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

Aaron Berman Nathaniel Hickok*
(Job Market Paper)

January 23, 2026

Please click here for the latest version.

Abstract

Many of the world's major aquifers are being rapidly depleted from agricultural irrigation, generating dynamic common-pool externalities by raising future extraction costs. Entry restrictions are commonly used to limit depletion because well drilling is easily monitored, but they are second-best compared to Pigouvian taxes that directly target the intensive margin of water use. When policies cannot be tailored to heterogeneous users, however, the relative effectiveness and political feasibility of entry fees and water-use taxes become theoretically ambiguous, depending crucially on the correlation between water users' productivity and externalities. To study this question, we develop a dynamic model of farmers' joint well-drilling and water-use decisions, integrated with a physically realistic model of groundwater flows, and estimate it using field-level data on aquifer levels, water use, and crop production in the Kansas High Plains Aquifer from 1959 to 2022. We find that field-level productivity and water-use externalities are strongly positively correlated due to the spatial concentration of high-productivity fields, leading uniform taxes to outperform entry fees in terms of aggregate welfare. Nevertheless, entry fees are preferred by most users because the optimal uniform tax exceeds the marginal social cost of water use for all but the most productive fields. However, driven by irreversible well investments that lock in depletion from high-externality early entrants, the effectiveness and popularity of entry fees decline rapidly over time. These findings highlight how heterogeneity and irreversibility jointly shape the efficiency and political feasibility of environmental regulation.

* 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, the Martin Family Society of Fellows for Sustainability, and the National Science Foundation Graduate Research Fellowship Program under Grant Nos. 1745302 and 2141064.

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 possible through 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). Yet extensive agricultural irrigation is now depleting many aquifers that were formed over millennia of precipitation and runoff accumulation (Jasechko et al. 2024; Carleton, Crews, and Nath 2025).

Aquifer depletion is 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 for all future users, imposing a dynamic common-pool externality. In principle, the first-best outcome can be restored through corrective water-use taxes (Pigou 1932).[1] In practice, however, Pigouvian taxes are rarely used to govern commons resources, and especially groundwater (Blomquist 1987; Ostrom 1990; Edwards and Guilfoos 2021). First, heterogeneity in hydrological characteristics and the density and productivity of farmers means that pumping externalities vary significantly across space and time.[2] This heterogeneity makes it technologically difficult to implement a first-best water-use tax unless productivity is observable and taxes can be highly personalized (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 water-use taxes, extensive margin restrictions on new entrants are simple to monitor and may face fewer political constraints since they do not impose costs on existing users. In the case of groundwater extraction, these policies are particularly straightforward to monitor since extraction requires long-lived, easily observed investments in wells and other irrigation equipment. 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 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).

1. An alternative solution is to establish clear property rights (Coase 1960). However, the complexity of hydrological flows makes it difficult for markets over these property rights to arise, heavily limiting their effectiveness (Libecap 2010; Ayres, Meng, and Plantinga 2021; Sears et al. 2023).

2. For instance, in highly congested areas with many nearby farmers using substantial amounts of water, externalities can be large, while in relatively sparse areas externalities may be close to zero.

In this paper, we study the relative effectiveness and political feasibility of extensive-margin entry fees compared to intensive-margin water-use taxes for limiting aquifer depletion. When policies are uniform despite heterogeneous externalities, their relative welfare impact is theoretically ambiguous. Since entry fees screen out low-productivity farmers, it depends critically on the correlation between productivity and externalities. And while entry fees may be more politically feasible than water-use taxes because they are costless for incumbents, even incumbents may prefer taxes when facing high external costs from farmers who would not be screened out by the entry fee. Since well investments are irreversible, both the relative effectiveness and the political support for each policy can also change significantly over time.

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 1950s and continued to expand in the ensuing decades. However, projections now 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 (KDHE) 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 U.S. Censuses of Agriculture and National Agricultural Statistics Service (NASS) surveys. Our dataset is uniquely well suited to answer our questions of interest because it covers both the extensive margin of entry and the intensive margin of water usage at a spatially granular level over a long time period. These features allow 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. Our model of groundwater flows allows us to capture the large

heterogeneity in externalities imposed by pumping over space and time, while our economic model allows us to capture the heterogeneity in productivity and entry dynamics that are crucial to the policy trade-offs that are the focus of this paper. On the model's 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 whether and when 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 wells have no scrap value. Instead, exit from irrigated 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 our 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 large number of fields in our setting, we model this equilibrium dynamic using a dynamic competitive 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 observed levels of water usage. Weather shocks provide an additional source of identifying variation. Our estimation results suggest that water usage is moderately elastic to water costs, with an average elasticity of about -1.

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. Moreover, it requires only weak assumptions about the beliefs of potential entrants, unlike conventional full-solution approaches which require fully specifying long-run beliefs (Rust 1987, 1994). 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. By contrast, we estimate that the existing entry permit regulations in Kansas impose a comparatively small cost of approximately $31,000.

With parameter estimates for both the intensive and extensive margins in hand, we then apply our model to understand the welfare costs of aquifer depletion, compare the relative effectiveness and political feasibility of entry fees and water-use taxes, and examine how the effects of each policy evolve over time due to irreversible investments. We first consider how aquifer depletion will continue under the status quo, using our combined economic and hydrologic 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 all fields will have fully depleted the underlying aquifer and be unable to continue irrigation by 2050. The fraction of depleted fields rises further to 55% by the end of the century, with depletion rates slowing as the aquifer becomes increasingly depleted and overall water use declines. 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.

Importantly, we also find that productivity and externalities are strongly correlated under the status quo. This correlation is driven by spatial concentration in productivity, which means that highly productive fields tend to be located in crowded areas of the aquifer that later suffer from severe depletion rates. Moreover, we find that externalities are strongly negatively selected over time, with early entrants having the highest externalities. This negative selection worsens the lock-in problem faced by entry fees since any delay in implementation results in irreversible entry by high-externality fields. Together, the positive correlation in productivity and externalities and the negative selection of externalities over time mean that we should expect entry fees to perform worse than even uniform water-use taxes.

To understand the quantitative trade-off between these two policy instruments, as well as how their relative effectiveness changes over time, we apply our model to find the welfare-maximizing uniform taxes and entry fees, setting each policy to begin in 1960. We find that both policies result in substantial welfare gains over the status quo, although water-use taxes do outperform entry fees. The welfare-maximizing uniform water-use tax is approximately $50 per acre-foot (AF) and results in an overall welfare of approximately $35.5B. For comparison, welfare under the status quo is $30.3B, and so in annuity-equivalent terms the tax results in an increase in welfare of nearly $110M per year. 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 permit 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 water-use 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.

However, even though the entry fee results in lower aggregate welfare than the intensive-margin tax, it is actually preferred by a large majority of fields, both with and without lump-sum rebates. The reason is that the optimal uniform tax is a weighted average of externalities (Diamond 1973), adjusted to account for the expected social surplus of the marginal entrant, and is heavily skewed by a long right-tail of high-externality fields. Hence, the level of the optimal uniform tax is above the marginal social cost of most fields and therefore distorts their water use below the welfare-maximizing level. Moreover, there are many low-productivity fields that are only marginal entrants under the status quo. These fields gain under the entry fee policy by foregoing entry and remaining rainfed producers while enjoying additional revenue from the large lump-sum payments the entry fee policy generates. This preference for entry fees over taxes is particularly striking because there are almost no incumbent irrigators when our counterfactual policies begin in 1960.

We also consider how the effects of each policy instrument evolve over time due to irreversible investments by delaying the start year of each policy and varying the length of this delay. The effectiveness of both water-use taxes and entry fees declines with delay as depletion becomes increasingly locked-in. For instance, the welfare gain of a water-use tax beginning in 1978, which is the same year that Kansas first implemented its entry permit 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. Driven by the negative selection in externalities over time, this 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 1978.

More surprisingly, the political popularity of entry fees also declines over time, in contrast with the intuition that entry permit policies should generally become more popular over time as more firms sink entry costs and therefore are no longer subject to the fees. The source of this decline in popularity is the rapid drop in effectiveness of the entry fee policy relative to water-use taxes. This decline in effectiveness dominates the decline in costs imposed on incumbents and results in a net gain in fields preferring the tax policy over time, even though the tax policy comes at the cost of large tax payments for water use. As a result,

the water-use tax overtakes the entry fee to become preferred by a majority of fields in the mid-1980s.

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; Timmins 2002; 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 control 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 by 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 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 water-use tax level 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 and our hydrologic

model of groundwater flows. Section 5 presents our estimation strategy. Section 6 shows the estimation results. Section 7 reports the results of our counterfactual policy experiments. Finally, section 8 concludes.

2 Background

The High Plains aquifer system is made up of a series of interconnected 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. 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 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 1950s 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 120 feet, a decline of nearly one foot per year.

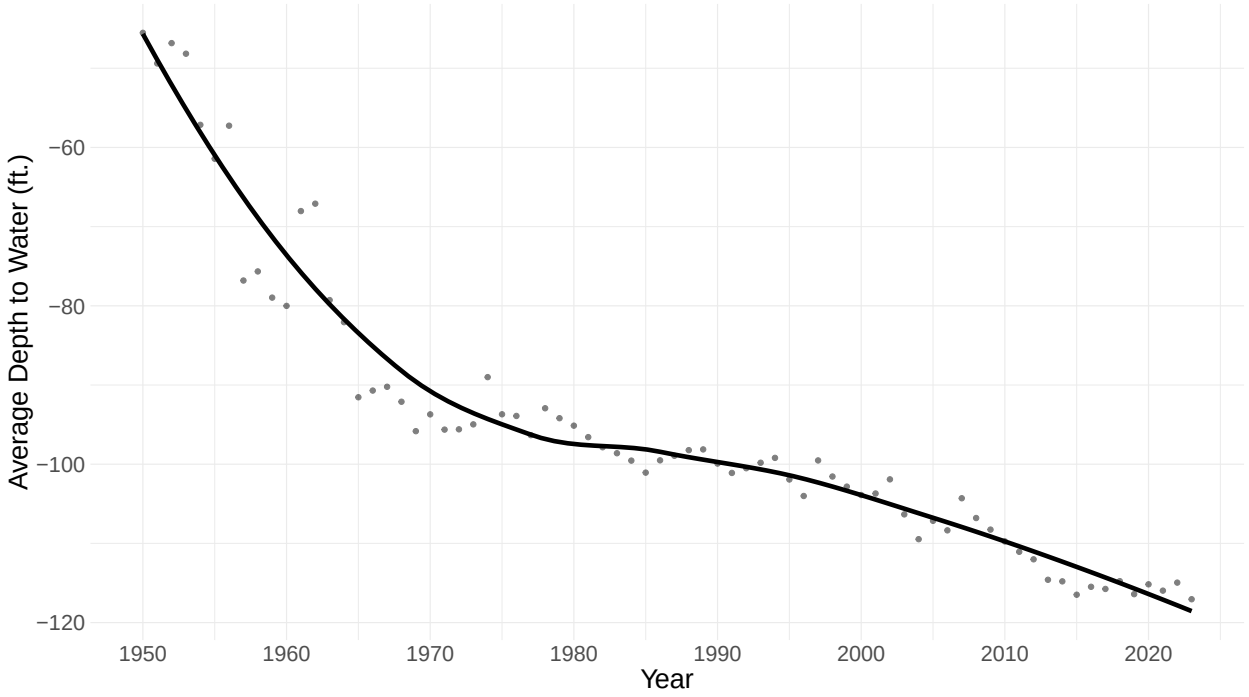


Figure 1: Average depth to water over time

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 line is a smoothed LOESS local regression.

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 1930s, 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.[3] 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. These permits, known as water rights, are granted by local groundwater management

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

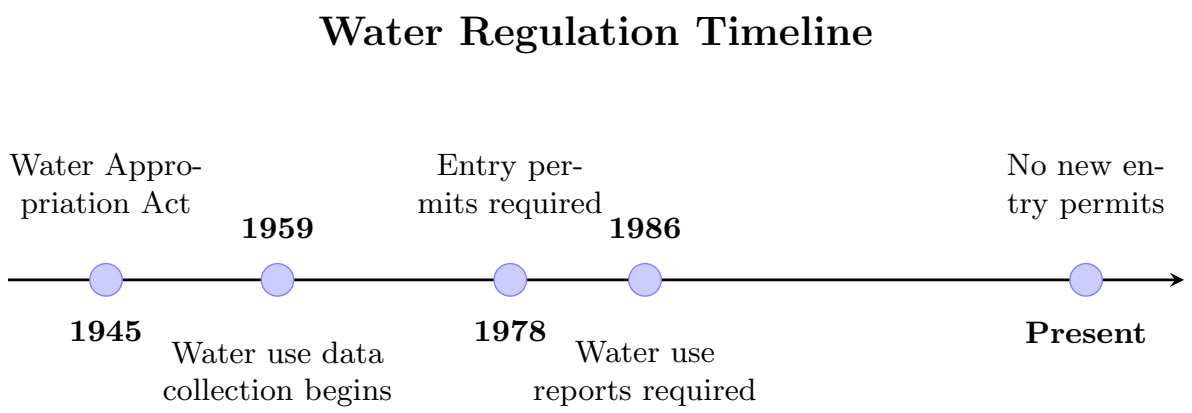


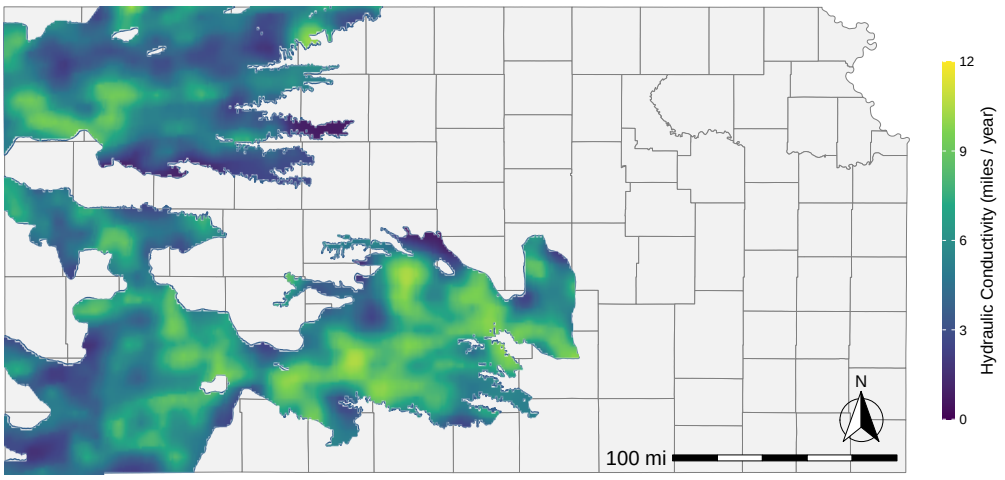
Figure 2: History of water regulation in Kansas

districts (GMDs) and were, until recently, easy to obtain, proving to be only a small barrier to entry. Today, in recognition of the fact that the aquifer is already vastly oversubscribed, Kansas GMDs have 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.

