Environmental Regulation with Irreversible Investments: Evidence from High Plains Aquifer Depletion

Aaron Berman Nathaniel Hickok*
(Job Market Paper)

September 18, 2025

Please click here for the latest version.

Abstract

Many of the world's major aquifers are rapidly depleting from agricultural irrigation, generating dynamic common-pool externalities by raising future extraction costs. While Pigouvian taxation can restore the first-best outcome, such policies are often infeasible due to political and technological constraints. By contrast, extensive-margin policies that regulate well entry are second-best but administratively simple. We study this tradeoff by developing and estimating a dynamic model of farmers' joint well-drilling and water-use decisions using a novel dataset of aquifer levels, agricultural water use, and crop production in the Kansas High Plains Aquifer from 1959–2022. We find that entry fees can create substantial welfare gains but that their effectiveness erodes rapidly over time. In 1960, entry fees would have captured three-quarters of the gains of Pigouvian taxation, by 1980 only one-quarter, and today virtually none. This rapid decline reflects negligible marginal extraction costs relative to fixed well drilling costs, ensuring depletion becomes largely locked in once wells are sunk. Together, these findings highlight the importance of irreversible investments in constraining second-best environmental regulations.

*Berman: MIT, bermana@mit.edu; Hickok: MIT, nhickok@mit.edu. We are grateful to Nikhil Agarwal, Ben Olken, and Tobias Salz for their invaluable advice. We thank Mert Demirer, Glenn Ellison, Jacob Moscona, Nancy Rose, Ben Vatter, and many seminar participants at the MIT Industrial Organization and Environmental lunches and workshops for their helpful comments and suggestions. We also thank Blake Brownie Wilson and Aaron Hrozencik for generous discussions about data and institutions. We acknowledge support from the George and Obie Shultz Fund and the Martin Family Society of Fellows for Sustainability. This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant Nos. 1745302 and 2141064. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

1 Introduction

Water is a critical resource for agricultural production. At the same time, in many arid and semi-arid regions of the world, rainfall is scarce and access to water is only made possible by groundwater resources stored in underground aquifers, which can then be pumped to the surface and used for irrigation. Increased access to these groundwater resources has enabled large advances in agricultural productivity over the past century (Murgai 2001; Hansen, Libecap, and Lowe 2009; Edwards and Smith 2018; Ryan and Sudarshan 2022; Libecap and Dinar 2023). But many aquifers, formed over millennia of precipitation and runoff accumulation, are now rapidly depleting as a result of extensive agricultural irrigation (Jasechko et al. 2024; Carleton, Crews, and Nath 2025).

In many respects, aquifers can be viewed as a classic example of a commons problem (Coman 1911; Hardin 1968; Gardner, Ostrom, and Walker 1990). As water is pumped from the aquifer, the water table declines and costs of extraction increase. Because individual users are small relative to the scale of the aquifer, they do not internalize the dynamic common pool externality that their pumping today imposes on future users and therefore over-extract relative to the first best. In this view, aquifer depletion is a natural consequence of a lack of regulation, and the first-best outcome can be restored through the use of simple corrective taxes or establishing clear property rights (Pigou 1932; Coase 1960).

In practice, neither property rights nor Pigouvian taxes are commonly used to govern commons resources, and especially groundwater (Blomquist 1987; Ostrom 1990; Edwards and Guilfoos 2021). For one, the complexity of hydrogeological flows makes it difficult to define and enforce property rights (Libecap 2010; Ayres, Meng, and Plantinga 2021; Sears et al. 2023). Similarly, heterogeneity in hydrogeological characteristics and aquifer depletion rates means that pumping externalities vary significantly across space and time, making it technologically difficult to implement a first-best Pigouvian tax (Brozović, Sunding, and Zilberman 2010). Moreover, even uniform Pigouvian taxes face steep monitoring requirements and are typically politically unpopular (Guilfoos, Khanna, and Peterson 2016; Boyd et al. 2018; Ryan and Sudarshan 2022).

Unlike both Pigouvian taxes and property rights, extensive margin restrictions on new entrants are simple to monitor and face fewer political constraints since they do not impose costs on existing users. This simplicity is especially true for groundwater extraction since it requires large, easily observed investments in wells and other irrigation technology. In addition, because groundwater wells are extremely long-lived investments and the fixed costs of well construction are large relative to the marginal cost of extracting water, depletion can become largely locked-in once wells are sunk, necessitating action before such investments

are made. As a result, policies that resemble entry restrictions have become a standard tool for managing aquifers, especially in the western U.S., which is the focus of this paper. In particular, new entrants are often required to obtain so-called water rights before drilling a well (Peck 2014). Since these rights have low transferability, they act as de facto entry permits (Brewer et al. 2007; Libecap 2010).

In this paper we study the efficiency trade-off inherent to extensive margin entry restrictions relative to a first-best Pigouvian benchmark. We focus on how this trade-off is influenced by technological irreversibility. In particular, since groundwater well drilling involves long-lived investments with large fixed costs, we study the extent to which depletion can become locked-in once investments are sunk, creating an opportunity cost of policy inaction. We then study how this opportunity cost can differentially affect policies that operate on the extensive and intensive margins.

Our empirical setting is the portion of the High Plains aquifer system underlying Kansas. This region is extremely agriculturally productive but also semi-arid and highly reliant on groundwater from the High Plains aquifer to provide sufficient water to meet crop water demands. Overall, irrigation accounts for nearly 95% of the water extracted from this aquifer, and extensive agricultural irrigation has created significant depletion that threatens future crop production (McGuire and Strauch 2024). Driven by post-WWII technological development, large-scale agricultural irrigation began in this region in the 1950's and continued to expand in the ensuing decades. As a result of extensive groundwater extraction for irrigation, current projections suggest that large portions of the aquifer will become depleted and unable to provide sufficient water for agricultural irrigation in the coming decades (Scanlon et al. 2012; Stover and Buchanan 2017).

To study these questions we construct a comprehensive, long-run panel dataset on agricultural irrigation in the Kansas High Plains aquifer at the field level spanning 1959-2022. Our main source of data is administrative microdata on water use, crop choices, and acres irrigated at the field level derived from water use reports submitted to the Kansas Division of Water Resources (KDWR). We similarly obtain data on well construction from driller reports submitted to the Kansas Department of Health and Environment and data on water levels in the aquifer from a network of observation wells maintained by the Kansas Geological Survey (KGS). We complement these granular data sources with county-level data on farm expenditures and production from 1959-2012 using data from historical Censuses of Agriculture. Our dataset is uniquely well-suited to answer our questions of interest because it is both extremely rich, covering both the extensive margin of entry and the intensive margin of water usage at a spatially granular level, but also has a long time-series, allowing us to study historical entry, which largely took place from 1960-1980, as well as the impacts of

long-run investments.

Our empirical approach combines a rich economic model of agricultural production and investment in groundwater wells with a physically realistic model of groundwater flows in a full equilibrium framework. On the intensive margin, production depends on water usage as well as labor and materials inputs, which are endogenously chosen by field owners to maximize profits. Fields are heterogeneous both in their productivity and also in their access to groundwater based on the depth to water and bedrock underlying the field. On the extensive margin, field owners decide when and if to invest in groundwater wells and other irrigation technology based on their expectations about future productivity, drilling costs, the climate, and the lifespan of the aquifer. Entry costs are determined by the timing of well drilling as well as the depth of the required well and the size of the field. Well drilling is irreversible in the sense that there is no scrap value. Instead, exit from agriculture arises endogenously through aquifer depletion, which increases the marginal cost of water extraction until irrigation is no longer profitable relative to rainfed agriculture.

Equilibrium in the model is driven by the path of aquifer depths. Increased entry results in faster aquifer depletion and thus a shorter usable lifetime for the aquifer, decreasing rents from irrigation. Equilibrium is reached when long-run expectations about the path of aquifer depths match actual depletion rates from realized entry. Given the number of fields in our setting, we model this equilibrium dynamic using an oblivious equilibrium approach to avoid a severe curse of dimensionality from a full Markov Perfect equilibrium approach (Hopenhayn 1992; Weintraub, Benkard, and Van Roy 2008).

To estimate the intensive margin production functions we take a two-step approach in the spirit of Olley and Pakes (1996) and Ackerberg, Caves, and Frazer (2015). Water input prices, which vary across the aquifer due to changes in depths, provide the main identifying variation. Although aquifer depths are endogenously related to productivity through aquifer depletion, we construct moments based on pre-determined lagged costs and implied productivity innovations that we recover by inverting from observed levels of water usage. Weather shocks provide an additional source of identifying variation. Our estimation results suggest that farmers are moderately elastic to water costs when they make their water usage decision, with an average elasticity of about -1. We also find that farmers practice moderate deficit irrigation, with yields reaching only about 90% of their maximum potential due to water limitations.

To estimate extensive margin entry costs we take a discrete Euler Equation approach, exploiting the fact that drilling a groundwater well is a terminal action that equalizes continuation values (Kalouptsidi, Scott, and Souza-Rodrigues 2021; Scott 2023; Hsiao 2025). Our approach is computationally light and does not require solving the full model for estimation.

The main identifying variation comes from variation in the timing of entry across areas of the aquifer combined with our estimates of productivity from the intensive margin. Intuitively, fields that delay entry despite high productivity must face high entry costs and likewise fields that enter early despite low productivity must face low entry costs. Our estimates imply an average entry cost of about $900,000 in 2017 dollars. We also find that the existing entry permitting regulations in Kansas impose only a moderate cost of approximately $31,000.

With parameter estimates for both the intensive and extensive margins in hand, we then consider counterfactual policies to understand the welfare costs of aquifer depletion, the relative effectiveness of extensive margin entry restrictions as a substitute for Pigouvian taxes on the intensive margin, and the opportunity cost of policy inaction due to irreversibility and lock-in. We first consider how aquifer depletion will continue under the status quo, using our combined economic and hydrogeologic model along with projections about future climate change to compute an equilibrium until the end of the century. We find that, under the status quo, nearly half of fields will have fully depleted the underlying aquifer and be unable to continue irrigation by 2050. However, the fraction of depleted fields only rises to 55% by the end of the century as overall water use declines as the aquifer becomes increasingly depleted. These predictions are only modestly impacted by alternative assumptions about future climate change since western Kansas is already an extremely dry region with a semi-arid climate.

We then consider counterfactual policies that either take the form of per unit intensive margin taxes on water use or extensive margin entry fees. We start by considering the welfare-maximizing uniform taxes and entry fees that begin near the start of our sample period in 1960. We find that both policies result in substantial welfare gains over the status quo. The welfare-maximizing uniform per unit tax is approximately $50 per AF and results in overall welfare of approximately $35.5B compared to $30.3B under the status quo, or an annualized welfare increase of nearly $110M per year in annuity-equivalent terms. Field depletion rates are also significantly lower, reaching only 35% by the end of the century. The welfare-maximizing entry fee is on the order of $1M, more than an order of magnitude larger than the implicit costs imposed by the existing entry permitting regulations, and results in overall welfare of approximately $34.3B, which is an annualized welfare increase of approximately $84M in annuity-equivalent terms. Similar to the per unit tax, field depletion rates shrink to only 37% by the end of the century as nearly half of the wells that are drilled under the status quo are never drilled under the counterfactual entry fee policy.

We also consider the opportunity cost of policy delay and lock-in by varying the start year of each policy. The effectiveness of both per unit intensive margin taxes and entry fees declines with delay as depletion becomes increasingly locked-in. For instance, the welfare

gains of an intensive margin tax beginning in 1978, which is the same year that Kansas first implemented its entry permitting process, is only $64M per year in annualized annuity-equivalent terms compared to the $110M in annualized welfare gains from an intensive margin tax starting in 1960. The decline in effectiveness is even steeper for entry fees. In 1960 entry fees captured more than 75% of the welfare gains achieved by intensive margin taxes, but this percentage declines to less than 25% by 1980 and virtually none today. A large part of this steep decline is driven by the simple fact that entry fees only affect fields that have not yet drilled a well, and the fraction of fields that have not yet entered shrinks rapidly from 1960-1980. An additional source of welfare losses is that the social surplus of fields is negatively selected over time. In particular, even though early entrants had high private surplus gains, spatial concentration in productivity means that they also have high externalities since they tend to be located in crowded areas of the aquifer that later suffer from severe depletion rates.

Related Literature We contribute to a large literature in economics studying the long-run management of natural resources, especially in common pool settings (Hotelling 1931; Dasgupta and Heal 1974; Solow 1974; Stiglitz 1974; Brown 1974; Ostrom 1990). Our primary contribution is to develop and estimate a dynamic model of water use and well drilling with rich heterogeneity in productivity to capture both short-run intensive margin changes in water use as well as long-run changes in entry into irrigation. We combine our empirical model of production and investment with a physical model of groundwater flows to capture the large variation in externalities that result from water extraction across the aquifer and quantify the welfare gains that result from counterfactual policies.

Previous work has largely studied the management of groundwater resources using an optimal control approach, abstracting from heterogeneity in productivity and physical realism in groundwater flows in order to maintain analytical tractability (Burt 1964; Gisser and Sánchez 1980; Kim et al. 1989; Negri 1989; Brill and Burness 1994; Koundouri 2004). More recent work has moved beyond the simplified "bathtub" aquifer framework to incorporate additional physical realism in groundwater flows, but due to data and computational limits have maintained relatively simple empirical models of agricultural irrigation (Brozović, Sunding, and Zilberman 2010; Guilfoos, Khanna, and Peterson 2016; Fenichel et al. 2016; Merrill and Guilfoos 2018; Sears et al. 2023). We address these limitations by first constructing a comprehensive dataset of irrigation water use, crop production, well drilling and aquifer conditions in an agriculturally important region over a long period. We then apply this data to estimate a rich empirical model of agricultural irrigation and investment in well drilling. Rather than directly solve the resulting computationally intractable optimal con-

trol problem, we use our empirical model, in combination with a physically realistic model of groundwater flows, to study the relative effectiveness of policies that operate on either the intensive or extensive margins.

Another strand of the literature has studied the value of water for agriculture either by directly estimating factor demand functions for water (Pfeiffer and Lin 2012; Hendricks and Peterson 2012; Pfeiffer and Lin 2014; Bruno and Jessoe 2021; Bruno, Jessoe, and Hanemann 2024; Hagerty 2024) or estimating agricultural production functions with irrigation water as an endogenously chosen input and then constructing the implied factor demand (Ryan and Sudarshan 2022; Rafey 2023; Ferguson 2024; Carleton, Crews, and Nath 2025). This strand of the literature has largely focused on the value of creating water markets for surface water through gains from trade and the losses from transaction costs that prevent those gains from being realized. Unlike surface water, difficulty in defining property rights has meant that there are essentially no water markets for groundwater, and in this paper we study how policy and regulation can encourage the efficient use of groundwater even in the absence of trade. Similar to Ryan and Sudarshan (2022) and Rafey (2023), we construct and estimate production functions for irrigated agriculture leveraging weather shocks and variation in costs of water due to heterogeneity in underlying aquifer characteristics to provide identifying variation. While Rafey (2023) focuses on the value of gains from trade in surface water markets, Ryan and Sudarshan (2022) similarly studies intensive margin Pigouvian water use taxes in the context of a heavily depleted aquifer in northern India. They focus on the short-run trade-off between equity and efficiency from water rationing relative to Pigouvian taxation. By contrast, we leverage our historical panel to study the long-run impacts of entry fee policies that operate on the extensive margin and how the relative effectiveness of these policies can decline over time due to lock-in of depletion.

More broadly, our paper contributes to a larger literature in environmental economics that studies regulation in settings with large, irreversible investments such as in oil wells (Anderson, Kellogg, and Salant 2018; Kellogg 2024), transportation (Barahona, Gallego, and Montero 2020; Donna 2021), paper mills (Gray and Shadbegian 1998), palm oil (Hsiao 2025), cement (Fowlie, Reguant, and Ryan 2016), and power plants (Fowlie 2010). Most closely related to this paper, Burlig, Preonas, and Woerman (2025) study how taxes on groundwater extraction can influence long-run investment decisions for perennial crop choices in California. They use their model to determine the level of tax required to reach sustainable pumping levels as set out by California's recently enacted Sustainable Groundwater Management Act. By contrast, farmers in our setting almost exclusively grow annual crops such as wheat and corn, eliminating investment dynamics in crop choices. Instead, investment dynamics arise through decisions to enter into irrigation by drilling a groundwater well, which is a margin

we are able to capture only because of the unique historical coverage of our data. Rather than determining the level of water use tax required for sustainable extraction, we leverage our hydrologic model of groundwater flows to solve for the welfare-maximizing taxes and entry fees and quantify the welfare gains from their implementation relative to the status quo. Moreover, we show how investment irreversibility can interact with the choice of policy instrument to generate a significant opportunity cost of policy inaction.

Roadmap The rest of the paper proceeds as follows. Section 2 provides details on our empirical setting, including the regulatory history and the hydrogeology of the High Plains aquifer. Section 3 details the data we compile on water use, crop production, and well drilling. Section 4 describes our empirical model of irrigated agriculture. Section 5 presents our estimation strategy. Section 6 shows the estimation results. Finally, section 7 reports the results of our counterfactual policy experiments.

2 Background

The High Plains aquifer system is made up of a series of interconnnected aquifers stretching across more than 100 million acres and eight states. It is one of the world's largest aquifer systems and supplies about a third of the water used annually for irrigation in the United States. The portion underlying Kansas, which is the focus of this paper, covers about 20 million acres in the western half of the state and consists of three separate aquifer formations: the Ogallala, the Equus Beds, and the Great Bend Prairie aquifers. The High Plains aquifer is the principal source of water in this region, as there are few surface water sources capable of providing sufficient water for large-scale agricultural irrigation. Moreover, because this region is predominantly agricultural, more than 90% of the water withdrawn each year is used for agricultural irrigation, with the rest being used for domestic, municipal, or industrial purposes.

Although the High Plains aquifer was first discovered in the late 19th century, large-scale agricultural irrigation did not begin until the 1950's when pump technology became sufficiently developed to allow for large volumes of water to be extracted at relatively low cost (Hornbeck and Keskin 2014). At the same time, the invention of the center pivot irrigation system drastically increased the extent of potentially irrigable land, resulting in widespread adoption of irrigation over the ensuing decades. Irrigation is especially valuable in western Kansas because the region lies in the rain shadow of the Rocky Mountains and consequently has a dry, semi-arid climate. Irrigation thus has the potential to greatly increase crop yields and also enables growing more water intensive crops, particularly corn. In large part due to

irrigation, corn has become the dominant crop of Kansas, replacing wheat, which has lower water requirements and can obtain relatively high yields even without irrigation.

