

UC Berkeley

UC Berkeley Previously Published Works

Title

Are We Consuming Too Much Groundwater?

Permalink

<https://escholarship.org/uc/item/1nr9t71t>

Authors

Greaney, Brian

Shapiro, Joseph

Wagner, Katherine

Publication Date

2026-10-06

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Are We Consuming Too Much Groundwater?*

Brian Greaney
Federal Reserve Bank
of Dallas

Joseph S. Shapiro
UC Berkeley
and NBER

Katherine R. H. Wagner
University of British
Columbia

October 2026

Abstract

We study the optimality of human groundwater extraction from many of the world's 3,400 aquifers, where institutional or market failures may produce non-optimal extraction decisions. We use remote sensing and economic data to estimate a dynamic model of water extraction for each aquifer, recover the discount factor that rationalizes observed groundwater extraction, and compare it against Ramsey and market benchmarks. Three- to four-fifths of aquifers are extracted more rapidly than is optimal, decreasing aquifer present value by several percent, and creating trillions of dollars in present value welfare losses. Policies that guarantee users constant water quantities indefinitely, through subsidies or prior appropriation, or policy rules that fully discount future extraction, generate larger welfare costs. A sustainable extraction path that maintains current fill forever generates meaningfully lower welfare than the optimal extraction path, because most aquifer stocks are already inefficiently low.

JEL Codes: E22, H23, H43, Q25

Keywords: Groundwater extraction, dynamic programming, renewable resource management, tragedy of the commons, agricultural production

*We thank numerous colleagues and participants in seminars at ASSA, Kassel, McMaster, the National University of Singapore, Paris School of Economics, Queen's, Simon Fraser, Stanford, Temple, UC Berkeley, UC Davis, U Chicago-Coase, U Ottawa, and Washington State for useful comments; Sheah Deilami, Hila Etkin, Josh Geselowitz, Yifei Liu, Zhehao (Colin) Lu, Kendra Marcoux, Brant Walker, Claude, and refine.ink for excellent research assistance and feedback; and the Giannini Foundation for generous support. The views expressed in this paper are those of the authors and do not necessarily reflect those of the Federal Reserve Bank of Dallas or the Federal Reserve System. Greaney: brian.greaney@dal.frb.org; Shapiro: joseph.shapiro@berkeley.edu; Wagner: katherine.wagner@ubc.ca.

1 Introduction

Humans are rapidly depleting many of Earth’s natural resources. [Hotelling \(1931, p. 137\)](#)’s description of the problem nearly a century ago still resonates:

Contemplation of the world’s disappearing supplies of minerals, forests, and other exhaustible assets has led to demands for regulation of their exploitation. The feeling that these products are now too cheap for the good of future generations, that they are being selfishly exploited at too rapid a rate, and that in consequence of their excessive cheapness they are being produced and consumed wastefully has given rise to the conservation movement.

Debates about the optimality of resource consumption arise in part because net depletion of some resources may be socially optimal.

We study these issues for groundwater, which supplies 40% of the water used for global agriculture ([Ojha et al. 2018](#); [Rodell et al. 2018](#)). In response to falling water tables in some aquifers, governments in Arizona, California, Texas, Australia, India, Iran, and elsewhere have begun debating or implementing laws to regulate groundwater extraction. Leading media and policymakers highlight risks of excess depletion ([Rojanasakul et al. 2023](#)).

Priors vary on the efficiency of prevailing extraction. Because formal groundwater institutions are uncommon, many observers believe that most aquifers are common pools whose many atomistic users act myopically, that groundwater levels are below the social optimum, and that extraction exceeds the optimal level ([Stavins 2011](#)).¹ We are unaware of systematic tests of this prior. Informal institutions that achieve non-myopic extraction may temper this prior ([Ostrom 1990](#); [Meinzen-Dick and Bruns 2024](#)).²

We begin with a classic dynamic model of renewable resource extraction. Groundwater produces output with a decreasing marginal product, and an agent chooses extraction to maximize the present discounted value of profits. Consuming water today increases current profits but lowers the water table, which raises future extraction costs.

Applying the model quantitatively requires measuring three main terms. We calculate groundwater extraction using remote sensing data. We estimate the marginal cost of groundwater extraction as the product of the height that groundwater must be lifted from the aquifer to the Earth’s surface and the energy cost per unit of lift height. We back out the

¹Formal institutions are rules that governments write and enforce, such as tradable water rights, volumetric prices or taxes, extraction permits, or pumping quotas. Informal institutions are norms that communities enforce themselves, such as rotation schedules for shared wells or voluntary agreements to limit pumping.

²These debates are longstanding; the first article in the first issue of the *American Economic Review* debated U.S. water use policy ([Coman 1911](#)), and the centennial of that article emphasized that these debates have persisted ([Libecap 2011](#); [Ostrom 2011](#); [Stavins 2011](#)).

marginal return to groundwater using production function estimates. Our remote sensing and economic data on production inputs and outputs, climate, and geology over a decade describe many of the world’s approximately 3,400 aquifers.

Data show that the marginal returns to groundwater typically exceed static marginal costs of extraction. This rejects the hypothesis of myopia and suggests that extraction in most aquifers implicitly considers future extraction costs.

We use the calibrated model for four purposes. First, we recover the “decision” discount factor that rationalizes observed groundwater extraction in each aquifer. This discount factor measures how impatiently the representative agent extracts water, so it provides a continuous measure of resource management. We compare it against two benchmarks—a Ramsey discount factor (Ramsey 1928; Addicott et al. 2020) and the long-term interest rate on low-risk government debt (Carleton and Greenstone 2022; Rennert et al. 2022).

This comparison assesses the optimality of current extraction. If institutional or market failures keep the representative agent from fully internalizing how extraction today affects future extraction costs, the decision discount factor that guides extraction may differ from the benchmark discount factor in the agent’s utility function. We find that 61% of the world’s aquifers have decision discount factors below a Ramsey benchmark and 82% below a market benchmark, implying that groundwater extraction in these aquifers exceeds optimal rates. Many aquifers have decision discount factors within a few percent of benchmarks, so while excessive extraction may be common, its welfare costs may not be extreme.

Although our model describes a representative agent for each aquifer, most actual aquifers have many individual farms, and different regions within an aquifer may have different local institutions. We ask what (potentially non-optimal) extraction decisions by a representative agent would reproduce the outcomes of actual decentralized management. We also show numerically that this discount factor interpretation relates monotonically to the classic property rights comparison of common-pool versus single-agent extraction.

Throughout the paper, we highlight four case study aquifer systems: the Central Valley Aquifer System in California, the Ogallala Aquifer in the U.S. Great Plains, the North China Aquifer System, and the Indus Basin in South Asia.³ For example, California’s Central Valley Aquifer System has a decision discount factor of 0.89. Atomistic water users who apply the Ramsey benchmark discount factor but act as if 41% of future extraction costs depend on others’ extraction would produce the same observed extraction.

Our second set of results compares the current extraction path against the path under the

³Individual aquifers in a system, such as the Ogallala, differ in extraction and agriculture. The case studies use systems because their larger area makes remotely sensed inputs more precise and because depletion in these regions is often documented and debated for the system as a whole (e.g., Richey et al. 2015b; Rojanasakul et al. 2023; USGS 2023). Statistics on the full sample include the individual component aquifers.

benchmark discount factor. The median aquifer is in a non-optimal steady state. Its current annual net extraction exceeds the optimal rate by 0.1% to 0.5% of aquifer capacity, a large change from its baseline net extraction of 0.0%. Although median net extraction is zero, most individual aquifers have non-zero net extraction. More importantly, most aquifers have inefficiently low fill: the median aquifer’s gross extraction equals recharge, but its extraction exceeds the optimal level because its stock is below the optimal level. Lower short-run extraction would reach a steady state with higher fill and greater aquifer present value.

Third, we measure the difference in the present value of social welfare between the current and benchmark extraction paths, which we interpret as the cost of non-optimal resource management. We also calculate the present value of the additional extraction costs that the current path incurs. We estimate that all global aquifers are worth \$103 trillion or more, comparable to the value of all publicly traded firms, at \$117 trillion (OECD 2025), or global GDP, at \$111 trillion (World Bank 2023). We interpret this far-out-of-sample statistic cautiously, since it compares against draining every aquifer, though its order of magnitude illustrates the potentially large value of natural capital in this setting. The global welfare loss from non-optimal extraction is 2.6% of aquifer present value. Three forces keep this percentage modest: dynamic feedback from higher future extraction costs, decision discount factors within a few percentage points of benchmarks, and the “tyranny of discounting” (Pearce et al. 1989) whereby large resource losses in the distant future have modest present values.

Fourth, we study three other non-optimal but potentially realistic extraction paths: maintaining current gross extraction forever (“constant extraction”), extracting myopically, or maintaining current fill indefinitely by consuming recharge each year (“sustainability”). Constant extraction lightly echoes the prior appropriation doctrine in Western U.S. surface water law (DWR 2019). Subsidies that fix real water prices could also produce constant extraction. Myopic extraction is equivalent to common-pool atomistic agents using a discount factor of zero. Zero net extraction resembles quantity restrictions such as California’s Sustainable Groundwater Management Act.

All three counterfactuals have high costs relative to the optimal extraction path, and the first two are costlier than observed extraction. Constant extraction eventually drains 31% of aquifers completely. Relative to the benchmark path, constant extraction has present value welfare costs of \$4 to \$13 trillion. Myopic extraction also imposes large welfare costs—in steady state, myopia completely drains 10% of aquifers. Myopic and constant extraction generate welfare costs of 3.4% to 12.9% of aquifer present value. Zero net extraction (sustainability) fully drains no aquifers. Sustainable extraction costs less than the other two counterfactuals, though it still falls short of the optimal path because it keeps aquifers from

reaching higher fill.

As one validation test, we calibrate the model to 2003-2009 and predict extraction in 2010-2016. Across aquifers, these predictions have high correlation with remotely-sensed extraction in 2010-2016. Discount factors, long-run fill, and long-run marginal extraction costs estimated separately for the two periods also have high correlation.

We contribute to several literatures. We provide the first estimates of actual versus optimal extraction paths and present discounted value for most of the world’s aquifers. Existing empirical work typically studies a single aquifer or region with sufficient data (e.g., [Timmins 2002](#); [Hornbeck and Keskin 2014](#); [Ayres et al. 2021](#); [Burlig et al. 2026](#); [Berman and Hickok 2025](#); [Hadachek et al. Forthcoming](#)). Many scholars conjecture from extraction trends or management regimes that most global groundwater stocks are below optimal levels ([Stavins 2011](#); [Lawell 2016](#); [Edwards and Guilfoos 2021](#); [Rodella et al. 2023](#)). We systematically test this hypothesis. Our approach—interpreting remote sensing and economic data for many endowments of a natural resource through a classic model of resource management—could apply to other resources, complementing approaches that measure net capital depletion for resources with market prices ([Arrow et al. 2004](#)). Our aquifer valuations can inform natural capital accounts ([Prabhakar et al. 2023](#)).

Like [Carleton et al. \(2025\)](#), we combine global spatial and remotely sensed data on water and agriculture, including GAEZ crop data and GRACE water storage. Both papers model extraction costs that rise as water tables fall, and calibrate models to observed aquifer levels and trends. [Carleton et al. \(2025\)](#) assume farmers extract water under open access, which makes each farmer’s problem static, and their spatial model quantifies how trade and agricultural policies reallocate water use across aquifers. We estimate the extraction behavior that characterizes each aquifer—the discount factor that rationalizes its observed extraction path—and test whether that path maximizes the aquifer’s present value. Their model quantifies the changes in food prices, production, and yields from moving water use across regions; ours quantifies the present value welfare effect of moving extraction within an aquifer across time.

We also relate a new empirical measure of resource management to a classic theoretical literature on extraction and the commons ([Gordon 1954](#); [Coase 1960](#); [Dasgupta and Heal 1979](#); [Ostrom 1990](#); [Provencher and Burt 1993](#)). Where theory typically contrasts two polar cases—common-pool versus single-agent extraction ([Levhari and Mirman 1980](#); [Karp and Sakamoto 2021](#))—we estimate the discount factor that rationalizes observed extraction, which provides a continuous measure of management, maps directly to dynamic economics, and relates monotonically to the property rights interpretation.⁴

⁴[Quaas et al. \(2012\)](#) estimate the shadow interest rate at which 13 European fisheries borrow from natural

Our findings recall [Gisser and Sanchez \(1980\)](#)’s classic argument that the welfare gap between common-pool and optimal groundwater extraction is small, though our framework includes depletion, exhaustion, and lift constraints, which can overturn their result ([Koundouri 2004](#); [Edwards and Guilfoos 2021](#)). Modest average losses here are therefore a result rather than a foregone conclusion. The low average cost of current practices also masks variation: some aquifers have losses exceeding 10% of their value, and myopic or constant extraction produces losses several times larger than the current path.

Our paper has important limitations. First, like a long resource literature since [Hotelling \(1931\)](#), we study single-agent dynamic optimization; the conclusion discusses linking intertemporal and interregional optimization. Second, like many hydrological models, we study a single-cell model of an aquifer with vertical sides and known depth. We allow recharge, lift height, specific yield, capacity, and most other inputs to differ across aquifers, which captures some of the variation that models of within-aquifer hydrodynamics generate, but we abstract from spatial forces within an aquifer such as lateral flow and local cones of depression, which are harder to address globally. Third, we focus on the classic stock externality, in which depleting a resource today raises future marginal extraction costs. In practice, aquifers can generate other externalities, such as interactions with surface water. We focus on the stock externality since it has been central to resource management debates for decades, ties closely to dynamic optimization, and has data for most of the world’s aquifers.


We proceed as follows. Section 2 provides background. Section 3 describes the model. Section 4 discusses data. Section 5 explains model estimation and calibration. Section 6 compares decision and benchmark discount factors. Section 7 studies resource management and welfare. Section 8 concludes.

2 Background on Water Resources

Groundwater fills underground aquifers, which can form over millennia. It is one component of total terrestrial water storage, which also includes moisture in soils and vegetation, snow, and surface waters such as rivers and lakes. Groundwater accounts for 20% of global fresh-water. Agricultural irrigation accounts for 70% of global groundwater extraction ([Ojha et al. 2018](#)). Groundwater extraction has grown as it increasingly substitutes for surface water, and analysts question this trend’s sustainability ([Siebert et al. 2010](#); [Rodell et al. 2018](#)).

Most of the Earth’s land surface stores groundwater. A group of hydrological agencies identifies 3,424 distinct global aquifers, ranging from major basins to small and shallow capital stocks, a concept related to our estimate of the discount factor that rationalizes observed extraction.

Figure 1: Aquifer Structure and Terminology



aquifers (BGR 2025). Aquifer systems group nearby aquifers that share geology and hydrology and may slowly exchange groundwater through lateral flow (Miller 2003).

We choose the four case study aquifer systems mentioned in the introduction (Central Valley, Ogallala, Indus, North China) given their importance by various metrics (Borunda 2022). For example, they are four of the 37 “major” aquifer systems that some hydrologists highlight (Richey et al. 2015a). Appendix A provides further background on them.

Figure 1 illustrates aquifer terminology. Aquifer dynamics depend on inflows (recharge) and outflows (extraction). An aquifer’s storage volume or reserves equal its current available water. Its capacity equals its maximum possible available water. An aquifer’s level is the top of its water table. When gross extraction exceeds recharge, reserves fall and water levels decline. Declining reserves raise the marginal cost of extraction, since users must lift groundwater further. Recharge depends on rainfall, surface water runoff, and geology. Rainfall recharges aquifers through porous rock but runs off non-porous rock as surface water (e.g., streams). Irrigation water can also recharge aquifers.

Figure 1 also clarifies other aquifer attributes. Groundwater fills only part of an aquifer; topsoil, rock, air, and other material fill the rest. Saturated thickness describes the vertical extent of groundwater mixed with rock, soil, and air. Groundwater extraction ultimately depends on the specific yield—the fraction of the saturated thickness that users can drain.⁵ Specific yield varies within an aquifer by rock type and is below 100%, partly since molec-

⁵We use the terms groundwater extraction, consumption, use, and depletion somewhat interchangeably.

ular forces attract water to soil and rock particles. Lower depths have lower specific yield (Sutanudjaja et al. 2018). Much of our analysis considers groundwater availability, which equals surface area times saturated thickness times specific yield (i.e., extractable volume or fill) (Faunt 2009; Richey et al. 2015a). Capacity equals maximum possible extractable fill. To measure capacity, we add the vertical distance from the bottom of the topsoil (just below the surface of the Earth) to the saturated thickness and multiply by surface area and specific yield, as above.⁶

Consider a simple example: a cubic aquifer 10 m on each side whose top lies 10 m below the ground. It occupies 1,000 m³; with a specific yield of 0.5, half of that volume could hold extractable water, so its capacity is 500 m³. If the saturated thickness is only 5 m—water fills only the bottom half of the aquifer—then groundwater availability is 250 m³.

Groundwater extraction requires drilling into the aquifer’s saturated thickness and pumping water to the surface, typically with diesel or electricity. Pumping groundwater can account for half of agricultural energy costs (Pradeleix et al. 2015). Pump costs depend on lift height, energy prices, and pump efficiency. An aquifer at capacity has positive lift height since users must pump groundwater through topsoil.

Groundwater institutions vary by region, though formal groundwater markets are rare (Carleton et al. 2025). California is unusual in recently beginning to enforce the Sustainable Groundwater Management Act, which uses price, quantity, and other tools to decrease extraction of the Central Valley Aquifer System. Much of India rations electricity as an indirect tool for groundwater management (Ryan and Sudarshan 2022).