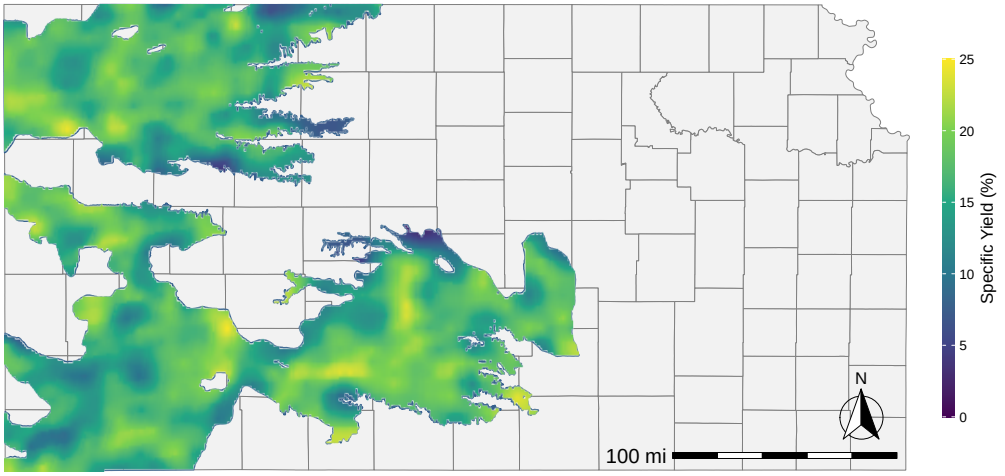
2.2 Aquifer Hydrogeology

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 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, showing how it varies across space. Hydraulic conductivity is highly variable across the aquifer, ranging from two miles per year at the



(a) Hydraulic conductivity



(b) Specific yield

Figure 3: Aquifer characteristics

Notes: This figure shows hydraulic conductivity across the aquifer in panel (a) and specific yields in panel (b). The values are interpolated from gridded data with a spatial resolution of five miles using ordinary kriging.

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.[4] 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.

Given hydraulic conductivity and specific yields, water flows laterally through the aquifer according to Darcy's Law (Woessner and Poeter 2024), which says that the flux of water through a cross-sectional area of the aquifer is equal to the product of hydraulic conductivity and the slope of water table elevation.[5] These lateral flows combine with vertical recharge from deep soil percolation and discharge from streamflows and groundwater extraction to jointly determine the law of motion for water volumes in the aquifer. Changes in water table depths are then proportional to these changes in volume and inversely proportional to specific yields.

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 datasets we use to construct this panel are administrative microdata derived from annual water use reports submitted by individual water users to the KDWR, data on well construction from driller reports submitted to the KDHE, and data on water levels in the aquifer from a network of observation wells maintained by the KGS. Water use reports are made at the level of an individual well and report the water pumped as well as the acreage and crop type irrigated in each year. We complement this field-level data with county-level data on farm expenditures and production using data from historical U.S. Censuses of Agriculture and various NASS surveys. Appendix A provides further details on data sources and data construction.

4. In particular, specific yields measure the fraction of water accessible for extraction per unit volume of the aquifer.

5. This slope is known as the hydraulic gradient.

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 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.[6] 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 affects 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
Well Construction Year	1978	16.8	1968	1975	1984
Water Use (AF)	201	168	98.0	162	249
Field Size (acres)	159	102	114	130	180
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
Depth (ft.)	108	62.8	49.4	107	149
Water Costs (2017$ / AF)	26.6	19.8	10.2	21.7	37.2
Irrigated Yield (2017$ / Acre)	530	226	379	469	623
Non-Irrigated Yield (2017$ / Acre)	222	113	151	187	267
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 A. Average water use is conditional on irrigating and effective rainfall and evapotranspiration are both measured for irrigated crops.

6. 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.

Table 1 presents a selection of basic descriptive statistics for our field panel. Conditional on irrigating, the average water use is approximately 200 acre-feet.[7] The average field size is approximately 160 acres.[8] 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 $26.6 per AF, implying that irrigating field owners spend on average about $5,350 per year on direct water costs relative to an average of about $100,000 per year in revenue. Non-irrigated field owners make significantly less revenue, about $45,000 per year on average.

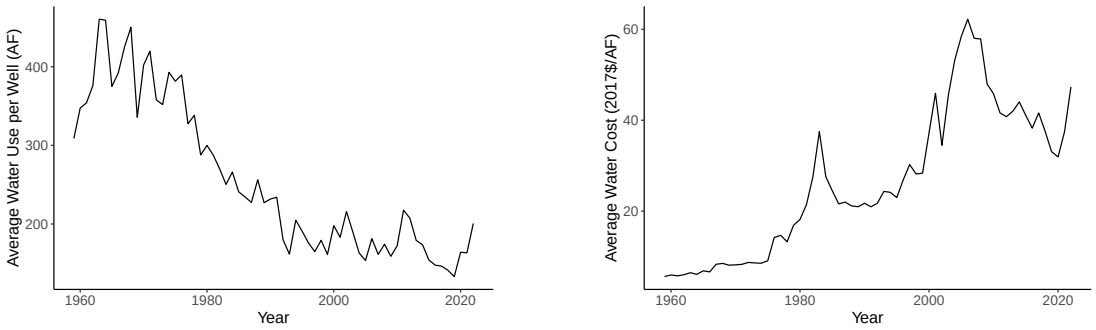
3.1 Water Use and Water Costs

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, with a high of more than 400AF per well in the 1960s and declining to a low of less than 200AF per well by 2020. This decline in water use reflects both changes in water costs, which have largely been increasing due to corresponding increases in aquifer depths, as well as a compositional shift in water users towards areas with lower water deficits and less water use on average. However, even though aquifer depths have been uniformly declining as seen in figure 1, water costs have not been uniformly increasing as they are also highly sensitive to the price of energy, particularly natural gas. Accordingly, water costs have experienced two spikes, once during the 1970s and early 1980s due to the energy crisis and again before the shale boom of the 2010s.

Similarly, figure 5 shows the average water use per well-year across the aquifer in panel (a) and the average water cost per well-year across the aquifer in panel (b). Driven by differential aquifer depths, water costs vary significantly across the aquifer and are highest in the southwest where the aquifer is deepest. Despite facing higher costs, wells in this southwest region also use the most water on average, which can be explained both by differences in climate—the southwest region of the aquifer is also the driest—and productivity—as we will show in section 7, the southwest region is also the most productive region with the highest marginal revenue from water use. This variation in water use conditional on climate and water costs will form a key part of our identification strategy detailed in section 5, allowing us to recover field-level productivity.

7. For comparison, a typical household uses approximately 0.5 AF of water each year.

8. 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 small differences in irrigation technology. For instance, fields using center pivots may or may not irrigate the corners of the field. Without corner irrigation, the effective field size of a 160 acre field is approximately 125 acres.



(a) Water use (b) Water costs

Figure 4: Water use and water costs over time

Notes: This figure shows average water use over time in panel (a) and average water costs in panel (b). Water use is averaged across constructed wells and measured in AF while water costs are averaged across all fields and measured in 2017$/AF.

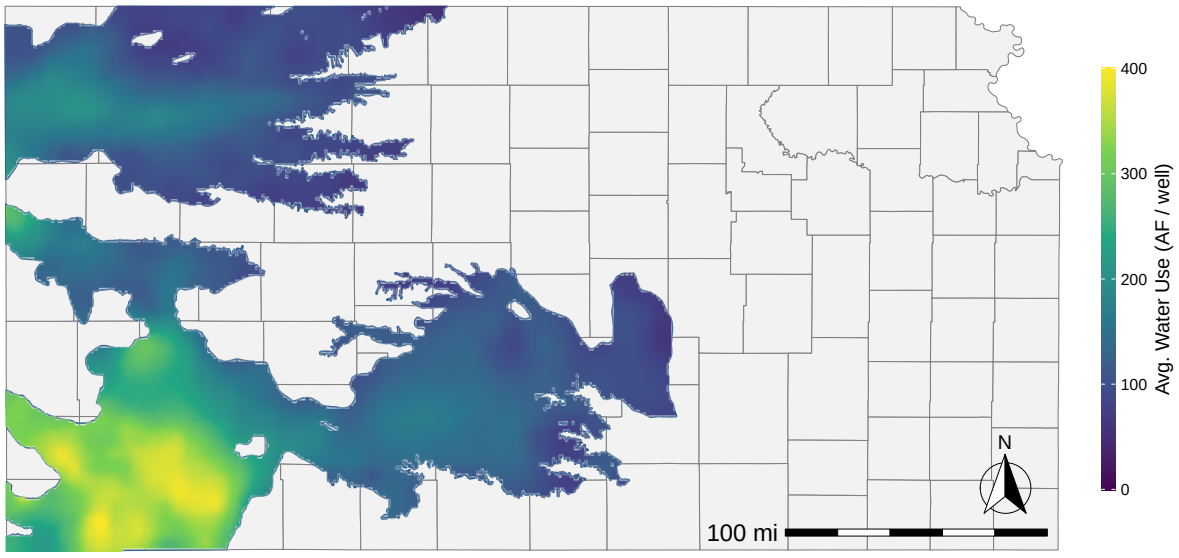
3.2 Well Construction

Figure 6 shows the cumulative number of wells constructed in or before each year from 1959-2022. In 1950 there were less than 5,000 wells constructed. The number of irrigation wells then increased rapidly in the 1960s and 1970s after the invention of the center pivot made large-scale agricultural irrigation feasible in the region. Irrigation well construction has slowed significantly in the decades since, both due to increased regulatory oversight and also increased aquifer depletion lowering the returns to irrigation. In total, there were more than 27,000 wells constructed by 2022.

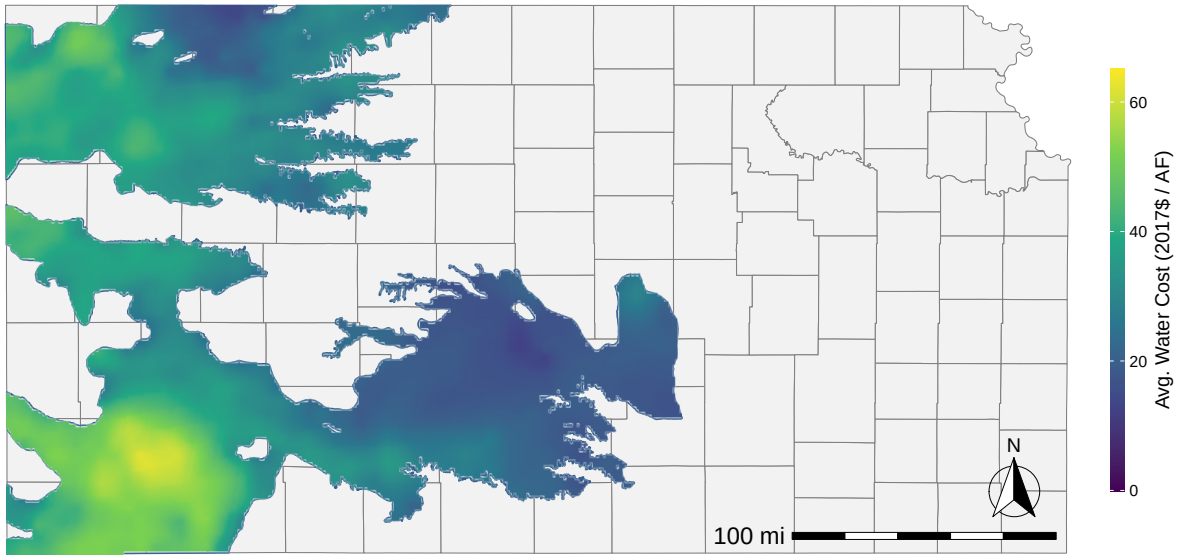
Figure 7 likewise shows spatial variation in the cumulative number of wells constructed by the end of our sample period in 2022 across the aquifer. The concentration of wells is highest in the southwest region of the aquifer, with some areas reaching a density greater than 4 wells per square mile. Due to the local nature of the aquifer, these areas also experience the largest external costs of water use, as depletion rates are highest and pumping from each well rapidly impacts a large number of other wells. As we will show in section 7, this variation in external costs of water use is key to the welfare impacts we recover for both the entry fee and water use tax policies we simulate in our counterfactuals.

4 Empirical Model

This section specifies an empirical model of supply for irrigated and rainfed agriculture, as well as a physically realistic hydrologic model of water flows in the aquifer. We divide land into atomistic fields that are independently managed by infinitely-lived owners. Field owners



(a) Water use



(b) Water costs

Figure 5: Water use and water costs over space

Notes: This figure shows water use per well-year in AF across the aquifer in panel (a) and water costs in 2017$/AF in panel (b). The values are interpolated from gridded data with a spatial resolution of five miles using ordinary kriging.

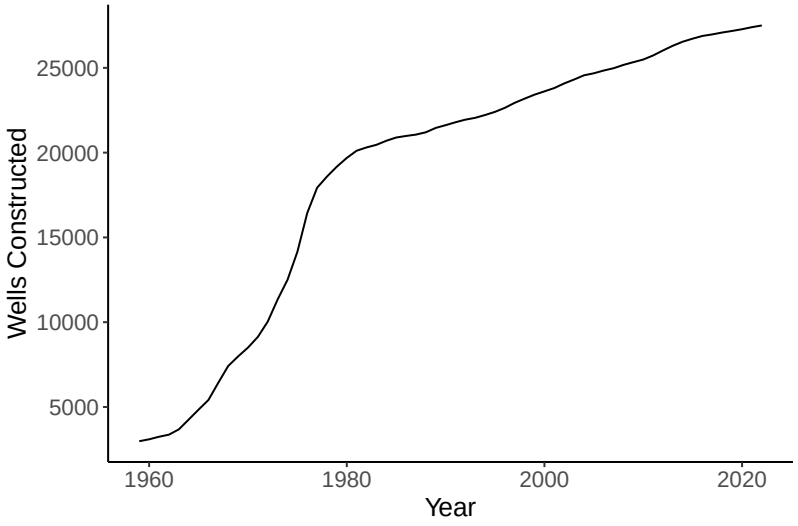


Figure 6: Wells constructed over time

Notes: This figure shows the cumulative number of wells constructed over time from 1959-2022.

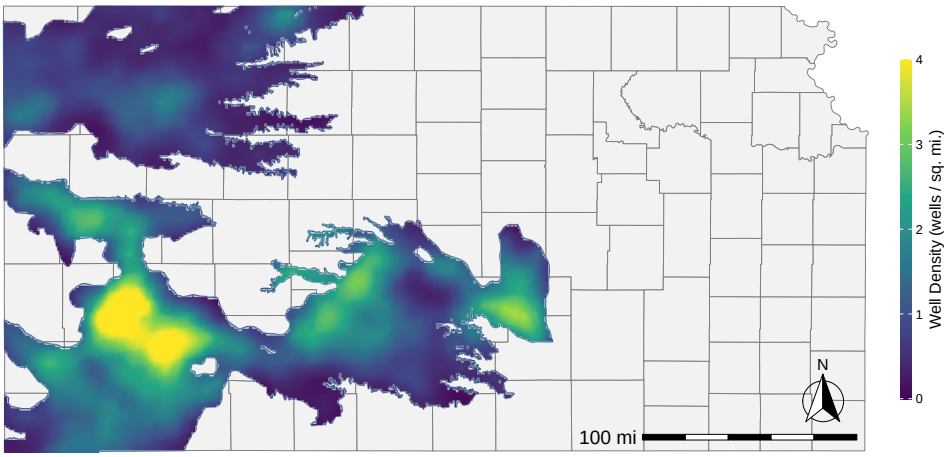


Figure 7: Well density over space

Notes: This figure shows the density of wells constructed across the aquifer by 2022 in units of wells per square mile. The values are interpolated from gridded data with a spatial resolution of five miles using ordinary kriging.

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. Finally, given realized pumping, the law of motion for water depths across the aquifer is determined by our hydrologic model.

4.1 Model Timing

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 water volumes 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}.

