Spatial Spillovers in the Mineral Rights Model: Evidence from the Gulf of Mexico

Arseniy Stolyarov

1 Introduction

The Gulf of Mexico represents one of the largest and most mature offshore petroleum provinces in the United States, yet firms' exploration outcomes exhibit strong spatial dependencies. In an R&D context, property rights over potential discoveries crucially shape the willingness to invest, often giving rise to free-riding concerns when new information gleaned from exploratory wells benefits adjacent tracts. While previous work has documented inefficiencies in bidding and drilling decisions due to such information externalities, few studies have measured the intensity of the spillovers.

This paper develops a joint spatial model of oil and gas discoveries using data from 1945–2024 to better quantify both the probability of a successful strike and the expected quantity extracted. By leveraging R-INLA and the Lindgren (2012) approach, we provide new insights into how leaseholders update their drilling decisions based on adjacent wells' outcomes. We then use these estimates to simulate alternative lease auction rules, highlighting implications for social welfare and regulatory policy.

2 Data

Our primary dataset encompasses all offshore well drilling and production activity in the Gulf of Mexico (GoM) from 1945 to 2024. The well-level records, obtained from the Bureau of Ocean Energy Management (BOEM), provide geographic identifiers, drilling dates, and production quantities (if any). We combine these records with tract-level auction data, including winning bids, lease terms, and royalty rates, to analyze how firms' bidding and drilling decisions evolve over time.

Each tract is associated with coordinates that allow us to measure spatial relationships across wells. We link drilling outcomes—whether a well struck oil/gas or was dry—to the tract's location and historical production data. The final dataset contains over 10.000 unique wells, distributed across more that 4000 federal tracts, enabling us to capture both cross-sectional and temporal variation in exploratory activities.

By structuring the data at the tract-well-year level, we can observe when and where new information becomes available (e.g., neighboring well outcomes) and track how that information influences subsequent drilling decisions.

3 Methodology

3.1 Spatial Model Setup

We analyze offshore exploration using a *joint spatial model* that simultaneously estimates the probability of striking oil or gas (hit probability) and the expected volume extracted (quantity). Let $z_i \in \{0, 1\}$ indicate whether the ith tract yields a successful discovery, and let y_i denote the realized production quantity. We adopt a Bernoulli likelihood for z_i and a Gamma likelihood for y_i. Both outcomes are linked to a latent Gaussian spatial field $X(s)$ via:

$$\operatorname{logit}(p_i) = \alpha_z + X(s_i), \quad z_i \sim \text{Bernoulli}(p_i), \tag{1}$$

$$\log(\mu_i) = \alpha_y + \beta\, X(s_i), \quad y_i \sim \Gamma\big(\cdot\big), \tag{2}$$

where s_i represents the coordinates of tract i.

3.2 Matérn Covariance Function

We specify $X(s)$ to follow a Gaussian random field with a *Matérn* covariance function:

$$\operatorname{Cov}\big(X(s_i), X(s_j)\big) = \sigma_x^2 \, \frac{2^{1-\nu}}{\Gamma(\nu)} \Big(\kappa \, \|s_i - s_j\|\Big)^{\nu} \mathcal{K}_\nu\Big(\kappa \, \|s_i - s_j\|\Big), \tag{3}$$

where σ_x^2 is the variance parameter, κ governs the spatial decay, ν is a smoothness parameter (often set or estimated), and $\mathcal{K}_\nu$ is the modified Bessel function of the second kind. We interpret

$$r = \frac{\sqrt{8\,\nu}}{\kappa}$$

as the approximate "range," beyond which spatial correlation becomes negligible.

