Research Statement—Eric Edwards

Summary

In natural resource economics, resource dynamics and human behavior interact to determine the use (or misuse) of a resource. A "social planner" is often posited to observe and control resource user behavior to maximize some objective function, often aggregate utility, or wealth. In the real world, social planners don't exist, at least not in this form. Instead, it is the institutions that govern the use of a natural resource—such as property rights, laws, and social norms—that determine the extent to which resource use maximizes resource value. My research is interested in when and which institutions are adopted or changed, how they operate, and what outcomes they achieve. This work builds on the pioneering papers of Ronald Coase, Harold Demsetz, Douglass North, and Elinor Ostrom, who provide a framework of property rights, transaction costs, and collective action, by exploiting new data and econometric methods in the context of natural resource governance.

My work focuses on a variety of natural resource settings, especially water, and falls into four broad categories. The first area is groundwater governance, where I have focused on elucidating and quantifying the benefits of groundwater management and the tradeoffs between more stringent management and the transaction costs required to implement it. This work has contributed to the literature in environmental economics on groundwater use and management. The second area is work on agricultural development in the 19th and 20th centuries, and contributes to a large literature on this topic in economic history. This work focuses on property rights to land and water, the role of these institutions in development, and the importance of non-property institutions in shaping outcomes. The third area is the role of insecure property right institutions to land and water on the development of Indigenous groups in the western US, contributing to economic literature on impediments to development of one of the most marginalized groups in the country. The fourth area examines how markets for water in arid regions differ from markets for other commodities as a result of transaction costs related to the laws and regulations surrounding water right trading, as well as the externalities caused by water transfers.

My work is distinctive in its focus on natural resource system dynamics and institutional details using quasi-experimental methods. Natural resource economists have traditionally used simulation models to understand the economics of complex natural system behavior, while institutions-oriented researchers have tended to describe, but not test, behavior. My research topic areas are linked by an approach that uses the details of resource behavior and the related institutions to create testable predictions and estimate causal relationships.

Groundwater Governance

The distribution of economic benefits from groundwater management is determined by underlying aquifer characteristics. While there is a long tradition of examining groundwater extraction in economics going back to the 1970s and 1980s (e.g. papers by Oscar Burt in *AER* and Micah Gisser in *JPE*) this heterogeneity of benefits had not been recognized until recently, and my work was the first to test empirically whether these effects are economically important. *"What Lies Beneath? Aquifer Heterogeneity and the Economics of Groundwater Management*" (*Journal of the Association of Environmental and Resource Economists*, 2016). The portions of an aquifer where water moves rapidly, those with high hydraulic conductivity, as well as those that receive less yearly recharge, face a more costly common-pool problem, and therefore receive higher benefits from management. The introduction of management districts in Kansas is used to test the effect of underlying aquifer characteristics; a landowner in a county with hydraulic conductivity one standard deviation higher sees a relative land value increase of 5%–8%.

Following this work, I co-authored a paper with Andrew Ayres and Gary Libecap explaining how the heterogeneity described in the *JAERE* article affects the choice of management institution. "*How Transaction Costs Obstruct Collective Action: Evidence from California's Groundwater"* (*Journal of Environmental Economics and Management*, 2018). We examine governance institutions in California's 445 groundwater basins to identify factors that influence the adoption of extraction controls. In most basins, unconstrained pumping occurs because the benefits of other approaches to management are small relative the benefits. While heterogeneity in benefits partially explains adoption of institutions that regulate extraction, this work also documents the critical role of the transaction costs associated with contracting in explaining why some basins with apparently high benefits from additional management interventions fail to adopt pumping controls.

After working on groundwater governance issues for many years, I collaborated with Todd Guilfoos to propose a unifying economic framework for understanding the emergence and purpose of groundwater governance. We examined ten basins located on six continents to show that groundwater governance addresses local externalities to balance the benefits of reducing common pool losses with the costs of doing so. "*The Economics of Groundwater Governance Institutions Across the Globe*" (*Applied Economics Perspectives and Policy*, 2021). While broad, basin-wide solutions to open access pumping are limited, spatially localized externality problems, like those described in the *JEEM* and *JAERE* papers, raise the benefits of management actions, allowing for the implementation of more stringent pumping controls in certain areas. This paper received an award from the journal for the best paper published in 2021. I believe our framework provides a starting point for all future work on the management of groundwater resources in the environmental economics literature.

Economic History of Agricultural Development and Institutions

There is a rich literature on the economic history of agricultural development in the US, (e.g. Allen and Lueck on why family farms dominate historically and Olmstead and Rhode on the development of the tractor). One aspect of development that has largely been ignored is the role of water. Today, over 50% of the value of agricultural production in the US is from irrigated farms, yet until Steven Smith and I worked on this topic, no estimate had been made of irrigation's role in US agricultural development. "*The Role of Irrigation in the Development of Agriculture in the United States*" (*Journal of Economic History*, 2018). We analyzed how productivity and farm values changed in the western United States as a result of technological and policy changes that expanded access to ground and surface water for irrigation. To statistically identify the effects, we compare counties based on their potential access to irrigation water defined by physical characteristics. We find areas with access to large streams or groundwater increase crop production relative to areas with only small streams by $19 billion annually, equivalent to 90 percent of the total annual increase in agricultural production in the western United States after 1940.