While widespread agricultural irrigation has transformed western Kansas into one of the most productive agricultural regions in the world, it has also led to significant depletion of water in the aquifer. A combination of low rainfall levels and low deep soil percolation means that the High Plains aquifer experiences extremely low levels of natural recharge. Recharge rates average less than one inch per year, about a quarter of annual withdrawals (Buchanan, Wilson, and Butler 2023). Figure 1 shows the average depth to water in the aquifer since the beginning of large-scale irrigation. In 1950, average depth to water was less than 50 feet, while today it is it 120 feet, a decline of nearly one foot per year, which is enough water to supply about 40 million households.[1]

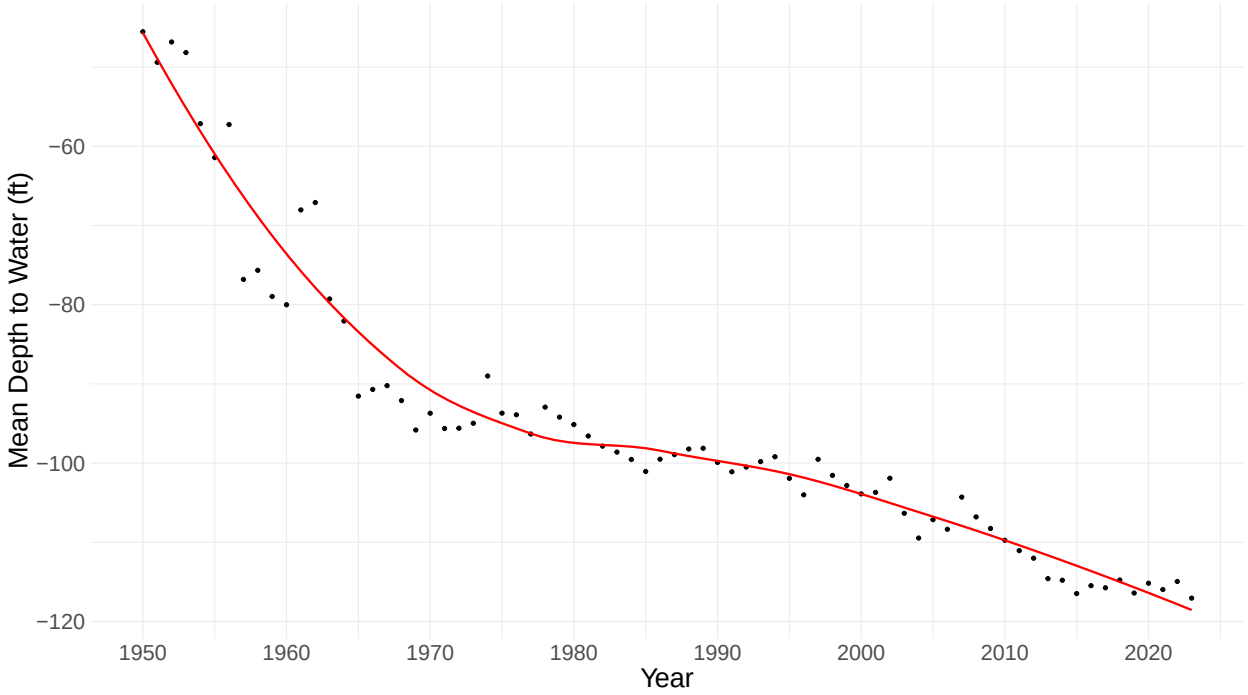

Figure 1: Average depth to water over time

Notes: This figure shows the average depth to water across the Kansas portion of the High Plains aquifer over time. Each point is an annual average and the solid red line is a smoothed LOESS local regression.

Groundwater use in the High Plains aquifer presents a classic dynamic common-pool externality. The costs of pumping groundwater increase with aquifer depths so that pumping by an individual user today increases costs for all future users. Because individual users are small relative to the scale of the aquifer, they do not internalize the full social cost of their

1. Household annual consumption averages approximately 0.5 acre-feet per year.

pumping and thus over-pump relative to the first-best. In the extreme, over-pumping can fully deplete the aquifer and result in infinite future extraction costs.

2.1 Regulatory History

Groundwater extraction in Kansas, as in much of the western US, has until recently been only loosely regulated (McGuire et al. 2003). Figure 2 shows the history of water regulation in Kansas, which largely evolved in response to a series of major droughts over the course of the twentieth century (Peck 2014). Following the Dust Bowl drought of the 1930's, the Water Appropriation Act of 1945 established state power to regulate water issues and unified groundwater and surface water regulation under a prior appropriation, or first-come first-served, framework.[2] From 1952 to 1957 Kansas experienced the most severe drought in the state's recorded history, which prompted the Kansas Division of Water Resources to begin collecting voluntary data on water use to better track statewide water demand (Layzell and Evans 2013).

Perhaps the largest change to groundwater regulation in Kansas took place in 1978 when the state moved from a free entry system to requiring permits for well drilling and water use. Known as water rights, these permits were, until recently, easy to obtain and proved to be only a small barrier to entry. Today, in recognition of the fact that the aquifer is already vastly oversubscribed, Kansas has largely stopped granting new entry permits. However, as we show in section 3, this change should largely be interpreted as fixing the set of potential, rather than actual, entrants to the aquifer, as new wells can still be constructed based on previously granted entry permits and there has been no discernible slowdown in new well construction as a result of this change. Moreover, the majority of well construction took place between 1960 and 1980 and so entry today is relatively unimportant compared to historical entry in the latter half of the twentieth century.

2.2 Aquifer Hydrogeology

Aquifers are formed by large areas of underground permeable rock or sediment such as silt, sand, gravel, or clay. The High Plains aquifer was formed several million years ago when large amounts of sediment eroded from the Rocky Mountains and settled on bedrock layers left behind by the ancient Western Interior Seaway. These bedrock layers shape the contours of the aquifer even today, determining the aquifer's saturated thickness. Since wells typically need to be drilled into the bedrock, depth to bedrock is also an important driver of entry

2. The main alternative approaches, the riparian and absolute ownership doctrines, tie the right to use water with property ownership.

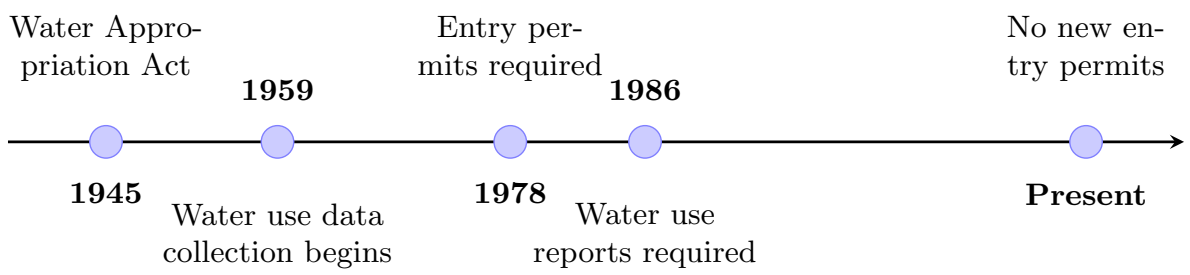

Figure 2: History of water regulation in Kansas

costs. In many areas of the aquifer, particularly the Ogallala formation, these bedrock layers can reach depths of more than five hundred feet.

Water flow in the aquifer, and by extension the degree of externality imposed on other users, is driven by two geologic features: hydraulic conductivity and specific yield. Hydraulic conductivity governs the speed at which water flows through the aquifer. Panel A of figure 3 plots the hydraulic conductivity of the aquifer at a resolution of five square miles. Hydraulic conductivity is highly variable across the aquifer, ranging from two miles per year at the lowest to more than twelve miles per year at the highest. By contrast, in the standard "bathtub" aquifer model used in most prior work in the economics literature, hydraulic conductivity is assumed to be essentially infinite—water extracted in one part of the aquifer causes immediate and homogeneous decreases in water table depth everywhere in the aquifer. In reality, an aquifer is better conceptualized as an overlapping set of local aquifers, with pumping in one area strongly affecting other nearby users but having only small and delayed effects on distant users.

The second geologic feature, specific yield, is related to the porosity of the aquifer and governs the rate at which water table depths decline as water is extracted from the aquifer. In areas with high specific yield the rock is relatively more porous and depths decrease more slowly per unit of volume withdrawn, while in areas with low specific yield depths decrease quickly. Since costs of extraction increase with depths rather than directly from volumes extracted, specific yields are an important determinant of the social cost of groundwater extraction. In the standard bathtub model, specific yields are also considered homogeneous and often implicitly assumed to be one, with volumes translating one-to-one into depth changes. By contrast, specific yields can vary from less than 5% to more than 25% across the aquifer, as shown in panel B of figure 3.

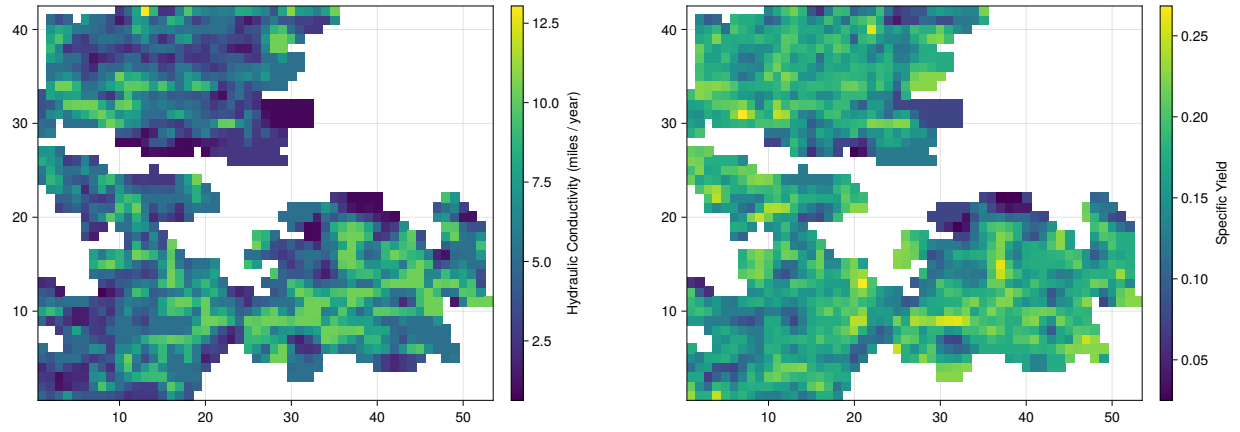

(a) Hydraulic Conductivity (b) Specific Yield

Figure 3: Aquifer Characteristics

Notes: This figure shows the average hydraulic conductivity across space in the aquifer in panel (a) and average specific yields in panel (b). Both panels have a spatial resolution five square miles.

3 Data

We construct an annual field-level panel of irrigation water use, well construction, water costs, aquifer depths, crop choices, and weather from 1959 to 2022. The main dataset we use to construct this panel is administrative microdata derived from annual water use reports submitted by individual water users to the Kansas Division of Water Resources (KDWR). We complement this field-level panel with county-level data on farm expenditures and production using data from historical Censuses of Agriculture and various NASS surveys. Appendix D provides further details on data sources and data construction.

We construct the marginal cost of water extraction in each year using the energy cost of water extraction and application. Energy costs increase with aquifer depth d_{it} according to (Rogers and Alam 2006):

$$c_{it} = P_t^{ng} \kappa^{ng} (d_{it} + \kappa) \tag{1}$$

where P_t^{ng} is the price of natural gas, which is the predominant fuel used for irrigation water extraction in Kansas, $\kappa^{ng} \approx 0.0233$ is a constant representing the amount of energy (in terms of Mcf of natural gas) required to lift one acre-foot of water by one foot, and $\kappa \approx 66.99$ is a constant equal to the additional lift required to maintain pump exit pressure. Equation 1 represents the main common pool externality imposed by irrigators: the more water they use today, the more the aquifer depletes and the higher costs are for all users tomorrow.

We also construct a measure of effective rainfall by summing evapotranspiration limited rainfall during the growing season (Allen et al. 1998). The intuition of this measure is

that any rainfall in excess of a crop's natural water demand, given by evapotranspiration, is not useful for increasing yields.[3] Likewise, rainfall outside the growing season does not increase yields beyond the relatively low amount of moisture that can be absorbed and stored in the soil. Evapotranspiration (ET) summarizes crop water demand as a function of humidity, temperature, wind speed, and other meteorological factors. Experimental evidence on agricultural yields shows that water effects crop yields only relative to ET and so we normalize water use by ET in our yield production functions in section 4.

Table 1: Descriptive Statistics

	Mean	SD	P25	P50	P75
Water Use (AF)	201.27	167.78	98.00	162.21	248.95
Depth (ft.)	108.18	62.75	49.36	107.18	148.90
Water Costs (2017$ / AF)	26.57	19.76	10.25	21.65	37.19
Effective Rainfall (ft.)	1.23	0.30	1.03	1.20	1.42
ET (ft.)	3.12	0.26	2.96	3.14	3.29
Field Size (acres)	159.22	101.78	114.00	130.00	179.61
Well Construction Year	1978.24	16.83	1968.00	1975.00	1984.00
Irrigated Yield (2017$ / Acre)	530.14	226.07	378.50	469.19	622.61
Non-Irrigated Yield (2017$ / Acre)	222.16	112.84	151.37	187.42	267.39
Materials Share	0.12	0.07	0.06	0.10	0.16
Labor Share	0.03	0.01	0.02	0.03	0.04

Notes: This table shows descriptive statistics for the field-level panel described in Appendix D. Average water use is conditional on irrigating and effective rainfall and evapotranspiration are both measured for irrigated crops.

Table 1 presents a selection of basic summary statistics for our field panel. Conditional on irrigating, the average water use is approximately 200 acre-feet (AF).[4] The average field size is approximately 160 acres.[5] The average effective rainfall is 1.23 feet while the average ET is 3.12 feet, resulting in a significant average irrigation deficit of 1.89 feet that must be made up by irrigation water to achieve full yields. Irrigation deficits in our setting are large due to the semi-arid climate of Western Kansas. Finally, average water costs are $26.57 per AF, implying farms spend on average about $5,350 per year on direct water costs relative to an average of about $60,000 per year in revenue.

3. In principle excess rainfall can actually be harmful to yields, but excess rainfall events are exceedingly rare in the semi-arid climate of Western Kansas.

4. A typical household uses approximately 0.5 AF of water each year.

5. A standard Kansas field encompasses a full PLSS quarter-section, which is a standard land unit in the Western U.S. measuring exactly 160 acres. Most of the variation we observe is due to differences in irrigation technology, for instance fields using center pivots often do not irrigate the corners, resulting in an effective field size of approximately 125 acres.

3.1 Water Use and Water Costs

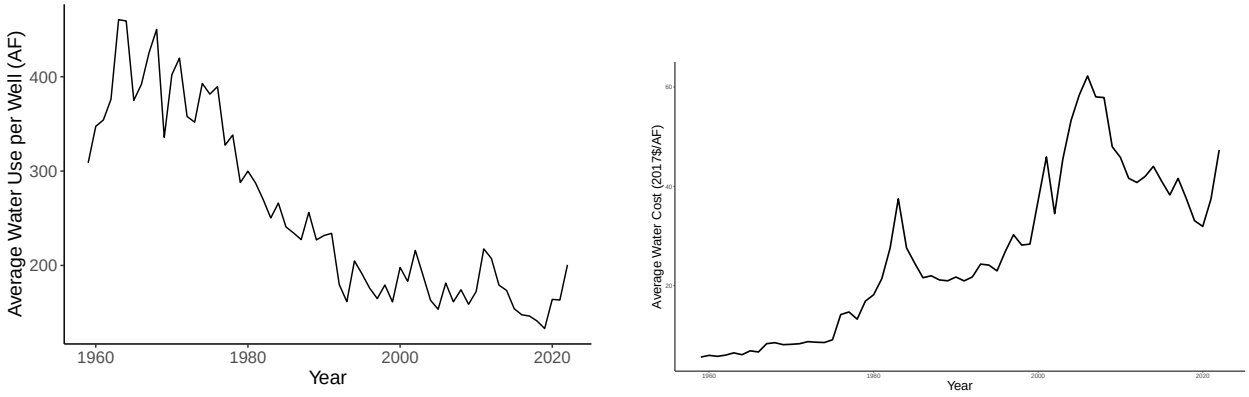

(a) Water Use (b) Water Costs

Figure 4: Water Use and Water Costs Over Time

Notes: This figure shows average water use per well in AF in panel (a) and average water costs in 2017$/AF in panel (b).

Figure 4 shows the average water use per well over time in panel (a) and the average water cost over time in panel (b). Average water use has largely been declining over the past century, reaching a high of nearly 500AF per well in the 1960's and declining to a low of almost 100AF per well by 2020. At the same time, water costs have largely been increasing due to corresponding increases in aquifer depths. However, even though aquifer depths have been uniformly declining, water costs have not been uniformly increasing as they are also highly sensitive to the price of natural gas. Accordingly, water costs have experienced two spikes, once during the 1970's and early 1980's due to the energy crisis and again before the shale boom of the 2010's.

3.2 Well Construction

Figure 5 shows the cumulative number of wells constructed in or before each year from 1959-2022. The number of irrigation wells increased rapidly in the 1960's and 1970's after the invention of the center pivot made large scale agricultural irrigation feasible in the region. Irrigation well construction has significantly slowed in the decades since, both due to increased regulatory oversight and also increased aquifer depletion lowering the returns to irrigation.

