

## Research Statement

My research agenda centers on the economics of energy, natural resources, and the environment. The questions I study address the effects of environmental objectives on competition in government auctions, distortionary contracting practices in the upstream oil and gas industry, methods for analyzing auction data (especially from natural resource sales), and external costs of pollution and the energy industry. Given the varied nature of these topics, I apply a diverse set of tools drawn from empirical industrial organization, mechanism design, and reduced-form econometrics.

### **External impacts of pollution and the energy industry**

In my job market paper, **“Air Pollution as a Cause of Violent Crime: Evidence from Los Angeles and Chicago”** (with Anthony Heyes, Erich Muehlegger, and Soodeh Saberian) my co-authors and I establish the first quasi-experimental evidence that exposure to pollution leads people to commit more violent crimes. A large body of research has quantified a range of adverse effects on human health and, more recently, productivity and cognition. Furthermore, a number of human and animal laboratory experiments suggest that pollution exposure could lead to aggression and loss of impulse control. To test for a causal link from pollution to violent crime, we use microdata on the commission of crimes in Chicago and Los Angeles. A simple approach using panel data would be vulnerable to omitted variables correlated with both pollution and crime, such as local fluctuations in economic activity.

We use distinct but related empirical strategies in the two cities, both of which exploit variation in pollution driven by day-by-day changes in the wind direction. Depending on the direction from which the wind is blowing, neighborhoods experience differential changes in air quality conditions. We compare relative levels of criminal activity in neighborhoods experiencing a wind-driven surge in pollution to those that do not in an approach akin to a difference-in-differences specification. In effect, we are using within-city, within-day variation in pollution exposure driven by the wind direction.

In Chicago, we compare crime on opposite sides of major interstates on days when the wind blows orthogonally to the direction of the interstate. The downwind side is exposed to elevated levels of mobile pollution, while the upwind side acts as a control that captures variation in day-to-day crime, ambient pollution and unobservables, such as neighborhood economic activity. In Los Angeles, ocean winds blowing on-shore push pollution from downtown to the northeast, into the foothills of the San Gabriel mountains, which trap the pollution. We compare criminal activity in areas where a sea-borne wind increases pollution to areas where it does not, on days with and without such a wind, in a difference-in-difference approach.

We find similar and robust results across settings. In both cities, we find that wind-driven changes in pollution increase violent crime (2.2% in Chicago and 6.1% in Los Angeles), but have a small, statistically insignificant effect on property crime. Consistent with a pollution-driven explanation, we find that the effect attenuates at higher wind speeds, and when we widen the interval of wind directions considered to be “treating” an area. In Chicago, there is evidence that pollution causes assaults to escalate to

batteries, which suggests that loss of impulse control is an important mechanism.

These results establish violent criminal activity as a new source of external costs associated with local pollutants. Back-of-the-envelope calculations show that the monetized cost associated with the additional crimes is economically significant and comparable to other outcomes of interest in the empirical environmental cost literature. Our results also imply that pollution has a subtler and more pervasive effect on decisionmaking and behavior than previously realized.

Hydraulic fracturing has immensely affected the scale and relative geography of natural gas supply, leading to a need for new pipeline capacity. These proposed pipeline expansions have often been met by strong local opposition partially motivated by concerns about safety and decreased property values. In **“Pipeline Awareness, Information Provision, and Local Housing Markets”**, Richard Sweeney and I test whether pipeline safety concerns are capitalized into nearby housing prices. Estimating pipeline capitalization is challenging for two reasons. First, pipelines are not randomly located: siting may be correlated with other disamenities that make for an attractive or low-cost route. Second, homeowners and prospective buyers seem to be poorly-informed about the location of pipelines.

To address these issues, we analyze housing prices around a deadly 2010 pipeline explosion in a suburb of San Francisco. First, we test for an effect of the events surrounding the explosion itself, which shocked pipeline awareness and salience. In the weeks following the accident, there were extensive and detailed media coverage and information releases regarding the location of natural gas transmission pipelines. Google search data from the weeks surrounding of the incident indicate a substantial spike in pipeline-related information acquisition. Second, we also examine the effects of a regulator-induced information campaign by PG&E, which targeted existing homeowners within 2000 feet of a transmission pipeline by sending them a disclosure letter.

Using data on housing transactions throughout Northern California and a detailed pipeline map, we employ a difference-in-differences strategy that compares house prices very close to a pipeline to those somewhat farther away. We find that prices homes within 1000 or 2000 feet of a natural gas pipeline did not change in the months following the explosion, relative to homes 2000-4000 feet from a pipeline. This result holds when we restrict our analysis to the Bay Area, and when we use a triple-difference approach incorporating data from Southern California. Finally, we evaluate the informational letter campaign using a regression discontinuity design. Receipt of the letter also had no effect on housing prices.