Figure 8 shows the timing of the model. Each year proceeds in three stages. First, field owners without a well decide whether or not to drill a well and become an irrigator. They do so before observing weather or their realized productivity shocks but after observing prices, well drilling costs, and water costs. If they choose not to drill a well and remain a rainfed producer, they grow rainfed crops and then face the same well-drilling decision in the following year. Second, fields with a well observe their productivity shocks and decide whether to grow rainfed or irrigated crops based on their expectations about the weather. Finally, fields growing irrigated crops observe the realized weather and choose water volumes for irrigation, along with their other flexible inputs, to maximize static profits. Fields growing rainfed crops likewise choose labor and materials but receive water only through exogenous rainfall levels.

The information set of each field owner i when deciding flexible inputs in period t, $\mathcal{I}_{it}$, as well as when deciding crop choice, $\mathcal{I}_{it}^{j}$, are recursively given by

$$\mathcal{I}_{it} = \{\vec{P}_t, A_{it}, c_{it}, \vec{\omega}_{it}, j_{it}, ER_{it}, ET_{it}, \mathcal{I}_{i,t-1}\} \tag{2}$$

$$\mathcal{I}_{it}^{j} = \{\vec{P}_t, A_{it}, c_{it}, \vec{\omega}_{it}, \mathcal{I}_{i,t-1}\} \tag{3}$$

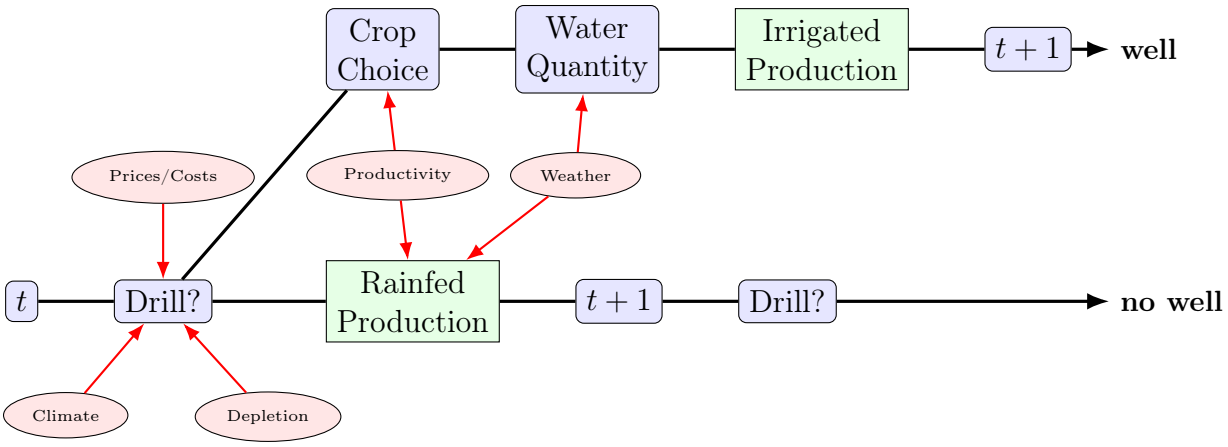


Figure 8: Model Timing

where $\vec{P}_t$ is the vector of prices, consisting of prices P_{jt} for each crop as well as natural gas prices P_t^{ng}, A_{it} is the size of the field in acres, $\vec{\omega}_{it}$ is the vector of productivities, consisting of both irrigated and rainfed productivity, j_{it} is the field owner's crop choice, ER_{it} is effective rainfall, and ET_{it} is evapotranspiration. The difference between $\mathcal{I}_{it}$ and $\mathcal{I}_{it}^j$ is that effective rainfall and evapotranspiration are part of the information set when deciding irrigation volumes but not crop choices. This difference approximates the timing of planting and irrigating decisions in the agricultural calendar. While planting decisions largely occur before rainfall and evapotranspiration are realized, irrigators do not need to commit to a particular level of water usage ex ante 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.

4.2 Intensive Margin

We model revenue yields as a function of water W_{it}, labor expenditures L_{it}, materials expenditures M_{it}, effective rainfall ER_{it}, and evapotranspiration ET_{it}

$$Y_{it} \equiv \frac{R_{it}}{A_{it}} = \exp(\omega_{it}^j) P_{jt} F_j(W_{it}, L_{it}, M_{it}, ER_{it}, ET_{it}) \tag{4}$$

where R_{it} is revenue and all inputs are measured in per acre terms. The assumption that productivity ω_{it}^j is Hicks neutral is standard in the literature and allows the marginal revenue product of water to vary flexibly across fields and over time (De Loecker and Syverson 2021). However, while it allows for water productivity to vary based on observables, it rules out any unobserved, factor-specific heterogeneity in productivity. We make this assumption

because we observe input usage at the individual level for only one flexible input, water, and therefore we are unable to separately identify factor-specific unobserved productivity from Hicks-neutral unobserved productivity.

Field owners make their crop choice decision to maximize expected static profits π_{it}^{j} before observing realized weather ER_{it} and ET_{it}

$$j_{it} = \arg\max_{j' \in \mathcal{J}_{it}} \mathbb{E}[\pi_{it}^{j'} \mid \mathcal{I}_{it}^{j}] \tag{5}$$

$$\pi_{it}^{\text{irr}} = \max_{W,L,M} A_{it} \left[\exp(\omega_{it}^{j}) P_{jt} F_j(\cdot) - c_{it} W - L - M\right] \tag{6}$$

$$\pi_{it}^{\text{rainfed}} = \max_{L,M} A_{it} \left[\exp(\omega_{it}^{j}) P_{jt} F_j(\cdot) - L - M\right] \tag{7}$$

where the crop choice set $\mathcal{J}_{it} = \{\text{irrigated, rainfed}\}$ if the field owner has an irrigation well and $\mathcal{J}_{it} = \{\text{rainfed}\}$ otherwise.

The assumption of static profit maximization rules 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 typical of the literature on common pool resources (Gisser and Sánchez 1980). It is also consistent with the assumption of atomistic fields, which implies that the individual decisions of fields affect the aquifer only in aggregate. 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.

4.2.1 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,

$$F_j = \left(1 - \exp(-\lambda_j EW_{it})\right) L_{it}^{\beta_{jL}} M_{it}^{\beta_{jM}}, \; EW_{it} := \frac{W_{it} + \nu_j ER_{it}}{ET_{it}} \tag{8}$$

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 8 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.[9] 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.

4.2.2 Productivity

We specify productivity for irrigated crops as a first-order Markov process

$$\omega_{it}^{\text{irr}} = \vec{x}_i \gamma^{\text{irr}} + t\gamma_t^{\text{irr}} + \mu_{it}^{\text{irr}} \tag{9}$$

$$\mu_{it}^{\text{irr}} = g(\mu_{i,t-1}^{\text{irr}}) + \xi_{it}^{\text{irr}} \tag{10}$$

where the productivity innovation $\xi_{it}^{\text{irr}} \equiv \omega_{it}^{j} - \mathbb{E}[\omega_{it}^{j} \mid \mathcal{I}_{i,t-1}]$. The vector of controls $\vec{x}_i$ includes a constant, latitude and longitude, elevation, slope, and measures of soil quality.[10] This specification decomposes field productivity into an individual-specific component that depends on field-specific characteristics, an aggregate component that evolves over time, and an idiosyncratic component. Similar to Ackerberg, Caves, and Frazer (2015), we assume that the idiosyncratic term μ_{it} follows a first-order Markov process so that the productivity innovation ξ_{it}^{j} is orthogonal to anything in the field owner's information set at time $t-1$.

For rainfed crops we specify productivity as

$$\omega_{it}^{\text{rainfed}} = \gamma_{c(i)}^{\text{rainfed}} + t\gamma_t^{\text{rainfed}} + \xi_{c(i)t}^{\text{rainfed}} \tag{11}$$

where $c(i)$ is the county of field i. The main restriction of equation 11 relative to equation 9 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.

9. 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 our effective rainfall measure that discounts rainfall in excess of crop water demand via evapotranspiration (Allen et al. 1998).

10. 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.

4.3 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 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^e_{it}, and we assume that wells have no scrap value.[11] 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 $s_{it} = \{c^e_{it}, D_{it}, \vec{P}_t, \vec{x}^e_i\}$ and privately observe a vector of shocks $\vec{\epsilon}^e_{it} = (\epsilon^e_{it0}, \epsilon^e_{it1})$, which we assume follow an i.i.d Type 1 extreme value distribution with scale parameter σ_e. The vector of field characteristics $\vec{x}^e_i$ includes the same field productivity characteristics as in section 4.2.2—soil quality, terrain, as well as latitude and longitude—along with the bedrock depth of the aquifer. We assume that fields do not know their productivity shocks when making their entry decisions but instead form beliefs about their productivity based on their field characteristics. We proxy for the local climate of each field with latitude and longitude, which provides a good approximation of weather processes in Kansas, where most of the cross-sectional variation in weather is spatial due to the rain shadow of the Rocky Mountains.

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

$$V(s_{it}) = \mathbb{E}[\max\{v(0; s_{it}) + \epsilon^e_{it0}, v(1; s_{it}) + \epsilon^e_{it1}\}] \tag{12}$$

$$v(0; s_{it}) = \mathbb{E}[\pi^0_{it} \mid s_{it}] + \beta E[V(s_{i,t+1}) \mid s_{it}] \tag{13}$$

$$v(1; s_{it}) = \mathbb{E}\left[\sum_{\tau=0}^{\infty} \beta^\tau \max\{\pi^0_{i,t+\tau}, \pi^1_{i,t+\tau}\} \mid s_{it}\right] - c^e_{it}(s_{it}) \tag{14}$$

where π^0_{it} and π^1_{it} give pre-entry and post-entry static profits respectively.

Equation 14 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 14 is that individual fields are atomistic, 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

11. 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.

endogenously affect the evolution of state variables but rules out any strategic interactions in order to maintain finite dependence (Arcidiacono and Miller 2011; Scott 2023). Finally, the assumption of logit shocks implies choice probabilities given by

$$p^e(s_{it}) = \frac{\exp(v(1; s_{it})/\sigma^e)}{\exp(v(0; s_{it})/\sigma^e) + \exp(v(1; s_{it})/\sigma^e)} \tag{15}$$

4.3.1 Entry Costs

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

$$c^e_{it} = \vec{x}^{e,c}_{it}\alpha^e + \mathbb{1}\{t \geq 1978\}\alpha^e_{GMD} + \alpha^e_t t + \alpha_{c(i)} \tag{16}$$

where the vector of controls $\vec{x}^{e,c}_{it}$ includes a constant, bedrock depths, and field size, $\mathbb{1}\{t \geq 1978\}$ is an indicator for active GMD entry permit regulations, and $c(i)$ gives the county of field i. The cost function specified in equation 16 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 entry permit regulation indicator allows costs to potentially increase after entry permit requirements began in 1978.[12] We also include a linear time trend to capture changing 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.

4.4 Hydrologic Model

As shown in figure 3, there are large differences in underlying aquifer characteristics that determine water flows across the aquifer. In turn, these differences drive significant heterogeneity in externalities imposed by pumping across the aquifer. To capture these differences, we adopt a physically realistic model of groundwater flows based on Darcy's Law (Woessner and Poeter 2024). We use this hydrologic model, together with vertical recharge and discharge, to compute the law of motion for water depths across the aquifer.

We conceptualize the aquifer as consisting of a set of cells $k \in \mathcal{A}$. Each cell k is a local aquifer with area $\tilde{A}$ characterized by it's hydraulic conductivity K_k, specific yield SY_k,

12. A restriction of this cost function is that is does not allow for the costs of regulatory oversight to change over time. The main concern is that obtaining new water rights has become more difficult over time, a process which is captured in equation 16 only through the linear time trend. However, this effect largely impacts the set of potential fields, which we take to be exogenous, as obtaining a water right is a prerequisite for drilling an irrigation well.

surface elevation E_k, water table depth D_{kt}, hydraulic head $H_{kt} \equiv E_k - D_{kt}$, bedrock depth $\underline{D}_k$, and saturated thickness $ST_{kt} \equiv \underline{D}_k - D_{kt}$. The total volume of water stored in a local aquifer cell k at time t, S_{kt}, is determined by its saturated thickness and specific yield

$$S_{kt} = SY_k ST_{kt} \tilde{A} \tag{17}$$

We assume infinite hydraulic conductivity within each cell, but we importantly allow for finite hydraulic conductivity across cells. In particular, the volume of water that flows from cell k to an adjacent cell k' is given by Darcy's Law

$$Q_{kk't} = \bar{K}_{kk'} \bar{ST}_{kk't} \Delta H_{kk't} \tag{18}$$

In terms of water volumes, the aquifer law of motion is then

$$S_{k,t+1} = S_{kt} - W_{kt} - \sum_{k' \in adj(k)} Q_{kk't} + NR_{kt} \tag{19}$$

where W_{kt} is the total volume of groundwater extracted by fields overlying aquifer cell k and NR_{kt} is net recharge.[13] Changes in volumes map into depths according to local specific yield through equation 17.

For our hydrologic model we discretize the aquifer into 2,226 cells with a spatial resolution of 5 miles.[14] Following the hydrologic literature, we calibrate net recharge amounts NR_{kt} to match observed changes in depths given the groundwater extraction levels we observe in the data. Appendix H provides details on this calibration process and the resulting model fit. Relative to state-of-the-art hydrologic models such as MODFLOW (Langevin et al. 2017), the main simplification in our model is that we consider only annual time steps rather than simulating the full differential equation implied by Darcy's Law. We take this approach to maintain computational tractability when embedding our hydrologic model within the full equilibrium framework of our empirical model and applying it at a regional scale. We show in appendix H that even with annual time steps our model is still able to fit the observed changes in depths, and moreover that it captures the local nature of groundwater extraction externalities implied by finite hydraulic conductivity.

13. Net recharge captures deep soil percolation from precipitation and past irrigation that flows into the aquifer net of any discharge either from natural sources such as streamflows or non-agricultural pumping, which occurs outside of our model.

14. We selected this spatial resolution because it is the most granular spatial resolution that satisfies the Courant condition for numerical stability given our annual time steps.

4.5 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 a dynamic competitive equilibrium approach in the spirit of Hopenhayn (1992) and Weintraub, Benkard, and Van Roy (2008). In particular, rather than tracking the full industry state,[15] fields instead track only their field-specific states. The impact of decisions by other fields enter only through field beliefs about future aquifer depletion. In equilibrium, as we detail in appendix C.2, fields have depletion expectations that are consistent with the realized path of aquifer depths.

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

1. Static profit maximization: Fields choose crop j_{it}, water usage W_{it}, labor L_{it}, and materials M_{it} to maximize static profits π_{it} according to equations 5, 6, and 7

2. Depletion expectations consistency: Field beliefs about the evolution of aquifer depths are consistent with the realized path of aquifer depths $\{\vec{D}_t\}_{t=0}^{\infty}$

3. Dynamic entry consistency: Field entry probabilities are given by equation 15

4. Hydrologic consistency: Aquifer depths $\vec{D}_t$ evolve according to equations 17 and 19 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 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 8. 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

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

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. This approach is computationally light since it does not require fully solving the model for estimation, and it also avoids the need to make potentially strong assumptions about long-run beliefs, requiring only significantly weaker assumptions on one-period ahead beliefs.

5.1 Intensive Margin

The main challenge associated with estimating the production function specified in equation 8, 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 $\theta_j \equiv (\lambda_j, \nu_j, \beta_{jL}, \beta_{jM})$.[16] 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 8, the implied inverse conditional factor demand for irrigation water is (see Appendix B for a derivation)

$$Y_{it} = c_{it} ET_{it} \Gamma_{it}, \ \Gamma_{it} := \frac{\widetilde{W}_{it}}{\lambda_j \exp(-\lambda_j EW_{it})} \tag{20}$$

16. 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.