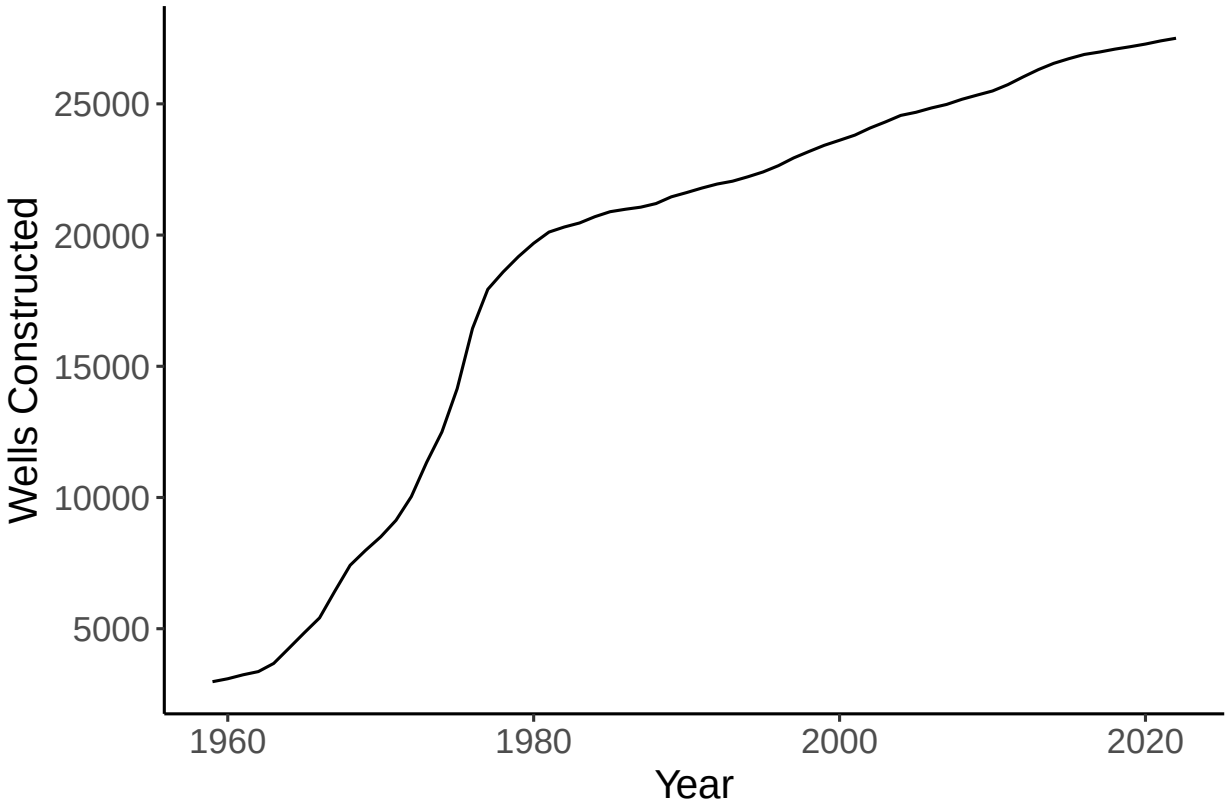

Figure 5: Wells Constructed Over Time

Notes: This figure shows the cumulative number of wells constructed in each year from 1959-2022.

4 Empirical Model

This section specifies an empirical model of supply for irrigated and rainfed agriculture. We divide land into atomistic fields that are independently managed by infinitely-lived owners. Field owners decide on the extensive margin whether or not to invest in irrigation technology. On the intensive margin, field owners combine water with other inputs to produce revenue according to an annual production technology. Input levels of water derive from exogenous precipitation for rainfed producers and endogenous groundwater extraction for irrigated producers. We abstract from demand by maintaining the assumption throughout that field owners are price takers on the world market and that their production does not impact world prices, which is plausible since Kansas agricultural production is small relative to the size of the world market.

4.1 Intensive Margin

We model production in each year t at the level of an individual field i. Consistent with the major crop types grown in Kansas, we consider two composite crop types, irrigated annuals and rainfed annuals. Field owners who have invested in irrigation choose between irrigated and rainfed crops depending on their water extraction costs c_{it}. Producers without irrigation wells grow rainfed crops. Then, given crop choice $j_{it} \in \{\text{irrigated, rainfed}\}$, field owners who have previously invested in irrigation choose their volume of water use W_{it}, and all producers choose their expenditures on other flexible inputs, labor L_{it} and materials M_{it}. Weather enters production through volumes of effective rainfall ER_{it} and evapotranspiration ET_{it}. Revenue yields are then given by

$$Y_{it} \equiv \frac{R_{it}}{A_{it}} = \exp(\omega_{it}^{j}) P_{jt} F_j(W_{it}, L_{it}, M_{it}, ER_{it}, ET_{it}) \tag{2}$$

where ω_{it}^{j} is a Hicks-neutral productivity term, P_{jt} is the price of crop j, A_{it} is the size of the field in acres, and all inputs are measured in per acre terms.

The assumption of Hicks neutrality is standard in the literature and allows the marginal revenue product of water to vary flexibly across fields and over time. However, it rules out any unobserved, factor-specific heterogeneity in productivity. Because we observe input usage at the individual level for only one flexible input, we are unable to separately identify factor-specific unobserved productivity from Hicks-neutral unobserved productivity. Instead, we allow for water productivity that varies based on observable weather, namely effective rainfall and evapotranspiration.

Field owners are assumed to choose flexible inputs to maximize static profits π_{it} given

their information set $\mathcal{I}_{it}$, which we recursively define as

$$\mathcal{I}_{it} = \{\vec{P}_t, A_{it}, c_{it}, \vec{\omega}_{it}, j_{it}, ER_{it}, ET_{it}, \mathcal{I}_{i,t-1}\},\ \mathcal{I}_{i0} = \emptyset \tag{3}$$

The assumption that effective rainfall and evapotranspiration are part of the information set when deciding irrigation volumes reflects the fact that irrigators do not need to commit to a particular level of water usage and can flexibly adjust their water use in response to weather shocks. As we describe further in section 5, the timing assumption on ω_{it}, along with static profit maximization, ensures that we can recover unobserved productivity by inverting observed factor demands.

The assumption of static profit maximization does, however, rule out any dynamics in irrigation water volumes and expenditures on labor and materials. The main restriction of this assumption is that field owners do not internalize the impact of their current water usage on future costs through aquifer depletion, which is consistent with the assumption of atomistic fields and is typical of the literature on common pool resources. We also assume that field owners are price takers in both output markets and the labor and materials input markets and that they face undifferentiated prices in these markets.

Finally, field owners make their crop choice decision to maximize expected static profits before observing realized weather ER_{it} and ET_{it}. Their information set when making crop choice decisions is given by

$$\mathcal{I}^j_{it} = \{\vec{P}_t, A_{it}, c_{it}, \vec{\omega}_{it}, \mathcal{I}_{i,t-1}\} \tag{4}$$

This assumption approximates the timing of planting decisions in the agricultural calendar, which largely occur before rainfall and evapotranspiration are realized. The assumption that crop prices are known when field owners make crop choice decisions reflects the fact that there are very active futures markets for most crops grown in Kansas. The price P_{jt} can be interpreted as the expected price, which we proxy for using the realized price at harvest time.

4.1.1 Production Technology

We follow the agronomic literature and assume that the water-yield relationship takes a Mitscherlich-Spillman form, which describes the physical relationship between water and crop yields (Hexem and Heady 1978). We further assume that F_j is Cobb-Douglas in labor and materials. In particular,

$$F_j = (1 - \exp(-\lambda_j EW_{it}))\, L_{it}^{\beta_{jL}} M_{it}^{\beta_{jM}},\ EW_{it} := \frac{W_{it} + \nu_j ER_{it}}{ET_{it}} \tag{5}$$

The parameter λ_j governs the water-yield relationship: a low λ_j implies a relatively larger amount of water is required to attain maximum yields, all else equal. The parameters β_{jL} and β_{jM} govern the elasticity of (potential) yield with respect to labor and materials.

The main economic assumptions embedded in equation 5 are as follows. First, consistent with a large body of experimental evidence on the effects of water on crop yields, we construct a measure of effective water EW_{it} that normalizes water and rainfall by evapotranspiration, which summarizes the natural crop water demand required to achieve maximum yields based on temperature, humidity, wind, and other weather factors (Steduto et al. 2012). Second, our measure of effective water assumes that irrigation water and rainfall are perfect substitutes up to the factor ν. The reason the two may not be perfect substitutes is that while irrigation water is typically applied at specific intervals designed to maximize yields, rainfall timing is stochastic and may over or under supply water throughout the growing season. This concern is mitigated by the ability of soil to store moisture and act as a buffer stock of water as well as the construction of an effective rainfall measure that discounts rainfall in excess of crop water demand via evapotranspiration (Allen et al. 1998). Finally, the Cobb-Douglas assumption imposes a unit elasticity of substitution between the Mitscherlich water aggregate $\widetilde{W}_{it} \equiv (1 - \exp(-\lambda_j EW_{it}))$ and the other factors, labor and materials. While this assumption restricts substitution patterns between water and other inputs, it still allows for a flexible factor demand for water, which is the main focus of this paper.

4.1.2 Productivity

My empirical strategy involves constructing moments based on the productivity innovation $\xi_{it}^j \equiv \omega_{it}^j - \mathbb{E}[\omega_{it}^j | \mathcal{I}_{i,t-1}]$. To recover this innovation term we first specify productivity for irrigated crops as

$$\omega_{it}^{\text{irr}} = x_i \gamma^{\text{irr}} + \gamma_t^{\text{irr}} t + \mu_{it}^{\text{irr}} \tag{6}$$

where the vector of controls x_i includes a constant, latitude and longitude, elevation, slope, and measures of soil quality.[6] This specification decomposes the field productivity into an individual-specific component that depends on field-specific characteristics, an aggregate component that evolves over time, and an idiosyncratic component. Second, we assume that the idiosyncratic term μ_{it} follows a first-order Markov process so that

$$\omega_{it}^{\text{irr}} = x_i \gamma^{\text{irr}} + \gamma_t^{\text{irr}} t + g(\mu_{i,t-1}^{\text{irr}}) + \xi_{it}^{\text{irr}} \tag{7}$$

6. We use two measures of soil quality, the concentration of organic carbon and the soil texture, as measured by the percentage of silt, sand, and clay.

Equation 7 is slightly more general than the typical first-order Markov assumption on ω_{it} as it places this restriction only on the idiosyncratic component ν_{it}^{irr}.

For rainfed crops we specify productivity as

$$\omega_{it}^{\text{rainfed}} = \omega_{c(i)}^{\text{rainfed}} + \omega_{t}^{\text{rainfed}} + \xi_{it}^{\text{rainfed}} \tag{8}$$

where $c(i)$ is the county of field i. The main restriction of equation 8 relative to equation 6 is that it assumes productivity is homogeneous within county-year rather than allowing for fully heterogeneous productivity at the individual-level. The reason for this is that we do not observe any individual-level input decisions for rainfed producers and so we can only recover aggregate, county-level productivity.

4.2 Extensive Margin: Entry

In each period t, any field i that has not previously adopted irrigation makes a binary entry decision e_{it} over whether or not to do so. Adopting irrigation entails drilling a well and constructing a means to apply extracted water onto the field, typically via a center pivot. Entrants must pay a sunk cost c_{it}^{e}, and we assume that there is no scrap value.[7] Instead, exit from irrigated agriculture arises endogenously in the model through aquifer depletion, which increases the marginal cost of water extraction.

Potential entrants are forward-looking and make decisions about whether or not to enter based on their expectations about future productivity, the climate, and the lifespan of the aquifer. They observe public state variables $s_{it} = \{c_{it}^{e}, d_{it}, \vec{P}_t, \vec{x}_i\}$ and privately observe a vector of shocks $\epsilon_{it}^{e} = (\epsilon_{it0}^{e}, \epsilon_{it1}^{e})$, which we assume follow an i.i.d Type 1 extreme value distribution with scale parameter σ_e. The vector of field characteristics $\vec{x}_i$ includes an ex-ante estimate of field productivity, the depth to bedrock, and the local climate of the field, which we proxy for with latitude and longitude. We do not directly include weather in the specification of the state s_{it}, which implicitly assumes that weather is not dynamic. This assumption improves tractability by limiting the dimensionality of the state space while still providing a good approximation of weather processes in Kansas, where most of the variation in weather is either spatial due to the rain shadow of the Rocky Mountains or temporal due to climate change and other aggregate weather phenomena.

The ex-ante value function $V(s_{it})$ and choice-specific value functions $v(0; s_{it})$ and $v(1; s_{it})$

7. In practice, irrigation wells are extremely long-lived and permanent exit from irrigation is rare. While some individual producers do exit, fields typically continue irrigated production under new ownership.

are given by

$$V(s_{it}) = \mathbb{E}[\max\{v(0; s_{it}) + \epsilon^e_{it0}, v(1; s_{it}) + \epsilon^e_{it1}\}] \tag{9}$$

$$v(0; s_{it}) = \mathbb{E}[\pi^0_{it} \mid s_{it}] + \beta E[V(s_{i,t+1}) \mid s_{it}] \tag{10}$$

$$v(1; s_{it}) = \mathbb{E}\left[\sum_{\tau=0}^{\infty} \beta^\tau \max\{\pi^0_{i,t+\tau}, \pi^1_{i,t+\tau}\} \mid s_{it}\right] - c^e_{it}(s_{it}) \tag{11}$$

where π^0_{it} and π^1_{it} give rainfed and irrigated profits respectively.

Equation 11 reflects the fact that irrigation investment is sunk, so the choice-specific value of entry is simply the expected discounted present value of flow profits net of the entry cost. Potential entrants who choose not to enter receive rainfed profits but pay no entry cost and face the same decision in the next period. An implicit assumption of equation 11 is that individual fields are atomic so that their individual decisions do not affect the evolution of the state variables. This assumption still allows for the aggregate decisions of many fields to endogenously affect the evolution of state variables but rules out any strategic interactions in order to maintain finite dependence (Arcidiacono and Miller 2011). Finally, the assumption of logit shocks implies choice probabilities given by

$$p^e(s_{it}) = \frac{\exp(v(1; s_{it})/\sigma^e)}{\exp(v(0; s_{it})/\sigma^e) + \exp(v(1; s_{it})/\sigma^e)} \tag{12}$$

4.2.1 Entry Costs

The main fixed cost associated with adopting irrigation is drilling an irrigation well. We specify entry costs as

$$c^e_{it} = x^e_{it}\alpha^e + \mathbb{1}\{t \geq 1978\}\alpha^e_{GMD} + \alpha^e_t t + \alpha_{c(i)} \tag{13}$$

where the vector of controls x^e_{it} includes a constant, bedrock depths, and field size, $\mathbb{1}\{t \geq 1978\}$ is an indicator for active GMD regulations, and $c(i)$ gives the county of field i. The cost function specified in equation 13 allows costs to increase with well depth, which is typically equal to the bedrock depth of the aquifer, as well as field size to account for increased costs from deeper wells and larger irrigation systems. The inclusion of the GMD regulation indicator allows costs to potentially increase due to increased regulatory oversight from GMDs, which began in 1978.[8] We also include a linear time trend to capture changing

8. A restriction of this cost function is that is does not allow for the costs of regulatory oversight to change over time. We hope to accommodate this feature in future work. The main concern is that obtaining new water rights has become more difficult over time, a process which is captured in equation 13 only through the linear time trend. However, this effect largely impacts the set of potential fields, which we take to be

costs over time and a county-level fixed effect to account for unobserved heterogeneity at the regional level, particularly regulatory costs such as filing and inspection fees.

4.3 Equilibrium

Equilibrium in our model is driven by the path of aquifer depths. As more fields drill wells, future potential entrants rationally expect the aquifer to deplete faster, which reduces future rents from irrigation and lowers entry probabilities. Entry occurs until the net present value of rents from irrigation equals the fixed cost of entry.

To model this equilibrium dynamic while avoiding the severe curse of dimensionality that would result from a full Markov Perfect Equilibrium, we adopt an oblivious equilibrium approach (Weintraub, Benkard, and Van Roy 2008). In particular, rather than tracking the full industry state,[9] fields instead track only their field-specific states. The impact of other fields decisions enter only through field beliefs about future aquifer depletion, and we assume fields have rational expectations that are consistent with the realized path of aquifer depths.

A full dynamic equilibrium is given by paths of crop choices, water usage, field entry probabilities, and aquifer depths $\{\vec{j}_t, \vec{W}_t, \vec{p}_t^{\,e}, \vec{d}_t\}_{t=0}^{\infty}$ that satisfy

1. Static profit maximization: Fields choose crop choice j_{it} and water usage W_{it} to maximize static profits π_{it}

2. Rational expectations: Fields beliefs about the evolution of aquifer depths d_{it} are consistent with the realized path of aquifer depths d_{it}[10]

3. Dynamic consistency: Field entry probabilities are given by equation 12

4. Hydrologic consistency: Aquifer depths $\vec{d}_t$ evolve according to the aquifer law of motion given realized water use $\vec{W}_t$

5 Estimation

In this section we describe how we estimate the parameters of the supply model specified in section 4. To estimate the intensive margin parameters we adopt a two-stage procedure

exogenous, as obtaining a water right is a prerequisite for drilling an irrigation well.

9. The full industry state includes the state vector for each field as well as the aquifer characteristics necessary to compute lateral flows of water.

10. For now we are assuming a particularly stark version of this: each field believes with certainty that aquifer depths will decrease by Δd_t, where Δd_t is the average decrease across the aquifer. This is purely to aid computation, we plan to relax this in future versions to capture local variation in aquifer depletion and non-degenerate beliefs.

similar to Rafey (2023). In the first stage we estimate labor and materials elasticities using moments derived from the first-order conditions for static profit maximization, which simplifies to a regression equation given the nested Cobb-Douglas functional form assumed in equation 5. Then in the second stage we estimate the parameters of the water-yield function separately for rainfed and irrigated crops using moment conditions derived from the implied productivity innovations. To estimate the extensive margin parameters we take a discrete Euler Equation approach similar to Kalouptsidi, Scott, and Souza-Rodrigues (2021) and Hsiao (2025), leveraging the one-period finite dependence structure of the entry decision model.

5.1 Intensive Margin

The main challenge associated with estimating the production function specified in equation 5, relative to the standard literature on production function estimation (see Olley and Pakes 1996 and the literature that followed), is that we only observe factor inputs at the field level for irrigation water usage W_{it}. We observe the other flexible factor inputs, labor and materials, as well as revenue, only at the county-level from historical Censuses of Agriculture and USDA survey data. However, a key advantage of our setting is that we observe field-specific variation in water input prices for irrigated crops due to variation in aquifer depths, which greatly aids identification (Doraszelski and Jaumandreu 2013; Gandhi, Navarro, and Rivers 2020). As we detail below, this variation allows us to invert observed irrigation water factor demands to recover implied field-level revenues, productivity, and factor inputs, conditional on a guess of the intensive margin parameters $\theta_j \equiv (\lambda_j, \nu_j, \beta_{jL}, \beta_{jM})$.[11] Weather shocks provide further plausibly exogenous variation in inputs, which we exploit to identify the parameters of the water-yield function for rainfed crops for which there is no endogenous choice of water usage.