3 Model

3.1 Model Description

A single agent manages each aquifer i . After observing the water available a at the beginning of a period, the agent chooses how much water w to consume. Each period, aquifer i receives a recharge quantity of available water r_i . Gross extraction of water and recharge occur at constant rates within each period. Without loss of generality, we measure a , r , and $w \in [0, 1]$ as a fraction of the aquifer’s capacity. Available water in aquifer i evolves according to the

⁶Water in topsoil is typically not included in groundwater, though is part of total terrestrial water storage. We treat the bottom of the topsoil as the upper limit of the aquifer. Some “confined” aquifers have impervious rock beneath the topsoil, which further limits the aquifer’s vertical extent. Global hydrological models typically do not distinguish between confined and unconfined aquifers due to lack of global data on these features and model only unconfined dynamics, a convention we adopt here. These models also typically omit lateral flow between aquifers, for similar data reasons (Richey et al. 2015b; de Graaf et al. 2015; Sutanudjaja et al. 2018; Müller Schmied et al. 2021).

law of motion

$$a' = \min\{a + r_i - w, 1\} \quad (1)$$

Gross extraction w is limited to available water reserves plus recharge $a + r_i$ since the agent cannot consume water from an empty aquifer. Recharge received when an aquifer is already full becomes surface water.

Agent i produces a numeraire agricultural good using the following technology:⁷

$$y_i = \mathcal{Z}_i k^\alpha l^\gamma w^{\phi_i} x_i^{1-\alpha-\gamma-\phi_i} \quad (2)$$

Here $\mathcal{Z}_i$ represents total factor productivity, k is capital, l is labor, and x is land. We use i subscripts for time-invariant and aquifer-specific variables and functions. Each period, the agent in each aquifer chooses k , l , and w . We assume that the price of capital p_i^k and the price of labor p_i^l are fixed, that the agent owns the land, and that current technology and preferences persist in future years.⁸ Sensitivity analyses include versions of w accounting for surface water and precipitation.⁹ The water share ϕ_i may vary across aquifers with characteristics that affect water demand. Appendix B shows that these assumptions imply the following profit-maximizing output for a given choice of gross water extraction w :

$$y_i(w) = z_i w^{\phi_i/\theta} \quad (3)$$

Transformed productivity z_i is then a function of terms the agent treats as fixed:

$$z_i = [\mathcal{Z}_i x_i^{1-\alpha-\gamma-\phi_i} (\alpha/p_i^k)^\alpha (\gamma/p_i^l)^\gamma]^{1/(1-\alpha-\gamma)} \quad (4)$$

The cost share of natural resources (water+land) is $\theta \equiv 1 - \alpha - \gamma$. Appendix B also shows that ignoring water costs, profit is

$$\tilde{\pi}_i(w) = \theta z_i w^{\phi_i/\theta} \quad (5)$$

⁷Much work assumes Cobb-Douglas agricultural technology, especially in micro settings (e.g., [Gollin et al. 2014](#); [Tombe 2015](#); [Donovan 2021](#); [Chen et al. 2023](#)), while some papers more focused on water assume a farm-level saturating response or Leontief technology for water within crops ([Berman and Hickok 2025](#); [Carleton et al. 2025](#)). Classic results identify settings under which Leontief production by individual agents produces aggregate Cobb-Douglas technology ([Houthakker 1955](#)).

⁸Sufficient conditions for fixed factor prices would be that each aquifer is a small open economy with frictionless intra-national factor mobility. Perfectly elastic labor supply is more plausible for the smooth, foreseeable changes along the balanced growth path we study than for a response to an unanticipated shock. Elastic labor could also reflect surplus workers in agricultural areas, particularly in developing countries.

⁹We abstract from backstop technologies like desalination, which is rarely used for irrigation due to its considerable cost ([Zolghadr-Asli et al. 2023](#); [Nargi 2024](#)).

The cost of water extraction reflects the energy needed to lift water to the surface. Let $\underline{h}_i$ and $\bar{h}_i$ denote the aquifer's minimum and maximum lift height, i.e., the lift height when the aquifer is full or empty (see Figure 1). Because recharge and extraction occur at constant rates throughout a period, for a period where initial water is a and gross extraction is w , mean lift height is

$$h_i(a, w) = \max \left\{ \underline{h}_i, \underline{h}_i + \left(1 - a + \frac{w - r_i}{2} \right) (\bar{h}_i - \underline{h}_i) \right\} \quad (6)$$

When an aquifer has initial fill a , extracting w units of water has total cost

$$\mathcal{C}_i(a, w) = c_i h_i(a, w) w \quad (7)$$

Here c_i represents an aquifer-specific marginal cost per unit of water $\times$ lift height, reflecting pump efficiency and energy prices. Given available water a and gross water extraction w , the representative agent earns profit

$$\pi_i(a, w) = \tilde{\pi}_i(w) - \mathcal{C}_i(a, w) \quad (8)$$

The representative agent solves the following Bellman equation:

$$V_i(a; \beta_i) = \max_{w \leq a + r_i} \pi_i(a, w) + \beta_i V_i(\min\{a + r_i - w, 1\}; \beta_i) \quad (9)$$

The discount factor β_i , central to our analysis, describes the representative agent's patience, given market or institutional failures like weak property rights. Our notation emphasizes the value function's dependence on β_i , which we assume is time-invariant for each aquifer, though we test whether it changes within our sample period. The second term on the right-hand side of (9) is the continuation value, which depends on current extraction w and recharge r_i . The agent faces a trade-off—increasing water consumption raises current profits but also raises future extraction costs, which decreases future profits. We let $w_i^*(a; \beta_i)$ denote the optimal policy function of (9).

We use the first-order condition of equation (9):

$$\frac{\partial \tilde{\pi}_i(w)}{\partial w} \geq \frac{\partial \mathcal{C}_i(a, w)}{\partial w} + \beta_i \frac{\partial V_i(\min\{a + r_i - w, 1\}; \beta_i)}{\partial a} \quad (10)$$

Equation (10) holds with equality for an interior solution where $w < a + r_i$. The left-hand side is the marginal revenue from extracting w units of water, and the right-hand side is the marginal cost of extraction. The marginal cost includes two terms—the current marginal

cost $\partial \mathcal{C}_i(\cdot)/\partial w$ and the discounted marginal dynamic cost $\beta_i \partial V_i(\cdot)/\partial a$. Our quantitative approach uses calibration and estimation to find values for all components of equation (10) besides the discount factor β_i , then infers the decision discount factor as the value that makes this first-order condition rationalize observed data.

While equation (9) characterizes the value of the optimal consumption path, we can also value arbitrary, non-optimal consumption paths. Let $\mathbf{w} = \{w_{it}\}_{t=0}^{\infty}$ represent an arbitrary water consumption path satisfying the feasibility constraint $w_{it} \leq a_{it} + r_i$, given initial water level a_0 . The value $\mathcal{V}_i$ of this arbitrary consumption path, discounted at β_i , equals the present discounted value of future profits:

$$\mathcal{V}_i(\mathbf{w}, a_0; \beta_i) = \sum_{t=0}^{\infty} \beta_i^t \pi_i(a_{it}, w_{it}) \quad (11)$$

Here the law of motion (1) determines the path of water availability $\{a_{it}\}_{t=0}^{\infty}$. Let $\mathbf{w}_i^*(a_0; \beta_i)$ denote the optimal gross water extraction path when initial fill is a_0 and the discount factor is β_i . The value of the optimal water extraction path equals the solution to (9):

$$\mathcal{V}_i(\mathbf{w}_i^*(a; \beta_i), a; \beta_i) = V_i(a; \beta_i)$$

Section 6.4 compares our discount factor interpretation of β_i to the literature's discussion of common-pool versus single-agent extraction.

3.2 Welfare Analysis

Welfare Loss from Non-Optimal Discounting

We calculate the welfare loss from the current extraction path as the difference in the present value of aquifer agricultural profits between the extraction paths under the decision discount factor β_i^d and under the benchmark discount factor β_i^* , which the representative agent would use absent institutional or market failures.¹⁰ Failures of property rights, common-pool management, or political discount rates may lead to inefficiently rapid discounting ($\beta_i^d < \beta_i^*$).

The welfare cost of the potentially non-optimal current extraction path, expressed as a fraction of the maximum aquifer value, is

$$\omega_i = \frac{V_i(a_{i0}; \beta_i^*) - \mathcal{V}_i(\mathbf{w}_i^*(a_{i0}; \beta_i^d), a_{i0}; \beta_i^*)}{V_i(a_{i0}; \beta_i^*)} \quad (12)$$

¹⁰We call this the benchmark rather than the optimal discount factor partly because we solve for the optimal consumption path given any discount factor, and partly because preference parameters are not typically objects to optimize. Our comparison of decision and benchmark discount factors parallels behavioral analysis of decision versus experienced utility (Kahneman and Sugden 2005).

Later we estimate ω_i for each aquifer.¹¹ To calculate ω_i , we solve the agent’s problem (9) to obtain the maximum value of an aquifer $V_i(a_{i0}; \beta_i^*)$ under the benchmark discount factor β_i^* . We also solve (9) with the decision discount factor to obtain the current extraction path $\mathbf{w}_i^*(a_{i0}; \beta_i^d)$. Evaluated using the benchmark discount factor β_i^* , the current water extraction path generates value $\mathcal{V}_i(\mathbf{w}_i^*(a_{i0}; \beta_i^d), a_{i0}; \beta_i^*)$, which is less than the maximum possible value: $\mathcal{V}_i(\mathbf{w}_i^*(a_{i0}; \beta_i^d), a_{i0}; \beta_i^*) < V_i(a_{i0}; \beta_i^*)$.¹²

3.3 Recovering Discount Factors

We use value function iteration to solve the agent’s problem, recover the decision discount factor β_i^d for each aquifer i , and calculate welfare consequences of non-optimal discounting; Appendix C.1 details the algorithm. Given an arbitrary candidate discount factor β_i , we solve the agent’s dynamic water extraction problem (9). This solution gives the agent’s extraction at the aquifer’s observed water level, $w_i^*(a_{i0}; \beta_i)$. We find the decision discount factor $\beta_i^d \in [0, 1)$ such that the model-implied water extraction $w_i^*(a_{i0}; \beta_i^d)$ equals current observed water extraction w_i . If no feasible $\beta_i^d \in [0, 1)$ matches observed extraction, we exclude the aquifer, though a sensitivity analysis includes it with the closest feasible discount factor.

The estimated policy function $w_i^*(a; \beta_i^d)$ describes extraction in aquifer i given beginning-of-period stock a . Iterating forward on the law of motion (1) from the observed stock a_{i0} gives model-predicted paths of extraction, aquifer level, profit, and extraction costs T periods ahead: $\{w_{it}, a_{it}, \pi_{it}, ec_{it}\}_{t=0}^T$. We iterate to $T = 1000$ periods and verify ex post that water levels in all aquifers converge by then.

4 Data

Applying the model requires measuring three groups of objects for each aquifer: the physical state and flow of groundwater; the marginal return to groundwater in production; and the marginal cost of extraction, plus benchmark discount factors. Sections 4.1-4.4 describe the data behind each; Section 5 converts them into model inputs. Appendix D provides further details and Appendix Table A.1 summarizes data sources.

¹¹Welfare in our model is the present discounted value of agricultural profits. These quantities are equivalent under the model’s assumptions (i.e., the prices of agricultural output, labor, and capital are fixed and the representative agent owns the land) with two additional assumptions: the benchmark discount factors are what the planner would choose, and there are no market failures other than the choice of discount factor.

¹²The representative agent could have chosen the estimated water consumption path, since it is feasible. Evaluated using the benchmark discount factor β_i^* , no feasible water consumption path can yield value above $V_i(a_{i0}; \beta_i^*)$.

Several measurement conventions apply across our datasets. For gridded datasets, we take an unweighted average of grid points within each aquifer $\times$ year. We deflate values to 2023 dollars using the U.S. GDP implicit price deflator (FRED 2024). For any variable, q_i denotes its value as a share of aquifer capacity κ_i and $\tilde{q} \equiv q_i \kappa_i$ its level in m^3 ; for example, w is gross extraction as a share of capacity, and $\tilde{w}$ is gross extraction in m^3 .¹³ We analyze annual data, since most of our data lack sub-annual frequency. For data reported by country, for an aquifer that spans borders we take the mean of the data across countries, weighted by the share of the aquifer’s area in each country. For the case study aquifer systems, we average component aquifers, weighted by their shares of each system’s area.

The final analysis dataset has 38,822 observations on 2,773 aquifers over the years 2003 to 2016.¹⁴ Our main results use 2003-2016 means of all variables to reduce noise in annual data, some of which have high variance or are interpolated between collection years. Many analyses exclude aquifers with limited agricultural activity (i.e., less than 1% cropped area). Our unit of analysis is the aquifer; some results distinguish the four case study aquifer systems or the 37 major aquifer systems, whose large area makes remotely sensed inputs more precise (Richey et al. 2015a).¹⁵

4.1 Data: Groundwater

To measure the annual change in the height that water from each aquifer must be lifted Δh_{it} , we use the Gravity Recovery and Climate Experiment (GRACE). Appendix D.1 describes data cleaning, including a scaling factor to reduce noise, calculating annual changes, separating groundwater from total water, and linking grid points to aquifers.

To measure mean annual recharge $\tilde{r}_i$ and surface area s_i , we use BGR (2025), a repository that harmonizes groundwater data from several hydrological research groups. Lacking annual recharge data, we assume that an aquifer’s annual recharge equals its historical mean, as some hydrological models do (Siebert et al. 2010; de Graaf et al. 2015). A sensitivity analysis lets recharge vary over time in proportion to observed precipitation, an idea related to Rodell et al. (2018).¹⁶

To measure the saturated thickness $\mathcal{T}_{i,2013}$ of each aquifer in the year 2013, we draw

¹³We use this notation only for variables measured in m^3 . Height h and surface area s are not scaled by capacity (h is in m and s in m^2).

¹⁴We therefore refer to the 2003-2016 data as either the current or baseline level.

¹⁵These analyses matter because our gridded water data measure water-height changes most accurately at scales above 300 km (including the case study systems); for smaller aquifers, measured storage change averages the surrounding region (Rodell et al. 2019). Appendix Figure A.5, Panel F shows the distribution of log aquifer land areas (in m^2).

¹⁶This sensitivity analysis reflects how floods, droughts, and other weather drive variation in groundwater recharge and depth.

on a widely used global hydrological model (WaterGAP, from [de Graaf et al. 2015](#)). To calculate the saturated thickness of each global aquifer, WaterGAP uses a global hydrogeological model parameterized using terrain geometry and lithology. Where observational data exist, model-derived estimates have a correlation of 0.85-0.87 with them for the preferred parameterization. Basing saturated thickness, and thus capacity, on observed aquifer properties improves on capacity estimates from earlier literature that assume constant specific yield and other inputs across space ([de Graaf et al. 2015](#); [Richey et al. 2015a](#)). Specific yield varies widely across aquifers (Appendix Figure A.5, Panels A and B). We assume that topsoil thickness $\mathcal{H}_i$, specific yield $\mathcal{Y}_i$, surface area s_i , and the maximum lift height $\bar{h}_i$ are time-invariant.

We use estimates of lift height $h_{i,2008}$ for each aquifer in the year 2008 from [Fan et al. \(2013\)](#). They calculate lift height data from 1.6 million global observations of water table depth, supplemented with a groundwater model. We assume the minimum lift distance $\underline{h}_i$ equals topsoil thickness $\mathcal{H}_i$.

One summary statistic previews a central result of the paper: median and mean net extraction as a share of aquifer capacity equal 0.0% and 0.1% (Table 4, Panel A), consistent with the near-zero mean overall net trend in freshwater availability in arable land ([Carleton et al. 2024](#)). The distribution is somewhat symmetric around zero (Appendix Figure A.7, Panel B). Since any long-run path ends in a steady state with zero net extraction, the average aquifer is already in a steady state, though perhaps not the optimal one.

4.2 Data: Production

We use proxies for output y_{it} to estimate production function parameters and productivity. We use gridded nighttime luminosity data from [Li et al. \(2020\)](#), in units of an index they construct. We report sensitivity analyses that separate records from two satellites they combine; we also use gridded data on purchasing power-adjusted GDP covering a subset of our time period ([Kummu et al. 2018](#)). These data rescale official subnational GDP per capita statistics by gridded population, so they approximate spatial variation in GDP within a region. These measures all have the advantage of proxying all economic activity that may benefit from water, not only crop yields.

We combine two sets of records to measure $y'_i(\tilde{w}_i)$, the marginal return to an additional unit of groundwater in agricultural production, which is isomorphic to the willingness to pay for groundwater. We obtain gridded data on unit production of 26 common crops from the Global Agro-Ecological Zones (GAEZ) project ([FAO and IIASA 2024](#)). GAEZ uses agronomic models and high-resolution data on hydrogeological characteristics to predict

yields under different growing conditions; much economic research uses it (e.g., [Alesina et al. 2013](#); [Costinot et al. 2016](#); [Montero and Yang 2022](#)).

We link the shares of crops in total production from GAEZ to estimates of the marginal return to an additional unit of groundwater in agricultural production for individual crops from [D’Odorico et al. \(2020\)](#). They obtain these estimates from agronomy models on increased production in irrigated land relative to rain-fed conditions, for 16 crops.