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

$$Y_{it} = \frac{L_{it}}{\beta_{jL}},\ Y_{it} = \frac{M_{it}}{\beta_{jM}} \tag{21}$$

Finally, productivities are given by the residual

$$\omega_{it}^{j} = \ln(Y_{it}) - \ln(P_{jt}) - \ln(\widetilde{W}_{it}) - \beta_{jL}\ln(L_{it}) - \beta_{jM}\ln(M_{it}) \tag{22}$$

Together, equations 20, 21, and 22 allow us 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 B we show that, under standard assumptions, this key inversion step is possible for any parametrically specified homothetically separable production function and does not rely on the specific functional form of equation 8. 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 yields $\hat{Y}_{it}^{\text{irr}}$ as described in section 5.1.1. We then estimate equation 9 by OLS to recover the productivity

residuals μ_{it}^{irr}.[17] 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

$$\mathbb{E}[Z_{it}\xi_{it}^{\text{irr}}] = 0 \tag{23}$$

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.[18] Equation 26 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

$$\mathbb{E}[Y_{ct} - \hat{Y}_{ct}] = 0 \tag{24}$$

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 of equation 22 together with the Cobb-Douglas factor share equations 21:

$$\begin{aligned} \omega_{ct}^{j} &= (1-b^{j})\ln(Y_{ct}^{j}) - \ln(P_{jt}) - \ln(B^{j}) - \ln(\widetilde{W}_{ct}) \\ b^{j} &:= \beta_{jL} + \beta_{jM}, B^{j} := \beta_{jL}^{\beta_{jL}}\beta_{jM}^{\beta_{jM}} \end{aligned} \tag{25}$$

17. Even though our data is selected on irrigation, we can still consistently estimate the parameters of equation 9 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 9 with a selection correction as in Olley and Pakes (1996).

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

Equation 25 relies on our assumption that rainfed productivity is homogeneous within county-year.[19] 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 fixed effects and a linear time trend to construct moments of the form

$$\mathbb{E}[Z_{ct}\xi_{ct}^{\text{rainfed}}] = 0 \tag{26}$$

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 21, which are implied by static profit maximization. Because of the Cobb-Douglas functional form assumption in equation 8, 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

$$\beta_{jL}R_{ct}^{j} = X_{ct}^{j},\ X_{ct}^{j} \in \{L_{ct}^{j}, M_{ct}^{j}\} \tag{27}$$

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

$$X_{ct} = \sum_{j} \beta_{jL}R_{ct}^{j} \tag{28}$$

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

$$X_{ct} = \beta_X^{\text{irr}}R_{ct}^{\text{irr}} + \beta_X^{\text{rainfed}}R_{ct}^{\text{rainfed}} + \beta_X^{\text{other}}R_{ct}^{\text{other}} + \epsilon_{ct} \tag{29}$$

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 29 is that these farms also produce according to a production function that is

19. 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.

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 29 and the micromoments that provide the main identifying variation for the parameters of the water-yield function. Similar to Rafey (2023), equation 29 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 as well as variation in rainfall and evapotranspiration 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 predetermined 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 Arcidiacono and Miller (2011) gives that

$$V(s_{it}) = v(1; s_{it}) - \sigma^e \ln(p^e(s_{it})) \tag{30}$$

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

$$\begin{aligned}\sigma^e \ln\left(\frac{p^e(s_{it})}{1-p^e(s_{it})}\right) &= v(1; s_{it}) - v(0; s_{it}) \\ &= \mathbb{E}[\pi^1_{it} - \pi^0_{it} \mid s_{it}] - c^e_{it} + \beta\mathbb{E}[c^e_{i,t+1} + \sigma^e \ln(p^e(s_{i,t+1})) \mid s_{it}]\end{aligned} \tag{31}$$

where the second row follows because finite dependence implies that continuation values difference out. Equation 31 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.[20]

Applying expectational errors and substituting estimated values to equation 31 gives

$$\ln\left(\frac{\hat{p}^e(s_{it})}{1-\hat{p}^e(s_{it})}\right) - \beta \ln(\hat{p}^e(s_{i,t+1})) = \left(\hat{\pi}^1_{it} - \hat{\pi}^0_{it} - c^e_{it} + \beta c^e_{i,t+1} + \eta^e_{it}\right)/\sigma^e \tag{32}$$

where η^e_{it} is an expectational error that gives the difference between realized and expected values.[21] Substituting the functional form for entry costs from equation 16 into 32 gives the regression equation

$$\ln\left(\frac{\hat{p}^e(s_{it})}{1-\hat{p}^e(s_{it})}\right) - \beta \ln(\hat{p}^e(s_{i,t+1})) = \frac{1}{\sigma^e}\Delta\pi_{it} + x^e_i\tilde{\alpha}^e + \mathbb{1}\{t \geq 1978\}\tilde{\alpha}^e_{GMD} + \tilde{\alpha}^e_t t + \tilde{\alpha}_{c(i)} + \frac{\eta^e_{it}}{\sigma^e} \tag{33}$$

where $\tilde{\alpha} \equiv \frac{(\beta-1)}{\sigma^e}\alpha$. To estimate equation 33 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, and field sizes as basis terms. Since discount factors are typically unidentified in dynamic

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

21. To be precise, $\eta^e_{it} \equiv \mathbb{E}[\pi^1_{it} - \pi^0_{it} + \beta(c^e_{i,t+1} + \sigma^e \ln(p^e(s_{i,t+1}))) \mid s_{it}] - (\pi^1_{it} - \pi^0_{it} + \beta(c^e_{i,t+1} + \sigma^e \ln(p^e(s_{i,t+1}))))$

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 one-period ahead 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, nor do we need to make any assumptions about long-run beliefs. 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 because, as we verify in appendix E, fields are indeed small relative to both their local aquifer and 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,[22] of 0.80. In specification (2) we additionally include the average yield moment and estimate that the 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

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

Table 2: Estimation Results: Production Functions

Panel A: Irrigated			
	(1)	(2)	(3)
Water (λ)	3.24	4.44	4.42
	(0.042)	(0.008)	(0.008)
Rainfall (ν)	0.36	0.34	0.35
	(0.005)	(0.005)	(0.004)
Labor (β_L)	0.04	0.04	0.04
	(0.004)	(0.004)	(0.004)
Materials (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)
Water (λ)	6.14	5.84	5.38
	(0.796)	(0.681)	(0.544)
Labor (β_L)	0.03	0.03	0.03
	(0.002)	(0.002)	(0.002)
Materials (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. Standard errors are computed via the delta method. Yield gaps are measured as the ratio between estimated realized yields and maximum 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.

-0.93, which is similar to existing estimates of groundwater input demand elasticities in other 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 \$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 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

We consider three main types of counterfactual policies: intensive-margin taxes, extensive-margin entry fees, and hybrid policies that combine both taxes and entry fees. We focus on uniform policies that do not vary across space or time, reflecting administrative and monitoring constraints. To understand the opportunity cost of policy delay and lock-in we also vary the start year of each policy. Our baseline case considers policies that begin in 1960, and as a primary alternative we consider policies that begin in 1978, which is the year Kansas implemented their entry permitting reform as described in section 2. For robustness,

in appendix **??** we also consider more granular policies that vary across spatial areas or across time to reflect changing conditions in the aquifer and the climate. In appendix G we consider the sensitivity of our results to alternative assumptions about the discount rate and future climate change.

7.1 Counterfactual Equilibria

For each counterfactual we initialize the model in 1960, taking our initial conditions from the data in 1959, and compute an equilibrium of the model under the counterfactual policy until 2099.[23] As described in section 4.5, 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 C.

Fully solving for the model equilibrium also requires specifying the evolution of the exogenous variables—prices $\vec{P}_t$, productivities $\vec{\omega}_{it}$, effective rainfall ER_{it}, and evapotranspiration ET_{it}—following the sample period, 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. 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 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, along with beliefs about productivity, prices, and aquifer depletion rates, are provided in Appendix C.

For each counterfactual policy τ we consider as our measure of social welfare the net

23. 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.

present value of profits and government revenue

$$V_t^\tau = \sum_{t' \in T, t' \geq t} \left(\frac{1}{1+r} \right)^{t'-t} \left(\sum_{i \in \mathcal{F}} \pi_{it'}^\tau + G_{t'}^\tau \right) \tag{34}$$

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 each counterfactual policy τ to maximize social welfare V_t^τ over the appropriate space of potential policies.

7.2 Results

We consider two main sets of counterfactuals. First, we implement the welfare-maximizing uniform taxes, entry fees, and hybrid policies starting in 1960, the initial year of our model. Second, we implement the welfare-maximizing uniform taxes, entry fees, and hybrid policies that are delayed until 1978. These counterfactuals show both the opportunity cost of policy inaction through increased lock-in, and also provide a natural point of comparison with the entry permit 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 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 a total welfare 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 9 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 10 shows both the average water use externality percentile across space in panel (a) as well as the average productivity percentile in panel (b) under the status quo.[24] Both high externality and high productivity fields are heavily concentrated in the southwest portion of the aquifer, resulting in a strong correlation between externalities and productivity as shown in panel (c) of figure 10. As a result, we should expect the intensive-margin tax to be more

24. In particular, figure 10 shows externalities as of 1960. While absolute externality levels do vary over time, percentiles are highly stable as the relative externality levels are largely driven by fixed aquifer characteristics and the eventual density of nearby wells.

effective than the extensive-margin entry fee, since the entry fee will only screen out low productivity fields who also have low externalities.

Figure 11 likewise shows the average water use externality percentile of entrants over time in the status quo in panel (a) as well as the average productivity percentile of entrants over time in panel (b). Both the average externality and the average productivity of new entrants are declining over time. This feature suggests that entry fees will suffer heavily from delays in implementation, as most of the high externality fields will have already entered and will thus be unaffected by the delayed entry fees.

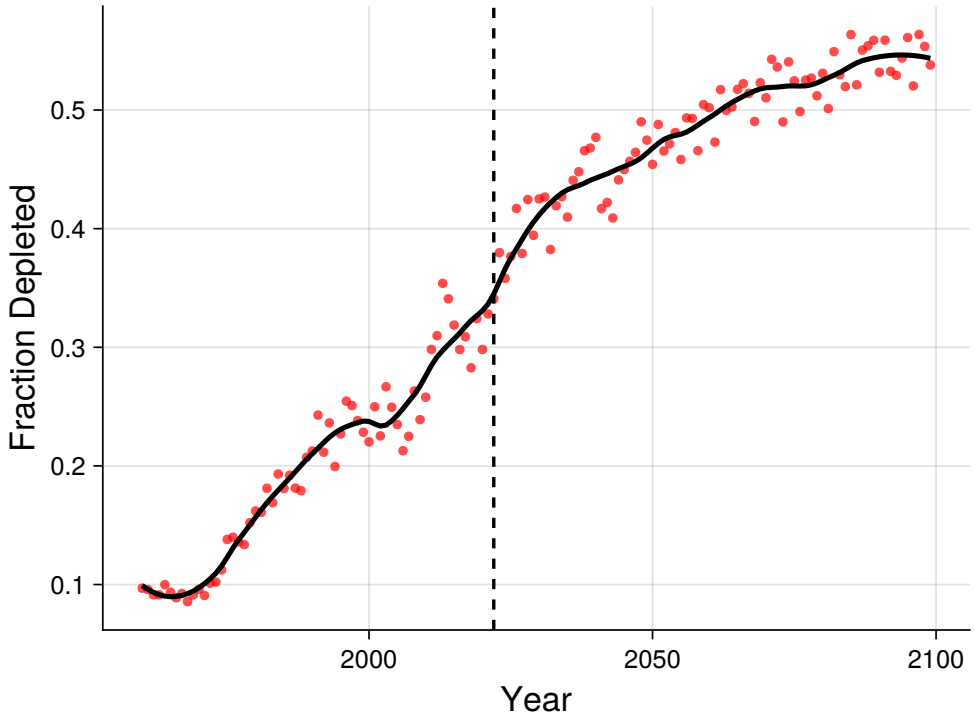
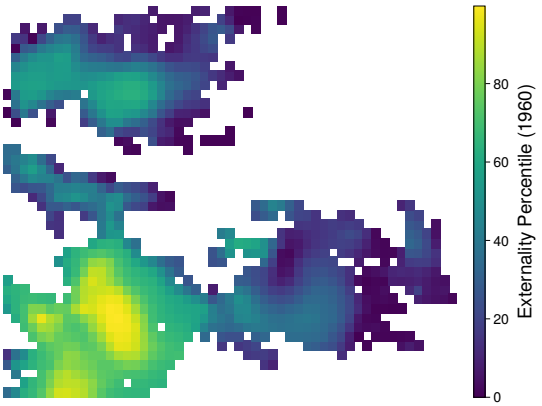


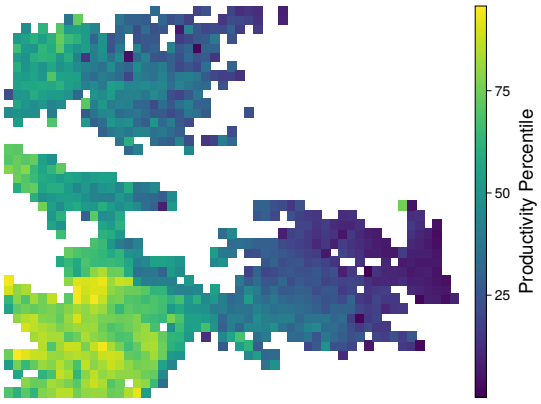
Figure 9: Time path of field depletion

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).

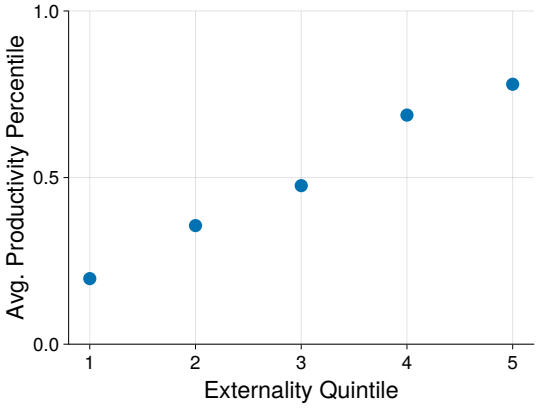
Columns (2) and (3) solve for the welfare-maximizing uniform water-use tax and entry fee respectively. We find an optimal tax of \$51.54/AF, resulting in a social welfare of \$35.48B and approximately 400 fewer wells. In annualized, annuity-equivalent terms, the change in welfare amounts to an increase of \$107.92M per year. We find an optimal entry fee of \$1M,



(a) Water use externality percentile



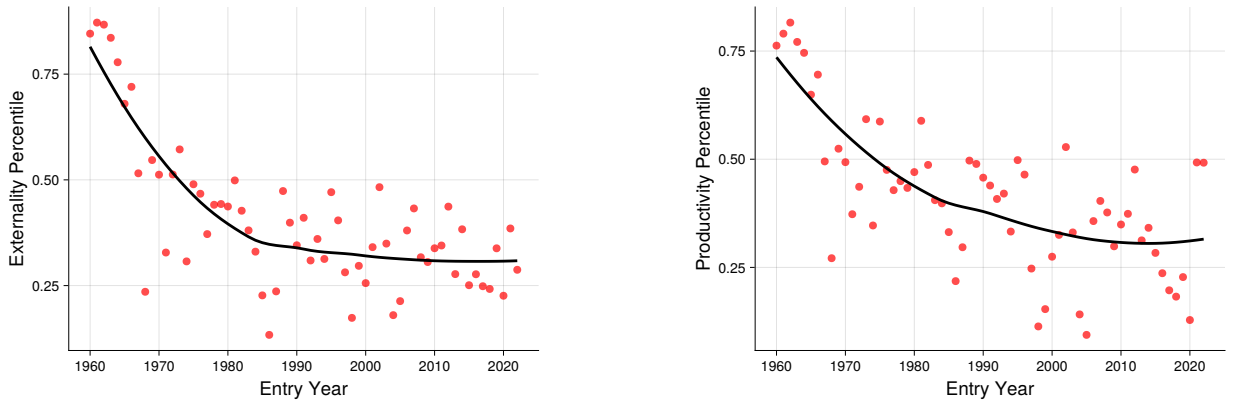
(b) Productivity percentile



(c) Correlation between externalities and productivity

Figure 10: Water use externality and productivity correlation

Notes: This figure shows the average externality percentile across grid cells of the aquifer in panel (a), the average productivity percentile in panel (b), and the average productivity percentile by externality quintile in panel (c). We measure externalities for each field as of 1960 using the shadow value of water, $\frac{dV_{t+1}}{dW_{it}}$. We measure productivity for each field using the average irrigated productivity $\bar{\omega}_i^{\text{irr}}$ across all periods of the model.