5.1.1 Inverting Water Demand

Fields that have invested in irrigation choose irrigation water as a flexible input to maximize static profits, setting marginal revenues equal to marginal costs. Given the production function specified by equation 5, the implied inverse conditional factor demand for irrigation

11. Exploiting duality in this way has a long history in production function estimation dating back to Marschak and Andrews (1944). More recent papers include Doraszelski and Jaumandreu 2013, Grieco, Li, and Zhang (2016), and Doraszelski and Jaumandreu (2018), which similarly exploit explicit inversions of parametric factor demand equations to estimate production functions in various contexts.

water is (see Appendix A for a derivation)

$$Y_{it} = c_{it} ET_{it} \Gamma_{it},\ \Gamma_{it} := \frac{\widetilde{W}_{it}}{\lambda_j \exp(-\lambda_j EW_{it})} \tag{14}$$

Since labor and materials are flexible inputs as well, we also have the normal Cobb-Douglas factor share equations

$$Y_{it} = \frac{L_{it}}{\beta_{jL}},\ Y_{it} = \frac{M_{it}}{\beta_{jM}} \tag{15}$$

Finally, productivities are given by the residual

$$\omega_{it}^{j} = \ln(Y_{it}) - \ln(P_{jt}) - \ln(\widetilde{W}_{it}) - \beta_{jL} \ln(L_{it}) - \beta_{jM} \ln(M_{it}) \tag{16}$$

Together, equations 14, 15, and 16 allow me to invert from observed irrigation water use and water extraction costs to recover $\hat{R}_{it}$, $\hat{L}_{it}$, $\hat{M}_{it}$, and $\hat{\omega}_{it}$ as a function of the parameters θ_j. This procedure is conceptually similar to the first-stage of Olley and Pakes (1996) and related estimation methods in that it inverts from factor demands to recover unobserved productivities, with two key differences. First, similar to Doraszelski and Jaumandreu (2013), we fully exploit the parametric structure of the model to recover productivity rather than relying on a non-parametric first stage. Second, we work with conditional factor demand functions rather than unconditional factor demand functions, leveraging the fact that the conditional demand functions can be directly inverted to recover field-specific revenues without the need for additional data beyond water usage and water costs. In Appendix A we show that, under standard assumptions, this key inversion step is possible for any homothetically separable production function and does not rely on the specific functional form of equation 5. The class of homothetically separable production functions includes most commonly specified production functions used in empirical work, including Cobb-Douglas, CES, and translog.

5.1.2 Water-Yield Estimation

We estimate the parameters of the water-yield function via two-step GMM, constructing micromoments based on productivity innovations in the spirit of Ackerberg, Caves, and Frazer (2015) and aggregate moments based on our county-level data. In this step we take as given estimates of the labor and materials elasticities $\vec{\beta}_j$, which we estimate in a first stage detailed below. We also estimate the irrigated and rainfed water-yield functions separately. Estimating separately provides computational advantages with no loss of efficiency because the moment conditions are uncorrelated.

Irrigated Water-Yield Function To estimate the parameters of the irrigated water-yield function, $\theta_{\text{irr}} = (\lambda_{\text{irr}}, \nu_{\text{irr}})$, we first recover the implied productivities $\hat{\omega}_{it}^{\text{irr}}$ and revenues $\hat{R}_{it}^{\text{irr}}$ as described in section 5.1.1. We then estimate equation 6 by OLS to recover the productivity residuals μ_{it}^{irr}.[12] Finally, we construct the productivity innovation $\hat{\xi}_{it}^{\text{irr}}$ as the residual of a regression of μ_{it}^{irr} on $\mu_{i,t-1}^{\text{irr}}$.

The micromoments constructed from the recovered productivity innovation that we use in estimation are of the form

$$\mathbb{E}[Z_{it}\xi_{it}^{\text{irr}}] = 0 \tag{17}$$

taking as instruments lagged water costs $c_{i,t-1}$, effective rainfall shocks ΔER_{it}, and inverse evapotranspiration shocks ΔET_{it}^{-1}, as well as their interactions.[13] Equation 20 allows the instruments to be correlated with the levels of productivity, requiring only that they are uncorrelated with the productivity innovation term ξ_{it}. This restriction is satisfied because weather shocks, which we construct by regressing effective rainfall and inverse ET on individual and year fixed effects, are exogenous and lagged water costs are pre-determined—that is, they are deterministic functions of $\mathcal{I}_{i,t-1}$ and by construction $\mathbb{E}[\xi_{it}|\mathcal{I}_{i,t-1}] = 0$.

In addition to these micromoments, we also construct an aggregate moment based on observed county-level revenues. In particular, we take observed county-level revenue yields as well as model-implied county-level revenue yields $\hat{Y}_{ct} = \frac{\sum_{i \in c} \hat{R}_{it}}{\sum_{i \in c} A_{it}}$ and form the moment

$$\mathbb{E}[Y_{ct} - \hat{Y}_{ct}] = 0 \tag{18}$$

This aggregate moment aids identification because it directly incorporates our county-level output data, while the micromoments are constructed entirely from our field-level data on irrigation water use and water costs.

Rainfed Water-Yield Function We take a similar approach to estimate the parameters of the rainfed water-yield function $\theta_{\text{rainfed}} = (\lambda_{\text{rainfed}})$. Since we only observe county-level data for rainfed crops, we recover implied rainfed productivities using a county-level version

12. Even though our data is selected on irrigation, we can still consistently estimate the parameters of equation 6 via OLS because the assumption that irrigation decisions are made before observing productivity innovations implies that this selection is not correlated with the idiosyncratic productivity term. While in principle there may also be selection stemming from exiting irrigated agriculture, in practice almost all producers with irrigation wells in our data choose to irrigate their crops, as it provides a substantial increase in yields at little cost. An alternative approach would be supplement equation 6 with a selection correction as in Olley and Pakes (1996).

13. These moment conditions correspond to the second-stage moments used in Olley and Pakes (1996) and the literature that follows.

of equation 16 together with the Cobb-Douglas factor share equations 15:

$$\omega_{ct}^{j} = (1 - b^{j}) \ln(Y_{ct}^{j}) - \ln(P_{jt}) - \ln(B^{j}) - \ln(\widetilde{W}_{ct}) \tag{19}$$
$$b^{j} := \beta_{jL} + \beta_{jM}, B^{j} := \beta_{jL}^{\beta_{jL}} \beta_{jM}^{\beta_{jM}}$$

Equation 19 relies on our assumption that rainfed productivity is homogeneous within county-year.[14] Estimation then proceeds analogous to the irrigated water-yield function. We recover the implied productivity innovations $\xi_{ct}^{\text{rainfed}}$ by residualizing ω^{rainfed} on county and year fixed effects and construct moments of the form

$$\mathbb{E}[Z_{ct}\xi_{ct}^{\text{rainfed}}] = 0 \tag{20}$$

taking as instruments county-level effective rainfall shocks ΔER_{ct} and inverse ET shocks ΔET_{ct}^{-1} as well as their squares. The key identification assumption is that the productivity innovations $\xi_{ct}^{\text{rainfed}}$ are uncorrelated with the weather shocks ΔER_{ct} and ΔET_{ct}^{-1}.

5.1.3 Labor and Materials Estimation

We estimate the labor and materials elasticities $\vec{\beta} = (\beta_{jL}, \beta_{jM})_{j \in \{\text{irrigated,rainfed}\}}$ using the factor share equations given by equation 15, which are implied by static profit maximization. Because of the Cobb-Douglas functional form assumption in equation 5, this amounts to a simple regression equation. Since our data on revenues is aggregated to the county-level and our data on expenditures is aggregated both to the county-level and across all crops, we construct a corresponding county-level estimating equation. In particular, we have that total revenue for crop j in county c in year t is given by

$$\beta_{jL} R_{ct}^{j} = X_{ct}^{j},\ X_{ct}^{j} \in \{L_{ct}^{j}, M_{ct}^{j}\} \tag{21}$$

where in a slight abuse of notation L_{ct}^{j} and M_{ct}^{j} are total expenditures rather than per acre expenditures on labor and materials for crop j in county c in year t. Summing over crops gives

$$X_{ct} = \sum_{j} \beta_{jL} R_{ct}^{j} \tag{22}$$

We therefore estimate the parameters $\vec{\beta}_X$ using the following estimating equation

$$X_{ct} = \beta_X^{\text{irr}} R_{ct}^{\text{irr}} + \beta_X^{\text{rainfed}} R_{ct}^{\text{rainfed}} + \beta_X^{\text{other}} R_{ct}^{\text{other}} + \epsilon_{ct} \tag{23}$$

14. For simplicity we also assume that rainfed effective rainfall and evapotranspiration are homogeneous within county-year by taking the average across all fields in the county.

where the covariate R_{ct}^{other} captures the revenue from farms not growing field crops (either irrigated or rainfed) that are nonetheless present in our expenditure data from the Census of Agriculture. These farm operations are typically livestock farms—the implicit assumption in equation 23 is that these farms also produce according to a production function that is Cobb-Douglas in labor and materials and moreover that these factors are also flexible inputs.

Relative to estimating jointly with the second-stage parameters, estimating the labor and materials output elasticities in a separate first stage provides a computationally simple estimator with minimal loss of efficiency since there is no covariance between the county-level moments of equation 23 and the micromoments that provide the main identifying variation for the parameters of the water-yield function. Similar to Rafey (2023), equation 23 also fully exploits our assumption of static profit maximization, avoiding the need for an instrument.

5.1.4 Identification

Our estimation approach uses data on a subset of inputs and their costs plus aggregate data on output and the other inputs to estimate the parameters of the intensive margin production functions. The key step is that, conditional on the parameters, we are able to invert from observed irrigation water use and water costs to recover implied field-level revenues and productivity as well as other input decisions. The intuition is simple: for any homothetically separable production function the ratio of factor demands is a function only of the input costs while the unobserved productivity ω determines only the scale of production. Since field owners face homogeneous labor and materials costs, variation in field-level water costs and water usage along with weather shocks is enough to separately identify the parameters of the model and the unobserved productivities. Intuitively, if a field owner faces high water costs or already has high effective rainfall and still uses large amounts of irrigation water, they must also have high productivity.

Our estimation procedure is closely related to the seminal approach of Olley and Pakes (1996) and the literature that followed. We avoid problems of functional dependence because we have variation in water costs due to differences in aquifer depths and variation in "effective" water (i.e. net of rainfall and relative to evapotranspiration) costs due to differences in climate and weather. As in Ackerberg, Caves, and Frazer (2015), we concentrate out the productivity residuals to recover the productivity innovation and construct moments based on that innovation term. We use as instruments weather shocks and lagged water costs. Although water costs are endogenously related to levels in productivity because of aquifer depletion, lagged water costs are pre-determined and thus uncorrelated with innovations in productivity, which is the key source of variation that we use in our estimation procedure.

5.2 Extensive Margin

To estimate the extensive margin we exploit the fact that entry into irrigation is a terminal action to generate a regression equation based on a discrete version of an Euler Equation (Kalouptsidi, Scott, and Souza-Rodrigues (2021); Hsiao (2025)). In particular, lemma 1 in Arcidiacano and Miller (2011) gives that

$$V(s_{it}) = v(1; s_{it}) - \sigma^e \ln(p^e(s_{it})) \tag{24}$$

Then from the Hotz and Miller (1993) inversion, we have

$$\begin{aligned} \sigma^e \ln\left(\frac{p^e(s_{it})}{1-p^e(s_{it})}\right) &= v(1; s_{it}) - v(0; s_{it}) \\ &= \mathbb{E}[\pi^1_{it} - \pi^0_{it} \mid s_{it}] - c^e_{it} + \beta \mathbb{E}[c^e_{i,t+1} + \sigma^e \ln(p^e(s_{i,t+1}))] \end{aligned} \tag{25}$$

where the second row follows because finite dependence implies that continuation values difference out. Equation 25 says that the log-difference in current period conditional choice probabilities can be constructed from current payoffs, one-period ahead costs, and one-period ahead conditional choice probabilities. This particularly simple structure results from the fact that entry into irrigation is a terminal action, implying one-period finite dependence since a field that enters in t and a field that enters in $t+1$ will be in the same state in $t+2$. This finite dependence structure also requires the assumption of independent, atomistic fields as in Scott (2023) since equilibrium dynamics can prevent future continuation values from canceling.[15]

Applying expectational errors and substituting estimated values to equation 25 gives

$$\ln\left(\frac{\hat{p}^e(s_{it})}{1-\hat{p}^e(s_{it})}\right) - \beta \ln(\hat{p}^e(s_{i,t+1})) = \left(\hat{\pi}^1_{it} - \hat{\pi}^0_{it} - c^e_{it} + \beta c^e_{i,t+1} + \eta^e\right)/\sigma^e \tag{26}$$

where η^e is an expectational error. Substituting the functional form for entry costs from equation 13 into 26 gives the regression equation

$$\ln\left(\frac{\hat{p}^e(s_{it})}{1-\hat{p}^e(s_{it})}\right) - \beta \ln(\hat{p}^e(s_{i,t+1})) = \frac{1}{\sigma^e}\Delta\pi_{it} + x^e_i\tilde{\alpha}^e + \mathbb{1}\{t \geq 1978\}\tilde{\alpha}^e_{GMD} + \tilde{\alpha}^e_t t + \tilde{\alpha}_{c(i)} + \frac{\eta^e}{\sigma^e} \tag{27}$$

where $\tilde{\alpha} \equiv \frac{(\beta-1)}{\sigma^e}\alpha$. To estimate equation 27 we flexibly compute entry probabilities in a first-stage using a probit regression of entry on a cubic monomial basis, using differences between irrigated and rainfed profits $\Delta\pi \equiv \pi^1 - \pi^0$, year t, well depths, latitude, longitude,

15. For instance, if one field non-atomistically affected prices then their entry would change the future state for all other fields, ruling out finite dependence.

and field sizes as basis terms. Since discount factors are typically unidentified in dynamic models, we also assume a discount factor of $\beta = 0.96$, corresponding to an interest rate of approximately 4% (Magnac and Thesmar 2002).

5.2.1 Discussion

Estimating via a discrete Euler Equation approach has a number of advantages relative to standard full-solution approaches based on nested fixed point methods (Rust 1987; Rust 1994). First, our approach is computationally light and does not require fully solving the model for estimation. Second, while we do have to assume rational expectations so that the expectational error η^e is uncorrelated with realized values, we do not need to make any additional, potentially strong assumptions about what those beliefs may be. Relative to full solution approaches, however, we cannot accommodate non-atomistic fields without breaking finite dependence. We believe the assumption of atomistic fields is reasonable in our setting since fields are indeed small relative to the size of the large agricultural markets in which they compete. Another potential disadvantage is that we cannot accommodate biased beliefs. For instance, if fields were overly optimistic about future aquifer depletion then this would load onto our estimated costs, resulting in an underestimate of entry costs.

6 Estimation Results

In this section we describe the main estimation results for both the intensive margin production function parameters and the extensive margin cost of entry parameters.

6.1 Intensive Margin

Table 2 shows estimation results for the parameters of the production function. Panel A describes the results for the irrigated crop production function. In specification (1) we include only the lagged water costs and weather shock instruments without any interactions or the aggregate yield moment. We find an average input demand elasticity of water of -1.50 and an average yield elasticity of 0.25. We also find an average yield gap, which is the ratio between realized and potential yields,[16] of 0.80. In specification (2) we additionally include the average yield moment and estimate that average yield gap is slightly higher at 0.88. Finally, in specification (3), which is our preferred specification, we additionally include all second-order instrument interactions. We find an average water input demand elasticity of -0.93, which is similar to existing estimates of groundwater input demand elasticities in other

16. Potential yields are the maximum yields a field could produce if their crops were not water-limited.

settings (Burlig, Preonas, and Woerman 2025). We also find an average yield elasticity of 0.17, suggesting that a 1% increase in applied irrigation water would increase yields by 0.17% on average given observed irrigation levels, as well as an average yield gap of 0.88. Together, these estimates imply that fields are practicing only limited deficit irrigation. That is, they are irrigating close to, although slightly below, the level required for reaching maximum crop yields. Finally, we find that the irrigation-rainfall substitutability coefficient is 0.35, suggesting that, in terms of yield production, approximately three feet of rainfall is equivalent to one foot of irrigation water.

Panel B of Table 2 describes the results for the rainfed crop production function. Similar to the irrigated crop production function, in specification (1) we include only a combined effective water shock instrument. In specification (2) we separate the effective water shock into an effective rainfall shock and an inverse evapotranspiration shock. Finally, in specification (3), which is our preferred specification, we include as instruments the effective rainfall and inverse evapotranspiration shocks as well as their squares. We find similar results across all specifications, and in our preferred specification we find an average water yield elasticity of 0.25 and an average yield gap of 0.89. Since rainfed crops, by definition, receive only rainfall and no irrigation water but still obtain similar yield gaps compared to irrigated crops, we see that rainfed crops have significantly lower water requirements to achieve high yields. This feature can also be seen by comparing the estimated effective water coefficient across rainfed and irrigated crops. For irrigated crops, we estimate an effective water coefficient of 4.42, or 1.55 in rainfall-equivalent terms. By contrast, for rainfed crops we estimate an effective water coefficient of 5.38, suggesting that the yield gap decreases about three times faster for rainfed crops than irrigated crops.

6.2 Extensive Margin

Table 3 shows the entry cost parameter estimation results. Specification (1) excludes the indicator for GMD regulations and finds an average entry cost of approximately $0.96M. Specification (2) includes this indicator and finds a similar average entry cost of $0.95M. In this specification we estimate that the GMD regulations impose an entry cost of approximately $35,000. Specification (3), which is our main specification, again includes an indicator for GMD regulations and also includes county fixed effects to account for unobserved heterogeneity at the regional level. We find average entry costs of $0.89M, with costs increasing by approximately $1,580/ft of well depth and and $1,578/acre of field size. At the same time, costs are falling by almost $7,000 per year. These costs are somewhat larger than engineering estimates for well drilling costs, which are in the range of a few hundred

Table 2: Estimation Results: Production Functions

Panel A: Irrigated			
	(1)	(2)	(3)
λ	3.24	4.44	4.42
	(0.042)	(0.008)	(0.008)
ν	0.36	0.34	0.35
	(0.005)	(0.005)	(0.004)
β_L	0.04	0.04	0.04
	(0.004)	(0.004)	(0.004)
β_M	0.22	0.22	0.22
	(0.004)	(0.004)	(0.004)
Avg. Water Elast.	-1.50	-0.89	-0.93
Avg. Water-Yield Elast.	0.25	0.18	0.17
Avg. Yield Gap	0.80	0.88	0.88
Observations	707,505	707,505	707,505
Fields	27,510	27,510	27,510
Years	64	64	64

Panel B: Rainfed			
	(1)	(2)	(3)
λ	6.14	5.84	5.38
	(0.796)	(0.681)	(0.544)
β_L	0.03	0.03	0.03
	(0.002)	(0.002)	(0.002)
β_M	0.22	0.22	0.22
	(0.013)	(0.013)	(0.013)
Avg. Water-Yield Elast.	0.20	0.22	0.25
Avg. Yield Gap	0.92	0.91	0.89
Observations	2,203	2,203	2,203
Counties	53	53	53
Years	47	47	47

Notes: This table shows the estimation results for the production function parameters. Standards errors are computed via the delta method. Yield gaps are measured as the ratio between estimated realized yields and maximium potential yields, which are the yields that would be realized if infinite water was applied so that the crop was not water limited. Panel A describes the estimates for the irrigated production function. Specification (1) gives results using lagged water costs and weather shocks as instruments. Specification (2) additionally includes the aggregate yield moment. Specification (3) is our preferred specification and also includes the second-order interactions as instruments. Panel B describes the estimates for the rainfed production function. Specification (1) gives results using effective water shocks as instruments. Specification (2) separates effective water shocks into effective rainfall and inverse evapotranspiration shocks. Specification (3) is our preferred specification and includes as instruments both effective rainfall and inverse evapotranspiration shocks as well as their squares.

thousand dollars (Tsoodle 2019), but we view this as reasonable since our entry costs are most accurately interpreted as the net present value of the initial fixed costs of irrigation, which includes both well drilling and the purchase of additional irrigation equipment, as well as all future flow costs of repairs and maintenance. Similar to specification (2), we estimate that GMD regulations impose an entry cost of approximately $31,000.