One aspect that this paper did not address, and the aspect most relevant to contemporary discussions of climate change, was the role of irrigation in agricultural resiliency. We followed up using 100 years of agricultural production and weather data across the United States to identify the extent to which access to stored water—distinguished by its source and location—affected drought resiliency. "*Water Storage and Agricultural Resilience to Drought: Historical Evidence of the Capacity and Institutional Limits in the United States"* (*Environmental Research Letters,* 2021). Arid regions with access to stored water avoided the 13% losses in crop value experienced in irrigated areas with more limited storage during droughts. This work also linked the use of these resources to their underlying governance regimes, suggesting that the incomplete governance of groundwater withdrawals in many areas allows resiliency in the near-term, but potentially at the expense of future water availability.

Climate resiliency, especially in historical settings, continues to be one of my primary research interests. Agricultural adaptation to climate change will involve water as a key input. Wally Thurman and I

examine agricultural drainage as adaptation to excess water, which occurs in many areas of the eastern United States. "*The Economics of Climatic Adaptation: Agricultural Drainage in the United States*" (*NBER Volume on American Agriculture, Water Resources, and Climate Change*, forthcoming) Although shifting growing regions and increased precipitation in the eastern US are already leading to the northern expansion of drainage, economic analysis has largely ignored this important issue.

We also have ongoing work looking at drainage from an historical perspective (*"The Development of Agricultural Drainage in the United States 1850-1969").* Of the 215 million acres of wetlands estimated to have existed in the contiguous United States at colonization, 124 million have been drained today, 80-87% for agricultural purposes. We argue that a key institutional innovation, the drainage management district, facilitated local investment in drainage. States in our sample adopted drainage laws between 1857 and 1932, and after adoption each state saw an increase in improved agricultural land in counties with poorly drained soils relative to well-drained counties. We estimate artificial drainage increased the value of agricultural land in each of the worst-drained counties of the eastern United States by $3.7-$6.5 billion (2020 dollars).

My work has also used agricultural climate suitability to do comparative institutional analysis in the early 20th century between the Argentine Pampas and the US Midwest. *Property Rights to Land and Agricultural Organization: An Argentina-United States Comparison* (*Journal of Law and Economics*, 2022). In the United States, land was distributed in small parcels and actively traded. In the Pampas, land was distributed in large plots, and trade was limited because land was a social and political asset, as well as a commercial one. We analyze why the absence of trade led to persistently larger farms, specialization in ranching, and peculiar tenancy contracts in Argentina relative to the United States. Our empirical analysis, based on county-level data for both regions, shows that geoclimatic factors cannot explain the observed differences in agricultural organization, lending credibility to an institutional explanation. These results as well as those on climate resiliency, while based on historical data, are relevant to contemporary discussions, as the property right institutions associated with land and water tend to persist.

<u>Indigenous Property Rights and Agricultural Development</u>

American Indian reservations, primarily located in the western US, have low incomes and high rates of poverty relative to adjacent communities. In a series of papers, we examine how land and water right insecurity plays a role in the economic development of these populations. In a paper with Muynag Ge, Reza Oladi, and Sherzod Akhundjanov, we examine the extent to which a lack of access to capital might explain differences in capital investments using irrigation systems as a proxy for on-farm investment around the Uintah-Ouray Indian Reservation in eastern Utah. "*Irrigation Investment on an American Indian Reservation" (American Journal of Agricultural Economics, 2020).* Uintah land is held in trust by the US government, and farmers on this land face significant barriers to acquiring capital to invest in irrigation equipment and infrastructure. We use the boundaries from a 1905 land allotment as a natural experiment, employing both sharp and fuzzy regression discontinuity designs to explore whether agricultural land use, irrigation levels, irrigation investment, and crop choice differ across the boundary. The original allocation provided similar land in the immediate neighborhood around its borders, and our results suggest that today tribal trust land is farmed and irrigated at rates similar to adjacent land. Conditional on being irrigated, however, tribal trust land is around thirty-two percentage points less likely to utilize capital-intensive sprinkler irrigation, and up to ten percentage points less likely to grow high-value crops. Trust ownership, which is characterized by cumbersome bureaucratic processes, limits on agricultural lease flexibility, and the inability to use land as collateral to acquire loans, is a likely explanation for the observed differences.

One aspect not addressed in the Uintah setting, where farmers on- and off-reservation have similar access to water, is that many reservations have faced considerable difficulties in acquiring legal rights to water.