(a) Average entrant externality over time

(b) Average entrant productivity over time

Figure 11: Externalities and productivity over time

Notes: This figure shows the average externality percentile of entrants over time in panel (a) and the average productivity percentile in panel (b). Each point is an average percentile for a given year and and the solid black line is a smoothed LOESS local regression. We measure externalities for each field as of 1960 using the shadow value of water, $\frac{dV_{t+1}}{dW_{it}}$. We measure productivity for each field using the average irrigated productivity $\bar{\omega}_i^{\text{irr}}$ across all periods of the model.

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.

As expected given the positive correlation between externalities and productivity seen in figure 10, which implies that the entry fee will be unable to screen out high externality fields, the entry fee results in lower welfare than the water-use tax. However, it still achieves three-quarters of the gains in welfare from the water-use tax. Hence, when undertaken early, we see that intensive-margin taxes and extensive-margin entry fees can both 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 fields 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 the local aquifer is not yet depleted.

The mechanisms behind each policy, however, differ significantly. Under the water-use tax, well entry is relatively unchanged 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. We will return to these mechanisms and what they imply for the distributional consequences and political feasibility of entry fee and water-use tax policies in section 7.2.1.

Table 4: Counterfactual Results

		Uniform			Delayed		
	Status Quo (1)	Tax (2)	Entry Fee (3)	Hybrid (4)	Tax (5)	Entry Fee (6)	Hybrid (7)
Policy							
Tax ($/AF)	0.00	51.54	0.00	51.45	66.44	0.00	86.70
Entry Fee ($M)	0.00	0.00	1.00	0.56	0.00	0.98	1.03
Welfare							
Welfare ($B)	30.32	35.48	34.33	37.22	33.40	31.05	34.83
Annualized Δ Welfare ($M)	0.00	107.92	83.95	144.31	64.55	15.37	94.47
Num. Wells	27,060	26,667	13,267	20,695	26,786	18,048	18,926
Δ Num. Wells	0	-393	-13,793	-6,365	-274	-9,012	-8,134
Aquifer							
Remaining Storage (MAF)	174.37	241.26	240.22	255.17	242.66	212.81	259.83
Depleted Fields (%)	53.79	34.77	37.30	31.68	32.69	46.23	26.24

Notes: This table shows the results of model simulations under the status quo (column 1), a uniform water-use tax beginning in 1960 (column 2), a uniform entry fee beginning in 1960 (column 3), a hybrid entry fee and tax policy beginning in 1960 (column 4), a uniform water-use tax beginning in 1978 (column 4), a uniform entry fee beginning in 1978 (column 6), and a hybrid entry fee and tax policy 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.

Column (4) solves for the welfare-maximizing hybrid policy, combining both a uniform water-use tax with a uniform entry fee. We find a similar optimal tax of \$51.45/AF but a significantly lower optimal entry fee of only \$0.56M. The reason for the lower entry fee is that the intensive margin tax is able to partially correct the over-pumping that occurs in the absence of policy, meaning that many more fields have positive social surplus in expectation. Indeed, more than 7,000 additional wells enter under the hybrid policy compared to the uniform entry fee policy in column (3). Moreover, the resulting change in welfare increases to nearly \$150M per year in annuity-equivalent terms, and the terminal field depletion rate decreases to 32%.

Columns (5) and (6) again solve for the welfare-maximizing uniform water-use 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 \$66.44/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, highlighting 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 a worse policy instrument but implementing it immediately. We discuss this trade-off further in section 7.2.2.

We find the welfare-maximizing entry fee to be \$0.98M, 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 fee 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. Lock-in and negative selection drive this rapid decline. 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%. More than just lock-in—the delayed entry fee still results in 9,000 fewer wells constructed, which is a significant fraction of the nearly 14,000 fewer wells constructed under the uniform entry fee policy implemented in 1960—the delayed entry fee also suffers from negative selection due to the spatial correlation in productivities and externalities shown in figure 10, which introduces a strong negative correlation between private surplus and externality levels. That is, many fields with high gains from irrigating are located close together and therefore impose large externalities due to the local nature of the aquifer. As shown in figure 11, these fields also enter earliest, implying that the delayed entry fee therefore targets not only fields that have low private surplus but also low externalities, as most of the high externality fields have already entered.

Finally, column (7) solves for the welfare-maximizing hybrid entry fee and water-use tax policy implemented in 1978. When used together, the entry fee and tax combine to increase welfare by \$94.5M per year in annualized terms, which is more than the sum of welfare increases when each policy is used alone, implying that the two policies can be complementary. The key reason for this is that when the uniform tax is implemented alone it must correct for two distortions: over-pumping on the intensive margin but also over-entry on the extensive margin. However, increasing the tax to further reduce entry comes at the cost of distorting the production of entrants. Combining the tax with an entry fee can relax this constraint, as the tax only needs to correct for over-pumping on the intensive margin while the entry fee can correct for over-entry, selecting only fields that are positive social surplus in expectation. Indeed, we can see in column (7) that unlike in 1960, when the entry fee under the hybrid policy was lower than with the entry fee alone, the optimal entry fee of \$1.03M is on par with the entry fee of \$0.98M that is optimal when only using entry fees. Moreover, the optimal tax increases to \$86.70/AF, compared to \$66.44/AF under the tax alone. The resulting tax is larger because of the correlation between externalities and productivity as seen in figure 10, which implies that the remaining fields that enter even with the entry fee have higher externalities on average.

7.2.1 Winners and Losers

Figure 12 shows the fraction of fields that gain relative to the status quo both with and without lump sum rebating of government revenues. Without rebates, a majority of fields lose compared to the status quo under both the tax and hybrid policies. By contrast, under the entry fee policy the opposite is true—a majority of fields gain even without rebates. Rebating closes the gap between policies and approximately 70% of fields gain under all three policies.

The aggregate similarity in the fraction of fields that gain under each policy conceals significant heterogeneity in distributional impacts. Figure 13 shows the fraction of fields that gain relative to the status quo across the aquifer under both the tax in panel (a) and the entry fee in panel (b) with lump sum rebating. The gains of the tax policy are concentrated largely in the southwest region of the aquifer while the gains from the entry fee policy are concentrated in the eastern and northern portions of the aquifer. The key mechanisms driving this pattern are two-fold. First, as we showed in figure 10, most of the high productivity and high externality fields are concentrated in the southwest portion of the aquifer. These high-productivity, high-externality fields drive the tax above the marginal externality of fields in the northern and eastern portion of the aquifer, which are less concentrated and therefore have lower externalities. Second, almost none of the highly productive fields in the southwest

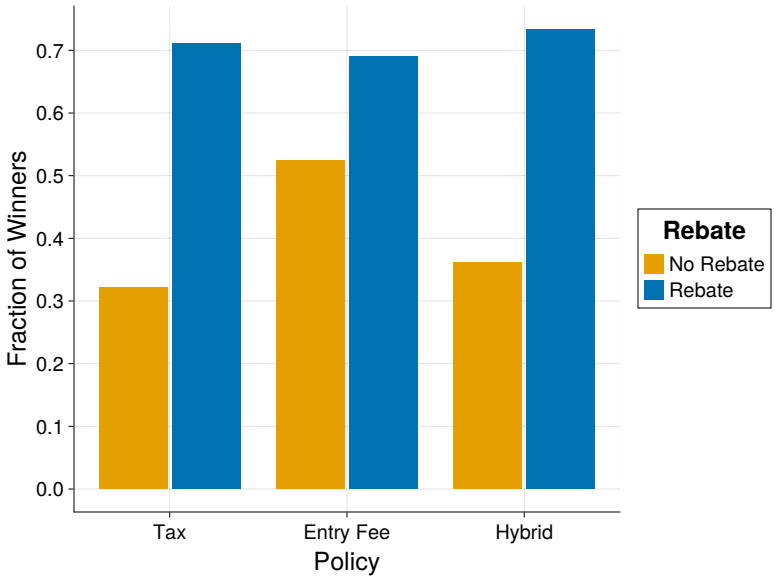
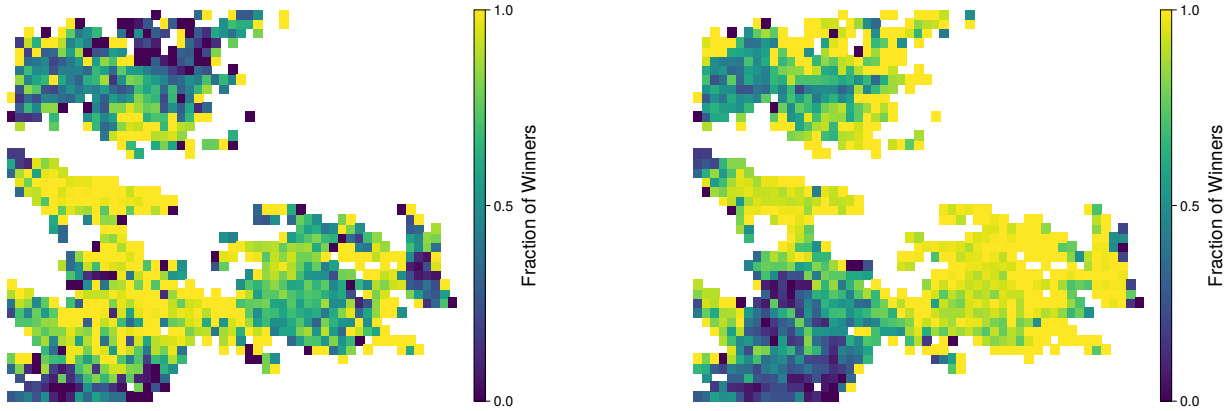


Figure 12: Fraction of winners by policy

Notes: This figure shows the fraction of fields that gain relative to the status quo under each uniform policy implemented in 1960, with and without lump sum rebates. The lump sum rebate amount is proportional to the size of the field.

portion are marginal to the entry fee policy while there are many marginal entrants in the northern and eastern portions of the aquifer. This pattern can be clearly seen in figure 14, which shows the average delay in entry across the aquifer from the tax policy in panel (a) and the entry fee policy in panel (b). While the tax policy has only marginal impacts on entry, the entry fee policy significantly delays, and often prevents entirely, entry in the less productive areas of the aquifer. Fields in these areas either gain as a direct result of improved aquifer conditions or, if they choose not to enter, from being compensated by the lump sum rebate as they are often close to indifferent between entering and not. By contrast, fields in the southwest portion realize few gains from the entry fee policy since most nearby fields still enter and the entry fee does not constrain pumping conditional on entry. These fields thus pay a large additional fixed cost for small gain and therefore often prefer the tax policy.

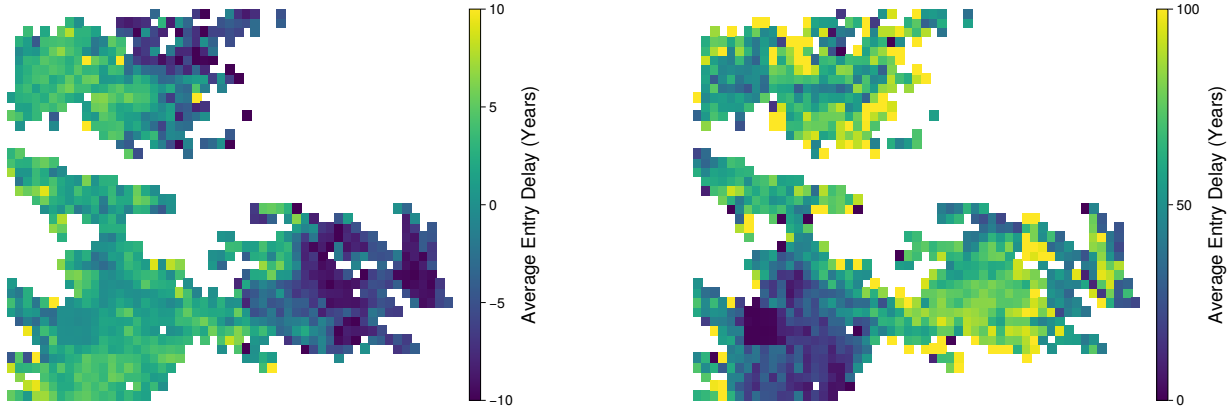
Figure 15 correspondingly shows the average gain in surplus across policies by both quintiles of productivity in panel (a) and externalities in panel (b). The tax and entry fee policies exhibit starkly different distributional patterns. Under the tax policy the gains are highest for high-productivity, high-externality fields while exactly the opposite is true for the entry fee policy. The difference is particularly stark for the highest productivity and externality quintiles which realize almost no surplus gains on average from the entry fee policy. The reason for this is that these fields are highly concentrated and almost never



(a) Tax (b) Entry Fee

Figure 13: Fraction of winners across space

Notes: This figure shows the fraction of fields that gain relative to the status quo under the uniform tax in panel (a) and under the uniform entry fee in panel (b) across grid cells of the aquifer with lump sum rebates. The lump sum rebate amount is proportional to the size of the field.



(a) Tax (b) Entry Fee

Figure 14: Entry delay across space

Notes: This figure shows the average delay in entry relative to the status quo under the uniform tax in panel (a) and under the uniform entry fee in panel (b) across grid cells of the aquifer. Fields that never enter are treated as if they entered in the year following the terminal year of the model, resulting in a maximum possible entry delay of 140 years.

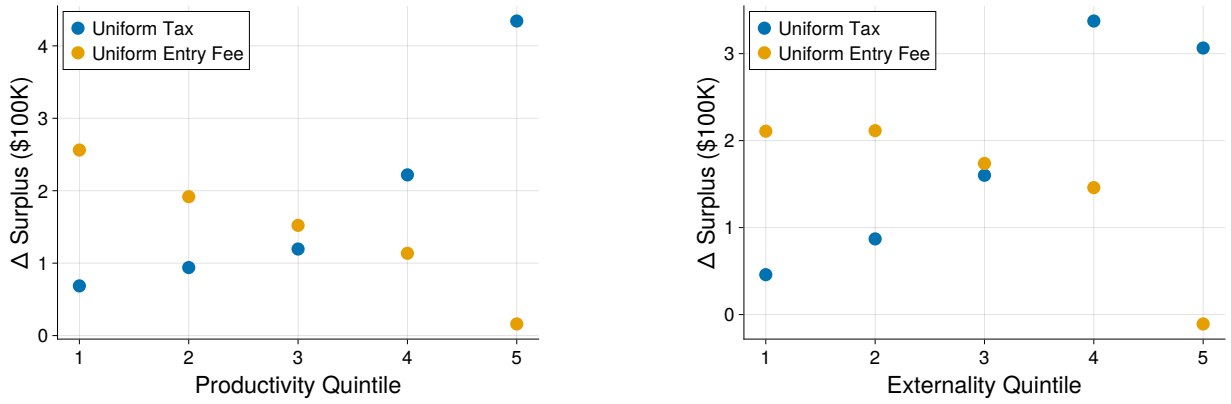
marginal to the entry fee. Because of the local nature of the aquifer, these fields therefore experience relatively little relief in depletion from entry declines in other areas of the aquifer but must still pay a large additional fixed cost of entry.