Table 3: Estimation Results: Entry Costs

	(1)	(2)	(3)
Constant	11,512,810	13,173,365	
	(4,294,916)	(4,292,715)	
σ^e	297,276	295,580	262,740
	(14,471)	(14,313)	(12,112)
Well Depth (ft.)	2,375	2,367	1,580
	(276)	(275)	(564)
Year	-5,757	-6,607	-6,942
	(2,205)	(2,204)	(1,999)
Acres	1,742	1,760	1,578
	(483)	(480)	(439)
GMD Active		34,521	30,996
		(8,751)	(7,780)
N	568,726	568,726	568,726
Average Entry Cost ($)	957,954	948,323	893,561
County Fixed Effects	No	No	Yes

Notes: This table shows the estimation results for the entry cost parameters. Standards errors are computed via the delta method. Specification (1) excludes the indicator for GMD regulations, specification (2) includes this indicator, and specification (3) both includes this indicator and county fixed effects.

7 Counterfactuals

In this section we describe the main counterfactuals that we simulate using our empirical model. We focus on two main types of counterfactual policies: intensive margin taxes and extensive margin entry fees. We consider both uniform policies that do not vary across space or time as well as more granular policies that can change over time to reflect changing conditions in the aquifer and the climate. We also consider the opportunity cost of policy delay and the costs of technological lock-in by varying the start year of each policy.

7.1 Solving the Model

For each counterfactual we initialize the model in the year 1960 and compute an equilibrium of the model under the counterfactual policy from 1960-2099.[17] As described in section 4.3, the equilibrium in our model is driven by paths of aquifer depths. To compute the equilibrium we start with a guess of the average depletion rate, forward simulate the model until 2099, and finally compute the realized average depletion rate. We then restart the procedure using this realized rate as our new guess of the average depletion rate and iterate until convergence. The full details of our equilibrium computation algorithm are provided in Appendix B.[18]

Fully solving for the model equilibrium also requires specifying the evolution of the exogenous variables following the sample period, which are prices $\vec{P}_t$, productivities $\vec{\omega}_{it}$, and weather ER_{it} and ET_{it}, as well as the corresponding expectations about these variables. We hold prices and productivities fixed at their observed levels in 2022, which is the final year of our sample period.[19] We assume field owners have myopic beliefs over prices and that they

17. We choose this terminal year to balance two potential sources of bias in our results. On the one hand, aquifer depletion is a long-run phenomenon, and even by 2099 we find that there are still some areas in the aquifer with a significant amount of remaining water. On the other hand, modeling long-run expectations involves significant additional assumptions necessary to extrapolate the exogenous variables of our model outside of our sample period.

18. One caveat here is that for the moment we still use a simplified crop choice decision that compares profits given realized weather rather than forming expectations over these variables. The key feature missed by this simplification is the insurance value of irrigation—if field owners knew in advance that the growing season will be rainy they would be more likely to grow rainfed crops. However, we believe this simplification is a reasonable approximation because even when comparing realized profits, we rarely observe fields with irrigation wells that choose to grow non-irrigated crops, reflecting the fact that the western portion of Kansas which overlies the High Plains Aquifer is semi-arid and rarely has rainy growing seasons. The main reason fields with irrigation wells switch from irrigated to rainfed crops is that they have depleted their groundwater resources, which occurs when saturated thickness ST_{it} falls below a minimum threshold $\underline{ST}_{it}$ that we set equal to 50 feet (Hecox, Macfarlane, and Wilson 2002).

19. Equations 6 and 8 already specify laws of motion for irrigated and rainfed productivities respectively, and we plan to use them to extend productivity levels into the future in future iterations. We similarly plan to implement a simple law of motion for prices.

assume productivities evolve according to the fitted values of equations 6 and 8 for irrigated and rainfed productivities respectively. For weather we use simulated weather outcomes from CMIP5 global climate model runs, in particular the CCSM4 model. Our baseline climate scenario is RCP4.5, which corresponds to approximately 2.5 to 3°C of warming, and we use RCP8.5, which corresponds to approximately 5°C of warming, for sensitivity analysis. Since we assume weather is not dynamic and is unobserved when making entry decisions, field owners do not directly account for realized weather when making their entry decisions. Instead, we specify a simple model for weather expectations based on a field's latitude, longitude, and the year, which accounts for both spatial variation in weather as well as temporal variation due to climate change. Details of our expected weather model are provided in Appendix B.

For each counterfactual policy τ we consider the social welfare

$$V_t^\tau = \sum_{t' \in T, t' \geq t} \left(\frac{1}{1+r} \right)^{t'-t} \left(\sum_{i \in \mathcal{F}} \pi_{it'}^\tau + G_{t'}^\tau \right)$$

where r is the social discount rate and G_t^τ gives the government revenue, if any, raised by the policy. We consider a range of discount rates $r \in [0.01, 0.05]$ and take as a baseline a discount rate of 2% (Carleton and Greenstone 2021). We choose the policy τ to maximize social welfare over the entirety of the model so that $\tau \in \arg\max_\tau V_{1960}^\tau$.

7.2 Results

We consider three main counterfactuals. First, we implement the welfare-maximizing uniform taxes and entry fees starting in 1960, the initial year of our model. Second, we implement the welfare-maximizing time-varying taxes and entry entry fees starting in 1960. We allow these time-varying policies to be adjusted once per decade from 1960-2020 to reflect changing conditions in the aquifer and climate change over the course of our sample period.[20] These counterfactuals demonstrate the value of more granular policy conditional on choice of policy instrument relative to the value of choosing a more complex policy instrument to begin with. Finally, we implement welfare-maximizing uniform taxes and entry fees that begin in 1978 rather than 1960. These counterfactuals firstly show the opportunity cost of policy inaction through increased lock-in, and secondly provide a natural point of comparison with the permit entry reforms described in section 2 that factually took place in 1978.

Table 4 gives an overview of the main results from each counterfactual, as well as the

20. We limit adjustments to once every decade to limit the dimensionality of the corresponding optimization problem, since evaluating the objective function requires solving for the full model equilibrium and is thus computationally expensive.

status quo outcome in column (1), which we take as a baseline. In the status quo there are a total of 27,060 wells constructed, resulting in discounted present value of $30.32B. Total water in the aquifer is reduced by more than half, with 174MAF remaining out of an initial total of almost 400MAF. In addition, more than half of the fields have locally depleted their groundwater resources, and as a result can no longer produce irrigated crops. Figure 6 shows the fraction of fields that are depleted in each year under our computed status quo equilibrium. At the end of our sample period approximately one-third of fields are depleted, and this number rises to nearly half by 2050. Field depletion slows in the second half of the century as increased costs drive lower water use, resulting in a final field depletion rate of 54% by the terminal year of 2099.

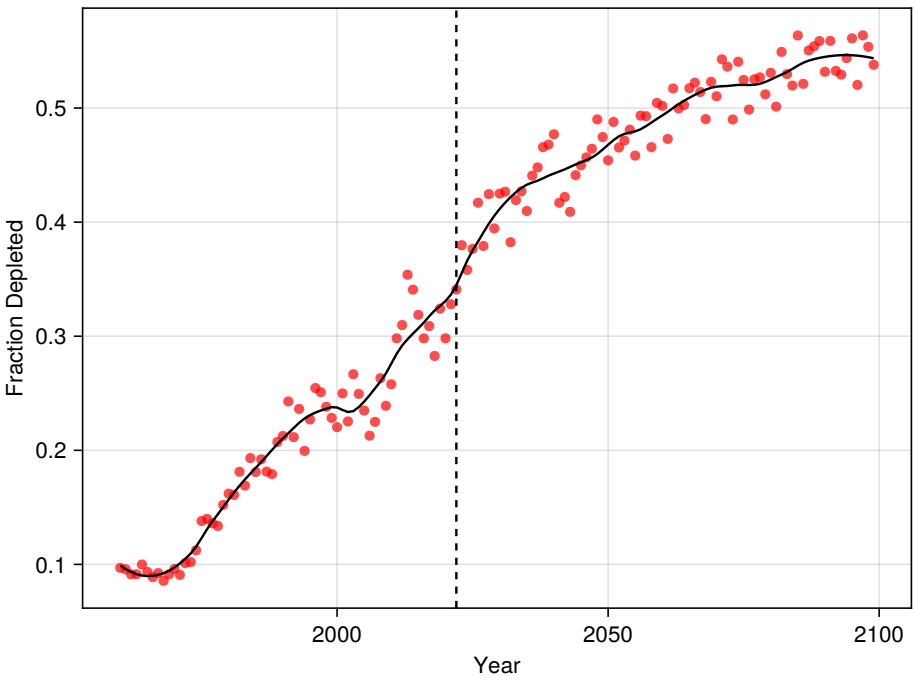

Figure 6: Time Path of Field Depletion

Notes: This figure shows the fraction of fields that are depleted by year under the computed status quo equilibrium. Each point is the fraction of depleted fields in a particular year and the solid black line is a smoothed LOESS local regression. The dashed black line is at the year 2022, which is the final year of our data. We measure field depletion by remaining saturated thickness in each year and consider a field to be depleted when saturated thickness falls below 50 feet, which is the minimum saturated thickness required to maintain pump output (Hecox, Macfarlane, and Wilson 2002).

Columns (2) and (3) solve for the welfare-maximizing uniform per AF tax and entry fee respectively. We find an optimal tax of $51.54/AF, resulting in a social welfare of $35.48B and approximately 400 fewer wells. In annualized, annuity-equivalent terms, the change in

welfare amounts to an increase of $107.92M per year. We find an optimal entry fee of $1M, resulting in a social welfare of $34.33B and approximately 14,000 fewer wells. The change in welfare amounts to an increase of approximately $84M per year, which is more than three-quarters of the change in welfare achieved by the per AF tax. Hence, when undertaken early we see that both per unit taxes and entry fees can significantly increase welfare. Indeed, both policies result in similar terminal field depletion rates of 35% and 37% respectively. Field depletion is critical for welfare because it forces field to produce rainfed crops at discretely lower yields than irrigated crops rather than to simply pay more for their irrigation water as is the case when aquifer depths increase but are not yet depleted. However, the mechanisms behind each policy differ significantly. Under the intensive margin tax, similar numbers of wells are drilled as under the status quo but significantly less water is used per well as fields now internalize the social cost of their aquifer depletion. By contrast, under the entry fee policy average water use per well actually increases because only the high-productivity wells still enter, but the number of wells drilled decreases by more than half.

Columns (4) and (5) solve for the welfare-maximizing time-varying per AF tax and entry fee respectively. We find an average optimal tax of $71.75/AF compared to $51.54/AF in the uniform case. The average tax increases under the time-varying scenario because externalities largely increase over time as the aquifer becomes more depleted and so the tax correspondingly increases. The resulting welfare is $36.24B, an increase of $124M/year in annualized terms, and the field depletion rate reduces to 28.75%. For comparison, we find an average optimal entry fee of $1.33M which results in a welfare of $34.9B, an increase of $96M/year in annualized terms. Unlike the time-varying tax, the optimal time-varying entry fee is actually U-shaped, with higher entry fees imposed at the beginning and the end. The reason for this is that it is valuable to both delay the entry of early entrants who have both high private surplus but also high externalities and also to prevent the entry of late entrants who have both low private surplus and low externalities. However, the entry fee still obtains the same relative gain in social welfare compared to the intensive margin tax as in the uniform case, showing that even though increased policy granularity is important for increasing welfare it does not change the comparison between entry fees and intensive margin taxes.

Columns (6) and (7) again solve for the welfare-maximizing uniform per AF tax and entry fee respectively. Unlike columns (1) and (2) which begin immediately, these policies are instead delayed until 1978 to match the timing of the existing 1978 reforms. We find an optimal tax of $57.90/AF, resulting in an overall welfare of $33.40B. Hence, relative to the tax beginning in 1960, the tax beginning in 1978 achieves only 60% of the gains. Moreover, the tax beginning in 1978 performs worse than the entry fee beginning in 1960, achieving only

77% of the gains of the latter policy, suggesting the importance of policy timing in settings with irreversibility. In the end, the opportunity cost of delaying policy by 18 years is larger than the losses from choosing an imperfect, second-best policy instrument but implementing it immediately. We discuss this trade-off further in section 7.2.1.

Finally, we find the welfare-maximizing entry fee to be \$0.86M, which is more than an order of magnitude larger than the actual entry cost impact of \$30,000 that we estimated for the existing permitting regulation in section 5. Welfare under the delayed entry permit is only \$31.05B, an annualized increase of \$15M/year relative to the status quo and less than 20% of the gains achieved by the immediately implemented entry fee. The reason for this is two-fold. First, by 1978 nearly two-thirds of potential entrants have already entered, locking in a large fraction of the status quo depletion. As a result, the terminal field depletion rate under the delayed entry fee reaches 46%. However, this is not the full story, as the delayed entry fee still results in 9,000 fewer wells constructed, similar to the 10,000 fewer wells constructed under the time-varying entry fee that results in \$34.90B in welfare. The second key reason that the delayed entry fee performs poorly is that spatial correlation in productivities and weather introduces a strong negative correlation between private surplus and externality levels. That is, many fields with high gains from irrigating, who thus enter early, are located close together are therefore impose large externalities due to the local nature of the aquifer. The delayed entry fee therefore targets not only fields who have low private surplus but also low externalities.

7.2.1 Policy Lock-In

Since recharge amounts in the High Plains Aquifer are low, averaging less than an inch per year, depletion is largely irreversible on reasonable timescales (Buchanan, Wilson, and Butler 2023). As a consequence, policy inaction carries a severe opportunity cost as depletion becomes "locked in" over time as additional fields enter. We investigate this possibility, as well as how it differentially affects taxes and entry fees by considering the welfare gains of optimal uniform taxes and entry fees as the start year of the policy becomes increasingly delayed.

Figure 7 shows the welfare gains by policy start year separately for taxes and entry fees. Unsurprisingly, taxes always provide larger welfare gains than entry fees. Indeed, by 2000 entry fees are essentially useless as most fields have already entered. And by 1978 entry fees are already worse than even extremely delayed taxes that start in 2010. Nonetheless, entry fees can still be effective policy instruments if they begin early—an entry fee beginning in 1960 is roughly equivalent to a tax beginning in 1975 and achieves more than 75% of the gains of a tax beginning in 1960.

Table 4: Counterfactual Results

		Uniform		Time-Varying		Delayed	
	Status Quo	Tax	Entry Fee	Tax	Entry Fee	Tax	Entry Fee
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Policy							
Avg. Tax ($/AF)	0.00	51.54	0.00	71.75	0.00	57.90	0.00
Avg. Entry Fee ($M)	0.00	0.00	1.00	0.00	1.33	0.00	0.86
Welfare							
Welfare ($B)	30.32	35.48	34.33	36.24	34.90	33.40	31.05
Annualized Δ Welfare ($M)	0.00	107.92	83.95	123.93	95.78	64.55	15.37
Num. Wells	27,060	26,667	13,267	27,251	16,708	26,786	18,048
Δ Num. Wells	0	-393	-13,793	191	-10,352	-274	-9,012
Aquifer							
Remaining Storage (MAF)	174.37	241.26	240.22	261.32	221.60	242.66	212.81
Depleted Fields (%)	53.79	34.77	37.30	28.75	41.95	32.69	46.23

Notes: This table shows the results of model simulations under the status quo (column 1), a uniform per AF tax beginning in 1960 (column 2), a uniform entry fee beginning in 1960 (column 3), a time-varying per AF tax beginning in 1960 (column 4), a time-varying entry fee beginning in 1960 (column 5), a uniform per AF tax beginning in 1978 (column 6), and a uniform entry fee beginning in 1978 (column 7). Average policy values are taken across all fields and years over the time horizon of the model. Annualized welfare changes are computed by converting the change in welfare relative to the status quo into an annuity equivalent.