170 of 226 American Indian reservations have unresolved water claims that potentially exceed the region's hydrological capacity, generating uncertainty for tribes and off-reservation water users. Leslie Sanchez, Bryan Leonard, and I constructed a complete and novel dataset on Indian water settlements and reservation characteristics which we then analyzed using a bargaining framework from the *JEEM* paper discussed above. *The Economics of Indigenous Water Claim Settlements in the American West (Environmental Research Letters, 2020).* We find that rapid off-reservation population growth, water scarcity, and large anticipated water entitlements catalyze disputes. When more users are involved in the negotiations, transaction costs delay settlement, increasing water insecurity. We use our findings to predict allocations for 25 ongoing water right negotiations, an exercise that helps bound the uncertainty facing water managers throughout the US West.

Building on this paper, ongoing work examines outcomes using satellite data and robust difference-in-difference methods to estimate the causal effect of water right settlements from 1974 to 2012 (*"The Long-term Outcomes of Restoring Indigenous Property Rights to Water").* Settlements increase agricultural land use by 8% and have no effect on developed land use. Our estimates of tribal water use indicate that many tribes are utilizing only a fraction of their entitlements, forgoing as much as $938M to $1.8B in revenue. We provide evidence that this gap is driven by land tenure constraints and a lack of irrigation infrastructure, consistent with earlier work from the paper in *AJAE*.

Water Markets and Transaction Costs

Across the western US, growing populations and urbanization along with environmental demands and a changing climate have strained water allocation mechanisms originally designed to provide water to agriculture. Water markets offer a reallocation mechanism that can provide both buyer and seller surplus, but as many economists have noted, water markets are not working as we might expect. In a paper with Quentin Grafton, Gary Libecap, Robert Young, and Clay Landry, we provide evidence of price differential between sectors in the US, indicating how much transaction costs limit trade. "*Comparative Assessment of Water Markets: Insights from the Murray-Darling Basin of Australia and the Western USA*" (*Water Policy*, 2012). These results are contrasted with Australia, where (at the time) transfers were successfully reallocating water from agricultural to urban uses, with large gains from trade.

In a paper with Oscar Cristi, Gonzalo Edwards, and Gary Libecap, we find similar patterns of sectoral price differentials in northern Chile's Atacama Desert, the driest region in the world. "*An Illiquid Market in the Desert: Estimating the Cost of Water Trade Restrictions in Northern Chile*" (*Environment and Development Economics*, 2018). This paper uses these prices to estimate the cost of the current policy which limits trades from agricultural users to mining firms. In response to the policy, mining firms have developed high-cost desalination and pumping facilities to secure adequate water supplies. We estimate the cost of market transactions that fail to occur due to the policy at US$52 million per year. Without trade restrictions, around 86 per cent of the remaining agricultural water in the region would be transferred to mining. The policy is designed to protect riparian ecosystems and indigenous agriculture, suggesting the policy provides both benefits and costs and demonstrating that lost environmental quality is a potential justification for a policy limiting water markets.

Along with Muyang Ge, Sherzod Akhundjanov, and Reza Oladi, I pursue an examination of this type of environmental cost in ongoing work ("*Left in the Dust? Environmental and Labor Effects of Rural-Urban Water Sales*"). We examine the United States' largest ever agriculture-to-urban water transfer, from Imperial County to San Diego County, California. Using synthetic control and event study approaches, we show that the increasing intensity of agricultural water use after water is transferred decreases inflows to the Salton Sea, exposing areas of fine-silted lakebed and creating additional dust. Dust-related air pollutants, PM10 and PM2.5, increase during the relevant period while placebo non-dust pollutants, Ozone and NO2, do not. Taken together, my work on water markets suggests that while restrictions on

market transfers have large deadweight losses, market transfers without adequate environmental protections can create large externalities.

These results raise questions about whether markets could be designed to mitigate externality problems associated with human water diversions, as occurred in Imperial County. The world's saline lakes (like the Great Salt Lake in Utah or the Salton Sea discussed above) are shrinking and human water diversions are the key contributor. In a paper with Sarah Null, we consider how the introduction of markets might affect the cost of reducing flows to the Great Salt Lake. "*The Cost of Addressing Saline Lake Decline and the Potential for Water Conservation Markets*" (*Science of the Total Environment*, 2019). We construct cost estimates of increasing water inflows to a level consistent with the production of key ecosystem services, like dust prevention, by comparing the cost of uniform cutbacks to cap-and-trade systems which allow intra- and inter-basin trading. The cost to permanently implement uniform water right cutbacks is $37.4 million under the ecosystem service-preservation criteria. Significant cost reductions occur from intra-basin water conservation markets (5–54% cost decrease) and inter-basin water conservation markets (22–57% cost decrease). This study is the first to compare the cost savings from inter- and intra-basin conservation markets and suggests that allowing marginal abatement costs to equalize within basins provides a significant fraction of the benefits of full inter-basin reallocations. As suggested by the work in San Diego, out-of-basin transfers are more controversial and prone to externalities. Follow-on work is needed to establish more broadly whether within-basin reallocation can provide a large fraction of the gains from trade without the same external costs.