Our findings suggest that housing prices do not respond to pipeline proximity after an event. In addition, informing current homeowners of the location of pipelines does not affect prices for those homes. This is an important finding: people are typically unaware of the 300,000 miles of transmission pipelines currently run underfoot in the U.S., unless there is a major accident. It suggests that opposition to natural gas pipelines is driven primarily by broader concerns about climate change and fracking, or construction disamenities. However, it is not consistent with recent state legislation that bans siting pipelines near schools, nursing homes, and hospitals. We plan to investigate these factors next using data on pipeline expansion project timing and location. Overall, this research is essential for distinguishing opposition due to real disamenities from NIMBYism and opposition to indirectly related environmental issues (such as fracking) in making long-run infrastructure decisions.

Finally, in **“Weather, Salience of Climate Change, and Congressional Voting”** (*Journal of Environmental Economics and Management*), Erich Muehlegger and I investigate the effect of local weather on the salience of climate change and voting on environmental legislation. Given the importance and complexity of climate change, it is crucial to understand how the public and elected officials develop their views and policy priorities regarding greenhouse gas emissions. Short-run weather phenomena are often cited to illustrate the presence or absence of climate change, but there is no existing evidence

connecting this relationship to real policy action.

We use data on Google search intensity and congressional roll call voting to assess the effect of unusual weather on the salience of climate change and environmental policy. We find that salience, as measured by the rate of Google searches for “global warming” and “climate change”, is significantly affected by recent unusual weather events. Furthermore, we provide evidence that members of the U.S. Congress and the Senate are more likely to cast environmentally-friendly votes after their home state experiences particularly unusual weather. Specifically, we find that a one standard deviation change in search intensity induced by unusual weather is associated with an 8.4% greater likelihood of casting an environmentally-friendly vote.

## Economics of upstream energy and resource markets

Much of my research deals directly with firm decision-making and regulation in upstream energy and natural resource markets. The very first step in many of these markets is the delegation of extraction duties from the owner of the resource to a firm. This principal-agent problem means that asymmetric information and moral hazard are prominent features in these settings. These frictions (and the mechanisms used to address them) are important drivers of private and external costs, supply behavior, and impacts of environmental regulation. I explore these themes theoretically and empirically in a series of three papers.

Government contracts with private firms increasingly incorporate environmental objectives or preferences for “green” products and producers. In **“Conservation Versus Competition? Environmental Restrictions in Government Contracting”** (Revisions requested, *AEJ: Microeconomics*), I estimate the costs and competitive effects of a suite of conservation restrictions imposed on Michigan timber harvesting contracts. Theoretically, such environmental objectives could undermine or bolster the primary goal of government contracting: competitive sourcing. Although a reduced-form analysis of bids reveals the effect of these objectives on government revenues, it cannot identify the overall cost of the policy or the effect on bidder surplus; lower bids could simply reflect weaker competition. I show how restrictions can influence equilibrium bids through changes in firm costs, strategic bidding behavior, and bidder participation.

I compile data from Michigan Department of Natural Resources (DNR) timber contracts, including a measure of seasonal operating restrictions and auction outcomes. I show that highly-restricted contracts receive 35% fewer bids and bring in 17% less revenue. To recover the magnitude and incidence of the cost of the objectives, I estimate a structural auction model that allows the restrictions to flexibly affect the loggers’ demand for the contracts. The structural estimates imply that the most-stringent restrictions impose compliance costs amounting to 15% of government revenue or 54% of firm surplus. However, these costs are borne almost entirely by the government in the form of lower equilibrium bids. Full passthrough occurs because the DNR is inelastic in supplying timber and compliance costs do not affect the dispersion of firm valuations. Loggers capture a larger share of total auction surplus for more restricted contracts, suggesting that the auction is less effective at fostering competition in these cases. Finally, I show that setting reserve prices optimally can partially mitigate the revenue disparity between more- and less-restricted contracts. Taken together, the results suggest that governments should be aware of strategic firm behavior when considering the impact of environmental objectives.