Figure 15 additionally shows that even though the tax increases welfare more than the entry fee policy, it does not dominate the entry fee policy. Indeed, for fields in the first three quintiles of productivity and externalities, the entry fee provides larger welfare gains on average than the tax policy. Many of these fields are either marginal entrants who benefit from large lump sum rebate payments while forgoing irrigation to continue rainfed production or are in relatively uncongested areas of the aquifer where the marginal externality is below the level of the tax. Moreover, spatial correlation in productivity means these fields are more likely to be located near other marginal entrants and therefore experience larger gains from the entry reductions of the entry fee policy.

A surprising consequence is that even though the tax policy provides more aggregate welfare than the entry fee policy, the entry fee policy is significantly more popular at the field-level, even with lump sum rebating. As shown in figure 16, more than 60% of fields prefer the entry fee policy to the tax policy. Almost all of the fields preferring the tax policy are located in the either the southwest area of the aquifer or small portions of the northern area where externalities are more closely aligned to the level of the tax. Low-productivity fields, especially in the eastern region of the aquifer, are typically close to indifferent between irrigated and rainfed agriculture in the status quo and can therefore benefit from the large lump sum payments they receive under the entry fee policy while remaining rainfed producers. Medium-productivity fields likewise are often located near many marginal entrants and thus can gain significantly from lower depletion rates due to fewer entrants under the entry fee policy. By contrast, while these fields also gain from lower levels of water use under the tax, the level of the tax is typically above their marginal externality and therefore reduces water use below the welfare-maximizing level. In some cases, this reduction can even be a larger distortion than the over-use under the status quo.

7.2.2 Policy Lock-In

There are two key sources of irreversibility that can entrench the effects of policy. First, well drilling investments are one-time sunk costs and enable pumping at relatively low marginal costs. Second, 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



(a) Productivity (b) Externality

Figure 15: Distributional consequences

Notes: This figure shows the average change in surplus under each policy by quintile of productivity in panel (a) and quintile of externality in panel (b). We measure externalities for each field in the status quo as of 1960 using the shadow value of water, $\frac{dV_{t+1}}{dW_{it}}$. We measure productivity for each field using the average irrigated productivity $\bar{\omega}_i^{\text{irr}}$ across all periods of the model.

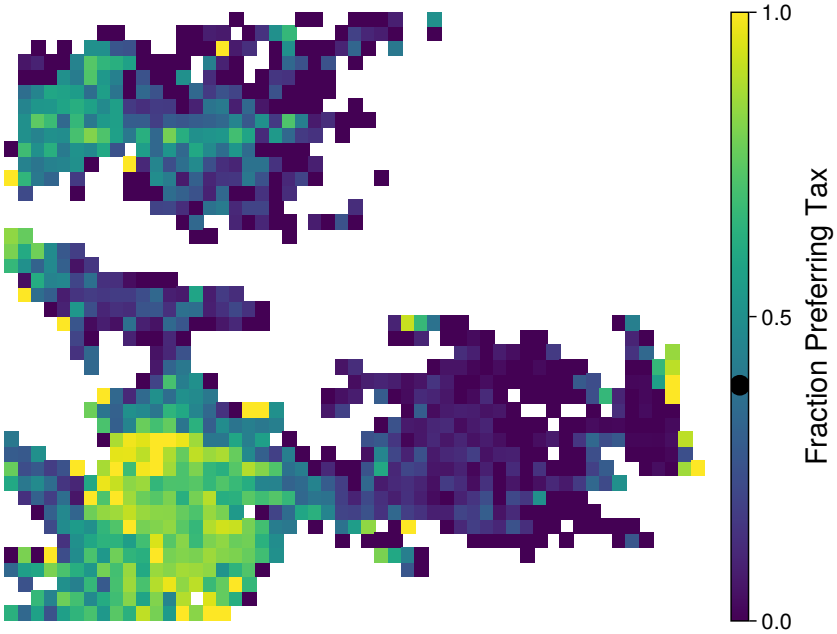


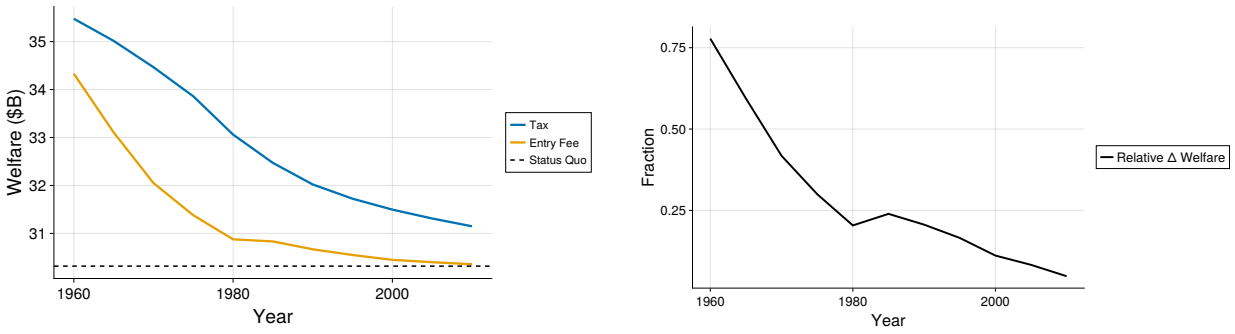
Figure 16: Policy preferences over space

Notes: This figure shows fraction of fields preferring the tax policy to the entry fee policy across grid cells of the aquifer. The solid circle represents the aggregate fraction preferring the tax policy to the entry fee policy across all fields, which is approximately 40%.

the welfare gains of optimal uniform taxes and entry fees as the start year of the policy becomes increasingly delayed. While we expect lock-in effects to be more severe for entry fees than taxes since entry fees can only affect fields that have not yet sunk wells, the question of political feasibility is ambiguous. On the one hand, entry fee policies are costless for incumbent fields that have already sunk wells. Thus, as more fields become incumbents the entry fee policy may become more popular over time relative to a tax policy. On the other hand, if the entry fee policy becomes relatively less effective compared to the tax policy over time, even incumbent fields may be willing to pay the cost of the tax to reduce the externalities imposed on them by other fields.

Figure 17 shows the welfare gains by policy start year separately for taxes and entry fees. Panel (a) shows the absolute welfare achieved by each policy compared to the status quo while panel (b) shows the relative welfare gains achieved by the entry fee policy compared to the tax policy. As we saw in table 4, entry fees beginning in 1960 are comparable to tax policies beginning in the same year, achieving more than 75% of the gains of the tax policy. However, while the effectiveness of both entry fee and tax policies declines over time, the decline is much more rapid for the entry fee as the locked-in fraction of depletion increases over time. This effect is particularly severe because, as we saw in figures 10 and 11, the correlation between productivity and externalities means that early entrants are also those with the largest externalities. This fact implies that the timing of policy is particularly important for entry fees in this setting to prevent high externality fields from entering.

Figure 18 shows the fraction of fields that prefer the entry fee policy to the tax policy by policy start year. As we saw in figure 16, more than 60% of fields initially prefer the entry fee policy to the tax policy, despite the fact that the tax policy results in larger welfare gains, and moreover that most fields have not yet entered and thus are subject to increased entry costs from the entry fee. However, the popularity of the entry fee actually declines over time, with the tax becoming preferred by a majority of fields in the mid-1980s. This finding runs counter to the common intuition that entry permitting policies should generally become more popular over time as more firms sink entry costs and therefore are no longer subject to the fee. The key source of this disparity is the rapid decline in effectiveness of the entry fee policy as seen in figure 17. This decline in effectiveness dominates the decline in costs imposed on incumbents and results in a net gain in fields preferring the tax policy over time. In other words, fields become increasingly willing to bear the costs of tax payments in exchange for a reduction in externalities as depletion becomes increasingly locked-in over time under an entry fee policy regime.



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

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

Figure 17: Policy Lock-In

Notes: This figure shows the welfare resulting from uniform water-use 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.

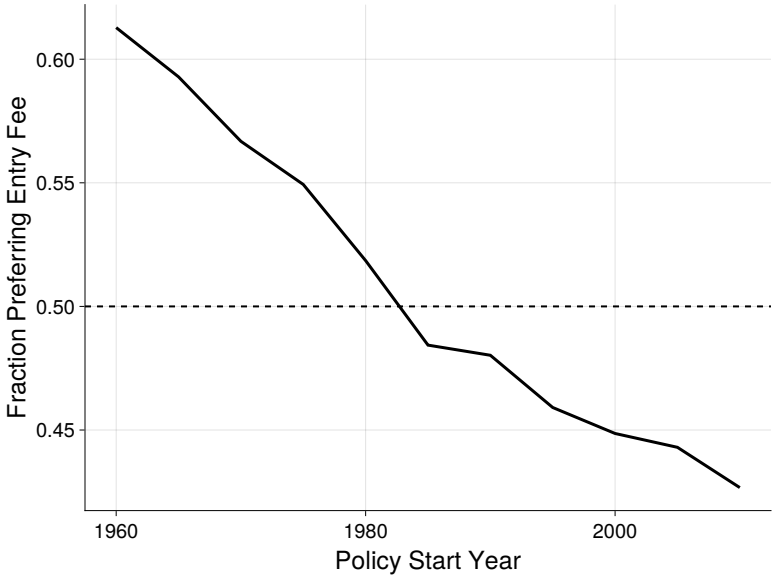


Figure 18: Fraction preferring entry fee over time

Notes: This figure shows the fraction of fields that prefer the entry fee policy to the tax policy over time with lump sum rebating compared to a benchmark of equal preference given by the dashed line.

8 Conclusion

This paper studies the relative effectiveness and political feasibility of entry fees and water-use taxes for limiting welfare losses from aquifer depletion in the Kansas High Plains aquifer. To do so, we develop a rich dynamic model of agricultural production and investment in groundwater wells and combine it with a physically realistic model of groundwater flows in a full equilibrium framework. We estimate the model using a field-level panel dataset on agricultural irrigation, well investment, and aquifer conditions in the Kansas High Plains aquifer spanning more than half a century. We show that externalities from pumping are highly heterogeneous across the aquifer and over time, driven by underlying heterogeneity in aquifer characteristics and the varying spatial density and productivity of wells and fields.

With uniform policies but heterogeneous externalities, the standard intuition that a water-use tax will outperform an entry fee can reverse. The key determinant of the relative welfare of the two policies is the correlation between productivity and externalities, as entry fees screen out low-productivity potential entrants without changing incentives for water use conditional on entry. Entry fees and water-use taxes may also differ significantly in the political support they engender, since entry fees do not impose costs on incumbents.

Under the status quo, we find that productivity and externalities are strongly positively correlated as a result of spatial correlations in productivity that lead to highly productive fields being concentrated in congested, severely depleted areas of the aquifer. As a consequence, we find that entry fees result in lower welfare than uniform water-use taxes. However, they still result in substantially increased welfare compared to the status quo, and moreover are preferred by a large majority of fields. This political preference is not driven by incumbency, as most fields have yet to invest in wells at the beginning of our sample period. Instead, it is driven by heterogeneity in externalities, as we find the level of the optimal uniform tax is above the marginal social cost of most fields.

We also find that the effectiveness of entry fees decreases rapidly over time because depletion becomes increasingly locked-in as irreversible well investments are sunk. This decrease in effectiveness is accelerated by negative selection in externalities over time, since high-externality entrants enter early and are thus unaffected by delayed entry fees. Contrary to the intuition that entry fee policies should generally become more popular over time because they do not burden incumbents, we also find that the popularity of entry fees falls over time, since the decline in costs imposed on incumbents is dominated by the decline in effectiveness relative to water-use taxes.

Looking forward, this paper leaves open multiple fruitful directions for future research. First, because almost all crops grown in Kansas are annuals such as corn or wheat, we ab-

stract from dynamics in crop choice that arise with perennial crops such as fruit and nut trees that are prevalent in many other settings. Extending our framework to incorporate crop choice as a dynamic investment decision would allow future work to study how water policy interacts with long-lived but reversible capital investments. Similarly, because water use in the Kansas High Plains aquifer is almost entirely agricultural, we focus on water use for irrigation rather than other sources of water use, such as municipal or industrial use. Inter-sectoral conflict in water allocation is likely to become increasingly important as climate change constrains surface water supplies and the expansion of water-intensive industries, including data centers, increases non-agricultural demand. Finally, the paper highlights that the political feasibility and efficiency of policies need not coincide. Future research could model the political dynamics of resource regulation explicitly, linking heterogeneous preferences over policy instruments to collective decision-making and long-run policy persistence.

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.

De Loecker, Jan, and Chad Syverson. 2021. "An industrial organization perspective on productivity." In *Handbook of Industrial Organization,* 4:141–223. Elsevier.

Diamond, Peter A. 1973. "Consumption Externalities and Imperfect Corrective Pricing." *The Bell Journal of Economics and Management Science* 4 (2): 526.

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.

Hotz, V. Joseph, and Robert A. Miller. 1993. "Conditional Choice Probabilities and the Estimation of Dynamic Models." *The Review of Economic Studies* 60, no. 3 (July): 497.

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).

Langevin, Christian D, Joseph D Hughes, Edward R. Banta, Alden Provost, Richard Niswonger, and Sorab Panday. 2017. *MODFLOW 6, the U.S. Geological Survey Modular Hydrologic Model,* September.

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.*

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.

Timmins, Christopher. 2002. "Measuring the Dynamic Efficiency Costs of Regulators' Preferences: Municipal Water Utilities in the Arid West." *Econometrica* 70 (2): 603–629.

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.

Woessner, William W., and Eileen P. Poeter. 2024. *Hydrogeologic properties of earth materials and principles of groundwater flow.* Guelph, Ontario: The Groundwater Project.

Appendix

A Data Construction

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 A1 overviews the full set of data sources we use for our analysis.

A.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%,

Table A1: 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

and since 1989 reporting has been mandatory and estimated compliance rates have risen to 93% (*Water Use Reporting* 2025). Figure A1 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 A2 shows the fraction of reports received from metered wells in each year. The fraction of metered wells reporting declined over the course of the 1960s and 1970s 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.

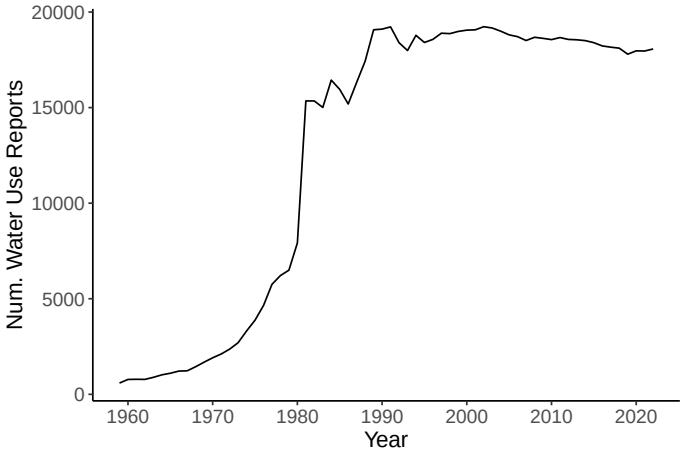


Figure A1: Number of water use reports by year

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

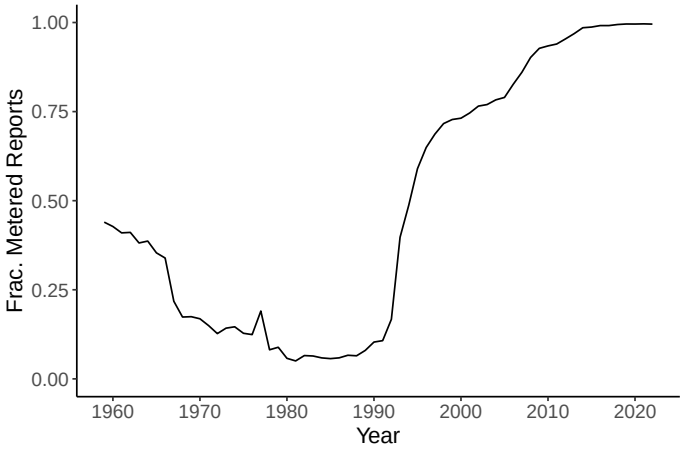


Figure A2: Fraction of metered water use reports by year

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

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 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).[25]

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.

