–691.

Broder, John M. and Clifford Krauss, "Judge Tells Government to Resume Permits for Drilling," *New York Times, February 11*, 2011.

Buchholz, Nicholas, "Spatial Equilibrium, Search Frictions, and Dynamic Efficiency in the Taxi Industry," *The Review of Economic Studies*, March 2022, *89.*

Collard-Wexler, Allan and Jan De Loecker, "Reallocation and Technology: Evidence from the US Steel Industry," *American Economic Review*, 2015, *105* (1), 131–71.

Copeland, Brian R. and M. Scott Taylor, "Trade and the Environment: Theory and Evidence," *Priceton Series in International Economics*, 2003.

Corts, Kenneth S, "Stacking the Deck: Idling and Reactivation of Capacity in Offshore Drilling," *Journal of Economics and 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, pp. 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.

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.

Frechette, Guillaume, Alessandro Lizzeri, and Tobias Salz, "Frictions in a Competitive Regulated Market: Evidence from Taxis," *American Economic Review*, August 2019, *109* (8), 2954–2992.

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, (1).

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 Cost Estimation for Hydrocarbon Wells," *Journal of Sustainable Energy Engineering*, 2015, *3* (1).

Kaiser, Mark 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.

Ludden, Jennifer, "Many Factors Cloud Cost Of Oil-Drilling Moratorium," *NPR*, May 2010.

Magnac, Thierry and David Thesmar, "Identifying Dynamic Discrete Decision Processes," *Econometrica*, 2002, *70*, 801–816.

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 CO2 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 the Association of Environmental and Resource Economists*, July 2023, *120.*

Ramey, Valerie A. and Matthew D. Shapiro, "Displaced Capital: A Study of Aerospace Plant Closings," *Journal of Political Economy*, 2021, *109* (5).

Terpstra, T., E. Hellingaand, 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," `https://nvreug.github.io/paper/bbm.pdf` 2023.