During the last decade, the fracking boom has substantially increased U.S. oil and natural gas production. These resources are often owned by private landowners, who sign lease contracts with firms to extract the oil and gas. In **“Asymmetric Information, Drilling Distortions, and Oil and Gas Leases”**,

Ryan Kellogg, Eric Lewis, and I seek to estimate the impacts and explain the presence of two pervasive features of these contracts: the royalty and the primary term. The royalty is a percentage of hydrocarbon revenue (not profit) that is paid to the landowner, driving a wedge between the landowner's and firm's incentives. The primary term is the period of time granted to the firm to exercise its option to drill a well on the lease. If the firm does not begin production by the end of the primary term, it loses the lease, and the landowner is then free to sign a new contract with another firm. If, however, the firm does commence production, the lease enters a secondary term, which lasts until the firm ceases operations. This creates a strong, discontinuous incentive to drill at least one well before the deadline, thus distorting the firm drilling decision.

Using detailed data on lease contracts and the timing of drilling, we show empirically that primary term expiration dates have an economically significant impact on firms' drilling decisions: a large share of wells are drilled just prior to expiration. This systematic pattern is difficult to fully explain with other factors such as information or common pool externalities. We then develop a model to explain why primary terms and royalties can help maximize the landowner's expected revenue from a lease, despite the distortions they generate. In our model, firm has private information about the expected productivity of the well and the inputs used in drilling. Intuitively, the optimal contract strikes a balance between extracting the firm's private information and distorting its incentive to choose the timing of and inputs into the well-drilling process. To achieve this balance, payments are tied to observable production outcomes and drilling timing decisions.

We have assembled a dataset of lease parameters, drilling and fracking inputs, realized natural gas production, and the timing of drilling for the Haynesville Shale in Louisiana. We are developing a structural model of firms' drilling decisions to quantify the welfare impacts of these contractual provisions and evaluate alternative lease designs. We also plan to estimate the extent to which primary term-driven production may have exacerbated the recent prolonged U.S. natural gas price slump.

In another ongoing project, "**Seller Commitment and the Empirical Analysis of First-Price Auctions**", I extend techniques for the empirical analysis of auctions to a setting in which the seller cannot commit to keeping an item off the market. This is a very relevant feature in natural resource auctions. For example, Porter (1995) notes that 47% of federal offshore oil and gas leases that do not receive a sufficient bid are eventually re-auctioned. Further, according to McAfee and Vincent (1997), the US Forest Service responds to a no-bid auction by re-auctioning the tract and reducing the reserve price by 10 percent. The possibility of a future re-auction introduces a dynamic component which will affect estimates of bidder valuations, division of surplus, and optimal seller policies such as reserve prices. The topic of optimal seller behavior with limited commitment has been analyzed in the theory literature (e.g., Burguet and Sákovics, 1996; McAfee and Vincent, 1997; Vartiainen, 2011; Skreta, 2015). However, to my knowledge, I am the first to incorporate this institutional feature into a structural empirical framework.

The workhorse approach to empirically analyzing first-price auctions, introduced by Guerre et al. (2000), essentially inverts the firm's first-order condition to recover valuations directly from bid data. In a context with re-auctioning, however, the researcher must account for three additional considerations in that first-order condition: (1) In all auctions, including the first round, the bidders anticipate that future rounds may be more favorable to them due to (revealed) weaker competition and a lower reserve price; (2) assuming some persistence in valuations, bidders are aware that later rounds are contested by bidders with an adversely-selected distribution of valuations; and (3) if there are value shocks in between auctions, the uncertainty creates option value associated with waiting.

Restricting the model to two rounds, I derive the unique symmetric Bayes Nash Equilibrium of this auction, which similar to that of a first-price auction with bid preparation costs. Drawing on related empirical approaches, I propose a estimator consistent with my model. These considerations prove important

in Monte Carlo simulations: my estimator performs well, while a naive static approach fails to recover the correct parameters. Finally, I discuss strategies for addressing unobserved auction heterogeneity in this setting. Future work will relax current model assumptions and apply the estimator to timber auction data. I also plan to extend the approach to accommodate interdependent values, which could capture interesting dynamics as bidders update their information set each period; it would also better characterize settings such as oil lease auctions.

## **References**

- Burguet, Roberto and József Sákovics**, “Reserve Prices without Commitment,” *Games and Economic Behavior*, 1996.
- Guerre, Emmanuel, Isabelle Perrigne, and Quang Vuong**, “Optimal Nonparametric Estimation of First-Price Auctions,” *Econometrica*, 2000.
- McAfee, R. Preston and Daniel Vincent**, “Sequentially Optimal Auctions,” *Games and Economic Behavior*, 1997.
- Porter, Robert**, “The Role of Information in U.S. Offshore Oil and Gas Lease Auctions,” *Econometrica*, 1995.
- Skreta, Vasiliki**, “Optimal Auction Design Under Non-Commitment,” *Journal of Economic Theory*, 2015.
- Vartiainen, Hannu**, “Auction Design without Commitment,” *Journal of the European Economic Association*, 2011.