4.3 Data: Inputs and Costs

To measure and instrument for capital and labor in production function estimates, we use factor quantities and prices. We use population and the employment share in agriculture from the [World Bank \(2023\)](#) to measure labor l_{it} . We measure capital stocks k_{it} from the Penn World Table ([GGDC 2022](#)). We measure wages from the International Labor Organization ([ILO 2023](#)). To adjust for differences in human capital, we multiply wages by $e^{-0.06H_c}$, where H_c is mean years of schooling in country c , measured from the World Bank’s World Development Indicators, as in [Eaton and Kortum \(2002\)](#). We obtain data on interest rates used in monetary policy from the International Monetary Fund ([IMF 2023](#)). We observe these variables by country and re-weight them to the aquifer level. They matter less for our welfare and counterfactual analyses than annual extraction does, since they enter only the production function estimates while extraction affects most parts of the model.

To calibrate groundwater extraction costs, we use data on electricity prices e_i from [IEA \(2021\)](#) and pump efficiency ξ . We obtain estimates of annual industrial electricity prices by country, averaged over the years 2003-2016 (Appendix Figure A.6). Industrial prices are the closest available proxy for agricultural electricity prices available in all countries.¹⁷ Engineering estimates of pump efficiency range from 35% to 85% and center around 60% ([Pradeleix et al. 2015](#); [Martin-Candilejo et al. 2020](#)). Our main results use the estimate of 54% from [Burlig et al. \(2026\)](#) based on groundwater pump microdata from California. Since most countries lack data on the energy source of pumps, our main results use electricity prices, though a sensitivity analysis uses country-specific pump energy sources and source-specific pump efficiencies ([Qin et al. 2024](#)).

As weather and geology controls and instruments, we use gridded data on climate and geology. We measure precipitation and temperature X_{it} from the Climatic Research Unit at the University of East Anglia ([Harris and Jones 2019](#)). We measure rock porosity and

¹⁷In the U.S. at least, some farmers pay industrial electricity rates ([EIA 2014](#)). Elsewhere, agricultural electricity prices can differ from other sectors’ prices because of subsidies, time-of-use or seasonal pricing, or nonlinear pricing ([Badiani-Magnusson and Jessoe 2018](#)). Industrial prices partly proxy for marginal cost, which matters to the planner more than subsidized prices do.

permeability from the high-resolution Global Hydrogeology Maps (Gleeson et al. 2014) and follow their assumption that soil porosity and permeability are fixed (see also de Graaf et al. 2015 and Richey et al. 2015a).

We use gridded cropland data from Ramankutty et al. (2008) to identify aquifers with limited agricultural activity. An aquifer’s cropland share is the aquifer average of each of its pixels’ cropland share.

4.4 Data: Benchmark Discount Factors and Others

The benchmark discount factor β_i^* is the value the representative agent would use absent institutional or market failures, and that it experiences in utility terms. Ramsey (1928) proposed the normative discount factor $\beta_i^* = 1/(1 + d_i + \eta_i g_i)$. Here d_i represents the pure rate of time preference (the discount rate on utility), η_i is the elasticity of the marginal utility of consumption, and g_i is the growth rate of consumption per capita.

We use estimates of β_i^* from Addicott et al. (2020), who calculate the utility discount rates implied by data on age-specific mortality rates and life expectancy for most countries. They combine these with estimates of the elasticity of the marginal utility of consumption and the rate of consumption growth from Drupp et al. (2018). They estimate Ramsey discount factors between 0.93 and 0.97 for most countries (Appendix Figure A.6, Panel B). The mean aquifer has a Ramsey discount factor of 0.95 (Table 1).

We also use a benchmark market-based discount factor of 0.98. This reflects long-term rates of low-risk government debt and other low-risk assets and prevailing views on the risk-free rate (Giglio et al. 2015; Drupp et al. 2018; Carleton and Greenstone 2022; Rennert et al. 2022; Bauer and Rudebusch 2023).

We present most results for both the Ramsey and market discount factors, since they have different advantages. Some maps show only the Ramsey benchmark, which varies by aquifer.¹⁸ The Ramsey discount factor also has an explicit welfare foundation, though it requires careful calibration and has faced longstanding debates. The market benchmark is simple, uniform across countries, and increasingly used in climate change economics.

As correlates of the decision discount factors, we use GDP per capita, formal or informal water markets, democracy, and other variables (Appendix D.5).

¹⁸Long-term government interest rates vary by country due to risk, but we follow the standard assumption that all countries share the same risk-free rate; Zimbabwe’s debt, for example, pays a higher rate than U.S. debt, but both countries face the same risk-free rate.

Table 1: Summary Statistics and Decision Discount Factors

	Median aquifer (1)	Mean aquifer (2)	Central Valley (3)	Ogallala (4)	Indus (5)	North China (6)
<i>Panel A: Model Inputs</i>						
Capacity (billion m ³)	44.13	284.25	327.01	737.05	2661.58	1434.39
Water: gross annual use (% capacity)	2.11%	3.06%	1.85%	1.42%	1.60%	3.33%
Annual recharge (% capacity)	2.06%	2.98%	1.62%	1.73%	1.17%	2.97%
Marginal return to groundwater (\$/m ³)	0.20	0.21	0.22	0.17	0.17	0.22
Marginal extraction cost (\$/m ³)	0.02	0.02	0.04	0.03	0.06	0.01
Current lift height (m)	18.08	20.92	41.88	20.31	45.23	26.34
Electricity price (\$/kWh)	0.07	0.07	0.08	0.07	0.08	0.04
Ramsey discount factor	0.95	0.95	0.96	0.96	0.92	0.91
<i>Panel B: Decision Discount Factor</i>						
Decision discount factor	0.924	0.860	0.890	0.935	0.714	0.890
Decision minus Ramsey discount factor	-0.022	-0.086	-0.069	-0.023	-0.206	-0.022

Notes: Columns (1) and (2) describe the median and mean of all global aquifers with data. Columns (3) through (6) describe the four case study aquifer systems. Ramsey discount factors are from [Addicott et al. \(2020\)](#). Panel B reports decision discount factors, estimated with the model in Section 3.3 and Appendix C.1. Statistics are unweighted. Dollar values are in 2023 U.S. dollars.

5 Parameter Estimation and Model Calibration

This section explains how we use our data to apply the model quantitatively.

5.1 Calibration: Aquifer Characteristics

This subsection converts the hydrological data into the model’s state variables—fill a , capacity κ , and lift height h —using accounting identities implied by the model. The starting point is water available in each aquifer $\tilde{a}_{i,2013}$ in the year 2013, the product of saturated thickness, specific yield, and surface area:¹⁹

$$\tilde{a}_{i,2013} = \mathcal{T}_{i,2013} \mathcal{Y}_i s_i$$

¹⁹[Richey et al. \(2015a\)](#), who also estimate aquifer capacity, assume a uniform specific yield of 1%, based on prior estimates for subsurface rock. We let specific yield vary across aquifers, though its median is 1%.

The annual change in available water is the product of aquifer area, specific yield, and the change in aquifer height, where increasing lift height decreases available water:

$$\Delta\tilde{a}_{it} = -s_i\mathcal{Y}_i\Delta h_{it} \quad (13)$$

Equation (13) effectively assumes that aquifers have vertical sides, following some hydrological research (Rodell et al. 2007; Rodell et al. 2009; Richey et al. 2015a). Gross annual groundwater extraction $\tilde{w}_{it}$ from each aquifer in m^3 equals recharge minus the change in available water, reflecting the law of motion in (1):

$$\tilde{w}_{it} = \tilde{r}_i - \Delta\tilde{a}_{it} \quad (14)$$

Iterating forward on the law of motion (1) then yields fill $\tilde{a}_{it}$ in each year, with the annual change in available water $\Delta\tilde{a}_{it}$ taken from (13). The recursion starts from 2013 fill $\tilde{a}_{i,2013}$, which saturated thickness data identify directly:

$$\tilde{a}_{it+1} = \tilde{a}_{it} + \Delta\tilde{a}_{it} \quad (15)$$

An analogous recursion, beginning from observed baseline values $h_{i,2008}$, yields lift height in each aquifer $\times$ year:

$$h_{it+1} = h_{it} + \Delta h_{it} \quad (16)$$

Maximum possible lift height is the sum of current lift height and saturated thickness:

$$\bar{h}_i = h_{i,2013} + \mathcal{T}_{i,2013}$$

Finally, each aquifer's capacity is the product of the aquifer's maximum possible lift height, less topsoil thickness, with surface area and specific yield:

$$\kappa_i = (\bar{h}_i - \mathcal{H}_i)s_i\mathcal{Y}_i \quad (17)$$

Log capacity varies considerably across aquifers (Appendix Figure A.5, Panels C and D).²⁰

²⁰Capacity has more uncertainty than our other data inputs because it depends on saturated thickness, which is not directly observed. We instead follow recent work that combines hydrological models with observational data on well depths and rock characteristics (de Graaf et al. 2015).

5.2 Estimation: Production Functions

Methods

We use versions of the following equations to estimate production function parameters:²¹

$$y_{it} = \mathcal{Z}_{it} k_{it}^{\alpha} l_{it}^{\gamma} w_{it}^{\phi_i} \quad (18)$$

$$\mathcal{Z}_{it} = e^{\Gamma X_{it} + \varphi_i + \zeta_t + \varepsilon_{it}} \quad (19)$$

Equation (18) allows for annual input variation in the technology from (2), and lets time-varying technology $\mathcal{Z}_{it}$ subsume land x_i . In (19), productivity combines temperature and precipitation X_{it} , soil quality and other time-invariant characteristics φ_i (including and subsuming land x_i), a global productivity disturbance ζ_t , and an idiosyncratic shock ε_{it} .

Combining (18) and (19) then taking logs gives

$$\log y_{it} = \alpha \log k_{it} + \gamma \log l_{it} + \phi_i \log w_{it} + X'_{it} \Gamma + \varphi_i + \zeta_t + \varepsilon_{it} \quad (20)$$

The welfare and counterfactual analyses in Sections 6 and 7 mainly use two parameters from equation (20)—the water elasticity ϕ_i and the resource share $\theta \equiv 1 - \alpha - \gamma$.

Estimating equation (20) faces two main challenges. First, no data report aquifer-level output, so y_{it} is unobserved. We proxy for it using remotely sensed nighttime luminosity $\bar{y}_{it}$ and, for a subset of years, gridded GDP. In our model of final and no intermediate goods, GDP equals output. We assume that the log of true agricultural output equals an aquifer-specific multiplier χ_i of nighttime luminosity plus normally-distributed measurement error, $\log y_{it} = \log(\chi_i \bar{y}_{it}) + v_{it}$, so that the relationship between nighttime luminosity and output can vary across aquifers.²² Second, OLS estimates of the input elasticities are likely biased. Measurement error can arise in remotely sensed data, which are typically smoothed averages of multiple images (Donaldson and Storeygard 2016; Proctor et al. 2023). Omitted variables bias is a classic concern in production functions, since the agent may observe productivity shocks that the econometrician does not (Marschak and Andrews 1944).

We thus instrument for water, capital, and labor. The first-stage equation for (20) is

$$\log b_{it} = Z'_{it} \Lambda + X'_{it} \Psi + \mu_i + \delta_t + u_{it} \quad (21)$$

²¹Market prices of water could also measure its marginal value, but water markets are rare and have substantial frictions.

²²Night lights capture all local economic activity. Our assumption that agricultural output is a fixed multiple of luminosity within each aquifer implies that other sectors move in proportion with agriculture, so luminosity-based estimates recover agriculture’s water elasticity and z_i can be read as aggregate productivity in equation (3). Sensitivity analyses use aquifer agricultural GDP downscaled from national estimates.

Here $b_{it} \in (k_{it}, l_{it}, w_{it})$ is endogenous and Z_{it} are instruments. To instrument for water, we use two- and three-period lags of precipitation and temperature and their interactions with soil permeability and porosity: through the law of motion, lagged weather determines current lift height and thus current extraction costs, with larger effects in more porous aquifers (Kotchoni et al. 2019). For capital and labor, we use up to three-period lags of interest rates and wages, following common practice in production function estimation (Blundell and Bond 2000; Doraszelski and Jaumandreu 2018). Identification requires that lagged input prices and lagged weather shift demand for inputs without being correlated with current shocks to aquifer productivity.²³

Implementation follows standard practice. We select instruments and weather controls (from among precipitation, temperature, their squares, and logs) using post-double-selection lasso (Belloni et al. 2012). Extensions estimate heterogeneous returns $\phi_i = \phi + \psi\nu_i$ using aquifer characteristics ν_i such as GAEZ measures of the return to irrigation, and cluster standard errors by aquifer. Sensitivity analyses also measure w with groundwater alone or combined with surface water, and control for precipitation directly.

Results

Lasso selects lagged precipitation, temperature, wages, interest rates, and some interactions as instruments (Appendix Table A.2). The first-stage Kleibergen–Paap rk Wald F statistic exceeds the critical value for a reduction in bias of at least 90% in the instrumental variables estimates with our three endogenous variables and the twelve selected instruments (Stock and Yogo 2005). Given this instrument strength, we also report model results calibrated to production function values from the literature.

Table 2 shows that water is moderately important for output. Our preferred specification, in column (4), implies a groundwater elasticity ϕ of 0.10 and a resource share θ of 0.51. Although the model does not directly use the land share, we also report the land share implied by constant returns, $\theta - \phi$, as an informative benchmark. The OLS estimate of the water elasticity in column (1) is small, consistent with attenuation from measurement error and omitted variables bias.²⁴ Instrumenting for groundwater raises the elasticity (columns (2) and (3)); our preferred specification in column (4) also instruments for capital and labor. Column (5) shows that calibrating the land share from the literature (Gollin and Udry 2021;

²³The underlying assumption is that aquifer productivity follows a Markov process. Some papers use current input prices as instruments, which is valid if aquifer-specific productivity is uncorrelated with aggregate demand. Because many aquifers produce significant output of specific crops, potentially challenging this assumption, we use lags.

²⁴For example, regional adoption of drought-resistant seeds would raise output while reducing water extraction, biasing OLS downward.

Ryan and Sudarshan 2022), rather than estimating it, yields a capital + labor share of 0.40, similar to the 0.49 that we estimate from re-weighted national capital and labor data. Finally, the water elasticity varies with the value of irrigation: an aquifer at the top of the GAEZ irrigation-return distribution has a groundwater elasticity of 0.35, versus 0.11 at the median (column (6)).

Few production estimates account for water. One exception, Ryan and Sudarshan (2022)’s analysis of farmers in Rajasthan, India, estimates water input shares of 0.04 to 0.18 and labor + capital shares of 0.42 to 0.54; our preferred estimates of 0.10 and 0.49 lie in the middle of these ranges, though the settings differ.

Appendix E discusses alternative estimates. Higher cropland, irrigation, and agricultural dependence predict higher returns to water (Appendix Table A.3). Alternative measures of output obtain estimates of the parameters ϕ and θ loosely centered around the main estimates, as does including surface water in total water inputs (Appendix Table A.7).

5.3 Calibration: Productivity

Applying the model also requires transformed productivity z_i from equation (4). We assume nighttime luminosity is proportional to GDP but does not equal output in each aquifer, so the production function estimates alone cannot measure productivity z_i .

We instead calculate transformed productivity as follows. Differentiating the production technology (3) with respect to water and solving for productivity gives

$$z_i = \frac{\theta \kappa_i^{\frac{\phi_i}{\theta}} y'_i(\tilde{w}_i)}{\phi_i(\tilde{w}_i)^{\frac{\phi_i}{\theta}-1}} \quad (22)$$

To calculate z_i , we use empirical analogues to each term in (22). Data report capacity κ_i and gross water extraction $\tilde{w}_i$. Table 2 estimates the water and resource shares ϕ and θ . We calculate the marginal return to groundwater $y'_i(\tilde{w}_i)$ as the weighted sum of the crop-specific marginal returns from D’Odorico et al. (2020), weighted by each crop’s share of production from GAEZ.²⁵

$$y'_i(\tilde{w}_i) = \sum_c \sigma_{ic} \times v_c$$

Here σ_{ic} is the share of crop c in total agricultural production of aquifer i and v_c is marginal groundwater revenue for crop c .

The marginal returns to groundwater vary by aquifer (Figure 2, Panel C). Areas with little

²⁵D’Odorico et al. (2020) calculate the crop- and continent-specific value of water as the increase in crop production due to additional irrigation water $\Delta y_i / \Delta \tilde{w}_i$, which provides a secant approximation to our derivative. They also multiply unit output by prices, while we interpret the output as numeraire.