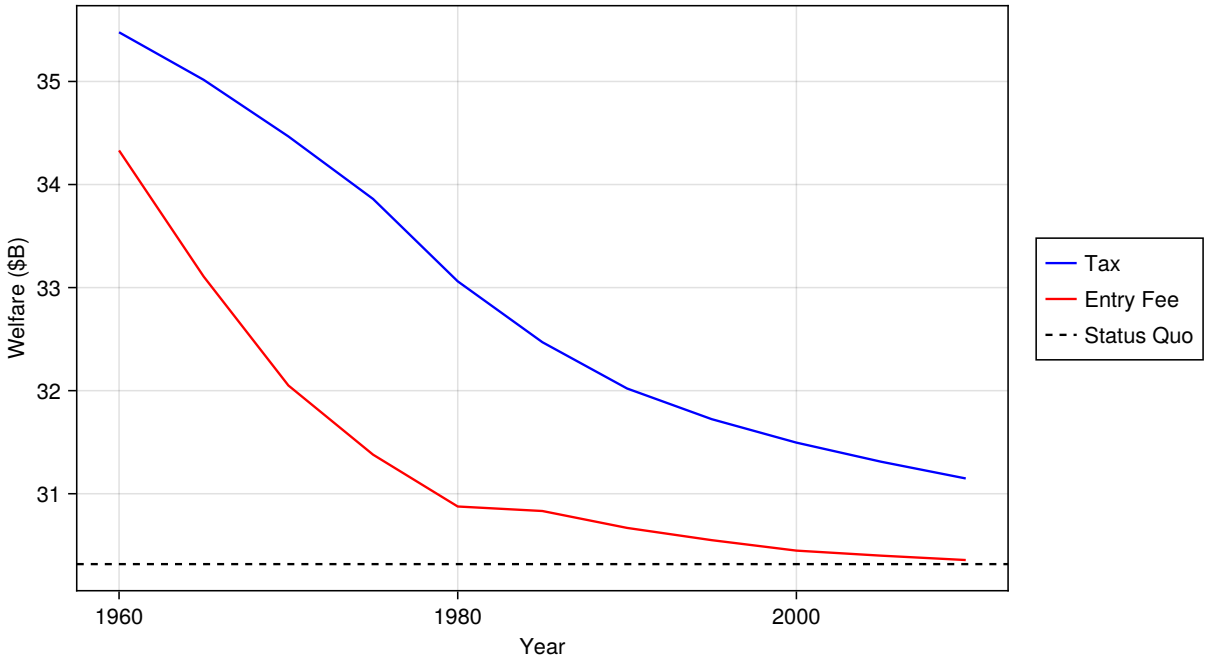

(a) Welfare from taxes and entry fees by start year

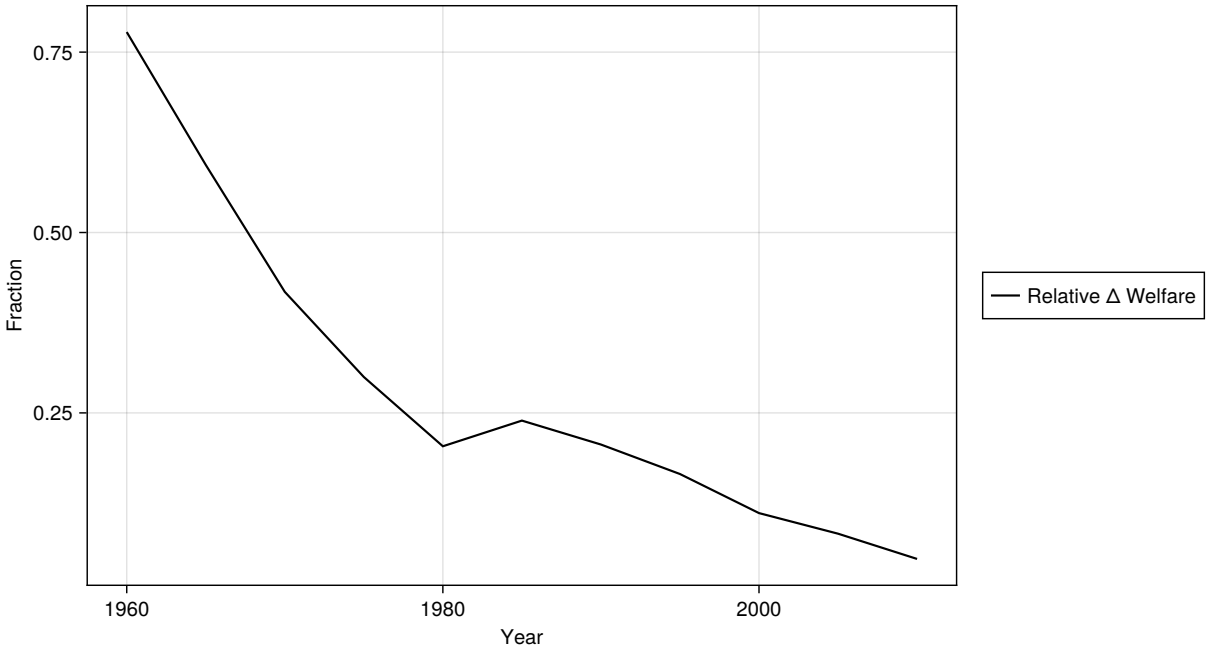

(b) Relative welfare gains from entry fees vs taxes by start year

Figure 7: Policy Lock-In

Notes: This figure shows the welfare resulting from uniform per AF taxes and entry fees by start year in panel (a) and the relative welfare gains in panel (b). The dashed line in panel (a) represents the status quo welfare.

8 Conclusion

This paper studies the welfare implications of second-best environmental regulation in the presence of technological irreversibility. To do so, we develop and estimate a dynamic structural model of agricultural irrigation that incorporates both intensive margin water-use decisions and extensive margin well-drilling decisions. We estimate the model using comprehensive administrative data from the Kansas High Plains Aquifer spanning 1959-2022. Counterfactual experiments demonstrate that extensive margin entry fees can achieve substantial welfare gains relative to the unregulated status quo, but their effectiveness relative to Pigouvian taxation declines sharply over time as depletion becomes locked in. In particular, we find that entry fees implemented in 1960 would have captured 75 percent of the welfare gains of Pigouvian taxation, declining to 25 percent by 1980 and approaching zero today. Our model predicts that under the status quo, more than half of fields will have depleted their groundwater resources by the end of the century. By contrast, even policy interventions beginning today can limit depletion to approximately one-third of fields, resulting in substantial welfare gains.

References

Abatzoglou, John T., Solomon Z. Dobrowski, Sean A. Parks, and Katherine C. Hegewisch. 2018. "TerraClimate, a high-resolution global dataset of monthly climate and climatic water balance from 1958–2015." *Scientific Data* 5, no. 1 (January): 170191.

Ackerberg, Daniel A., Kevin Caves, and Garth Frazer. 2015. "Identification Properties of Recent Production Function Estimators." *Econometrica* 83 (6): 2411–2451.

Allen, Richard G., Luis S. Pereira, Dirk Raes, and Martin Smith. 1998. *Crop evapotranspiration - Guidelines for computing crop water requirements.* Technical report 56. Rome: Food and Agriculture Organization of the United Nations.

Anderson, Soren T., Ryan Kellogg, and Stephen W. Salant. 2018. "Hotelling under Pressure." *Journal of Political Economy* 126, no. 3 (June): 984–1026.

Arcidiacono, Peter, and Robert A. Miller. 2011. "Conditional Choice Probability Estimation of Dynamic Discrete Choice Models With Unobserved Heterogeneity." *Econometrica* 79 (6): 1823–1867.

Ayres, Andrew B., Kyle C. Meng, and Andrew J. Plantinga. 2021. "Do Environmental Markets Improve on Open Access? Evidence from California Groundwater Rights." *Journal of Political Economy* 129, no. 10 (October): 2817–2860.

Barahona, Nano, Francisco A Gallego, and Juan-Pablo Montero. 2020. "Vintage-Specific Driving Restrictions." *The Review of Economic Studies* 87, no. 4 (July): 1646–1682.

Blomquist, William. 1987. "Getting Out of the Commons Trap: Variables, Process, and Results in Four Groundwater Basins," 70. Bloomington, Indiana: Indiana University.

Boyd, Robert, Peter J. Richerson, Ruth Meinzen-Dick, Tine De Moor, Matthew O. Jackson, Kristina M. Gjerde, Harriet Harden-Davies, et al. 2018. "Tragedy revisited." *Science* 362, no. 6420 (December): 1236–1241.

Brewer, Jebidiah, Robert Glennon, Alan P. Ker, and Gary D. Libecap. 2007. *Water Markets in the West: Prices, Trading, and Contractual Forms.* NBER Working Paper 13002. March.

Brill, Thomas C., and H. Stuart Burness. 1994. "Planning versus competitive rates of groundwater pumping." *Water Resources Research* 30 (6): 1873–1880.

Brown, Gardner. 1974. "An Optimal Program for Managing Common Property Resources with Congestion Externalities." *Journal of Political Economy* 82, no. 1 (January): 163–173.

Brozović, Nicholas, David L. Sunding, and David Zilberman. 2010. "On the spatial nature of the groundwater pumping externality." *Resource and Energy Economics,* Spatial Natural Resource and Environmental Economics, 32, no. 2 (April): 154–164.

Bruno, Ellen M., and Katrina Jessoe. 2021. "Using Price Elasticities of Water Demand to Inform Policy." *Annual Review of Resource Economics* 13, no. Volume 13, 2021 (October): 427–441.

Bruno, Ellen M., Katrina K. Jessoe, and W. Michael Hanemann. 2024. "The Dynamic Impacts of Pricing Groundwater." *Journal of the Association of Environmental and Resource Economists* 11, no. 5 (September): 1201–1227.

Buchanan, Rex C., Brownie B. Wilson, and James J. Butler Jr. 2023. *The High Plains Aquifer.* Technical report 18. Lawrence, KS: Kansas Geological Survey, January.

Burlig, Fiona, Louis Preonas, and Matt Woerman. 2025. "Groundwater and Crop Choice in the Short and Long Run." February.

Burt, Oscar R. 1964. "Optimal Resource Use Over Time with an Application to Ground Water." *Management Science* 11, no. 1 (September): 80–93.

Carleton, Tamma, Levi Crews, and Ishan Nath. 2025. "Agriculture, Trade, and the Spatial Efficiency of Global Water Use." April.

Carleton, Tamma, and Michael Greenstone. 2021. "Updating the United States Government's Social Cost of Carbon." *SSRN Electronic Journal.*

Coase, R. H. 1960. "The Problem of Social Cost." *The Journal of Law and Economics* 3 (October): 1–44.

Coman, Katharine. 1911. "Some Unsettled Problems of Irrigation." *American Economic Review* 1, no. 1 (March): 1–19.

Dasgupta, Partha, and Geoffrey Heal. 1974. "The Optimal Depletion of Exhaustible Resources." *The Review of Economic Studies* 41 (Symposium on the Economics of Exhaustible Resources): 3–28.

Donna, Javier D. 2021. "Measuring long-run gasoline price elasticities in urban travel demand." *The RAND Journal of Economics* 52 (4): 945–994.

Doraszelski, Ulrich, and Jordi Jaumandreu. 2013. "R&D and Productivity: Estimating Endogenous Productivity." *The Review of Economic Studies* 80 (4 (285)): 1338–1383.

———. 2018. "Measuring the Bias of Technological Change." *Journal of Political Economy* 126, no. 3 (June): 1027–1084.

Edwards, Eric C., and Todd Guilfoos. 2021. "The Economics of Groundwater Governance Institutions across the Globe." *Applied Economic Perspectives and Policy* 43, no. 4 (December): 1571–1594.

Edwards, Eric C., and Steven M. Smith. 2018. "The Role of Irrigation in the Development of Agriculture in the United States." *The Journal of Economic History* 78, no. 4 (December): 1103–1141.

Fenichel, Eli P., Joshua K. Abbott, Jude Bayham, Whitney Boone, Erin M. K. Haacker, and Lisa Pfeiffer. 2016. "Measuring the value of groundwater and other forms of natural capital." *Proceedings of the National Academy of Sciences* 113, no. 9 (March): 2382–2387.

Ferguson, Billy. 2024. "Trade Frictions in Surface Water Markets." November.

Fowlie, Meredith. 2010. "Emissions Trading, Electricity Restructuring, and Investment in Pollution Abatement." *American Economic Review* 100, no. 3 (June): 837–869.

Fowlie, Meredith, Mar Reguant, and Stephen P. Ryan. 2016. "Market-Based Emissions Regulation and Industry Dynamics | Journal of Political Economy: Vol 124, No 1." *Journal of Political Economy* 124, no. 1 (February): 249–302.

Gandhi, Amit, Salvador Navarro, and David A. Rivers. 2020. "On the Identification of Gross Output Production Functions." *Journal of Political Economy* 128, no. 8 (August): 2973–3016.

Gardner, Roy, Elinor Ostrom, and James M. Walker. 1990. "The Nature of Common-Pool Resource Problems." *Rationality and Society* 2, no. 3 (July): 335–358.

Gisser, Micha, and David A. Sánchez. 1980. "Competition versus optimal control in groundwater pumping." *Water Resources Research* 16, no. 4 (August): 638–642.

Gray, Wayne B., and Ronald J. Shadbegian. 1998. "Environmental Regulation, Investment Timing, and Technology Choice." *The Journal of Industrial Economics* 46 (2): 235–256.

Grieco, Paul L. E., Shengyu Li, and Hongsong Zhang. 2016. "Production Function Estimation with Unobserved Input Price Dispersion." *International Economic Review* 57 (2): 665–689.

Guilfoos, Todd, Neha Khanna, and Jeffrey M. Peterson. 2016. "Efficiency of Viable Groundwater Management Policies." *Land Economics* 92, no. 4 (November): 618–640.

Gutentag, Edwin, Frederick Heimes, Richard Luckey, and John Weeks. 1984. "Geohydrology of the High Plains Aquifer In Parts of Colorado, Kansas, Nebraska, New Mexico, Oklahoma, South Dakota, Texas, and Wyoming." *United States Geological Survey: Staff Publications* (January).

Hagerty, Nick. 2024. "Transaction Costs and the Gains from Trade in Water Markets." December.

Haines, Michael, Price Fishback, and Paul Rhode. 2018. *United States Agriculture Data, 1840 - 2012,* August.

Hansen, Zeynep K., Gary D. Libecap, and Scott E. Lowe. 2009. *Climate Variability and Water Infrastructure: Historical Experience in the Western United States.* Working Paper, December.

Hardin, Garrett. 1968. "The Tragedy of the Commons." *Science* 162, no. 3859 (December): 1243–1248.

Hecox, G. R., P. A. Macfarlane, and B. B. Wilson. 2002. *Calculation of Yield for High Plains Wells: Relationship between saturated thickness and well yield.* Technical Report OFR 2002-25C. Kansas Geological Survey.

Hendricks, Nathan P., and Jeffrey M. Peterson. 2012. "Fixed Effects Estimation of the Intensive and Extensive Margins of Irrigation Water Demand." *Journal of Agricultural and Resource Economics* 37 (1): 1–19.

Hexem, Roger W., and Earl Orel Heady. 1978. *Water Production Functions for Irrigated Agriculture.* Iowa State University Press.

Hopenhayn, Hugo A. 1992. "Entry, Exit, and firm Dynamics in Long Run Equilibrium." *Econometrica* 60, no. 5 (September): 1127.

Hornbeck, Richard, and Pinar Keskin. 2014. "The Historically Evolving Impact of the Ogallala Aquifer: Agricultural Adaptation to Groundwater and Drought." *American Economic Journal: Applied Economics* 6, no. 1 (January): 190–219.

Hotelling, Harold. 1931. "The Economics of Exhaustible Resources." *Journal of Political Economy* 39, no. 2 (April): 137–175.

Hsiao, Allan. 2025. "Coordination and Commitment in International Climate Action: Evidence from Palm Oil." May.

Jasechko, Scott, Hansjörg Seybold, Debra Perrone, Ying Fan, Mohammad Shamsudduha, Richard G. Taylor, Othman Fallatah, and James W. Kirchner. 2024. "Rapid groundwater decline and some cases of recovery in aquifers globally." *Nature* 625, no. 7996 (January): 715–721.

Judd, Kenneth L., Lilia Maliar, and Serguei Maliar. 2011. "Numerically stable and accurate stochastic simulation approaches for solving dynamic economic models: Approaches for solving dynamic models." *Quantitative Economics* 2, no. 2 (July): 173–210.

Judd, Kenneth L., Lilia Maliar, Serguei Maliar, and Inna Tsener. 2017. "How to solve dynamic stochastic models computing expectations just once: How to solve dynamic stochastic models." *Quantitative Economics* 8, no. 3 (November): 851–893.

Judd, Kenneth L., Lilia Maliar, Serguei Maliar, and Rafael Valero. 2014. "Smolyak method for solving dynamic economic models: Lagrange interpolation, anisotropic grid and adaptive domain." *Journal of Economic Dynamics and Control* 44 (July): 92–123.

Kalouptsidi, Myrto, Paul T. Scott, and Eduardo Souza-Rodrigues. 2021. "Linear IV regression estimators for structural dynamic discrete choice models." *Journal of Econometrics* 222, no. 1, Part C (May): 778–804.

Kansas Farm Statistics, 1920-2023, Selected Years. 2024. Technical report 58. Lawrence, KS: Institute for Policy & Social Research, September.

Kansas Geological Survey. 2025. *Master Well Inventory,* January.

KDWR. 2024. *Water Information Management and Analysis System (WIMAS).*

Kellogg, Ryan. 2024. *The End of Oil.* Working Paper, November.

Kim, C. S, Michael R Moore, John J Hanchar, and Michael Nieswiadomy. 1989. "A dynamic model of adaptation to resource depletion: theory and an application to groundwater mining." *Journal of Environmental Economics and Management* 17, no. 1 (July): 66–82.

Koundouri, Phoebe. 2004. "Potential for groundwater management: Gisser-Sanchez effect reconsidered." *Water Resources Research* 40 (6).

Layzell, Anthony L., and Catherine S. Evans. 2013. *Kansas Droughts: Climatic Trends Over 1,000 Years.* Technical report 35. Lawrence, KS: Kansas Geological Survey, August.

Libecap, Gary D. 2010. "Water Rights and Markets in the U.S. Semi Arid West: Efficiency and Equity Issues." *SSRN Electronic Journal.*

Libecap, Gary D., and Ariel Dinar, eds. 2023. *American Agriculture, Water Resources, and Climate Change.* National Bureau of Economic Research Conference Report. Chicago, IL: University of Chicago Press.

Magnac, Thierry, and David Thesmar. 2002. "Identifying Dynamic Discrete Decision Processes." *Econometrica* 70, no. 2 (March): 801–816.

Marschak, Jacob, and William H. Andrews. 1944. "Random Simultaneous Equations and the Theory of Production." *Econometrica* 12, nos. 3/4 (July): 143.

McGuire, V. L., M. R. Johnson, R. L. Schieffer, J. S. Stanton, S. K. Sebree, and I. M. erstraeten. 2003. *Water in Storage and Approaches to GroundWater Management, High Plains Aquifer, 2000.* Technical report 1243. Denver, CO: U.S. Geological Survey.

McGuire, Virginia L., and Kellan R. Strauch. 2024. *Water-Level and Recoverable Water in Storage Changes, High Plains Aquifer, Predevelopment to 2019 and 2017 to 2019.* Scientific Investigations Report 2023-5143. Reston, VA: U.S. Geological Survey.

Merrill, Nathaniel H., and Todd Guilfoos. 2018. "Optimal Groundwater Extraction under Uncertainty and a Spatial Stock Externality." *American journal of agricultural economics* 100, no. 1 (January): 220–238.