25. 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.

A.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 expenditures as the sum of expenditures on fuel, seeds, fertilizer, and chemicals.[26] 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.[27] 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.

A.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 A3 plots the gridded recharge amounts across the aquifer. The western half of the aquifer receives almost no recharge. Recharge rates in the eastern half are also low on average but can reach up to 10 inches per year.

26. 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.

27. 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.

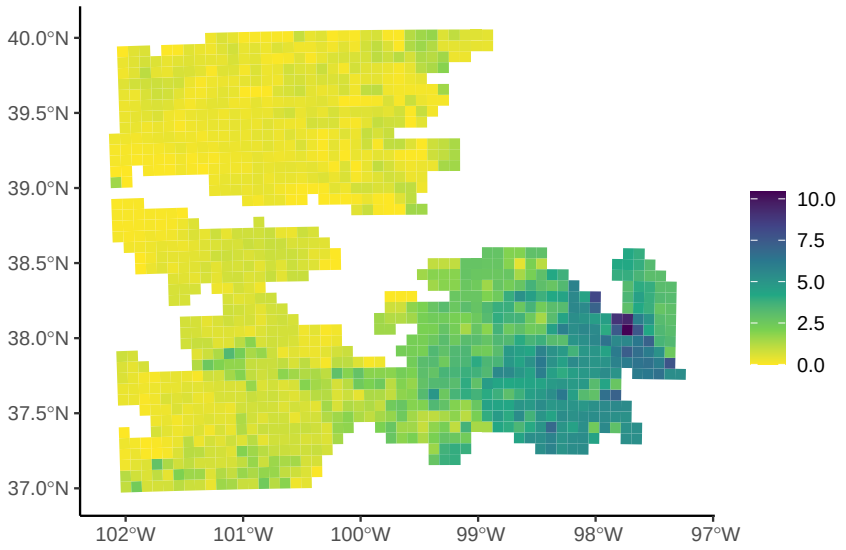


Figure A3: Potential recharge (in./year)

Notes: This figure shows the average potential recharge of the High Plains aquifer at a 5 mile resolution. Recharge is measured in inches per year.

B Proofs

B.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 8 yields

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

which is the expression in the text.

B.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 for R 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

$$\Omega\Phi(X_1\Sigma)\Sigma\Phi'(X_1\Sigma) = C\Phi(X_1\Sigma)$$
$$R\Sigma\Phi'(X_1\Sigma) = C\Phi(X_1\Sigma)$$
$$R = \frac{C\Phi(X_1\Sigma)}{\Sigma\Phi'(X_1\Sigma)}$$

□

C Computation

In this section we provide the computational details required for solving the model and constructing counterfactual equilibria. We also specify the additional long-run beliefs over weather, productivity, prices, and aquifer depletion that are necessary to compute these equilibria.

C.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[28] 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,[29] B^e is a matrix of weight adjustments,[30] 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

$$\hat{V}^e(s_{it}) = \mathbb{E}[\max\{\pi^0, \pi^1\}] + \beta \mathbb{E}[\hat{V}^e(s_{i,t+1})] \tag{35}$$

28. 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.

29. 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 $\mu \equiv \mathbb{E}[\vec{\epsilon}]$.

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

Equation 35 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 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

$$\hat{V}(z) = \sigma^e \ln(\exp(\hat{v}(0;z)/\sigma^e) + \exp(\hat{v}(1;z)/\sigma^e)) \tag{36}$$

$$\begin{aligned}\hat{v}(0;z) &= \mathbb{E}[\pi^0_{it} \mid s_{it} = z] + \beta E[V(s_{i,t+1}) \mid s_{it} = z] \\ &= \mathbb{E}[\pi^0_{it} \mid s_{it} = z] + (\tilde{\Phi} \odot B)\vec{w}\end{aligned} \tag{37}$$

$$\begin{aligned}\hat{v}(1;z) &= \mathbb{E}[\max\{\pi^0, \pi^1\}] - c^e(z) + \beta\mathbb{E}[\hat{V}^e(s_{i,t+1})] \\ &= V^e(s_{it}) - c^e(z)\end{aligned} \tag{38}$$

where the first line applies the closed-form for the expectation over type 1 extreme value random draws in equation 12.

C.2 Beliefs

Weather We specify beliefs over weather ER^j_{it} and ET^j_{it} as a function of location, particularly latitude and longitude, and year. In particular, we estimate the following regressions

$$ER^{\text{irr}}_{it} = \beta^{ER,\text{irr}}\vec{X}_{it} + \epsilon^{ER,\text{irr}}_{it} \tag{39}$$

$$ET^{\text{irr}}_{it} = \beta^{ET,\text{irr}}\vec{X}_{it} + \epsilon^{ET,\text{irr}}_{it} \tag{40}$$

$$ER^{\text{rainfed}}_{it} = \beta^{ER,\text{rainfed}}\vec{X}_{it} + \epsilon^{ER,\text{rainfed}}_{it} \tag{41}$$

$$ET^{\text{rainfed}}_{it} = \beta^{ET,\text{rainfed}}\vec{X}_{it} + \epsilon^{ET,\text{rainfed}}_{it} \tag{42}$$

where the vector $\vec{X}$ contains a constant, latitude, longitude, and the year. Latitude and longitude capture cross-sectional spatial variation in weather due to the rain shadow of the Rocky Mountains while the linear time trend captures temporal variation due to climate change. We include only first-order terms for simplicity, as higher-order terms do not significantly improve fit since most of the cross-sectional variation is directly east-west due to the north-south orientation of the Rocky Mountains. We assume field beliefs over weather correspond to the fitted values of these regressions.

Productivity We similarly specify beliefs over productivity as a function of location-specific characteristics—latitude and longitude, terrain, and soil quality—as well as the year. In particular, we estimate equations 9 and 11 and assume field beliefs about irrigated and rainfed productivity, respectively, correspond to the fitted values of these regressions. These

beliefs capture the expected ex-ante productivity of a field based on their field-specific characteristics and also allows these beliefs to evolve over time.

Prices We assume that fields treat prices as martingales and believe the price tomorrow will be the same as the price today. This martingale assumption reflects the fact that there are very active spot and futures commodity markets for both crops and energy in our setting.

Depletion We assume field beliefs about aquifer depletion correspond to the long-run average depletion rate θ^{depth}. In particular, let $\Delta D_{it} \equiv D_{it} - D_{i,t-1}$ be the realized depletion at field i from time $t-1$ to time t. We compute $\theta^{\text{depth}} = \frac{1}{NT}\sum_{t=1}^{T}\sum_{i=1}^{N}\Delta D_{it}$ and assume that fields believe that $D_{i,t+1} = D_{it} + \theta^{\text{depth}}$. We adopt this relatively simple belief structure for two reasons. First, forming more sophisticated beliefs over local depletion rates requires detailed knowledge of local hydrologic conditions and past local depletion rates, both of which were typically not easily observable by fields, especially in the period from 1960-1980 when most entry occurred. As we show in section D, our relatively simple belief model provides a good fit for observed entry rates. Second, as we detail below, constructing counterfactual equilibria requires finding a fixed point in a space with the same dimensionality as depletion beliefs. There are therefore significant computational advantages to a low-dimensional model of beliefs.

C.3 Computing Equilibria

To compute counterfactual equilibria corresponding to a counterfactual policy τ we first guess a path of aquifer depths and a corresponding value of the parameters governing beliefs over aquifer depletion, θ^{depth}.[31] We then compute the corresponding value functions V^{τ} and $V^{e,\tau}$ and 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 we select the n fields with the highest entry probabilities 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 large number of fields.

31. In practice we start at the path of depths realized under the status quo.

D Model Fit

Figure D1 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. While the model produces a smoother transition from the early period of rapid entry before 1980 to the more recent period of relatively slow entry, it is still able to closely replicate the patterns of entry in the data, with significant entry from 1960-1980 followed by a decline in entry in the ensuing decades.

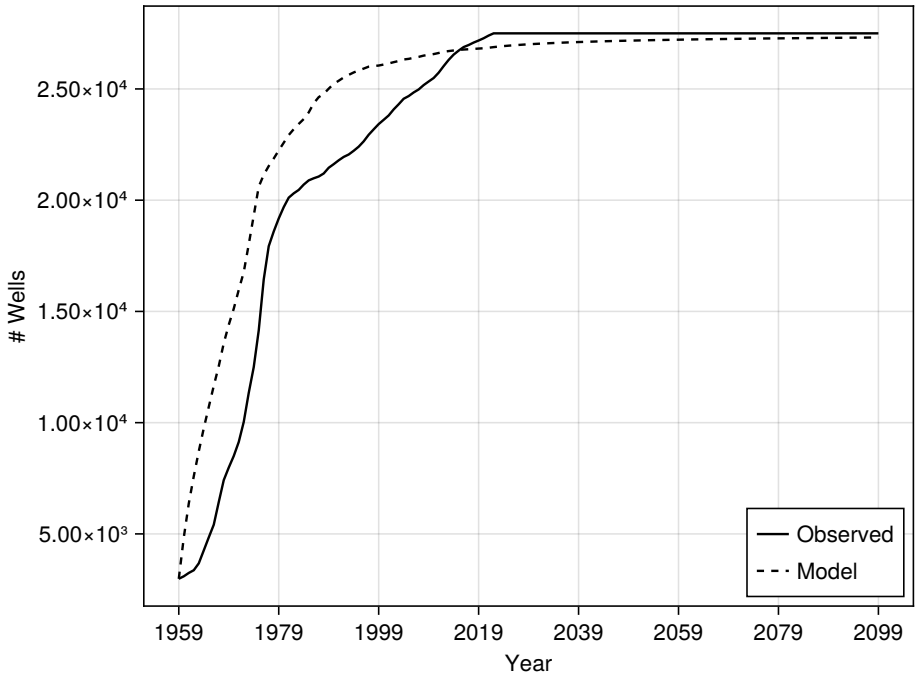


Figure D1: Cumulative number of wells by year

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).

E Atomistic Fields

A key assumption of our model is that fields are atomistic 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 E1 shows this change in costs averaged across 100 randomly chosen grid cells of the aquifer to account for differences in aquifer characteristics.

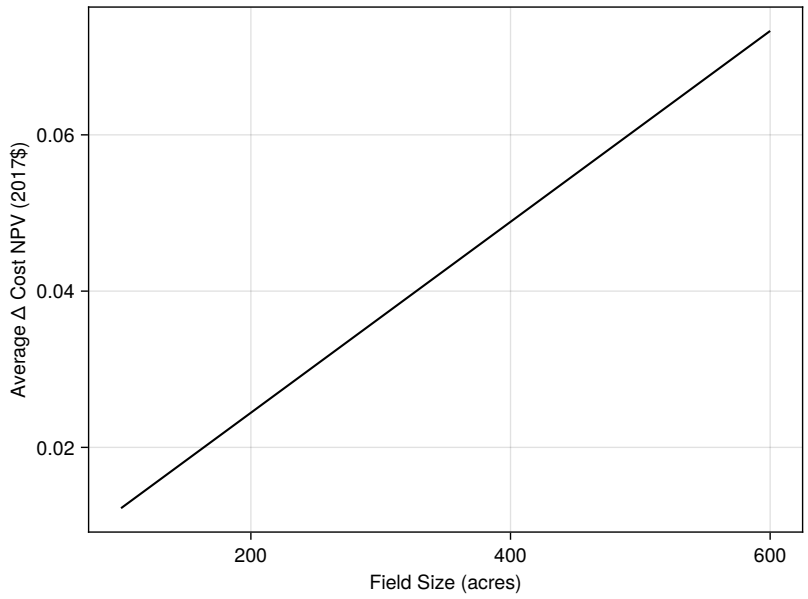


Figure E1: Changes in cost from field-specific pumping

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.

F Granular Counterfactuals

We consider two types of granular counterfactuals. First, we consider policies that vary spatially at the level of a GMD, which is the main spatial unit of groundwater regulation in Kansas. Second, we consider policies that vary over time. 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.[32] Table F1 overviews the results from each granular counterfactual policy we consider. As in table 4, column (1) shows the status quo outcome, which we take as a baseline.

Columns (2) and (3) solve for the welfare-maximizing spatially-varying water-use taxes and entry fees respectively. We find an average optimal tax of \$40.71/AF, compared to the optimal uniform tax of \$51.54/AF, and an average optimal entry fee of \$0.93M, compared to the optimal uniform entry fee of \$1M. Welfare increases to \$36.19B under the spatially-varying tax policy and \$34.45B under the spatially-varying entry fee policy, a relatively minor increase compared to \$35.48B and \$34.33B under the uniform tax and entry fee policies respectively.

Figure F1 shows the average tax level across space in panel (a) and the average entry fee level across space in panel (b). Reflecting the significant variation in productivity and externalities seen in figure 10, there is significant variation in both the tax and entry fee levels across space, with the tax increasing from less than \$20/AF to more than \$60/AF and the entry fee increasing from less than \$0.8M to more than \$1.2M. Figure F2 likewise shows the average change in welfare within each GMD. The entry fee increases welfare more than the tax policy in GMDs 1, 2, and 5 while the tax policy increases welfare more in GMDs 3 and 4, where both productivity and externalities are highest and thus the entry fee is unable to screen out high-productivity but high-externality fields.

Columns (4) and (5) solve for the welfare-maximizing time-varying water-use 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, as can be seen in panel (a) of figure F3 which plots the optimal tax by year. The resulting welfare is \$36.24B, an increase of \$124M/year in annualized terms, and the field depletion rate reduces to 28.75%.

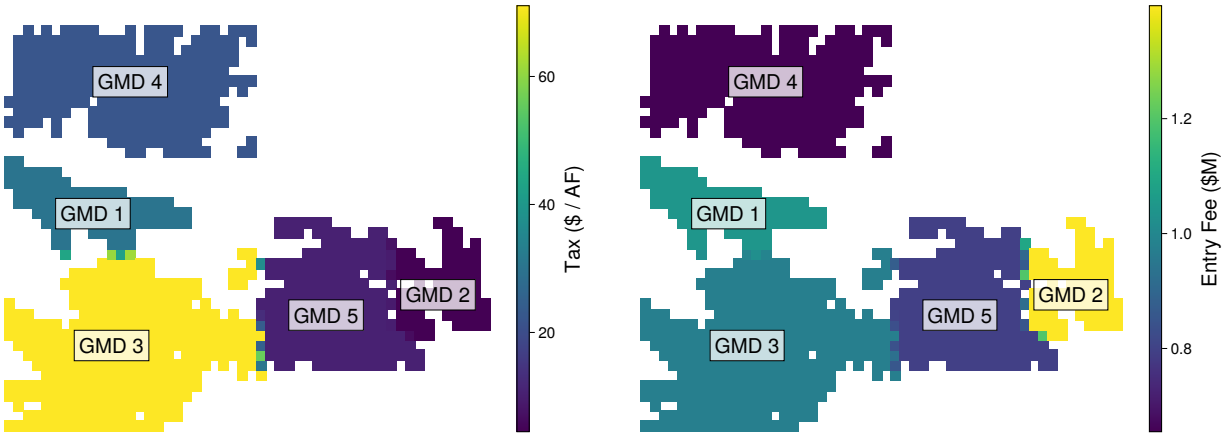
For comparison, we find an average optimal entry fee of \$1.33M which results in a welfare of \$34.9B, an increase of \$96M/year in annualized terms. Unlike the time-varying tax, the

32. 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.