Table 2: Estimated Production Function Parameters

	OLS (1)	IV, water only (2)	IV, water only (3)	IV, all (4)	IV, calibrated (5)	IV, GAEZ (6)
$\log(\text{Groundwater}) (\phi)$	0.004 (0.003)	0.227*** (0.077)	0.253*** (0.079)	0.103** (0.044)	0.103** (0.044)	-0.130 (0.095)
$\log(\text{Groundwater}) \times \text{Value} (\phi_i)$	—	—	—	—	—	0.479** (0.194)
$1 - \log(\text{Capital}) - \log(\text{Labor}) (\theta)$	0.863*** (0.049)	—	0.893*** (0.051)	0.514*** (0.155)	0.603*** (0.044)	0.405** (0.162)
Implied $\log(\text{Land}) (\theta - \phi)$	0.858*** (0.049)	—	0.639*** (0.085)	0.411** (0.173)	0.500 —	0.470*** (0.158)
Aquifer FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Weather controls	Yes	Yes	Yes	Yes	Yes	Yes
F statistic	—	21.608	21.105	10.401	10.401	5.082
Observations	29,854	29,854	29,854	29,854	29,854	29,854

Notes: Columns (2)-(4) and (6) instrument for $\log(\text{Groundwater})$, where groundwater represents gross extraction in m^3 . Columns (4) and (6) instrument for $\log(\text{Labor})$ and $\log(\text{Capital})$. Section 5.2 describes the process for choosing instruments and weather controls. Column (5) recovers θ implied using the groundwater estimate from column (4) and the land share from Ryan and Sudarshan (2022). Column (6) uses the aquifer’s percentile in the distribution of irrigation’s effect on crop yield from GAEZ to measure water sensitivity. Capital is in millions of 2023 U.S. dollars. Labor is population employed in agriculture. Implied $\log(\text{Land})$ is the implied value of the land share under the assumption of constant returns to scale. The dependent variable is $\log(\text{Night Lights Luminosity})$ by aquifer-year. F statistic is the first-stage Kleibergen–Paap rk Wald F statistic. Parentheses show standard errors clustered by aquifer. Asterisks denote p-value less than 0.01 (***), 0.05 (**), 0.10 (*).

agriculture, like the Sahara, have low marginal returns. Regions with productive staples, like the Midwestern U.S., have moderate marginal returns. Areas with high-value crops, like almonds and fruit in California’s Central Valley, have high marginal returns.

5.4 Calibration: Cost Function

We measure the marginal cost of groundwater extraction by calibrating the cost function (7). Following engineering practice (Robinson 2002; Pradeleix et al. 2015; Martin-Candilejo et al. 2020), the cost in (7) of lifting w_i cubic meters of groundwater is lift height h_{it} (in meters) times an aquifer-specific engineering cost c_i times extraction w_i . This cost depends on pump efficiency ξ , electricity prices e_i , and a constant ρ reflecting flow rate, gravitational acceleration, and water density:²⁶

$$c_i = \frac{e_i \rho}{\xi} \quad (23)$$

As discussed above, we calibrate pump efficiency and the physical constant to Burlig et al. (2026). A sensitivity analysis allows fuel type, price, and pump efficiency to vary by region.²⁷

The marginal cost of lifting one cubic meter of groundwater to the surface, $c_i \times h_{it}$, varies by region (Figure 2, Panel E). The Western U.S. and Western Australia have relatively high lift costs; regions with cheap energy like the Middle East and Russia have lower costs. Panel F shows that most lift costs range between \$0.01 and \$0.04/ m^3 . Lift height is far more skewed than electricity prices and so drives more of the dispersion in marginal extraction costs (Appendix Figure A.6, Panels D and F).


Appendix F.1 reports one model validation test. A model fitted to 2003-2009 data predicts 2010-2016 extraction well, and the two periods imply similar decision discount factors, welfare losses, and long-run fill. This test is informative, though limited by its seven-year horizon.

²⁶Engineering models have a mixed record in predicting realized costs in complex settings like residential energy efficiency (Fowlie et al. 2018; Burlig et al. 2020; Davis et al. 2020). The physics of lifting water is far simpler, however, and more amenable to calculation; results are also similar under region-specific fuel types and pump efficiencies (Appendix Table A.5, column (5)).


²⁷Since in the model $\mathcal{C}_i(a, w)$ is the extraction cost per unit of aquifer capacity, rather than m^3 , we multiply (23) by aquifer capacity κ_i to link the model with data.

Figure 2: Aquifer Characteristics


(a) Spatial Patterns of Gross Annual Extraction




(b) Distribution of Gross Annual Extraction




(c) Spatial Patterns of Marginal Return to Groundwater




(d) Distribution of Marginal Return to Groundwater



(e) Spatial Patterns of Marginal Extraction Costs



(f) Distribution of Marginal Extraction Costs



Notes: Panels A and B show mean gross annual groundwater extraction as a percent of aquifer capacity from GRACE (Rodell et al. 2019) and WHYMAP (BGR 2025). Panels C and D show marginal returns in \$/m³ of groundwater calibrated using data from GAEZ and D’Odorico et al. (2020). Panels E and F show the marginal cost of extracting groundwater in \$/m³, calibrated using data from Fan et al. (2013) and the International Energy Agency (IEA 2021). Each aquifer-level estimate averages the years 2003-2016. We calculate marginal returns and marginal cost following Sections 5.3 and 5.4.

6 Results: Discount Factors

6.1 Graphical Explanation of Decision Discount Factors

Figure 3 shows how the model estimates the decision discount factor, using the Central Valley Aquifer System as an example. The horizontal axis shows gross extraction as a share of capacity (w), and the vertical axis shows marginal revenue or cost in $\$/m^3$. The thick downward-sloping marginal revenue curve, $\partial\tilde{\pi}_i(w)/\partial w$, shows how current groundwater extraction affects current profits. Its positive value reflects the return to water extraction, and its downward slope reflects decreasing returns. The dashed upward-sloping curves show marginal extraction costs under different discount factors, including the stock externality that extraction today raises future extraction costs. Orange crosses show market discounting, yellow triangles Ramsey discounting, purple squares the decision discount factor β_i^d , and green asterisks myopia ($\beta_i = 0$); the vertical dashed line shows observed extraction w .

Graphically, our estimation algorithm recovers the decision discount factor β_i^d as the value which makes the marginal revenue and cost curves intersect at observed water consumption w (Figure 3). Myopic extraction has the lowest marginal cost curve because it ignores future costs. Larger discount factors shift marginal costs up, since they put more weight on the greater lift height required for future extraction. Policy functions then map aquifer fill to water consumption (see Appendix F.4 and Appendix Figure A.11).

In the Central Valley, the triangles and crosses lie above the squares (Figure 3): the Ramsey and market benchmarks exceed the decision discount factor, so current extraction exceeds the optimal level.


6.2 Describing Decision Discount Factors

Estimated decision discount factors vary widely (Figure 4, Panel A). Most North American aquifers have relatively low decision discount factors, though the Ogallala does not; so do aquifers in South Asia and Australia, but not parts of the Middle East and West Africa.

Mean and median decision discount factors are 0.86 and 0.92 (Table 1, Panel B). Most aquifers have decision discount factors well above zero, so they are far from myopia. Because the market discount factor is 0.98 and most Ramsey discount factors are between 0.93 and 0.97, aquifers are split between discount factors well below these benchmarks and those somewhat near them. Few aquifers have a decision discount factor below 0.50. Aquifers with decision discount factors near these benchmarks help explain why, for the mean aquifer, welfare losses from excess extraction are modest as a percent of aquifer present value.

We highlight unweighted averages because management typically involves aquifer-specific

Figure 3: Marginal Cost and Marginal Revenue Curves for the Central Valley Aquifer System




Notes: The horizontal axis shows gross extraction as a share of aquifer capacity. The dashed vertical line shows gross groundwater extraction based on GRACE (Rodell et al. 2019). The upward-sloping lines show marginal costs for four discount factors: complete myopia, the decision discount factor (i.e., the estimated marginal cost), the Ramsey benchmark discount factor, and the market benchmark discount factor. With complete myopia, $\beta = 0$. The market discount factor is $\beta = 0.98$. In the Central Valley Aquifer System, the decision discount factor is $\beta^d = 0.89$ and the Ramsey discount factor is $\beta^* = 0.96$.

policies. Weighting across aquifers by crop value, GDP, or population gives a similar range (Appendix Table A.4). The median crop- and population-weighted discount factors are 5%-9% lower than the unweighted median, so aquifers with high crop value and population are extracted more myopically; the corresponding weighted net water use measures are lower. The GDP-weighted median is about 2% above the unweighted median.

Few prior estimates of decision discount factors exist for comparison. In the Indian state of Rajasthan, Ryan and Sudarshan (2022) estimate that a discount factor of 0.82 would rationalize an observed ration of electricity used to pump groundwater; this value is near our global mean estimate of 0.86, though their discount factor would make the level of the ration efficient and may not explain observed pumping behavior.²⁸ Burlig et al. (2026) use a


²⁸Our decision discount factor for the nearby but non-overlapping Indus Basin is lower, at 0.71 (Table 1,

Figure 4: Decision Discount Factors, by Aquifer



Notes: The decision discount factor represents the value that rationalizes observed extraction. Panel A maps the decision discount factor (β^d) for each aquifer. Panel B shows the distribution of decision discount factors across aquifers. Section 3.3 and Appendix C.1 explain how we estimate decision discount factors.

Figure 5: Decision Minus Ramsey Discount Factors, by Aquifer



Notes: The decision discount factor represents the value that rationalizes observed extraction. Panel A maps the difference between the decision discount factor that we estimate and the Ramsey benchmark discount factor from Addicott et al. (2020). Panel B graphs the distribution of this difference across aquifers. Section 3.3 and Appendix C.1 explain how we estimate decision discount factors.

calibrated discount factor of 0.90 to rationalize groundwater pumping in the Central Valley Aquifer System, close to our estimate of 0.89 for that system (Table 1, Panel B).

Our decision discount factors have sensible correlations with other variables (Appendix Table A.8). Aquifers with higher marginal returns have higher decision discount factors; richer and more democratic countries have lower discount factors, consistent with decentralized common-pool extraction; and some evidence suggests aquifers in countries with formal water markets have higher discount factors.

6.3 Comparing Decision and Benchmark Discount Factors

Most aquifers have decision discount factors 2.5 percentage points or more below their Ramsey benchmark (Figure 5, Panel A). Parts of East Asia, the Middle East, and Sub-Saharan Africa do not. The similar maps in Panel A of Figures 4 and 5 suggest that these patterns reflect variation in decision rather than Ramsey benchmark discount factors. This finding echoes the low correlation between decision and Ramsey discount factors (Appendix Table A.8). In Sub-Saharan Africa, for example, extraction costs are lower (Figure 2, Panel E), mostly because of cheaper electricity rather than lift height (Appendix Figure A.6), so rapid extraction generates a smaller externality. The higher decision discount factors in poorer regions such as Sub-Saharan Africa may also reflect less infrastructure.

While the modal gap between decision and benchmark discount factors is near zero, many aquifers have decision discount factors far below benchmarks (Figure 5, Panel B). A few aquifers also have decision discount factors slightly above benchmarks, potentially due to factor or other input market distortions. Discount factors fall below the Ramsey benchmark in 60.8% of aquifers and below the market benchmark in 81.5% (Table 3). These shares are moderately higher when weighted by aquifer capacity. Even aquifers that are refilling on net can have decision discount factors below benchmarks if they would optimally refill faster.

Decision discount factors are 2 versus 9 percentage points below the Ramsey discount factor for the median versus mean aquifer (Table 1, Panel B), reflecting the skewness of decision discount factors. The Ramsey and market benchmark discount factors exceed the decision discount factor for all four case studies. Some of these patterns persist when weighting across aquifers by crops, GDP, or population (Appendix Table A.4).

6.4 Discount Factor Versus Property Rights Interpretation of β_i

We primarily interpret resource management through the gap between the decision discount factor and the benchmark discount factor. A more standard approach to analyzing effi-

Panel B).

Table 3: Welfare Gains from Following Benchmark Discount Factors

	Median aquifer (1)	Mean aquifer (2)	Global total (3)	Central Valley (4)	Ogallala (5)	Indus (6)	North China (7)
<i>Panel A: Welfare gain (% of aquifer value) from following...</i>							
Ramsey discounting	1.35%	2.55%	—	2.75%	0.48%	2.58%	0.18%
Market discounting	1.69%	2.90%	—	5.54%	1.24%	10.83%	3.37%
<i>Panel B: Welfare gain (billion \$) from following...</i>							
Ramsey discounting	0.07	1.88	2,575	2.12	0.54	5.13	0.54
Market discounting	0.20	6.11	8,362	9.01	2.93	91.87	44.88
<i>Panel C: PDV extraction cost savings (billion \$) from following...</i>							
Ramsey discounting	0.01	2.15	2,939	5.39	2.35	17.52	5.24
Market discounting	0.33	7.54	10,325	13.22	5.29	126.84	57.38
<i>Panel D: Share of aquifers over-extracted</i>							
Ramsey discounting	—	—	0.608	—	—	—	—
Market discounting	—	—	0.815	—	—	—	—
Ramsey, capacity weighted	—	—	0.717	—	—	—	—
Market, capacity weighted	—	—	0.890	—	—	—	—

Notes: Columns (1) and (2) describe the median and mean of all global aquifers with data. Columns (4) through (7) describe the four case study aquifer systems. “Ramsey discounting” and “market discounting” describe counterfactuals where the planner chooses the extraction path while using a Ramsey or a market discount factor. Ramsey benchmark discount factors are from [Addicott et al. \(2020\)](#). The market discount factor is 0.98. Estimates use the model in Section 3.3 and Appendix C.1 and are unweighted unless otherwise noted. Dollar values are in 2023 U.S. dollars.

ciency, which directly represents property rights, compares whether common-pool or single-agent decision-making better characterizes extraction (Levhari and Mirman 1980; Karp and Sakamoto 2021). This subsection relates the two interpretations. Some research equates single-agent extraction with optimal resource management; we instead allow a single representative agent to manage an aquifer non-optimally, by extracting under a decision discount factor that differs from the benchmark.

Suppose each aquifer has a unit mass of atomistic agents who extract water from private wells on a shared aquifer. The law of motion for a well’s water level then depends on the owner’s gross extraction w_{it} and on aquifer-wide average extraction $\bar{w}_{it}$, which the atomistic agents take as given:

$$a_{it+1} = \min\{a_{it} + r_i - \lambda_i w_{it} - (1 - \lambda_i)\bar{w}_{it}, 1\}$$

The parameter $\lambda_i \in [0, 1]$ represents the extent of a well owner’s control over the water level. The value $\lambda_i = 0$ is a common pool, where agents take the water level as given; $\lambda_i = 1$ is a single agent who fully internalizes her own extraction. The agent’s problem is then

$$V_i(a; \beta_i, \lambda_i) = \max_{w \leq a+r_i} \pi_i(a, w) + \beta_i V_i(\min\{a + r_i - \lambda_i w - (1 - \lambda_i)\bar{w}_i(a), 1\}; \beta_i, \lambda_i) \quad (24)$$

The problem in (24) has the optimal solution $w_i^*(a; \beta_i, \lambda_i)$. We focus on equilibria with symmetry across atomistic agents, so that $\bar{w}_i(a) = w_i^*(a; \beta_i, \lambda_i)$. This equilibrium is a fixed point of (24). Appendix C.2 describes our computational algorithm to solve for this fixed point.²⁹

Our discount factor interpretation is analytically equivalent to this property rights interpretation for the polar cases. The common pool ($\lambda_i = 0$) corresponds to myopia ($\beta_i = 0$), and the single-agent optimum ($\lambda_i = 1$) to discounting at the benchmark rate ($\beta_i = \beta_i^*$), the two cases that many resource models describe.³⁰

Between these polar cases, we compute (β_i^d, λ_i) pairs such that the main model with decision discount factor β_i^d induces the same extraction as the property rights model with parameter λ_i . For the property rights model, we set β_i equal to either the Ramsey or market discount factor.

For the Central Valley Aquifer System, we trace the relationship between the decision

²⁹Simpler settings have simpler solutions—in a two-period model of resource extraction, the equilibrium with N agents that each have discount factor β_i is equivalent to the choice of a planner with discount factor β_i/N (Karp 2017).

³⁰In both the baseline and property rights models, agents fully internalize effects of their water consumption on *within-period* water level. In the property rights model, aquifer-wide extraction only affects next period’s water level.

discount factor β_i^d and the property rights parameter λ_i across a wide range of water consumption levels (Panel A of Appendix Figure A.1). Either the property rights or discount factor interpretation can rationalize observed levels of water extraction. The relationship between β_i^d and λ_i is strictly increasing and nonlinear. The decision discount factor is consistent with a value of $\lambda_i = 0.59$, implying that users act as if 41% of extraction costs depend on others' extraction choices. We find similar patterns for the three other case study aquifer systems (Panels B-D of Appendix Figure A.1).³¹ All these patterns hold whether using the Ramsey or market benchmark. These one-to-one functions suggest that the discount factor in most of this paper is analogous to the longer-standing property rights interpretation. The standard comparison of open access to single-agent extraction characterizes property rights only as binary, while this analysis provides a continuous measure.

7 Results: Resource Management and Welfare

7.1 Results: Resource Management

Under either benchmark discount factor, the Central Valley Aquifer System would cut extraction nearly in half to refill the aquifer (Figure 6, Panel A). In the baseline year, extraction of the Central Valley Aquifer System equals 1.8% of aquifer capacity, and under the decision discount factor it gradually decreases to 1.6% of aquifer capacity in steady state. Under the benchmark discount factors, extraction instead drops to about 1% of capacity in the baseline period, then rises to the same steady state of 1.6% by 2040. Halving extraction would have large short-run consequences for farmers.³² Steady state extraction is the same under the decision and benchmark discount factors because extraction equals recharge in any steady state.³³


The same benchmark paths that refill the aquifer also reduce extraction costs and eventually raise annual profit (Figure 6). Under the decision discount factor, the Central Valley Aquifer System fill declines from about 95% of capacity today to 90% by the year 2100. Under either benchmark, the aquifer instead fills completely by 2040. Extraction costs rise by about 40% over the 21st century under the decision discount factor but fall to nearly zero within a few decades under either benchmark. The current extraction path increases short-run profit. Within a decade, however, annual profit under either benchmark path exceeds

³¹We calculate these relationships only for the case study aquifer systems because the calculation is computationally intensive.

³²As in other analyses of optimal policy, we study the policies that maximize social welfare, abstracting from political economy, which would matter for pursuing the optimal extraction path in practice.

³³At capacity ($a = 1$), an aquifer can stay 100% full as long as extraction does not exceed recharge.