3.3 Estimation via R-INLA

We estimate the latent Gaussian model using *R-INLA* (Rue et al. (2009), Simpson et al. (2011)) an integrated nested Laplace approximation method that offers a fast and accurate alternative to traditional Markov Chain Monte Carlo for high-dimensional spatial models. R-INLA approximates the posterior distributions of the parameters $\{\alpha_z, \alpha_y, \beta, \sigma_x^2, \kappa, \dots\}$ and the latent field $X(s)$ without exhaustive simulation. Our model is an extension to Hodgson (2022) binary outcome model and Nguyen (2024) mixed model that only considers productive wells in the later period of Gulf development.

3.4 Dynamic Drilling and Counterfactual Analysis

Once the spatial model is fit, we use its output—specifically, the tract-level probabilities of a successful strike and the expected quantities—to inform a *dynamic structural model* of drilling decisions. Firms update their beliefs about potential payoffs as new information from neighboring tracts becomes available. We then simulate *counterfactual auction rules* (e.g., alternative royalty rates or information restrictions) to assess how these changes might alter exploration timing, speed of region development, reduce duplicated drilling and increase overall social surplus.

4 Results

4.1 Parameter Estimates and Spatial Correlation

Table 1 presents the posterior means and standard deviations for the key parameters in our joint spatial model. We observe an unconditional probability of discovering oil/gas of around 10%, which is consistent with observed drilling outcomes. The range parameter r suggests that spatial correlation declines significantly after approximately 19 nautical miles, indicating that successful (or failed) wells have the most pronounced impact on adjacent tracts within this radius.

Table 1: Spatial Correlation Estimates From The Drilling Outcomes (1959-2024)

	oil	oil	gas	gas
	mean	sd	mean	sd
α_y	8.469	1.743	16.081	0.448
α_z	-1.249	0.515	-1.256	0.582
ϕ	0.364	0.008	0.636	0.015
β	3.329	0.171	0.760	0.050
$r = \frac{\sqrt{8}}{\kappa}$	1.655	1.244	1.338	1.216
σ_x^2	4.185	1.461	3.695	1.317

In both the oil and gas models, the coefficient β is statistically and economically significant, implying that once a tract is revealed to have positive resources, neighboring tracts not only update their probability of discovery but also their expected quantity. This joint estimation approach thus delivers more precise inferences compared to modeling the probability of a hit independently of the volume extracted.

4.2 Impact on Drilling and Bidding Decisions

Our initial regressions indicate that a discovery on a nearby tract raises the probability of subsequent drilling by 34%, while a dry hole reduces it by 30.3%. These results corroborate the notion of strong *information externalities* in the offshore context. Firms appear to use neighboring outcomes as a key input to their drilling decisions, which can generate free-riding behavior and potential inefficiencies when adjacent tracts benefit from privately funded exploration efforts.

4.3 Implications for Lease Auctions

To gauge the real-world impact of these externalities, we incorporate the estimated spatial probabilities into a dynamic drilling decision model. Preliminary simulations suggest that the current first-price auction format may lead to suboptimal lease allocations, as the high-risk tracts with positive spatial signals could be undervalued when drilling information is dispersed among multiple market participants. These findings motivate further analysis of alternative lease rules—such as different royalty rates, joint-bidding restrictions, or partial information disclosure—to potentially increase social welfare.

5 Extending the Analysis: Common-Value Auction Framework

Incorporating observable bids into our model provides an additional way to identify the private signals that firms receive before participating in an auction and correct the priors of the tract value after the release of the auction results. Following (Hendricks, Porter (2003)), oil and gas auctions can be viewed as a common-value setting in which tracts ultimately have the same true value (i.e., the oil and gas beneath them), but individual bidders receive imperfect private signals. Under this framework, signals and valuations can be inferred by observing both the bids and the realized outcomes once a tract is drilled and produced.

However, a key complication arises because not all auctioned tracts are eventually drilled, leaving some true values unobserved. To identify the joint distribution of signals and valuations in these cases, we require a *cost-shifter* that influences the drilling decisions of the auction winners. Several papers employ external cost or price variables to achieve identification. For instance, (Nguyen 2021) uses average rig rental rates, while Bhattacharya, Ordin, Roberts (2022) rely on unexpected changes in the oil prices.