Murgai, Rinku. 2001. "The Green Revolution and the productivity paradox: evidence from the Indian Punjab." *Agricultural Economics* 25, nos. 2-3 (September): 199–209.

Negri, Donald H. 1989. "The common property aquifer as a differential game." *Water Resources Research* 25 (1): 9–15.

Olley, G. Steven, and Ariel Pakes. 1996. "The Dynamics of Productivity in the Telecommunications Equipment Industry." *Econometrica* 64 (6): 1263–1297.

Ostrom, Elinor. 1990. *Governing the Commons: The Evolution of Institutions for Collective Action.* Political Economy of Institutions and Decisions. Cambridge: Cambridge University Press.

Peck, John C. 2014. "Legal Responses to Drought in Kansas." *Kansas Law Review* 62 (May): 1141–1167.

Pfeiffer, Lisa, and C.-Y. Cynthia Lin. 2012. "Groundwater pumping and spatial externalities in agriculture." *Journal of Environmental Economics and Management* 64, no. 1 (July): 16–30.

———. 2014. "The Effects of Energy Prices on Agricultural Groundwater Extraction from the High Plains Aquifer." *American Journal of Agricultural Economics* 96, no. 5 (October): 1349–1362.

Pigou, Arthur. 1932. *The Economics of Welfare.* 4th. London: Macmillan.

Rafey, Will. 2023. "Droughts, Deluges, and (River) Diversions: Valuing Market-Based Water Reallocation." *American Economic Review* 113, no. 2 (February): 430–471.

Rogers, Danny H., and Mahbub Alam. 2006. *Comparing Irrigation Energy Costs.* Technical report MF2360. Kansas State University, August.

Rust, John. 1987. "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher." *Econometrica* 55 (5): 999–1033.

———. 1994. "Chapter 51 Structural estimation of markov decision processes." In *Handbook of Econometrics,* 4:3081–3143. Elsevier, January.

Ryan, Nicholas, and Anant Sudarshan. 2022. "Rationing the Commons." *Journal of Political Economy* 130, no. 1 (January): 210–257.

Scanlon, Bridget R., Claudia C. Faunt, Laurent Longuevergne, Robert C. Reedy, William M. Alley, Virginia L. McGuire, and Peter B. McMahon. 2012. "Groundwater depletion and sustainability of irrigation in the US High Plains and Central Valley." *Proceedings of the National Academy of Sciences* 109, no. 24 (June): 9320–9325.

Scott, Paul T. 2023. "Dynamic Discrete Choice Estimation of Agricultural Land Use." December.

Sears, Louis, C.-Y. Cynthia Lin Lawell, Gerald Torres, and M. Todd Walter. 2023. "Managing Common Pool Resources: Lessons from Groundwater Resource Extraction in California." June.

Soil Survey Staff. 2023. *Gridded Soil Survey Geographic (gSSURGO) Database for Kansas.* Technical report. United States Department of Agriculture, Natural Resources Conservation Service, November.

Solow, Robert M. 1974. "The Economics of Resources or the Resources of Economics." *American Economic Review* 64, no. 2 (May): 1–14.

Stanton, Jennifer S., Sharon L. Qi, Derek W. Ryter, Sarah E. Falk, Natalie A. Houston, Steven M. Peterson, Stephen M. Westenbroek, and Scott C. Christenson. 2011. *Selected approaches to estimate water-budget components of the High Plains, 1940 through 1949 and 2000 through 2009.* Technical report 2011-5183. U.S. Geological Survey.

Steduto, Pasquale, Theodore C. Hsiao, Elias Fereres, and Dirk Raes. 2012. *Crop yield response to water.* Technical report 66. Rome: Food and Agriculture Organization of the United Nations.

Stiglitz, Joseph. 1974. "Growth with Exhaustible Natural Resources: Efficient and Optimal Growth Paths." *The Review of Economic Studies* 41:123.

Stover, Susan, and Rex Buchanan. 2017. "The High Plains Aquifer: Can We Make It Last?" *GSA Today* 27:44–45.

Tsoodle, Leah J. 2019. *2017 Irrigation Equipment Cost Survey in Kansas.* Technical report. Manhattan, KS: Kansas State University, January.

USDA NASS. 2024. *Quick Stats.*

Water Use Reporting. 2025.

Weintraub, Gabriel Y., C. Lanier Benkard, and Benjamin Van Roy. 2008. "Markov Perfect Industry Dynamics With Many Firms." *Econometrica* 76 (6): 1375–1411.

Wilson, Brownie. 2024. *WIZARD Water Well Levels Database,* April.

A Proofs

A.1 Derivation of Inverse Conditional Factor Demand

Consider a farmer growing irrigated crops. They solve

$$\max_{W,L,M} A_{it}\left[\exp(\omega_{it}^{j})P_{jt}L^{\beta_{jL}}M^{\beta_{jM}}(1-\exp(-\lambda_j EW_{it}))-c_{it}W-L-M\right]$$

where $EW_{it} \equiv \frac{W+\nu_j ER_{it}}{ET_{it}}$. The first order condition with respect to W is then given by

$$\exp(\omega_{it}^{j})P_{jt}L^{\beta_{jL}}M^{\beta_{jM}}\lambda_j \exp(-\lambda_j EW_{it}) = c_{it}ET_{it}$$

Multiplying both sides by $\widetilde{W}_{it} \equiv 1-\exp(-\lambda_j EW_{it})$ and substituting using equation 5 yields

$$\begin{aligned} Y_{it} &= c_{it}ET_{it}\Gamma_{it} \\ \Gamma_{it} &:= \frac{\widetilde{W}_{it}}{\lambda_j \exp(-\lambda_j EW_{it})} \end{aligned}$$

which is the expression in the text.

A.2 Proof of Invertibility

Lemma 1. *Consider the revenue production function $R = \Omega F(\vec{X})$ and the associated conditional factor demand function $X_i^* = g_i(R, \vec{c}\,; \Omega)$. If F is homothetic and X_i is a flexible input then g_i is invertible and g_i^{-1} is independent of Ω.*

Proof. Our proof is constructive. Shephard (1953, p. 47) shows that the conditional factor demand function factors

$$g_i(R, \Omega, \vec{c}) = f(R, \Omega)\Gamma_i(\vec{c})$$

and the relative factor demands are therefore given by

$$\frac{X_i^*}{X_j^*} = \frac{\Gamma_i}{\Gamma_j}$$

Moreover, by definition of homotheticity, we can write

$$R = \Omega F(\vec{X}) = \Omega\Phi(\sigma(\vec{X}))$$

for some positive function $\sigma(\cdot)$ that is homogeneous of degree one and some positive, strictly increasing function Φ. Because relative factor demands are constant, we can rewrite the

firm's profit maximization problem as

$$\max_{X_1} \Omega F\left(\sigma\left(X_1, X_1\frac{\Gamma_2}{\Gamma_1}, ..., X_1\frac{\Gamma_N}{\Gamma_1}\right)\right) - X_1\sum_i c_i\left(\frac{\Gamma_i}{\Gamma_1}\right)$$
$$\max_{X_1} \Omega F\left(X_1\sigma\left(1, \frac{\Gamma_2}{\Gamma_1}, ..., \frac{\Gamma_N}{\Gamma_1}\right)\right) - X_1\sum_i c_i\left(\frac{\Gamma_i}{\Gamma_1}\right)$$

The first order condition for this problem is

$$\Omega\Sigma\Phi'(X_1\Sigma) = C$$

where $\Sigma = \sigma\left(1, \frac{\Gamma_2}{\Gamma_1}, ..., \frac{\Gamma_N}{\Gamma_1}\right)$ and $C = \sum_i c_i\left(\frac{\Gamma_i}{\Gamma_1}\right)$. Multiplying both sides by $\Phi(X_1\Sigma)$ gives

$$\begin{aligned}
&\Omega\Phi(X_1\Sigma)\Sigma\Phi'(X_1\Sigma) = C\Phi(X_1\Sigma)\\
&R\Sigma\Phi'(X_1\Sigma) = C\Phi(X_1\Sigma)\\
&R = \frac{C\Phi(X_1\Sigma)}{\Sigma\Phi'(X_1\Sigma)}
\end{aligned}$$

□

B Computation

In this section we provide the computational details required for solving the model and constructing counterfactual equilibria.

B.1 Value Function Approximation

Because the state space is moderately high-dimensional, we adopt the Smolyak sparse grid approximation methods of Judd et al. (2014) and pre-compute integrals using monomial quadrature integration (Judd, Maliar, and Maliar 2011; Judd et al. 2017). We take advantage of the binary action space to increase our approximation accuracy by separately approximating the pre-entry value function $V(s_{it})$ and the post-entry value function, which we denote $V^e(s_{it})$.

In particular, to approximate each value function we form a third-order[21] basis of orthogonal Chebyshev polynomials $\{\phi_n\}_{n\leq N}$ and consider the approximations $\hat{V}(s_{it}) \approx \vec{\phi}(s_{it}) \cdot \vec{w}$ and $\hat{V}^e(s_{it}) \approx \vec{\phi}(s_{it}) \cdot \vec{w}^e$ for weight vectors $\vec{w}$ and $\vec{w}^e$. Evaluating $\vec{\phi}$ at each of the corresponding N grid points of Chebyshev extrema forms the $N \times N$ Vandermonde matrix Φ, whose rows index grid points and columns index basis functions.

To recover the post-entry value function V^e we apply a result from Proposition 1 of Judd et al. (2017), which in our setting implies that $E[\hat{V}^e(s_{i,t+1}) \mid s_{it}] = (\tilde{\Phi}^e \odot B^e)\vec{w}^e$, where $\tilde{\Phi}^e$ is the Vandermonde matrix evaluated at an appropriately adjusted set of grid points,[22] B^e is a matrix of weight adjustments,[23] and $\odot$ is the element-wise Hadamard product. The key benefit of this result is that both $\tilde{\Phi}^e$ and B^e can be pre-computed, which greatly reduces the computational burden of computing the high-dimensional expectation over future states. It also increases accuracy, as the expectations are exact rather than numerically approximated with quadrature or Monte Carlo integration rules. Finally, we compute the vector of weights $\vec{w}^e$ to solve the approximated Bellman equation

$$\hat{V}^e(s_{it}) = \mathbb{E}[\max\{\pi^0, \pi^1\}] + \beta\mathbb{E}[\hat{V}^e(s_{i,t+1)}] \tag{28}$$

Equation 28 has the closed-form solution $\vec{w}^e = (\Phi - \beta(\tilde{\Phi}^e \odot B^e))^{-1}\vec{\pi}^e$, where $\vec{\pi}^e \equiv \mathbb{E}[\max\{\pi^0, \pi^1\}]$.

We follow a similar procedure to solve for the pre-entry value function. We first form the

21. Third-order here refers to the approximation order of the Smolyak selection rule, which prunes basis functions from the full tensor product of uni-dimensional basis functions by their importance for the quality of approximation. Under this selection rule the number of basis polynomials grows polynomially rather than exponentially in the dimension of the state space.

22. In particular, let the evolution of the state be given by the process $s_{i,t+1} = Z(s_{it}, \vec{\epsilon})$, where $\vec{\epsilon}$ is a vector of shocks. Then for each grid point z the adjusted grid point is given by $Z(z, \mu)$, where $\mu \equiv \mathbb{E}[\vec{\epsilon}]$.

23. In particular, each row corresponds to a grid point z with representative element $b_n(z) = \frac{\mathbb{E}[\phi_n(Z(z,\vec{\epsilon})]}{\phi_n(Z(z,\mu))}$

analogous adjusted Vandermonde matrix $\tilde{\Phi}$ and matrix of weight adjustments B. We then iterate on the approximated Bellman equation for each grid point z until convergence

$$\hat{V}(z) = \sigma^e \ln(\exp(\hat{v}(0;z)/\sigma^e) + \exp(\hat{v}(1;z)/\sigma^e)) \tag{29}$$

$$\begin{aligned}\hat{v}(0;z) &= \mathbb{E}[\pi^0_{it} \mid s_{it} = z] + \beta E[V(s_{i,t+1}) \mid s_{it} = z] \\ &= \mathbb{E}[\pi^0_{it} \mid s_{it} = z] + (\tilde{\Phi} \odot B)\vec{w} \end{aligned} \tag{30}$$

$$\begin{aligned}\hat{v}(1;z) &= \mathbb{E}[\max\{\pi^0, \pi^1\}] - c^e(z) + \beta\mathbb{E}[\hat{V}^e(s_{i,t+1)}] \\ &= V^e(s_{it}) - c^e(z) \end{aligned} \tag{31}$$

where the first line applies the closed-form for the expectation over type 1 extreme value random draws in equation 9.

B.2 Computing Equilibria

To compute counterfactual equilibria corresponding to a counterfactual policy τ we first compute the corresponding value functions V^τ and $V^{e,\tau}$. To avoid recomputing the value functions for each potential equilibrium path of aquifer depths, we augment the state space with the parameters of the depth process θ^{depth}.[24] To compute equilibria we first guess a path of aquifer depths and a corresponding value of θ^{depth}.[25] We then forward simulate the model as described in section 7.1 to compute the realized path of aquifer depths, which we use to update θ^{depth}, repeating until convergence.

Since the model equilibrium only provides entry probabilities for each field, we map each model equilibria to site-specific entry using a greedy equilibrium selection rule. In particular, in each period the n fields with the highest entry probabilities are selected to enter, where n is chosen so that the aggregate realized entry rate equals the aggregate ex ante entry probability. Using this greedy equilibrium selection rule provides a significant computational advantage compared to modeling the full binomial distribution of entry due to the number of fields.

C Model Fit

Figure 8 shows the cumulative number of wells that have been drilled in each year under both the status quo equilibrium computed by the model and in the data. The model somewhat

24. For now we assume that field owners believe depths evolve according to an exogenous process such that $\Delta d_{it} = \rho$, where ρ is the average change in depths, so $\theta^{\text{depth}} \equiv \rho$.

25. In practice we start at the status quo.

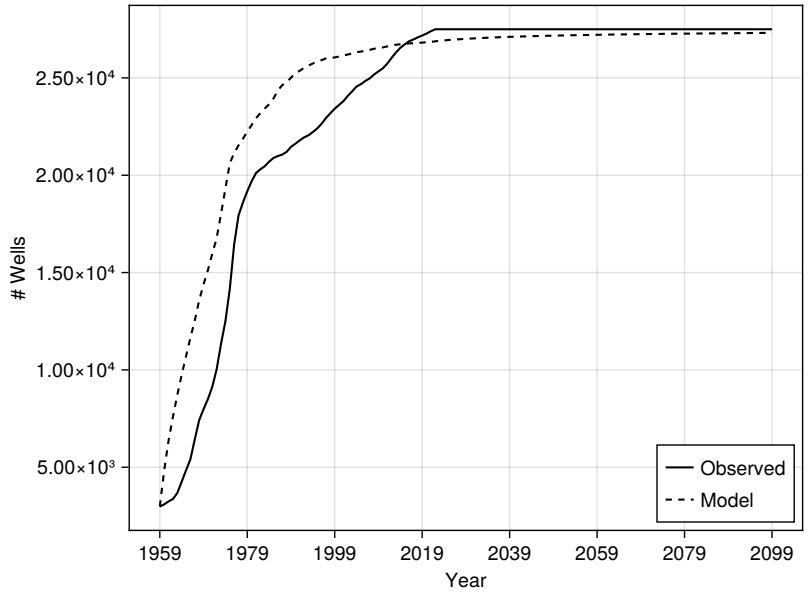

Figure 8: Number of Wells by Year

Notes: This figure shows the cumulative number of wells that have been drilled in each year of the model under the computed status quo equilibrium (dashed) and data (solid).

overestimates entry in the early periods of the model which we believe is largely due to the limited way in which we am currently modeling beliefs. In particular, beliefs about costs of entry are currently myopic (as in the field owners believe that they will have the same entry costs next year as they have this year) while our estimates suggests that costs of entry are dropping significantly, which would serve to delay entry.

D Data

We construct three main datasets. The first dataset is an annual field-level panel for the set of potentially irrigated fields overlying the High Plains aquifer from 1959-2022. For each field-year we observe whether or not the field has drilled an irrigation well as well as reported water usage and acres irrigated. We also construct measures of aquifer depths, water costs, weather, and prices as described below. We include as potentially irrigated fields all fields that we eventually observe irrigating in our reported water usage data, which we view as a weak restriction since most entry took place from 1960-1980 and the pace of new entry has slowed considerably in the decades since. This definition also implicitly assumes that each field that eventually enters is a potential entrant for each year of our sample. We believe this is a reasonable assumption since the extent of agricultural land in Kansas peaked in approximately 1950 with over 95% of state land used for agriculture. Land for agriculture has moderately declined in the decades since and now makes up about 85% of state land (*Kansas Farm Statistics, 1920-2023, Selected Years* 2024).

The second dataset we construct is an annual county-level panel of observed farm expenditures and production. We construct measures of farm expenditures on labor and materials using historical Censuses of Agriculture. We then combine this expenditure data with annual county-level survey data on irrigated and rainfed yields for the main crops grown in Kansas. The third and final dataset consists of hydrogeologic characteristics of the High Plains aquifer across space, which we use for our hydrogeologic model.

Table 5 overviews the full set of data sources we use for our analysis.

Table 5: Data Sources

Source	Period	Description
KDWR (2024)	1959-2024	Annual water usage and acres irrigated
USDA NASS (2024)	1972-2018	Annual county-level farm revenues
Haines, Fishback, and Rhode (2018)	1840-2012	Quinquennial county-level farm expenditures
Abatzoglou et al. (2018)	1958-2024	Monthly precipitation and evapotranspiration
Wilson (2024)	1900-2024	Annual High Plains aquifer water levels
Kansas Geological Survey (2025)	1894-2025	High plains aquifer well construction
Gutentag et al. (1984)	NA	High Plains aquifer characteristics
Stanton et al. (2011)	NA	High Plains aquifer recharge
Soil Survey Staff (2023)	NA	Kansas soil quality and texture

D.1 Field Panel

The main dataset we use to construct our field-level panel consists of water use reports made by individual water users to the Kansas Division of Water Resources (KDWR). For our main sample we consider reports made by groundwater irrigators in areas overlying the High Plains aquifer. Reports are made at the level of an individual well and report the water pumped and acreage irrigated in each year. Since 1990 reports have also included the crop irrigated by that well.