Figure 6: Resource Management Paths by Discount Factor, Central Valley Aquifer System



Notes: This figure shows the predicted gross groundwater extraction (% capacity), predicted fill levels (% capacity), predicted future water extraction costs (\$/m³), and predicted annual profit (billion \$) for the Central Valley Aquifer System. Each line represents a prediction under a different discount factor. The solid blue line represents the current extraction path.

profit under the current path, and it remains moderately higher in steady state.

Under the decision discount factor, mean and median net water extraction nearly equal zero (Table 4, Panel A). Zero net extraction may be excessive if the optimal fill level exceeds the baseline value. For example, a completely drained aquifer might have zero net extraction because extraction equals recharge, but the optimal path would likely cut short-term extraction to raise long-term fill. The Ogallala Aquifer is refilling, but the optimal path would decrease extraction further.³⁴ In all cases, Ramsey or market discounting would initially let recharge moderately exceed extraction, refilling the aquifer.

These water extraction plans affect long-run fill (Table 4, Panel B). Steady state median fill is only 94.1% under the current extraction path, versus 100% under Ramsey or market discounting. The benchmark discount factors also substantially lower steady state extraction costs as aquifers reach higher fill (Panel C). Weighting by crop value, GDP, or population gives similar patterns (Appendix Table A.4). These medians mask variation: 46% of aquifers exceed 95% fill under the decision discount factor (Appendix Figure A.10), while 61% to 84% fill completely under the Ramsey and market benchmarks.

7.2 Results: Social Welfare

Aquifer Present Value

All global aquifers together have present value of \$103 trillion under Ramsey discounting and \$245 trillion under market discounting (Appendix Table A.5); sensitivity analyses vary but all exceed \$49 trillion. Valuing an aquifer in full effectively measures the value of draining it, far outside the moderate changes that most of our scenarios consider. We therefore read this order of magnitude—which, to our knowledge, no existing work has estimated—as informative but interpret it cautiously.


This magnitude may or may not be surprising. Water is critical for life, so a vast sum for Earth’s groundwater may seem reasonable. Agriculture, however, is also a modest share of global GDP, and water is only one input to agricultural production.

Costs of the Current Extraction Path

The costs of non-optimal groundwater extraction are large where decision discount factors are low, including in parts of the Western U.S. and Mexico (Figure 7, Panel A). The modal aquifer has a welfare cost of non-optimal extraction below 1.0% of aquifer present value,

³⁴This pattern represents aggregate extraction throughout this major aquifer system (Scanlon et al. 2012; Richey et al. 2015b); the southern portion of the system, beneath Texas and Kansas, is being depleted much more rapidly than the wetter and less cultivated northern portion (USGS 2023).

Figure 7: Welfare Loss Due to Deviation from Ramsey Benchmark Discount Factor, by Aquifer



Notes: Panel A maps the welfare loss in each aquifer due to the current extraction path deviating from the Ramsey benchmark. We measure the welfare loss as a percent of each aquifer’s present value. Panel B graphs the distribution of these losses across aquifers.

though the distribution has a long right tail (Panel B). Several percent of aquifers have losses above 5.0% of aquifer present value, a few exceed 10%, and none exceeds 16%. Results are comparable for the market benchmark (Appendix Figure A.2). The welfare loss has correlations of 0.59 and 0.71 with the absolute gap between the decision discount factor and the Ramsey and market benchmarks, respectively (Appendix Figure A.3).

The welfare costs of non-optimal aquifer depletion have large dollar though small percentage magnitudes (Table 3, Panels A and B). These welfare losses of current extraction relative to the benchmark path equal \$2.6 to \$8.4 trillion in aggregate, or 2.6% to 2.9% of aquifer value. Much higher extraction costs drive these losses: in aggregate, the present value of extraction costs on the observed path exceeds that under Ramsey and market discounting by \$2.9 to \$10.3 trillion (Panel C). In California’s Central Valley Aquifer System, for example, the loss is \$2.1 to \$9.0 billion in present value; following the Ramsey or market optimal extraction path would save \$5.4 to \$13.2 billion in present value extraction costs, respectively. Sensitivity analyses obtain qualitatively similar results (Appendix F.3).

Three reasons explain why non-optimal extraction creates modest percentage welfare costs. First, non-optimal extraction generates a natural dampening force: extracting faster than is optimal lowers the water table and raises extraction costs, which slows future extraction. Second, the benchmark path typically lowers short-run profits as the aquifer refills, and discounting shrinks the present value of subsequent profits. Discussions of modest to moderate early estimates of the welfare costs of climate change (Nordhaus 1977) sometimes attributed them to the “tyranny of discounting” (Pearce et al. 1989; Cline 1992)—the tendency of low discount factors to assign small present costs to large distant future damages.

Third, most aquifers have a decision discount factor above 0.90, so the typical gap from the benchmarks is modest.

For each case study aquifer, Appendix Figure A.4 illustrates how these welfare costs vary with the decision discount factor. When the decision discount factor equals the benchmark, extraction is optimal; smaller or larger decision discount factors lower welfare. The welfare loss is steepest near the benchmark and flatter further away. In each aquifer, the welfare cost of using a discount factor that is 10 percentage points too low is about 7% to 12% of aquifer value, but the welfare cost of using a discount factor that is 50 percentage points too low is only slightly higher. When the discount factor is small, future pumping costs and output have small present values, so changing the discount factor matters less for welfare. Small management improvements in aquifers with decision discount factors near the benchmark could therefore recover much of the welfare loss.

7.3 Counterfactual Extraction Patterns

This subsection analyzes three counterfactuals. Under myopia, the representative agent ignores future extraction costs ($\beta_i^d = 0$). Under constant extraction, current extraction persists in physical units forever, or equals recharge once an aquifer is drained. Under sustainability, the agent consumes recharge every period.

Each counterfactual has an institutional analogue: common-pool extraction can induce myopia; the prior appropriation doctrine of the Western U.S. or growing water subsidies could induce constant extraction; and sustainable extraction reflects the spirit of California’s Sustainable Groundwater Management Act. Since our main results find modest welfare costs of the current path, the first two counterfactuals ask whether weakening those results’ assumption of complete dynamic optimization implies larger costs. The third highlights classic trade-offs between sustainable extraction and extraction that maximizes the aquifer’s present value (Arrow et al. 2004).

Table 4 compares each counterfactual to the current path in baseline net extraction, long-run fill, long-run marginal extraction costs, and welfare losses, and Appendix Figure A.10 shows the distributions of long-run fill.

Counterfactual: Myopic Extraction

Myopia produces the largest departures from the current extraction path. In the mean aquifer, baseline net extraction rises from a tenth of a percent of aquifer capacity to 6.4%; if this rate of net extraction persisted, the mean aquifer would lose 95% of its capacity within a half century ($= 1 - 0.94^{50}$). Long-run fill in the mean aquifer falls from 90%

Table 4: Impact of Counterfactual Scenarios on Resource Management

	Median aquifer (1)	Mean aquifer (2)	Central Valley (3)	Ogallala (4)	Indus (5)	North China (6)
<i>Panel A: Baseline Net Water Use (% of aquifer capacity)</i>						
Observed extraction path	0.0	0.1	0.2	-0.3	0.4	0.4
Ramsey discounting	-0.1	-0.5	-0.5	-0.6	-0.2	-0.1
Market discounting	-0.5	-1.0	-0.7	-0.8	-0.6	-0.9
Myopic extraction	4.1	6.4	2.8	2.6	1.1	7.4
<i>Panel B: Long-Run Fill (% of aquifer capacity)</i>						
Observed extraction path	94.1	89.8	88.5	98.4	87.0	91.5
Ramsey discounting	100.0	86.5	100.0	100.0	97.5	98.1
Market discounting	100.0	94.9	100.0	100.0	100.0	100.0
Myopic extraction	71.6	61.4	75.5	73.4	84.1	67.7
Constant water extraction	91.1	60.7	0.0	100.0	0.0	0.0
Zero net extraction	94.8	93.2	92.4	90.5	91.9	96.6
<i>Panel C: Long-Run Marginal Extraction Cost (cents/m³)</i>						
Observed extraction path	1.8	3.1	5.7	0.5	8.9	3.2
Ramsey discounting	0.1	3.2	0.0	0.0	1.8	0.7
Market discounting	0.0	1.3	0.0	0.0	0.0	0.0
Myopic extraction	10.1	10.4	12.2	7.4	10.9	12.0
Constant water extraction	3.9	14.2	49.6	0.0	68.2	37.2
Zero net extraction	1.5	2.3	3.8	2.6	5.7	1.4
<i>Panel D: Welfare Loss (Ramsey) (% of aquifer value)</i>						
Observed extraction path	1.4	2.6	2.8	0.5	2.6	0.2
Myopic extraction	8.9	9.2	10.7	10.2	5.4	8.4
Constant water extraction	1.7	3.4	4.9	1.2	6.8	0.5
Zero net extraction	1.5	2.9	1.7	1.7	0.8	0.1
<i>Panel E: Welfare Loss (Market) (% of aquifer value)</i>						
Observed extraction path	1.7	2.9	5.5	1.2	10.8	3.4
Myopic extraction	14.3	12.9	14.6	14.4	14.3	16.0
Constant water extraction	2.2	5.4	12.3	2.4	38.9	11.3
Zero net extraction	1.6	2.9	3.7	4.1	6.6	1.6

Notes: Ramsey and market discounting describe the extraction paths a planner chooses with the Ramsey discount factor (from [Addicott et al. 2020](#)) or the market discount factor (0.98). Myopia describes extraction under a discount factor of $\beta = 0$. Constant water extraction repeats the baseline extraction of the current path in every future year. Zero net extraction extracts recharge in each year, so aquifer fill stays constant. In Panel A, constant water extraction equals the observed extraction path and zero net extraction equals zero by construction, so the panel omits both. All estimates are based on the model in Section 3.3 and Appendix C.1 and are unweighted. Dollar values are in 2023 U.S. dollars.

under the current path to 61%, and the share of aquifers nearly drained in the long run rises from about 1% to about 10%. Because myopic extraction lowers the water table, long-run marginal extraction costs more than triple. The resulting welfare losses are 9.2% to 12.9% of aquifer present value, roughly four times those of the current path. Intuitively, myopia front-loads extraction into the baseline and immediately subsequent years; this maximizes profits early, but profits then decline rapidly as the water table falls.

Counterfactual: Constant Extraction

Constant extraction produces smaller departures than myopia. By construction, baseline net extraction under constant extraction and under the current path are equal. Long-run fill in the mean aquifer is 61%, similar to myopia, but the distribution across aquifers is bimodal: constant extraction completely drains 30.8% of aquifers and completely fills 45.2%.³⁵ For example, constant extraction empties the Central Valley Aquifer System, where extraction exceeds recharge, but fills the Ogallala Aquifer. Long-run marginal extraction costs in the mean aquifer rise more than four-fold. The resulting welfare losses are 3.4% to 5.4% of aquifer present value—above the current path but well below myopia. Intuitively, constant extraction diverges from the current path only gradually, and the two paths differ most in far future years, when differences matter less in present value; the welfare cost relative to the current path arises because repeating current extraction forever ignores feedback from rising marginal extraction costs.

Counterfactual: Zero Net Extraction

Sustainability produces the smallest changes of the three counterfactuals. Baseline net extraction is zero by construction, close to the current path for the mean aquifer. Long-run fill in the mean aquifer is 93%, slightly above the current path. The fill distribution is compressed—no aquifer is completely drained, and none has long-run fill below 60%. Long-run marginal extraction costs are similar to the current path. The resulting welfare losses are 2.9% of aquifer present value under either benchmark, essentially the same as the current extraction path—an average that combines welfare gains for aquifers depleting on net (e.g., the Central Valley Aquifer System, the Indus Basin, and the North China Aquifer System) with welfare losses for aquifers that would refill under the current path. Intuitively, sustainable extraction equals optimal extraction in the long run, since both equal recharge; the welfare cost arises because freezing fill at its current, inefficiently low level prevents most aquifers

³⁵Steady state fill under constant extraction is mechanical: it fully drains aquifers where observed extraction exceeds recharge and completely fills aquifers where recharge exceeds observed extraction.

from reaching their value-maximizing fill. A sustainable policy can therefore still have large welfare costs relative to optimal extraction.

8 Conclusions

Economists have long worried that inefficiently rapid extraction of natural resources generates large social costs. To put numbers on these concerns, we apply a classic model of resource extraction to modern remote sensing and economic data.

This paper’s title describes its research question: are people consuming too much groundwater? Our answer is both no and yes. For the median global aquifer, extraction equals recharge today, as in any steady state, although most individual aquifers have non-zero net extraction. At the same time, we calculate that most aquifers have inefficiently low stocks. The optimal plan would cut short-run extraction to refill most aquifers, then reach a steady state in which extraction equals recharge. Lower short-run extraction reduces short-run profits, but the long-term benefit of a higher water table, which lowers extraction costs and raises profits, outweighs this cost.

We find that three-fifths to four-fifths of aquifers extract groundwater more rapidly than is optimal, potentially due to failures of institutions or property rights, generating global costs of several trillion dollars, or about 3% of aquifers’ global present value. Myopia or repeating current extraction indefinitely raises these costs, though they remain under 13% of present value for the average aquifer. A third counterfactual where users sustainably consume only recharge has a moderate welfare cost relative to the optimal extraction path.

Although natural resource theory focuses on the distinction between common-pool and single-agent extraction, we instead estimate the discount factor that best characterizes the current extraction path. We show that it relates monotonically to a property rights interpretation.

We leave several avenues for future work. Spatial analysis of the environment and natural resources has burgeoned ([Balboni and Shapiro 2025](#)), and future work could combine our intertemporal analysis with interregional (trade and spatial) optimization, a combination that has informed climate change analysis but remains unexplored for groundwater.

Future work could also study stochastic fluctuations around the long-run policy we analyze. One constraint is limited evidence on stochastic recharge for each global aquifer. Using aquifers as a buffer stock would raise their value beyond our nonstochastic estimates, amplifying our findings that most aquifers are extracted too rapidly and that their global present value is enormous. Stochastic recharge matters most when an aquifer is empty, however, which is rare in our scenarios, particularly in the short run. Future work could test whether

long-run extraction trajectories matter more for welfare than short-run fluctuations, a comparison that lightly echoes growth versus business cycle analysis in macroeconomics (Lucas 1988) and the focus of climate change research on long-term dynamics rather than short-run fluctuations. Some case study evidence finds that stochastic refill matters less than long-run extraction (e.g., Merrill and Guilfoos 2018), though whether this generalizes is unknown.

Aquifers also interact with other natural resources and ecosystem services, such as through saltwater intrusion or surface water. Data limit accounting for these in all global aquifers, but clarifying their importance for groundwater policy would be valuable.

Finally, classic literature and modern policy papers investigate how to account for environmental and natural resource goods in national accounts (Kokkelenberg and Nordhaus 1999; Prabhakar et al. 2023; UN 2024). Governments in the U.S., UK, and elsewhere have implemented national strategies to formalize this measurement. Our estimates suggest that the global value of aquifers is large. An important question is how statistical agencies could adapt the analysis from this paper to measure dynamic values of natural resources.

References

- Addicott, Ethan T, Eli P Fenichel, and Matthew J Kotchen, “Even the representative agent must die: using demographics to inform long-term social discount rates,” *Journal of the Association of Environmental and Resource Economists*, 2020, 7 (2), 379–415.
- Alesina, Alberto, Paola Giuliano, and Nathan Nunn, “On the Origins of Gender Roles: Women and the Plough,” *Quarterly Journal of Economics*, 2013, 128 (2), 469–530.
- Arrow, Kenneth, Partha Dasgupta, Lawrence Goulder, Gretchen Daily, Paul Ehrlich, Geoffrey Heal, Simon Levin, Karl-Göran Mäler, Stephen Schneider, David Starrett, and Brian Walker, “Are We Consuming Too Much?,” *Journal of Economic Perspectives*, 2004, 18 (3), 147–172.
- Ayres, Andrew B, Kyle C Meng, and Andrew J Plantinga, “Do environmental markets improve on open access? Evidence from California groundwater rights,” *Journal of Political Economy*, 2021, 129 (10), 2817–2860.
- Badiani-Magnusson, Reena and Katrina Jessoe, “Electricity prices, groundwater, and agriculture: The environmental and agricultural impacts of electricity subsidies in India,” in “Agricultural productivity and producer behavior,” University of Chicago Press, 2018, pp. 157–183.
- Balboni, Clare and Joseph S Shapiro, “Spatial environmental economics,” in “Handbook of Regional and Urban Economics,” Vol. 6, Elsevier, 2025, pp. 585–652.
- Bauer, Michael D and Glenn D Rudebusch, “The rising cost of climate change: evidence from the bond market,” *Review of Economics and Statistics*, 2023, 105 (5), 1255–1270.
- Belloni, A., D. Chen, V. Chernozhukov, and C. Hansen, “Sparse Models and Methods for Optimal Instruments With an Application to Eminent Domain,” *Econometrica*, 2012, 80 (6), 2369–2429.