Table F1: Granular Counterfactual Results

		Spatial		Time-Varying	
	Status Quo	Tax	Entry Fee	Tax	Entry Fee
	(1)	(2)	(3)	(4)	(5)
Policy					
Avg. Tax ($/AF)	0.00	40.71	0.00	71.75	0.00
Avg. Entry Fee ($M)	0.00	0.00	0.93	0.00	1.33
Welfare					
Welfare ($B)	30.32	36.19	34.45	36.24	34.90
Annualized Δ Welfare ($M)	0.00	122.87	86.44	123.93	95.78
Num. Wells	27,060	26,423	14,393	27,251	16,708
Δ Num. Wells	0	-637	-12,667	191	-10,352
Aquifer					
Remaining Storage (MAF)	174.37	214.19	228.25	261.32	221.60
Depleted Fields (%)	53.79	39.13	38.13	28.75	41.95

Notes: This table shows the results of model simulations under the status quo (column 1), a spatially differentiated water-use tax beginning in 1960 (column 2), a spatially differentiated entry fee beginning in 1960 (column 3), a time-varying water-use tax beginning in 1960 (column 4), and a time-varying entry fee beginning in 1960 (column 5). 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) Tax (b) Entry Fee

Figure F1: Spatial policy

Notes: This figure shows the average optimal per-AF water-use tax across space in panel (a) and the average optimal entry fee across space in panel (b) at a spatial resolution of five miles. Labels indicate the location of each GMD.

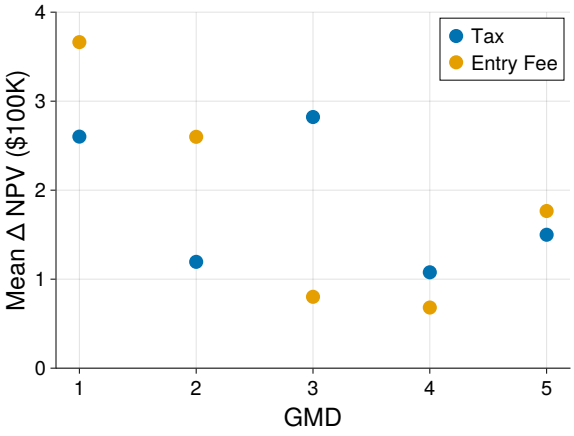
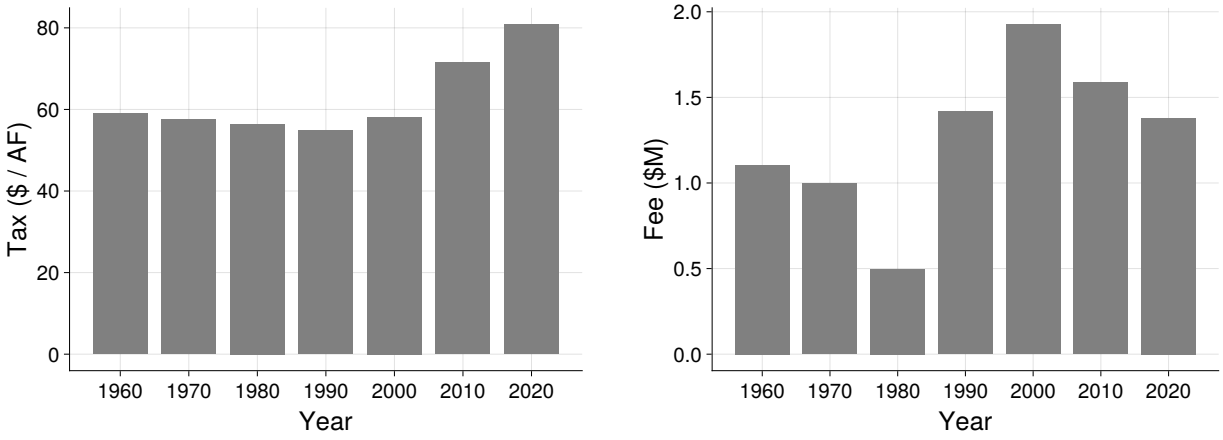


Figure F2: Welfare change from spatial policies

Notes: This figure shows the mean change in welfare relative to the status quo across GMDs for the spatial tax and entry fee policies.

optimal time-varying entry fee is actually U-shaped, with higher entry fees imposed at the beginning and the end. This can be seen in panel (b) of figure F3 which plots the optimal entry fee by year. 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 can be important for increasing welfare it does not change the comparison between entry fees and water-use taxes.



(a) Tax

(b) Entry Fee

Figure F3: Time-Varying policy

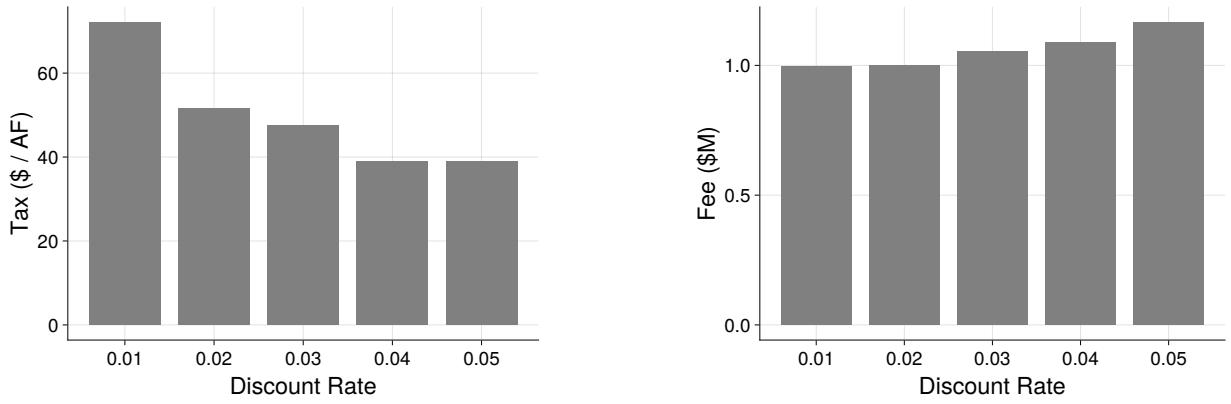
Notes: This figure shows the average optimal per-AF water-use tax over time in panel (a) and the average optimal entry fee over time in panel (b).

G Counterfactual Robustness

We consider the robustness of our counterfactuals along two dimensions, the assumed discount rate and assumptions about future climate change. Our baseline results assume a discount rate of 2% (Carleton and Greenstone 2021), and for robustness we consider a range of discount rates $r \in [0.01, 0.05]$. Our baseline results also assume future climate change proceeds according to the RCP4.5 scenario, which assumes approximately 2.5 to 3°C degrees of warming. For robustness we consider the RCP8.5 scenario, which assumes approximately 5°C of warming.

G.1 Discount Rate Robustness

We first consider the robustness of the optimal tax and entry fee to varying the discount rate. Panel (a) of figure G1 shows the optimal uniform tax starting in 1960 for discount rates from 0.01 to 0.05, and panel (b) shows the optimal uniform entry fee. With a discount rate of 1% the optimal tax is approximately \$70/AF, decreasing to approximately \$40/AF with a discount rate of 5%. This decline reflects smaller average externalities as future increases in pumping costs are weighted less with higher discount rates. The optimal uniform entry fee is less sensitive to the discount rate, increasing from approximately \$1M with a discount rate of 1% to approximately \$1.2M with a discount rate of 5%. The reason the entry fee increases with the discount rate is, as shown in section **??**, the optimal entry fee is U-shaped over time, with larger entry fees in both early and later years. As the discount rate increases more weight is put on the the middle period for which a lower entry fee is optimal.



(a) Tax

(b) Entry Fee

Figure G1: Robustness of optimal uniform tax and entry fee to discount rate

Notes: This figure shows the optimal uniform tax in panel (a) and entry fee in panel (b) starting in 1960 for discount rates from 0.01 to 0.05.

G.2 Future Climate Change

We next consider sensitivity to varying assumptions about future climate change. Figure G2 compares our baseline results under RCP4.5 to a more extreme climate scenario where we instead assume warming in line with RCP8.5.[33] For both the status quo and all counterfactual policies we consider, our welfare results are insensitive to the assumed climate scenario, with welfare decreasing only slightly under RCP8.5 compared to RCP4.5. We observe this pattern first because, due to discounting, most of our welfare results are driven by in-sample outcomes prior to 2022 for which the assumed climate scenario has no impact. Second, because western Kansas already has an exremely dry, semi-arid climate, the impacts of increased warming from climate change are small relative to the baseline rates of evapotranspiration.

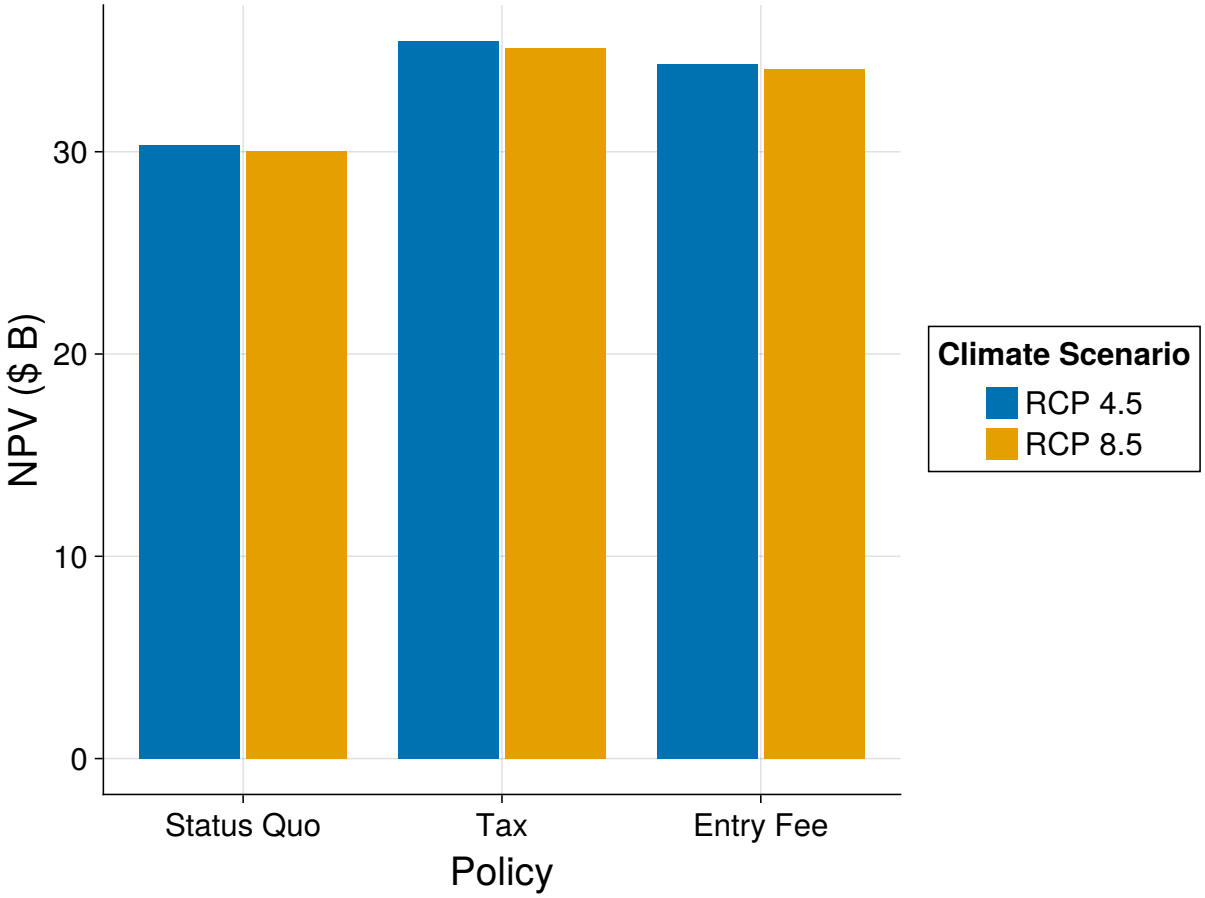


Figure G2: Welfare under RCP4.5 vs. RCP8.5

Notes: This figure shows welfare by climate scenario in the status quo, under the optimal uniform tax starting in 1960, and under the optimal uniform entry fee starting in 1960.

33. While climate change is projected to be good for crop yields in some areas, this is not the case in our setting, as climate change is projected to result in reduced rainfall and higher evapotranspiration.

H Hydrologic Model

We calibrate recharge amounts to match the observed path of aquifer depths in our data given water extraction amounts. Figure H1 shows the average path of aquifer depths in our model relative to the observed path of average aquifer depths in the data. The reason our model cannot exactly replicate aquifer depths in the data is that we impose constant recharge amounts over time. However, our model is still able to closely replicate the average path of depths. And while recharge amounts can in principle vary from year-to-year due to differences in precipitation and deep soil percolation, this variation is extremely small relative to water extraction levels from pumping.

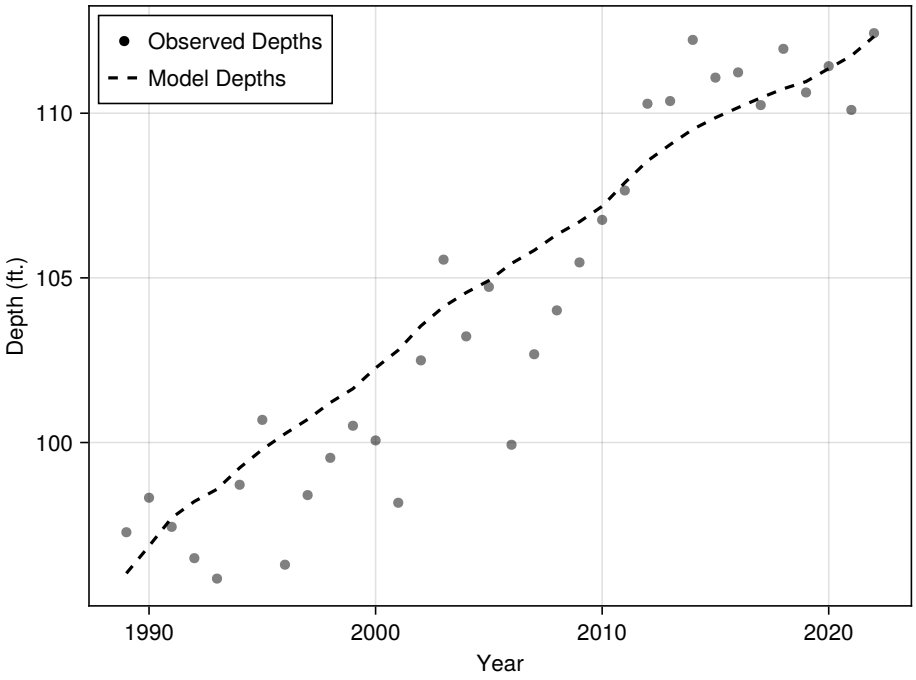


Figure H1: Hydrologic model fit

Notes: This figure shows the path of average aquifer depths under our calibrated model (dashed line) relative to the observed depths (gray dots).