The water use reporting system began in 1954, although 1959 is the first year with a substantial number of recorded reports and we consider it to be the first year of our sample. Before 1989 reporting was voluntary, with estimated compliance rates of approximately 60%, and since 1989 reporting has been mandatory and estimated compliance rates have risen to 93% (*Water Use Reporting* 2025). Figure 9 shows the number of reports received in each year from 1959-2022, which is the last year of our sample. There were just under 1,000 reports in 1959. The number of reports increased rapidly until 1980, mirroring a rapid increase in the number of irrigators. Reporting spiked again in the years leading up to the beginning of mandatory reporting and has leveled off since.

Most wells are now required to be equipped with flow meters to accurately track pumped volumes. Wells without flow meters typically track volumes by converting from hours of well operation using the flow rate of the well. Figure 10 shows the fraction of reports received from metered wells in each year. The fraction of metered wells reporting actually declined over the course of the 1960's and 1970's as most newly constructed wells did not have flow meters. Since 1990 the number of metered wells increased rapidly and they are now essentially ubiquitous.

To construct our field panel, we define a field to be the area irrigated by a single well and measure the field size by acres irrigated. We adopt this definition because we only observe well locations rather than directly observing field boundaries, but fields are unlikely to be irrigated by more than one well because wells that are located within close proximity can create local interference. Staff at the KDWR also manually confirm that reported water use at multiple wells owned by a single farmer are non-overlapping to prevent double-counting irrigated acreage. Since our data is largely self-reported, we also adopt two sets of filters to ensure data quality. First, we only include fields of at least 40 acres, which is the area of a quarter-quarter PLSS section and the smallest land unit typically used in Western Kansas. Second, we only include water use observations of more than one inch and less than twice the local evapotranspiration, which drops less than 1% of our sample.

We observe well construction dates from a well construction database maintained by the Kansas Geological Survey (KGS), and we determine the depth to water in each field-year

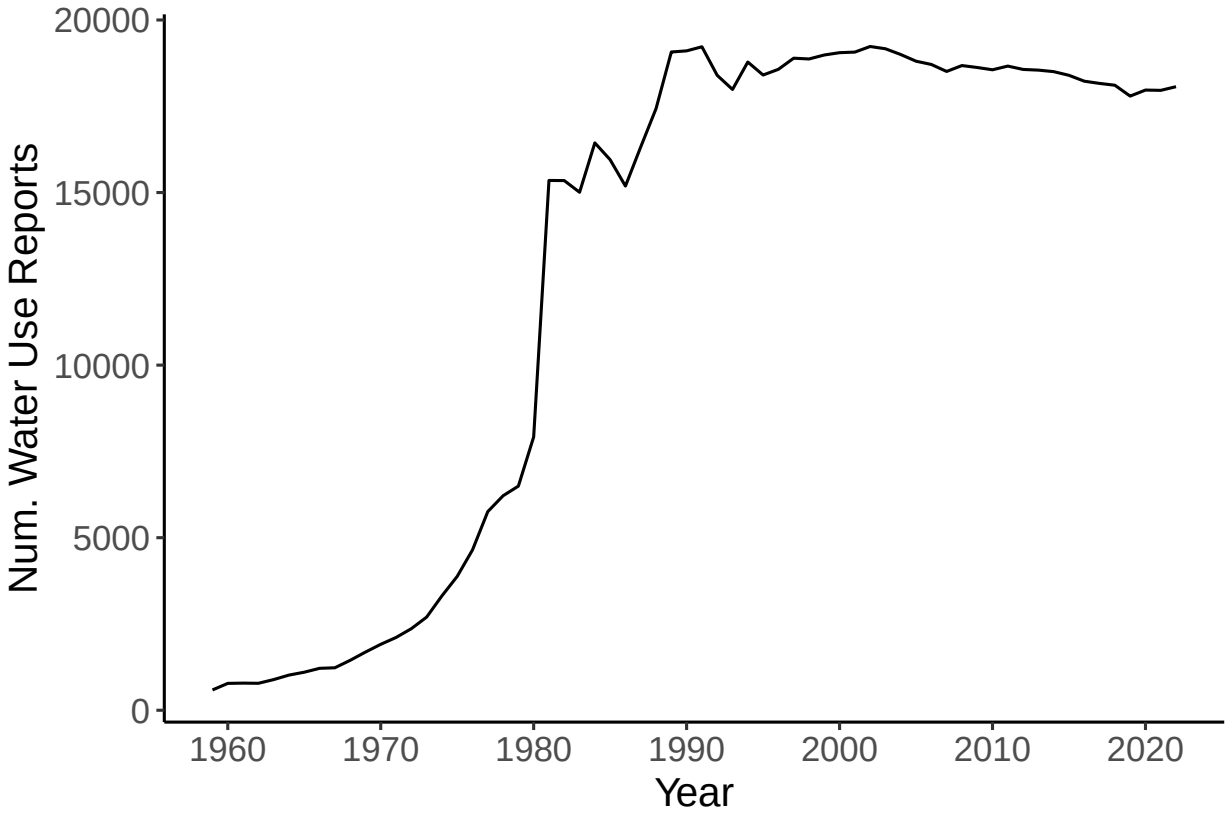

Figure 9: Number of Water Use Reports by Year

Notes: This figure shows the number of water use reports from 1959-2022.

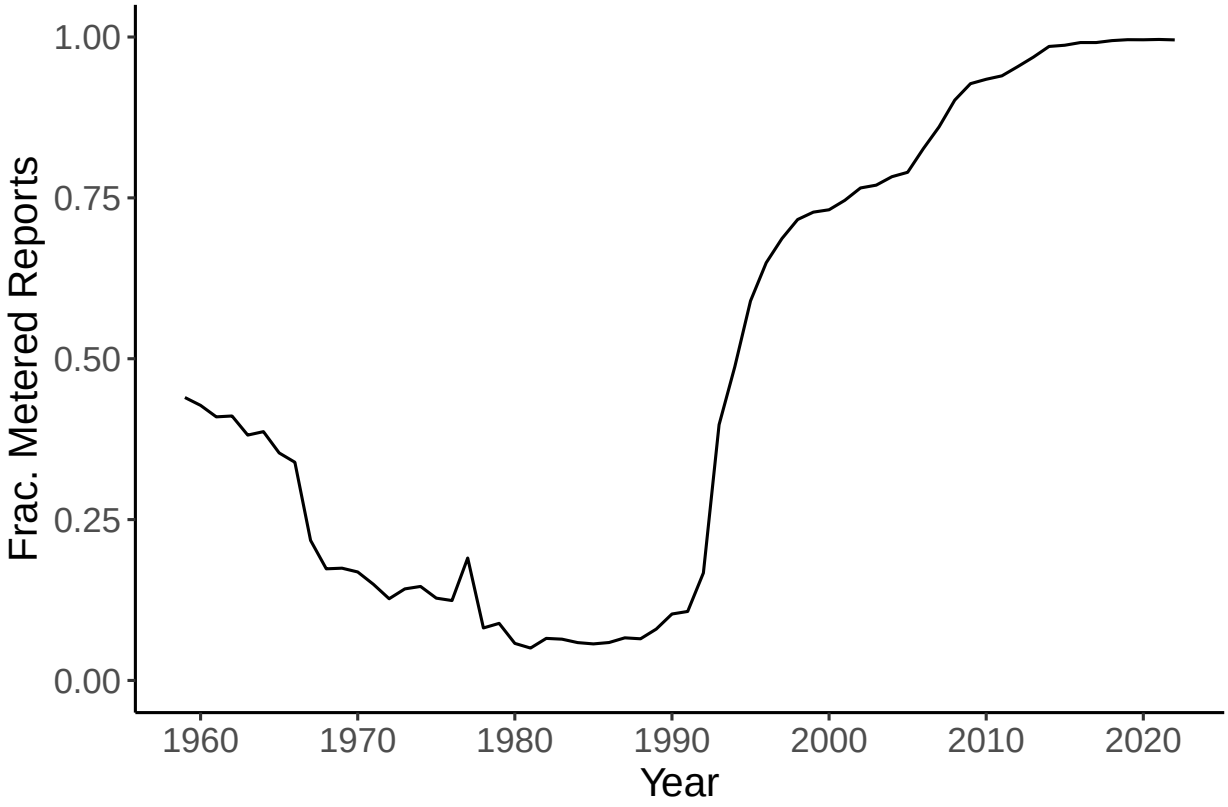

Figure 10: Fraction of Metered Water Use Reports by Year

Notes: This figure shows the fraction of water use reports from metered wells for the years 1959-2022.

using a network of aquifer observation wells maintained jointly by the KGS, KDWR, and various local GMDs. The observation wells track water levels throughout the High Plains aquifer, with annual measurements taken by staff at the KGS. For each year we restrict to measurements taken between October of the previous year and March of the current year so that measured depths are largely unaffected by interference from ongoing irrigation activity. We then impute the aquifer depth at each field using an inverse distance weighted average of observation well measurements.

We construct a measure of water marginal costs according to equation 1, obtaining natural gas prices from the U.S. Energy Information Association (EIA) Forms 857 and 910, which give the average commercial natural gas price in Kansas in each year from 1967 onwards. Prior to 1967 we use the national average price since state-specific prices are unavailable. We weight the national price by the average ratio of Kansas to U.S. prices to ensure consistency between the two series. We deflate all prices to 2017 dollars using the PCE.

We measure the weather for each field-year using historical meteorological data from TerraClimate, which has monthly precipitation and reference evapotranspiration totals at a 4km resolution since 1958. We construct three measures of weather in each year: growing season rainfall, growing season effective rainfall, and total evapotranspiration (Allen et al. 1998, Table 11). We convert from reference evapotranspiration to crop-specific evapotranspiration using monthly crop-specific ET coefficients (Table 12).[26]

We also include a variety of field-specific characteristics that do not vary over time. First, bedrock depths and surface elevation are obtained from digitized maps of the High Plains aquifer maintained by the KGS. Measures of soil quality, soil texture, and field slope are obtained from the Gridded Soil Survey Geographic (gSSURGO) Database maintained by the USDA Natural Resources Conservation Service (NRCS). We use soil organic carbon content as a proxy for soil quality and percentages of silt, sand, and clay as measures of soil texture.

D.2 County Panel

Similar to our field-level panel, we construct a county-level panel of farm expenditures and production using historical Censuses of Agriculture and NASS production surveys. In particular, we take county-level farm expenditure data from the 1959-2012 Censuses of Agriculture. We use expenditures on hired labor to measure labor expenses and construct materials ex-

26. For field-years in which we do not observe crop choices we use corn as a reference crop for irrigated production and winter wheat as a reference crop for rainfed production, which are the dominant crop choices in Kansas for irrigated and rainfed production respectively.

penditures as the sum of expenditures on fuel, seeds, fertilizer, and chemicals.[27] We again deflate all expenditures to 2017 dollars using the PCE.

From 1972-2018 the USDA also measured irrigated and rainfed production totals and acreage annually at the county-level for most major crops in Kansas through a series of surveys.[28] We use these measures to construct annual yields at the county-level separately for irrigated and rainfed crops. Finally, we construct a composite price index for irrigated and rainfed crops by creating a quantity-weighted average of crop prices in each year. For 1972-2018 we directly use the quantities measured by NASS survey data, aggregating across counties. We extend the series to our full sample of 1959-2022 by back- and forward-filling using the 1972 and 2018 weights respectively. We again deflate to 2017 dollars using the PCE.

D.3 Aquifer Characteristics

We obtain data on aquifer characteristics from two sources. First, we use digitized maps of the High Plains aquifer to obtain specific yields and hydraulic conductivity (Gutentag et al. 1984) at a 5 mile resolution. Second, we use estimated recharge amounts derived from a soil water balance approach, again at a 5 mile resolution (Stanton et al. 2011). Figure 11 plots the gridded specific yield and hydraulic conductivity data. Both specific yield and hydraulic conductivity vary significantly across the aquifer, with the southeastern portion having both the highest specific yields and the highest hydraulic conductivity.

27. Contract labor is another source of labor expenditures, but we exclude this category because it is inconsistently tracked across census cycles. We also subtract natural gas expenses from fuel expenditures since most natural gas is used for water pumping and we directly track water expenditures via water costs.

28. Corn and wheat production were tracked for the entire period, while sorghum production was tracked until 2009 and soybeans from 1984-2016. The USDA ceased annual production tracking after 2018 due to the passage of the 2018 Farm Bill and declining response rates.

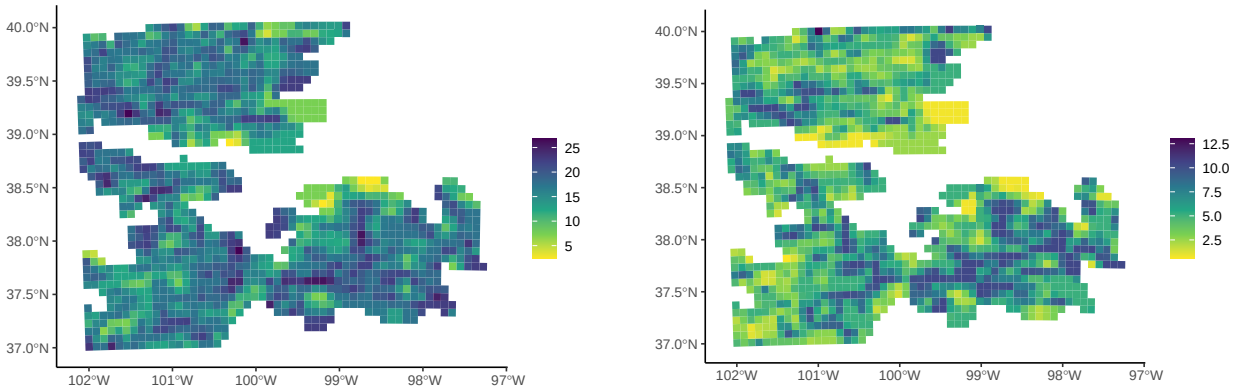

(a) Specific Yield (%)

(b) Hydraulic Conductivity (miles/year)

Figure 11: Aquifer Characteristics

Notes: This figure shows the specific yield (panel a) and hydraulic conductivity (panel b) of the High Plains aquifer at a 5 mile resolution. Specific yield is unitless and hydraulic conductivity is measured in miles per year.

E Model Assumption Analysis

A key assumption of our model is that fields are "small" and thus do not internalize the change in future costs resulting from their own pumping. In this section we assess the validity of that assumption by computing the net present value of the change in future marginal costs associated with one foot of pumping for fields of various sizes. Figure 12 shows this change in costs averaged across 100 randomly chosen grid cells of the aquifer to account for differences in aquifer characteristics.

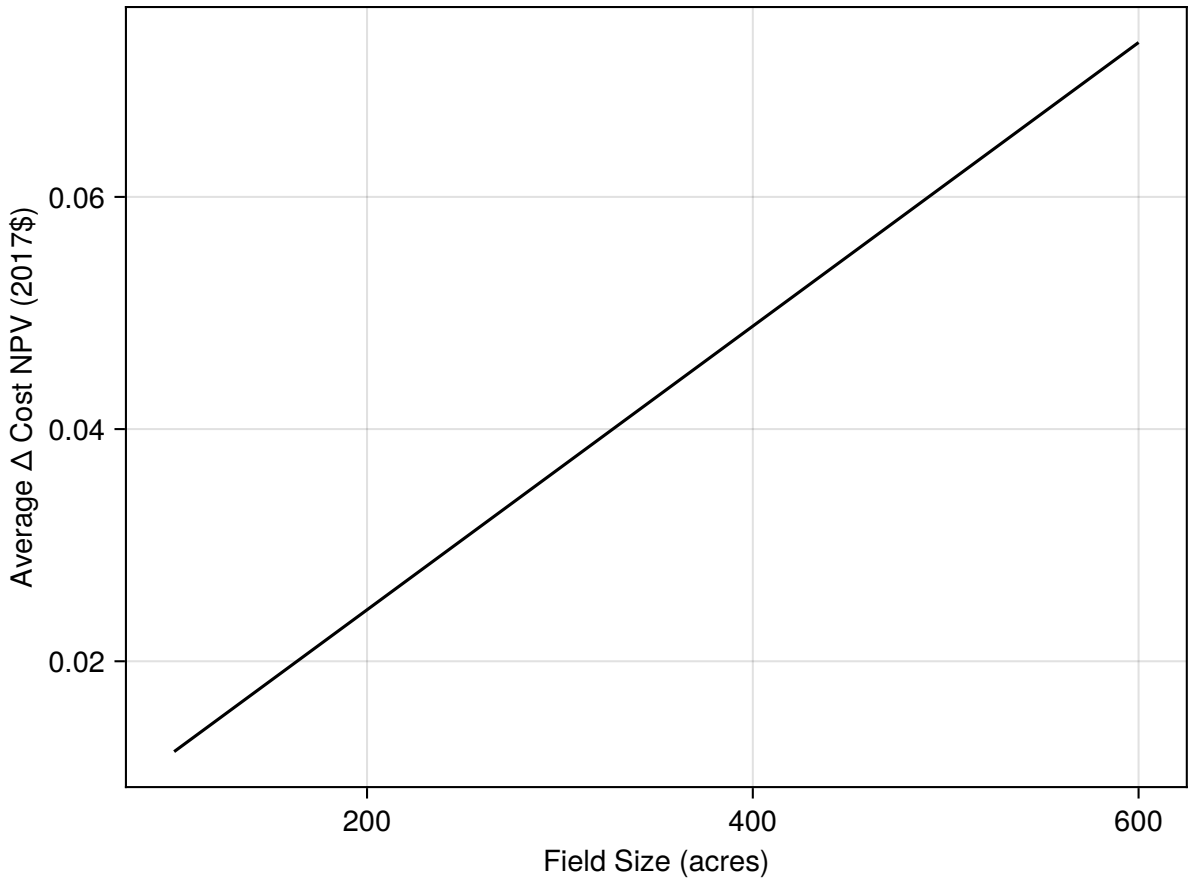

Figure 12: Changes in Cost from Field-Specific Pumping

Notes: This figure shows the average change in costs resulting from one foot of pumping across field sizes. Averages are taken over 100 randomly chosen grid cells of the aquifer to account for differences in aquifer characteristics such as hydraulic conductivity and recharge amounts. We present changes in costs in net present value terms by computing the change in the path of depths resulting from each pumping over 100 years and converting to costs by assuming a natural gas price of $9, which is the average value observed in our data. We assume a discount rate of 0.96.

As expected from the linear change in costs due to changes in depths as specified by equation 1, changes in cost grow linearly with changes in field size. Nevertheless, for all

field sizes the net present value of cost changes is only a few cents. For comparison, the average water marginal cost in our data is approximately \$27/AF, which corresponds to a net present value of approximately \$700 with a discount rate of 0.96. The change in costs from field-specific pumping is thus indeed small, representing only 0.01% of this value.