- Berman, Aaron and Nathaniel Hickok**, “Environmental Regulation with Irreversible Investments: Evidence from High Plains Aquifer Depletion,” 2025. Working Paper, MIT.
- BGR**, “World-wide Hydrogeological Mapping and Assessment Programme (WHYMAP),” Jan 2025. https://www.whymap.org/whymap/EN/Home/whymap_node.html, accessed 5/19/2021.
- Blundell, Richard and Stephen Bond**, “GMM Estimation with Persistent Panel Data: An Application to Production Functions,” *Econometric Reviews*, 2000, 19 (3), 321–340.
- Borunda, Alejandra**, “The deceptively simple plan to replenish California’s groundwater,” *National Geographic*, 2022, March 23.
- Burlig, Fiona, Christopher Knittel, David Rapson, Mar Reguant, and Catherine Wolfram**, “Machine learning from schools about energy efficiency,” *Journal of the Association of Environmental and Resource Economists*, 2020, 7 (6), 1181–1217.
- , Louis Preonas, and Matt Woerman**, “Groundwater and Crop Choice in the Short and Long Run,” 2026. NBER Working Paper 28706.
- Carleton, Tamma and Michael Greenstone**, “A Guide to Updating the US Government’s Social Cost of Carbon,” *Review of Environmental Economics and Policy*, 2022, 16 (2), 196–218.
- , Levi Crews, and Ishan Nath**, “Is the World Running Out of Fresh Water?,” *AEA Papers and Proceedings*, 2024, 114, 31–35.
- , -, and -**, “Agriculture, Trade, and Global Water Use,” 2025. Working Paper, UC Berkeley.
- Chen, Chaoran, Diego Restuccia, and Raül Santaeulàlia-Llopis**, “Land misallocation and productivity,” *American Economic Journal: Macroeconomics*, 2023, 15 (2), 441–465.
- Cline, William R.**, *The Economics of Global Warming*, Washington, DC: Institute for International Economics, 1992.
- Coase, R. H.**, “The Problem of Social Cost,” *Journal of Law and Economics*, 1960, 3, 1–44.
- Coman, Katharine**, “Some unsettled problems of irrigation,” *American Economic Review*, 1911, 1 (1), 1–19.
- Costinot, Arnaud, Dave Donaldson, and Cory Smith**, “Evolving Comparative Advantage and the Impact of Climate Change in Agricultural Markets: Evidence from 1.7 Million Fields around the World,” *Journal of Political Economy*, 2016, 124 (1), 205–248.
- Dasgupta, Partha S. and Geoffrey Heal**, *Economic Theory and Exhaustible Resources*, Cambridge University Press, 1979.
- Davis, Lucas W, Sebastian Martinez, and Bibiana Taboada**, “How effective is energy-efficient housing? Evidence from a field trial in Mexico,” *Journal of Development Economics*, 2020, 143, 102390.
- de Graaf, I. E. M., E. H. Sutanudjaja, L. P. H. van Beek, and M. F. P. Bierkens**, “A high-resolution global-scale groundwater model,” *Hydrology and Earth System Sciences*, 2015, 19, 823–837.
- D’Odorico, P., D. D. Chiarelli, L. Rosa, A. Bini, D. Zilberman, and M. C. Rulli**, “The Global Value of Water in Agriculture,” *Proceedings of the National Academy of Sciences*, 2020, 117 (36), 21985–21993.
- Donaldson, Dave and Adam Storeygard**, “The View from Above: Applications of Satellite Data in Economics,” *Journal of Economic Perspectives*, 2016, 30 (4), 171–198.

- Donovan, Kevin**, “The equilibrium impact of agricultural risk on intermediate inputs and aggregate productivity,” *Review of Economic Studies*, 2021, 88 (5), 2275–2307.
- Doraszelski, Ulrich and Jordi Jaumandreu**, “Measuring the bias of technological change,” *Journal of Political Economy*, 2018, 126 (3), 1027–1084.
- Drupp, Moritz A., Mark C. Freeman, Ben Groom, and Frikk Nesje**, “Discounting disentangled,” *American Economic Journal: Economic Policy*, 2018, 10 (4), 109–134.
- DWR**, “Technical Information for Preparing Water Transfer Proposals (Water Transfer White Paper),” Technical Report, California Department of Water Resources 2019.
- Eaton, Jonathan and Samuel Kortum**, “Technology, Geography, and Trade,” *Econometrica*, 2002, 70 (5), 1741–1779.
- Edwards, Eric C and Todd Guilfoos**, “The economics of groundwater governance institutions across the globe,” *Applied Economic Perspectives and Policy*, 2021, 43 (4), 1571–1594.
- EIA**, “Many industrial electricity customers are farmers,” 2014. <https://www.eia.gov/todayinenergy/detail.php?id=16231>, accessed October 13, 2025.
- Fan, Y., H. Li, and G. Miguez-Macho**, “Global patterns of groundwater table depth,” *Science*, 2013, 339 (6122), 940–943.
- FAO and IIASA**, “Global Agro-Ecological Zones version 4 (GAEZ v4),” Feb 2024. <http://www.fao.org/gaez/>, accessed 6/28/2022.
- Faunt, Claudia C.**, “Groundwater Availability of the Central Valley Aquifer, California,” Technical Report, U.S. Geological Survey 2009.
- Fowlie, Meredith, Michael Greenstone, and Catherine Wolfram**, “Do energy efficiency investments deliver? Evidence from the weatherization assistance program,” *Quarterly Journal of Economics*, 2018, 133 (3), 1597–1644.
- FRED**, “Gross Domestic Product: Implicit Price Deflator,” Oct 2024. <https://fred.stlouisfed.org/series/GDPDEF>, accessed 10/28/2024.
- GGDC**, “Penn World Table Version 10,” August 2022. <https://www.rug.nl/ggdc/productivity/pwt/>, accessed 8/31/2022.
- Giglio, Stefano, Matteo Maggiori, and Johannes Stroebel**, “Very Long-Run Discount Rates,” *Quarterly Journal of Economics*, 2015, 130 (1), 1–53.
- Gisser, Micha and David A Sanchez**, “Competition versus optimal control in groundwater pumping,” *Water Resources Research*, 1980, 16 (4), 638–642.
- Gleeson, T., N. Moosdorf, J. Hartmann, and L.P.H. van Beek**, “A glimpse beneath earth’s surface: GLobal HYdrogeology MaPS (GLHYMPS) of permeability and porosity,” *Geophysical Research Letters*, 2014, 41 (11), 3891–3898.
- Gollin, Douglas and Christopher Udry**, “Heterogeneity, measurement error, and misallocation: Evidence from African agriculture,” *Journal of Political Economy*, 2021, 129 (1), 1–80.
- ., David Lagakos, and Michael E Waugh**, “The agricultural productivity gap,” *Quarterly Journal of Economics*, 2014, 129 (2), 939–993.
- Gordon, H. Scott**, “The Economic Theory of a Common-Property Resource: The Fishery,” *Journal of Political Economy*, 1954, 62 (2), 124–142.
- Hadachek, Jeffrey, Ellen M Bruno, Nick Hagerty, and Katrina Jessoe**, “Externalities of Climate Adaptation in Common-Pool Groundwater Resources,” *Journal of Public Economics*, Forthcoming.

- Harris, I.C. and P.D. Jones**, “Climatic Research Unit (CRU) Time-Series (TS) Version 3.26 of High-Resolution Gridded Data of Month-by-month Variation in Climate (Jan. 1901- Dec. 2017),” Mar 2019. <http://dx.doi.org/10.5285/7ad889f2cc1647efba7e6a356098e4f3>, accessed 8/3/2019.
- Hornbeck, Richard and Pinar Keskin**, “The historically evolving impact of the Ogallala Aquifer: Agricultural adaptation to groundwater and drought,” *American Economic Journal: Applied Economics*, 2014, 6 (1), 190–219.
- Hotelling, Harold**, “The Economics of Exhaustible Resources,” *Journal of Political Economy*, 1931, 39, 137–175.
- Houthakker, HS**, “The Pareto distribution and the Cobb-Douglas production function in activity analysis,” *Review of Economic Studies*, 1955, 23 (1), 27–31.
- IEA**, “IEA World Energy Prices,” Sept 2021. <https://www.iea.org/data-and-statistics/data-product/energy-prices>.
- ILO**, “ILOSTAT Database,” June 2023. <https://ilostat.ilo.org/>, accessed 6/7/2023.
- IMF**, “Interest Rate Data: Selected Indicators,” June 2023. <https://data.imf.org/regular.aspx?key=63087881>, accessed 6/7/2023.
- Kahneman, Daniel and Robert Sugden**, “Experienced utility as a standard of policy evaluation,” *Environmental and Resource Economics*, 2005, 32, 161–181.
- Karp, Larry**, *Natural Resources as Capital*, MIT Press, 2017.
- and Hiroaki Sakamoto**, “Sober optimism and the formation of international environmental agreements,” *Journal of Economic Theory*, 2021, 197, 105321.
- Kokkelenberg, Edward C and William D Nordhaus**, *Nature’s Numbers: Expanding the National Economic Accounts to Include the Environment*, National Academies Press, 1999.
- Kotchoni, D. O. Valerie, Jean-Michel Vouillamoz, Fabrice M. A. Lawson, Philippe Adjomayi, Moussa Boukari, and Richard G. Taylor**, “Relationships between rainfall and groundwater recharge in seasonally humid Benin: a comparative analysis of long-term hydrographs in sedimentary and crystalline aquifers,” *Hydrogeology Journal*, 2019, 27, 447–457.
- Koundouri, Phoebe**, “Potential for groundwater management: Gisser-Sanchez effect reconsidered,” *Water Resources Research*, 2004, 40 (6).
- Kummu, Matti, Maija Taka, and Joseph H. A. Guillaume**, “Gridded global datasets for Gross Domestic Product and Human Development Index over 1990-2015,” *Scientific Data*, 2018, 5 (180004), 1–15.
- Lawell, C.-Y. Cynthia Lin**, “The Management of Groundwater,” *Annual Review of Resource Economics*, 2016, 8, 247–259.
- Levhari, David and Leonard Mirman**, “The Great Fish War: An Example Using a Dynamic Cournot-Nash Solution,” *Bell Journal of Economics*, 1980, 11 (1), 322–334.
- Li, Xuecao, Yuyu Zhou, Min Zhao, and Xia Zhao**, “A harmonized global nighttime light dataset 1992-2018,” *Scientific Data*, 2020, 7 (168).
- Libecap, Gary D**, “Institutional path dependence in climate adaptation: Coman’s ‘Some Unsettled Problems of Irrigation’,” *American Economic Review*, 2011, 101 (1), 64–80.
- Lucas, Robert E**, “On the mechanics of economic development,” *Journal of Monetary Economics*, 1988, 22 (1), 3–42.

- Marschak, Jacob and William H. Andrews**, “Random Simultaneous Equations and the Theory of Production,” *Econometrica*, 1944, 12 (3/4), 143–205.
- Martin-Candilejo, Araceli, David Santillan, and Luis Garrote**, “Pump Efficiency Analysis for Proper Energy Assessment in Optimization of Water Supply Systems,” *Water*, 2020, 12 (1), 132.
- Meinzen-Dick, Ruth and Bryan Bruns**, “Crafting combinations to govern groundwater,” *International Journal of the Commons*, 2024, 18 (1), 585–600.
- Merrill, Nathaniel H and Todd Guilfoos**, “Optimal groundwater extraction under uncertainty and a spatial stock externality,” *American Journal of Agricultural Economics*, 2018, 100 (1), 220–238.
- Miller, James A.**, “Principal Aquifers,” in “The National Atlas of the United States of America,” U.S. Geological Survey, 2003.
- Montero, Eduardo and Dean Yang**, “Religious festivals and economic development: Evidence from the timing of Mexican saint day festivals,” *American Economic Review*, 2022, 112 (10), 3176–3214.
- Müller Schmied, Hannes, Denise Cáceres, Stephanie Eisner, Martina Flörke, Claudia Herbert, Christoph Niemann, Thedini Asali Peiris, Eklavya Popat, Felix Theodor Portmann, Robert Reinecke, Maike Schumacher, Somayeh Shadkam, Camelia-Eliza Telteu, Tim Trautmann, and Petra Döll**, “The global water resources and use model WaterGAP v2.2d: model description and evaluation,” *Geoscientific Model Development*, 2021, 14 (2), 1037–1079.
- Nargi, Lela**, “Can desalination quench agriculture’s thirst?,” *Knowable Magazine*, 2024. November 26, <https://knowablemagazine.org/content/article/food-environment/2024/can-desalination-of-groundwater-grow-crops> (accessed March 3, 2026).
- Nordhaus, William D.**, “Economic growth and climate: the carbon dioxide problem,” *American Economic Review*, 1977, 67 (1), 341–346.
- OECD**, “Capital Markets,” 2025. <https://www.oecd.org/en/topics/sub-issues/capital-markets.html>, accessed November 20, 2025.
- Ojha, C., M. Shirzaei, S. Werth, D.F. Argus, and T.G. Farr**, “Sustained groundwater loss in California’s Central Valley exacerbated by intense drought periods,” *Water Resources Research*, 2018, 54, 4449–4460.
- Ostrom, Elinor**, *Governing the Commons*, Cambridge University Press, 1990.
- ..**, “Reflections on ‘Some Unsettled Problems of Irrigation’,” *American Economic Review*, 2011, 101 (1), 49–63.
- Pearce, David, Anil Markandya, and Edward Barbier**, *Blueprint for a Green Economy*, Earthscan, 1989.
- Prabhakar, Arati, Gina M. Raimondo, and Shalanda D. Young**, “National Strategy to Develop Statistics for Environmental-Economic Decisions: A U.S. System of Natural Capital Accounting and Associated Environmental-Economic Statistics,” Technical Report, Office of Science and Technology Policy 2023.
- Pradeleix, Ludivine, Philippe Roux, Sami Bouarfa, Bochra Jaouani, Zohra Lili-Chabaane, and Veronique Bellon-Maurel**, “Environmental Impacts of the Contrasted Groundwater Pumping Systems Assessed by Life Cycle Assessment Methodology: Contribution to the Water-Energy Nexus Study,” *Irrigation and Drainage*, 2015, 64, 124–138.

- Proctor, Jonathan, Tamma Carleton, and Sandy Sum**, “Parameter recovery using remotely sensed variables,” 2023. NBER Working Paper 30861.
- Provencher, Bill and Oscar Burt**, “The externalities associated with the common property exploitation of groundwater,” *Journal of Environmental Economics and Management*, 1993, *24* (2), 139–158.
- Qin, Jingxiu, Weili Duan, Shan Zou, Yaning Chen, Wenjing Huang, and Lorenzo Rosa**, “Global energy use and carbon emissions from irrigated agriculture,” *Nature Communications*, 2024, *15* (3084).
- Quaas, Martin F, Rainer Froese, Helmut Herwartz, Till Requate, Jörn O Schmidt, and Rüdiger Voss**, “Fishing industry borrows from natural capital at high shadow interest rates,” *Ecological Economics*, 2012, *82*, 45–52.
- Ramankutty, N., A.T. Evan, C. Monfreda, and J.A. Foley**, “Farming the planet: Geographic distribution of global agricultural lands in the year 2000,” *Global Biogeochemical Cycles*, 2008, *22*, GB1003.
- Ramsey, Frank P.**, “A Mathematical Theory of Saving,” *Economic Journal*, 1928, *38* (152), 543–559.
- Rennert, Kevin, Frank Errickson, Brian C. Prest, Lisa Rennels, Richard G. Newell, William Pizer, Cora Kingdon, Jordan Wingenroth, Roger Cooke, Bryan Parthum, David Smith, Kevin Cromar, Delavane Diaz, Frances C. Moore, Ulrich K. Muller, Richard J. Plevin, Adrian E. Raftery, Hana Sevcikova, Hannah Sheets, James H. Stock, Tammy Tan, Mark Watson, Tony E. Wong, and David Anthoff**, “Comprehensive evidence implies a higher social cost of CO₂,” *Nature*, 2022, *610*, 687–692.
- Richey, Alexandra S., Brian F. Thomas, Min-Hui Lo, James S. Famiglietti, Sean Swenson, and Matthew Rodell**, “Uncertainty in global groundwater storage estimates in a Total Groundwater Stress framework,” *Water Resources Research*, 2015a, *51*, 5198–5216.
- , —, —, John T. Reager, James S. Famiglietti, Katalin Voss, Sean Swenson, and Matthew Rodell**, “Quantifying renewable groundwater stress with GRACE,” *Water Resources Research*, 2015b, *51* (7), 5217–5238.
- Robinson, D.W.**, “Construction and Operating Costs of Groundwater Pumps for Irrigation in the Riverine Plain,” Technical Report, CSIRO Land and Water 2002.
- Rodell, M., J.S. Famiglietti, D.N. Wiese, J.T. Reager, H.K. Beaudoin, F.W. Landerer, and M.H. Lo**, “Emerging Trends in Global Freshwater Availability,” *Nature*, 2018, *557*, 651–659.
- Rodell, Matthew, Isabella Velicogna, and James S. Famiglietti**, “Satellite-based estimates of groundwater depletion in India,” *Nature*, 2009, *460*, 999–1002.
- , James S. Famiglietti, David N. Wiese, John T. Reager, Hiroko K. Beaudoin, Felix W. Landerer, and Min-Hui Lo**, “Trends in Global Freshwater Availability from the Gravity Recovery and Climate Experiment (GRACE), 2002–2016,” 2019. <https://doi.org/10.7927/H4TT4P2C>, accessed 8/22/2018.
- , Jianli Chen, Hiroko Kato, James S. Famiglietti, Joe Nigro, and Clark R. Wilson**, “Estimating Groundwater Storage Changes in the Mississippi River Basin (USA) using GRACE,” *Hydrogeology Journal*, 2007, *15* (1), 159–166.