In the Gulf of Mexico, we propose exploiting the evolving *spatial information* derived from neighboring tracts. Specifically, once a firm wins a lease, the updated expectations regarding oil or gas potential on adjacent tracts—inferred via our spatial model—can act as a tract-specific cost or value shifter. Unlike rig rental rates (which can be locked in through long-term contracts) or short-run oil price fluctuations (more relevant in onshore settings with shorter production horizons), the spatial signals in the Gulf of Mexico reflect geological updates and potential long-term returns on investment. Some leases in the Gulf produce for more than 20 years, underscoring the importance of forward-looking behavior that incorporates this spatially generated information.

With the joint distribution of bids and signals in hand, we can then compare alternative institutional arrangements. Once the joint distribution of bids and signals is identified, we can simulate different alternative scenarios. For instance, we can model a "Social Planner" who sequentially draws a signal for each tract and observes the spatial updates as wells are drilled.

This allows us to contrast a fully informed social planner's drilling decisions against those arising under a competitive market mechanism, where firms hold fragmented signals and face strategic bidding incentives. By quantitatively examining these scenarios, we aim to shed light on whether existing offshore lease auctions in the Gulf of Mexico align with socially optimal exploration and production, or whether alternative auction rules could improve welfare.

6 Conclusion

This paper investigates the role of spatial spillovers and private signals in shaping drilling and bidding decisions in the Gulf of Mexico. By estimating a joint spatial model for discovery probability and quantity, we find that firms strongly update their beliefs about nearby tracts based on observed drilling outcomes, highlighting the significance of information externalities. Our initial results suggest that standard first-price lease auctions may fail to allocate exploration rights optimally, as signals of competitive firms can lead to potential free-riding.

Building on this foundation, we plan to incorporate observable bids into a common-value framework, where spatially updated expectations serve as a tract-specific cost or value shifter. This approach allows for more sophisticated dynamic simulations.

In ongoing work, we expand the dynamic drilling model to simulate alternative policy scenarios, such as different royalty rates, partial information disclosure, or changes to joint-bidding rules. These counterfactuals not only clarify the economic trade-offs inherent in offshore leasing but also offer actionable guidance for policymakers seeking to balance private incentives with the efficient R&D policy.

References

1. Bhattacharya, Vivek, Andrey Ordin, and James W. Roberts. "Bidding and drilling under uncertainty: An empirical analysis of contingent payment auctions." Journal of Political Economy 130.5 (2022): 1319-1363.

2. Hendricks, Kenneth, and Robert H. Porter. "The timing and incidence of exploratory drilling on offshore wildcat tracts." The American Economic Review (1996): 388-407.

3. Hendricks, Kenneth, Joris Pinkse, and Robert H. Porter. "Empirical implications of equilibrium bidding in first-price, symmetric, common value auctions." The Review of Economic Studies 70.1 (2003): 115-145.

4. Hendricks, Kenneth, Robert H. Porter, and Charles A. Wilson. "Auctions for oil and gas leases with an informed bidder and a random reservation price." Econometrica: Journal of the Econometric Society (1994): 1415-1444

5. Hodgson, Charles. Information externalities, free riding, and optimal exploration in the uk oil industry. SIEPR, Stanford Institute for Economic Policy Research, 2019.

6. Hohn, Michael. Geostatistics and petroleum geology. Springer Science & Business Media, 1998.

7. Lindgren, Finn. "Continuous domain spatial models in R-INLA." The ISBA Bulletin 19.4 (2012): 14-20

8. Nguyen, Anh. Information Design in Common Value Auction with Moral Hazard: Application to OCS Leasing Auctions. Working Paper, 2024.

9. Simpson, Daniel, Finn Lindgren, and Håvard Rue. "Fast approximate inference with INLA: the past, the present and the future." arXiv preprint arXiv:1105.2982 (2011).