- Rodella, Aude-Sophie, Esha Zaveri, and Francois Bertone**, “The hidden wealth of nations: The economics of groundwater in times of climate change,” Technical Report, The World Bank Group 2023.
- Rojanasakul, Mira, Christopher Flavelle, Blacki Migliozi, and Eli Murray**, “America Is Using Up Its Groundwater Like There’s No Tomorrow,” *New York Times*, 2023. August 28.
- Ryan, Nicholas and Anant Sudarshan**, “Rationing the Commons,” *Journal of Political Economy*, 2022, 130 (1), 210–257.
- Scanlon, Bridget R., Claudia C. Faunt, Laurent Longuevergne, Robert C. Reedy, William M. Alley, Virginia L. McGuire, and Peter B. McMahon**, “Groundwater depletion and sustainability of irrigation in the US High Plains and Central Valley,” *Proceedings of the National Academy of Sciences*, 2012, 109 (24), 9320–9325.
- Siebert, S., J. Burke, J.M. Faures, K. Frenken, J. Hoogeveen, P. Döll, and F. T. Portmann**, “Groundwater use for irrigation – a global inventory,” *Hydrology and Earth System Sciences*, 2010, 14, 1863–1880.
- Stavins, Robert N.**, “The problem of the commons: Still unsettled after 100 years,” *American Economic Review*, 2011, 101 (1), 81–108.
- Stock, James H. and Motohiro Yogo**, “Testing for Weak Instruments in Linear IV Regression,” in Donald W. K. Andrews and James H. Stock, eds., *Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg*, Cambridge University Press, 2005, pp. 80–108.
- Sutanudjaja, Edwin H., Rens van Beek, Niko Wanders, Yoshihide Wada, Joyce H. C. Bosmans, Niels Drost, Ruud J. van der Ent, Inge E. M. de Graaf, Jan-nis M. Hoch, Kor de Jong, Derek Karssenberg, Patricia Lopez Lopez, Stefanie Pessenteiner, Oliver Schmitz, Menno W. Straatsma, Ekkamol Vannamettee, Dominik Wisser, and Marc F. P. Bierkens**, “PCR-GLOBWB 2: A 5 arcmin global hydrological and water resources model,” *Geoscientific Model Development*, 2018, 11, 2429–2453.
- Timmins, Christopher**, “Measuring the Dynamic Efficiency Costs of Regulators’ Preferences: Municipal Water Utilities in the Arid West,” *Econometrica*, 2002, 70 (2), 603–629.
- Tombe, Trevor**, “The missing food problem: Trade, agriculture, and international productivity differences,” *American Economic Journal: Macroeconomics*, 2015, 7 (3), 226–258.
- UN**, “System of Environmental-Economic Accounting: Ecosystem Accounting,” Technical Report, United Nations 2024.
- USGS**, “Water-Level and Recoverable Water in Storage Changes, High Plains Aquifer, Predevelopment to 2019 and 2017 to 2019,” Technical Report, United States Geological Survey Scientific Investigations Report 2023-5143 2023.
- World Bank**, “World Development Indicators,” Dec 2023. <https://www.worldbank.org/en/home>.
- Zolghadr-Asli, Babak, Neil McIntyre, Slobodan Djordjevic, Raziye Farmani, and Liliana Pagliero**, “The sustainability of desalination as a remedy to the water crisis in the agriculture sector: An analysis from the climate-water-energy-food nexus perspective,” *Agricultural Water Management*, 2023, 286, 108407.

Online Appendix

A Background on Case Study Aquifer Systems

This appendix provides additional background on the four case study aquifer systems.³⁶ We choose these four for their capacity, population, GDP, and prominence in public discussion, and because in the hydrological map data one aquifer almost completely contains each system (BGR 2025).

The Central Valley Aquifer System supports 40% of American fresh produce (Borunda 2022). California has historically not metered groundwater extraction, so farmers pump with little oversight. In response to a falling water table, California’s 2014 Sustainable Groundwater Management Act aims for sustainable extraction by 2042, using price and non-price policies (Burlig et al. 2026; Bruno et al. 2022).

The Central Valley Aquifer System and the Ogallala Aquifer have the greatest extraction among U.S. aquifers, together accounting for about half of U.S. groundwater extraction since 1900. The Ogallala is larger but less intensively cultivated than the Central Valley, and produces about 30% of U.S. agricultural output (Rhodes et al. 2023).

The Indus Basin straddles arid regions of India, Pakistan, China, and Afghanistan and may have the highest net depletion of any major aquifer system, driven by agricultural demand and low precipitation (Richey et al. 2015b).

The North China Aquifer System also has among the highest global depletion rates. The arid plain above it is China’s largest wheat and maize producing area, and agriculture accounts for most of its groundwater extraction. Between 2003 and 2010, it lost 50 km³ of groundwater, more than the capacity of China’s Three Gorges Dam (Feng et al. 2013).

B Derivation of Model Results

Given w , the agent faces the static optimization problem

$$\max_{k,l} \mathcal{Z}_i k^\alpha l^\gamma w^{\phi_i} x_i^{1-\alpha-\gamma-\phi_i} - p_i^k k - p_i^l l$$

³⁶Names for these aquifer systems vary across researchers; we follow the naming of the National Aeronautics and Space Administration (NASA) (Richey et al. 2015a). The Ogallala Aquifer is the largest aquifer in the High Plains Aquifer System, and NASA often names the whole system after this main unit. Some major aquifer systems, such as the Indus Basin, take their names from the surface water basins they almost perfectly overlap.

where p_i^k is the rental price of capital and p_i^l is the price of labor. The first-order conditions are

$$\begin{aligned}\frac{\alpha y_i(w)}{k} &= p_i^k \\ \frac{\gamma y_i(w)}{l} &= p_i^l\end{aligned}$$

Substituting $k = \alpha y_i(w)/p_i^k$ and $l = \gamma y_i(w)/p_i^l$ into the production function gives

$$\begin{aligned}y_i(w) &= \mathcal{Z}_i \left(\frac{\alpha y_i(w)}{p_i^k} \right)^\alpha \left(\frac{\gamma y_i(w)}{p_i^l} \right)^\gamma w^{\phi_i} x_i^{1-\alpha-\gamma-\phi_i} \\ \Rightarrow y_i(w) &= \left[\mathcal{Z}_i x_i^{1-\alpha-\gamma-\phi_i} \left(\frac{\alpha}{p_i^k} \right)^\alpha \left(\frac{\gamma}{p_i^l} \right)^\gamma \right]^{\frac{1}{1-\alpha-\gamma}} w^{\frac{\phi_i}{1-\alpha-\gamma}}\end{aligned}$$

which can be written as equation (3)

$$y_i(w) = z_i w^{\frac{\phi_i}{\theta}}$$

Here $z_i = [\mathcal{Z}_i x_i^{1-\alpha-\gamma-\phi_i} (\alpha/p_i^k)^\alpha (\gamma/p_i^l)^\gamma]^{1/(1-\alpha-\gamma)}$ and $\theta = 1 - \alpha - \gamma$ are functions of terms the agent takes as fixed. Substituting the first-order conditions into the objective function gives profit in equation (5), net of water costs:

$$\begin{aligned}\tilde{\pi}_i(w) &= y_i(w) - \alpha y_i(w) - \gamma y_i(w) \\ &= \theta y_i(w)\end{aligned}$$

C Computational Details

C.1 Solving the Water Use Problem

We solve the water extraction problem (9) by value function iteration on an evenly spaced grid of 100 nodes for aquifer fill $a \in [0, 1]$. We suppress the aquifer index i , write a_j for discretized fill levels, and proceed as follows:

1. Guess an initial discretized value function $V^0(a_j)$.³⁷ Set the iteration counter to $n = 0$.
2. Interpolate $V^n(a)$ and its derivative $\partial_a V^n(a)$ with cubic splines.
3. Compute the unconstrained extraction w_j^u that solves the unconstrained first-order

³⁷We implement this with an initial guess of $V^0(a_j)$ everywhere.

condition (10) at each a_j :

$$\phi z(w_j^u)^{\phi/\theta-1} - c[\underline{h} + (1 - a_j - r/2)(\bar{h} - \underline{h})] - c(\bar{h} - \underline{h})w_j^u - \beta \partial_a V^n(\min\{a_j - w_j^u + r, 1\}) = 0 \quad (25)$$

The optimal discretized policy function is $w_j = \min\{w_j^u, a_j + r\}$.

4. Update the value function by substituting the policy function into the Bellman equation (9):

$$V^{n+1}(a_j) = \pi(a_j, w_j) + \beta_i V^n(\min\{a_j - w_j + r, 1\})$$

Check for convergence. If

$$\max_j \left| \frac{V^{n+1}(a_j) - V^n(a_j)}{V^n(a_j)} \right| < \varepsilon_{tol}$$

for a numerical tolerance parameter $\varepsilon_{tol} > 0$, stop. Otherwise, set $n = n + 1$ and return to step (2).

As the main text explains, we exclude aquifers for which no $\beta_i^d \in [0, 1)$ matches observed extraction, though a sensitivity analysis adds them back. In 14% of aquifers, model-implied extraction is at the lower bound of the β_i^d values for which the model converges, and in 6% it is at the upper bound.

C.2 Computing the Property Rights Parameter λ_i

To calculate the property rights parameter λ_i of Section 6.4, we solve the agent's problem (24) for a given $\lambda_i \in [0, 1]$ as follows:


1. Guess the function $\bar{w}^0(a_j)$ and set the iteration counter to $n = 0$.
2. Given λ_i and $\bar{w}^n(a_j)$, solve (24) with the algorithm in Appendix C.1.
3. In a symmetric equilibrium, the policy function $w(a_j)$ obtained in step (2) equals $\bar{w}(a_j)$. Check for convergence: if

$$\max_j |w(a_j) - \bar{w}^n(a_j)| < \varepsilon_{tol}$$

for the numerical tolerance parameter $\varepsilon_{tol} > 0$, stop. Otherwise, update the guess for $\bar{w}(a_j)$ to

$$\bar{w}^{n+1}(a_j) = \bar{w}^n(a_j) + \varsigma[w(a_j) - \bar{w}^n(a_j)]$$

Figure A.1: Estimated Relationship Between Discount Factor and Property Rights Interpretations: Case Study Aquifer Systems



Notes: This figure shows the relationship between the property rights parameter λ and the discount factor β for the four case study aquifer systems, using the methods of Section 6.4 and Appendix C.2. We estimate λ assuming β^* is given by the Ramsey or market discount factor.


where $\varsigma \in (0, 1]$ is a numerical damping parameter. Set $n = n + 1$, and return to step (2).

We then search for the $\lambda_i \in [0, 1]$ at which model-predicted extraction equals observed extraction: $w(a^{\text{data}}; \lambda_i) = w^{\text{data}}$.

C.3 Welfare Analysis

To compute the welfare loss ω from excess extraction (equation (12)), we first use the algorithm of Appendix C.1 to compute the discretized policy function w_j implied by the decision discount factor β_i^d , and interpolate it with cubic splines to obtain $w(a)$, $a \in [0, 1]$. The law

Figure A.2: Welfare Losses Due to Difference Between Estimated and Market Discount Factors



Notes: Panel A maps each aquifer's welfare loss, as a percent of its present value, from the current extraction path deviating from the market benchmark. Panel B shows the distribution of these losses.

of motion (1) and policy function give the fill level at $t + 1$:

$$a_{t+1} = \min\{a_t - w(a_t) + r, 1\}.$$

Iterating forward from the observed fill level a_0 gives predicted paths of fill $\{a_t\}_{t=0}^T$ and extraction $\{w_t\}_{t=0}^T$. If T is large enough that a has approximately converged, we approximate the value of the current extraction path under the benchmark discount factor (equation (11)) by³⁸

$$\mathcal{V}(\mathbf{w}^*(a_0; \beta_i^d), a_0; \beta_i^*) \approx \sum_{t=0}^{T-1} (\beta_i^*)^t \pi(a_t, w_t) + \frac{(\beta_i^*)^T \pi(a_T, w_T)}{1 - \beta_i^*}.$$

We then compute the aquifer's maximum value under β_i^* , $V(a_0; \beta_i^*)$, using the algorithm in Appendix C.1. Equation (12) then gives ω .


D Data Details

D.1 Groundwater Storage and Recharge

GRACE satellites measure changes in Earth's gravity field over each aquifer and report the monthly net change in the level of total terrestrial water in vertical centimeters on a 1-by-1 degree grid from April 2002 to January 2017. We use annual changes for the full calendar years 2003 to 2016, dropping the partial years 2002 and 2017 since many of our other datasets

³⁸We set $T = 1000$ and verify ex post that the fill level of every aquifer converges over this horizon.

Figure A.3: Relationship Between Welfare Loss and Decision Discount Factor Minus Benchmarks



Notes: Panel A plots each aquifer’s welfare loss, as a percent of its present value, against the difference between its decision and Ramsey discount factors. Panel B uses the market discount factor. Each graph shows a scatterplot of the estimates and quantile bin means.

are annual, and calibrate the model to these years. Because the regressions use up to three lags, the production function estimates use 2006-2016.³⁹

GRACE reports the height of the change in water storage, rather than the change in the height that water must be lifted. The two differ because one meter of extractable groundwater may spread across many meters of porous rock. Dividing GRACE’s change in water thickness by specific yield converts it to the change in lift height.⁴⁰

The raw GRACE data measure total terrestrial water storage, which includes groundwater, surface water, snow, and vegetation and soil moisture. The Global Land Data Assimilation System (GLDAS) measures two other components, plant and soil moisture and snow water-equivalent (Beaudoing and Rodell 2016). Surface water, the remaining component, accounts for only about 8% of changes in global terrestrial water storage, while groundwater, plant and soil moisture, and snow water-equivalent account for 17%, 51%, and 23%, respectively (Rodell et al. 2007; Getirana et al. 2017).

To obtain each aquifer’s annual change in groundwater storage from the raw GRACE data, we follow Rodell et al. (2007), Rodell et al. (2009), and Richey et al. (2015a) in six steps. First, we multiply the GRACE data by the land-grid-scaling factor provided by


³⁹A separate GRACE Follow-On (GRACE-FO) mission, launched in May 2018, continues to track terrestrial water storage. Our data end with the first mission because of the two-year gap in the January-to-January panel, the pandemic and its economic fluctuations, and uncertain continuity between the two missions’ data.

⁴⁰For example, if GRACE reports a 2 m change in water thickness for an aquifer with specific yield of 0.1, lift height has changed by 20 m.

Table A.1: Data Sources

Data source	Variables	Unit	Grid resolution	Description
GRACE (Rodell et al. 2019)	Δh_{it}	cm/year	55×55 km	Height of the change in water storage; used to compute change in lift height (m)
WHYMAP (BGR 2025)	$\tilde{r}_i$	mm/year		Mean annual aquifer recharge; input to aquifer dynamics model
	s_i	m ²		Surface land area; used to derive aquifer characteristics such as capacity
WaterGAP (de Graaf et al. 2015)	$\mathcal{T}_i, \mathcal{H}_i, \mathcal{Y}_i$	m, unitless	11×11 km	Saturated thickness, topsoil thickness, specific yield; used to compute aquifer capacity $\kappa_i = (\bar{h}_i - \mathcal{H}_i)s_i\mathcal{Y}_i$
Fan et al. (2013)	$h_{i,2008}$	meters	1×1 km	Initial lift height; benchmark for water table depth in 2008
Li et al. (2020)	$\bar{y}_{it}$	Index	1×1 km	Nighttime light intensity; proxy for aquifer-level economic output
Kummu et al. (2018)		USD	9×9 km	Gridded purchasing power-adjusted GDP; alternative proxy for output
GAEZ (FAO and IIASA 2024)		tons, %	9×9 km	Yields of 26 crops under various conditions; combined with marginal returns to estimate $y'_i(\tilde{w}_i)$
D’Odorico et al. (2020)		\$/m ³	9×9 km	Marginal value of irrigation water by crop; used to estimate marginal returns to groundwater
World Bank WDI (World Bank 2023)		people		Population, employment share in agriculture; used to compute effective labor input
	H_c	years		Mean years of schooling; adjusts for differences in human capital
Penn World Table (GGDC 2022)	k_{it}	USD		Capital stock; capital input in the production function
ILO (ILO 2023)		USD		Wages; used in productivity estimation and calibration
IMF (IMF 2023)	p_i^k	%		Interest rates; used as proxy for cost of capital in production function
IEA (IEA 2021)	e_i	\$/kWh		Electricity prices; used to calculate marginal cost of groundwater extraction
Climatic Research Unit (Harris and Jones 2019)	X_{it}	mm, °C	55×55 km	Climate data; used as controls in production function and to calculate recharge
Global Hydrogeology Maps (Gleeson et al. 2014)				Soil characteristics; informs recharge and aquifer structure
Addicott et al. (2020)	β_i^*			Benchmark discount factor; used to compute welfare-optimal extraction path
HYDE / AquaKnow (Klein Goldewijk 2017)	IrrigArea _{<i>i</i>}	km ² %	9×9 km	Irrigated area; used to estimate groundwater dependence of agriculture

Figure A.4: Estimated Relationship Between the Discount Factor and Welfare Losses for Case Study Aquifer Systems



Notes: Each graph shows how a case study aquifer's welfare loss, as a percent of aquifer value, under Ramsey and market discounting varies with the decision discount factor β^d , setting beginning fill to its observed value. Dashed and dotted gray vertical lines show the Ramsey discount factor ($\beta^* = 0.96$ in the Central Valley and Ogallala, 0.92 in the Indus, and 0.91 in North China) and the market discount factor ($\beta = 0.98$).

NASA to reduce post-processing noise. Second, we calculate the January-to-January change in terrestrial water storage, averaging the datasets of the three agencies that process the data into raster format, in units of deviations from the 2004-2009 mean. Third, we isolate groundwater storage by subtracting plant and soil moisture and snow water-equivalent from the GLDAS (Beaudoing and Rodell 2016); we first convert the GLDAS data to the same units as GRACE. Fourth, we average the gridded changes within each aquifer, using the aquifer maps of the World-wide Hydrogeological Mapping and Assessment Programme (WHYMAP) (BGR 2025). Fifth, we exclude the top and bottom 1% of observations, since some are clear

outliers (e.g., changes in groundwater height of hundreds of meters). Finally, we exclude Antarctica, which several of our main datasets do not cover.

We measure recharge from the Water Global Assessment and Prognosis (WaterGAP) model, version 2, a global hydrological model that captures groundwater dynamics (de Graaf et al. 2015; Müller Schmied et al. 2021).⁴¹ These gridded recharge data account for differences in local precipitation, soil porosity, and inflow from returning surface water runoff. In Müller Schmied et al. (2021), recharge depends on daily effective precipitation (snow and rain), soil moisture, and soil porosity and saturation. Effective precipitation is precipitation net of plant canopy absorption, adjusted by a soil-specific runoff coefficient for how readily water permeates the soil. When daily effective precipitation that does not run off plus existing soil moisture exceeds the soil’s saturation threshold, the excess recharges the underlying aquifer. The model adds estimates of return flow from irrigation and inflow from surface water bodies to obtain total recharge.⁴² WHYMAP reports binned estimates of annual recharge based on this methodology (BGR 2025). We assign each aquifer the midpoint of its bin, except for the top bin (above 300 mm per year), where we assign the bin minimum, and the bottom bin (below 2 mm per year), where we assign the bin maximum. A sensitivity analysis scales each aquifer’s WHYMAP recharge, a 30-year average, by the percentage deviation of rainfall in our sample period from its 100-year average.


Calculating aquifer capacity in equation (17) uses data on specific yield, saturated thickness, topsoil thickness, lift height, and surface area. Since an aquifer’s rock may vary in specific yield, we follow Richey et al. (2015a) and use the minimum across its grid points, the most binding constraint on groundwater flow. We assume capacity is fixed over time because pollution and subsidence typically change it little (de Graaf et al. 2015; Ojha et al. 2018; Sutanudjaja et al. 2018); globally, subsidence causes a loss of less than 0.03% of aquifer capacity per decade, with heterogeneity across aquifers (Hasan et al. 2023).⁴³ Evidence on pollution’s effect on groundwater storage is limited; one study of an Indian aquifer finds that dilution from recharge may fully offset pollution incursion (Sarah et al. 2021).

⁴¹To assess this model, de Graaf et al. (2015) compare its predicted lift heights with observed well depths from the global map compiled by Fan et al. (2013). They find a correlation of 0.85-0.87 between predicted and observed data for their preferred parameterization.

⁴²This model accounts for flow between surface water and groundwater, but like other global hydrological models, it omits lateral flow between neighboring aquifers for lack of data and models recharge as a function of precipitation, land cover, and geology (e.g., de Graaf et al. 2015; Sutanudjaja et al. 2018).

⁴³Subsidence with small effects on storage capacity can have large above-ground infrastructure costs (Hasan et al. 2023).

Figure A.5: Patterns of Specific Yield, Capacity, and Area, by Aquifer



Notes: Panels A and B show aquifer-level specific yield (out of 1). Panels C and D show log capacity (in m^3). Panels E and F show log aquifer surface area (in m^2).

D.2 Measurement Error

Satellite- and model-derived data may have measurement error. We respond in two ways: we average data across years, and, apart from the four case studies, we emphasize moments and other aggregates across many aquifers rather than values for individual aquifers.

D.3 Production Proxies

In analyzing nighttime luminosity as a proxy for aquifer GDP, we use the temporally-harmonized version of the gridded remotely sensed data of global luminosity spanning 2003-2016. The source analysis, [Li et al. \(2020\)](#), harmonizes records from the two satellite programs, the Visible Infrared Imaging Radiometer Suite (VIIRS) years from 2012-2018 and the Defense Meteorological Satellite Program (DMSP) years from 1992-2013. We also analyze the two periods separately.

How well does nighttime luminosity proxy economic activity? Nighttime luminosity offers gridded resolution and global scope, while GDP data vary in quality across countries and are rarely available subnationally ([Chen and Nordhaus 2011](#); [Donaldson and Storeygard 2016](#)). Economists use night lights widely for these reasons: [Gibson et al. \(2020\)](#) estimate that over 150 economics papers since 2012 have used night lights to proxy local economic activity. [Henderson et al. \(2012\)](#) estimate a linear relationship between night lights and GDP of 0.28 to 0.53 in a panel of countries over the two decades before 2012, with a higher correlation in urban than rural areas. Changes in measurement and spatial resolution have since improved the data’s approximation of rural economic activity ([Perez-Sindin et al. 2021](#)). For the pre-2012 data, [Gibson et al. \(2020\)](#) report a statistically significant, positive relationship between night lights and GDP where agriculture is less than 20% of GDP; aquifer-level agricultural GDP data do not exist, but 80% of aquifers in our data are less than 20% cropped.

Production function estimates also use two alternative GDP outcomes. The first is country GDP for 2003 to 2016 from the World Bank’s World Development Indicators, multiplied by the share of the population employed in agriculture, from the same source, to proxy agricultural GDP. The second is gridded, purchasing power-adjusted GDP data from [Kummu et al. \(2018\)](#), available for 1990 to 2015. These gridded data scale official national statistics, so only approximate the spatial distribution of subnational economic activity.

D.4 Marginal Return to Groundwater and Costs Data

We draw irrigated and rain-fed crop production, and the marginal return to groundwater in irrigating these crops, from the GAEZ dataset of the Food and Agriculture Organization and

the International Institute for Applied Systems Analysis. GAEZ data on production under actual conditions cover banana, barley, cassava, cotton, groundnut, fodder, fruit, maize, millet, oil palm, olives, other cereals, potato, pulses, rapeseed, rice, stimulants, sorghum, soy, sugarbeet, sugarcane, sunflower, tobacco, vegetables, wheat, yams, and one “rest of crops” category. The marginal return to groundwater, from [D’Odorico et al. \(2020\)](#), varies by crop and, for some crops, by continent. They calculate it as the increase in crop production from irrigation divided by the associated increase in irrigation water.⁴⁴ Their data cover barley, cassava, groundnut, maize, millet, potatoes, oil palm, rapeseed, rice, rye, sorghum, soybean, sugarbeet, sugarcane, sunflower, and wheat. We use continent-specific prices where available (maize, rice, soybean, and wheat). For the 11 crops in GAEZ but not in [D’Odorico et al. \(2020\)](#), we average the marginal return to groundwater for the crops observed in both datasets, using the GAEZ shares as weights.

Section 5.3 calculates the marginal return to groundwater for each aquifer. To do so, we calculate the marginal return $y'_i(\tilde{w}_i) = \sum_{c=1}^{26} \sigma_{ic} \times v_c$ for each pixel, where σ_{ic} is the share of crop c in total agricultural production of each pixel within aquifer i and v_c is the value of groundwater in irrigation for crop c in that pixel. Pixels without crop production get a marginal return of 0, and the aquifer’s marginal return is the average across its pixels.

Country-level data differ somewhat in coverage. For missing years between 2003 and 2016, we interpolate linearly within each country. We impute missing country-level data using the annual continent mean.

Besides the World Development Indicators labor data in Section 4, we use gridded population. Its finer resolution approximates aquifer labor inputs more precisely than country statistics do, but it is available only every five years, so we interpolate linearly between them.

The main text calculates extraction costs from electricity prices and the efficiency of electric pumps. An alternative measure lets pump fuel type and efficiency differ across aquifers. We obtain the efficiency of electric and diesel pumps and each fuel’s share of groundwater extraction energy from [Qin et al. \(2024\)](#), and industrial electricity and diesel prices from [IEA \(2021\)](#). We re-weight these country-level variables to the aquifer level and calculate extraction cost as the average of electric and diesel pumping costs, weighted by each fuel’s share of extraction energy.

⁴⁴They estimate

$$y'_i(\tilde{w}_i) = \frac{P_i \Delta Y_i}{IWR_i}$$

where P_i is the unit price of agricultural output in location i , ΔY_i is the change in agricultural output due to irrigation, and IWR_i is the amount of water required to irrigate.

D.5 Other Data

Our analysis of heterogeneous returns to groundwater uses several measures of the value of water in irrigated agriculture. Our preferred measure is groundwater’s relative contribution to agricultural production, from GAEZ data. We measure irrigation’s effect on crop yields in each aquifer as the sum across crops of actual yield under irrigation less actual yield under rainfall, weighted by each crop’s share of total yields. We convert these effects to percentiles since yield returns in physical quantities are highly dispersed.

Many very small aquifers lack data on cropland shares; we do not classify such an aquifer as low agriculture if GAEZ reports non-zero crop value for it in 2010.

Other measures of groundwater’s local importance use gridded data on cropped land, irrigated land, and dominant land use. Irrigated land in 2005 is from the History Database of the Global Environment (HYDE) and the AquaKnow Historical Irrigation Dataset ([Klein Goldewijk 2017](#); [Siebert et al. 2015](#)). From these data we measure the cropped and irrigated shares of aquifer land area, and indicators for any cropland or irrigation above the aquifer. A separate GAEZ gridded dataset on dominant land use class indicates whether agriculture is the dominant land use above the aquifer.

To analyze correlates of the decision discount factors, we use non-agricultural aquifer characteristics: GDP per capita, democracy, and water markets. GDP per capita and democracy data are from [Kummu et al. \(2018\)](#) and [INSCR \(2022\)](#), respectively, re-weighted from countries to aquifers. We manually code water markets from [Carleton et al. \(2025\)](#) by country. International aquifers have less than 95% of the aquifer in a single country.

D.6 Model Inputs Discussion

Table 1 reports model inputs for the four case study aquifer systems and for the mean and median global aquifer.

Among the four case study aquifer systems (Table 1), the Indus Basin has the largest capacity, which contributes to its relatively low depletion as a share of capacity. All four are relatively high-value, with marginal returns ranging from $\$0.17/m^3$ for the Ogallala to $\$0.22/m^3$ for the Central Valley, in the upper half of the distribution (Figure 2). Marginal costs are 5% to 35% of marginal returns, with differences driven largely by lift height. Ramsey discount factors range from 0.91 to 0.96.

The case studies resemble the mean and median global aquifer, with a few differences. Mean and median gross extraction are about 2-3% of capacity, similar to the case studies. Mean and median marginal returns are weakly below those of the Central Valley and North China systems; weighting by aquifer crop value, GDP, or population increases marginal

returns by approximately 10%. Median and mean extraction costs, both $\$0.02/m^3$, are at the low end of the case studies, because of lift height rather than electricity prices. Ramsey discount factors are similar, also 0.91 to 0.96.

Figure 2, Panel A, maps the percent of each aquifer’s capacity extracted annually w_i (i.e., gross extraction), averaged across years. Most global land has some form of groundwater, though not all aquifers are widely used. Aquifers with hard-to-access groundwater, arid climates, or inhospitable soils, like the Australian Outback, the Arabian Peninsula, or the Sahara Desert, have limited extraction. More fertile areas like California’s Central Valley or the breadbaskets of Eastern India have higher extraction rates. The modal aquifer extracts less than 1.0% of capacity annually, though many extract up to 5.0% and some up to 10.0% (Panel B); the mean aquifer extracts 3.1% (Table 1).

Panel C of Appendix Figure A.7 maps annual recharge r_i across aquifers, equal to gross minus net extraction. Aquifers beneath the Amazon and Congolese rain forests and snow-covered Russian tundra have among the fastest recharge. Recharge correlates less clearly with agricultural production. Annual recharge in most aquifers is 0.5% to 1.5% of capacity (Panel D).⁴⁵ Because recharge has a long right tail, however, the mean across aquifers is larger, at 3.0% (Table 1).

E Heterogeneous Production Function Estimates

To examine heterogeneity in the return to groundwater, we add to equation (20) an interaction between groundwater extraction and a measure of groundwater’s importance in each aquifer. This regression estimates ϕ_i for each aquifer i :


$$\log y_{it} = \alpha \log k_{it} + \gamma \log l_{it} + \phi \log w_{it} + \psi \log w_{it} \times \nu_i + X'_{it}\Gamma + \varphi_i + \zeta_t + \varepsilon_{it} \quad (26)$$

where ν_i measures differences in the value of groundwater across aquifers and the aquifer-specific water input share is $\phi_i = \phi + \psi \times \nu_i$. We estimate it with the same instruments as equation (20). We also estimate it without the interaction, excluding low-agriculture aquifers (those with less than 1% cropland). Appendix Table A.3 reports the results. We measure the importance of groundwater by the aquifer’s cropland share (column (1)), irrigated share (column (2)), and indicators for any agriculture (column (3)), any irrigation (column (4)), or agriculture as the dominant land use (column (5)) above the aquifer. Columns (6) and (7) report OLS and IV results on the high-agriculture data subsample.


⁴⁵Most aquifers have some level of recharge, and only 6% of aquifers with data have annual recharge below 0.1% of aquifer capacity.

Figure A.6: Additional Aquifer Characteristics, by Aquifer


(a) Spatial Patterns of Ramsey Discount Factor




(b) Distribution of Ramsey Discount Factor




(c) Spatial Patterns of Mean Lift Height




(d) Distribution of Mean Lift Height



(e) Spatial Patterns of Electricity Prices



(f) Distribution of Electricity Prices



Notes: Panels A and B show Ramsey discount factors from [Addicott et al. \(2020\)](#). Panels C and D show mean lift height in meters from [Fan et al. \(2013\)](#). Panels E and F show electricity prices in \$/kWh from the International Energy Agency. The means are aquifer-level averages over all years between 2003 and 2016.


Table A.2: First Stage for Instrumental Variables Regressions

	(1) log(Groundwater)	(2) log(Capital)	(3) log(Labor)
Average Temperature, Lag 2	-0.008 (0.006)	-0.001 (0.001)	-0.002** (0.001)
Average Temperature, Lag 2 $\times$ Porosity	0.001 (0.042)	-0.017* (0.009)	-0.010* (0.005)
Average Temperature, Lag 3 $\times$ Porosity	-0.043* (0.024)	-0.028*** (0.005)	-0.006* (0.003)
asinh(Average Temperature, Lag 2)	-0.052*** (0.019)	-0.049*** (0.004)	-0.020*** (0.002)
log(Total Precipitation, Lag 2) $\times$ Porosity	-0.683*** (0.139)	0.276*** (0.046)	0.030 (0.018)
asinh(Average Temperature, Lag 3)	0.062*** (0.016)	-0.035*** (0.004)	-0.030*** (0.002)
log(Wage, Lag 1)	-0.002 (0.016)	-0.001 (0.007)	-0.069*** (0.006)
log(Wage, Lag 2)	0.048*** (0.017)	0.011*** (0.004)	-0.028*** (0.003)
log(Interest Rate, Lag 2)	0.094*** (0.014)	0.045*** (0.002)	-0.031*** (0.001)
log(Interest Rate, Lag 1)	-0.073*** (0.012)	0.096*** (0.004)	-0.032*** (0.002)
log(Wage, Lag 3)	-0.042*** (0.014)	0.025*** (0.005)	-0.014*** (0.004)
log(Interest Rate, Lag 3)	0.021* (0.012)	-0.001 (0.004)	0.004** (0.002)
Aquifer FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Weather controls	Yes	Yes	Yes
F statistic	11.664	116.508	256.495
Observations	29,854	29,854	29,854


Notes: This table shows the first-stage regressions for Table 2. The outcome variables are the log of groundwater extraction, capital, and labor by aquifer $\times$ year, respectively. Instruments and controls come from the two-step model selection process (see text). Log(Groundwater) is gross groundwater extraction in m³. Log(Capital) is the aquifer-specific value of capital in million 2023 U.S. dollars. Log(Labor) is aquifer population employed in agriculture. F statistic is the first-stage F statistic for each endogenous variable. Standard errors clustered by aquifer are in parentheses. Asterisks denote p-value less than 0.01 (***), 0.05 (**), 0.10 (*).

Figure A.7: Net Extraction and Recharge, by Aquifer


(a) Spatial Patterns of Net Annual Extraction




(b) Distribution of Net Annual Extraction



(c) Spatial Patterns of Annual Recharge



(d) Distribution of Annual Recharge



Notes: Panels A and B show aquifer-level net groundwater extraction as a share of aquifer capacity. Net extraction equals gross extraction minus recharge. Panels C and D show recharge as a share of aquifer capacity. Data are from GRACE and [de Graaf et al. \(2015\)](#). The means are aquifer-level averages over all years between 2003 and 2016.

Groundwater is more valuable where agriculture is more prevalent (Appendix Table A.3). For example, aquifers at the 90th percentile of the distributions of cropland and irrigation have elasticities of 0.41 and 0.19 (columns (1) and (2)), above the mean of 0.10 across all aquifers (Table 2). The high-agriculture sample yields an IV estimate of 0.274 (column (7)), again above the mean.

Appendix Table A.7 shows that the input shares are generally robust to different GDP outcomes, water definitions (including surface water in w_{it}), data subsamples, night lights satellites, and labor measures. Relative to our main estimate of $\phi = 0.10$, the groundwater input share is 0.12 with all water inputs (column (1)), 0.17 with agricultural GDP as the outcome (column (2)), 0.15 and 0.10 with the newer and older night lights satellites (columns (5) and (6)), 0.10 with log rather than inverse hyperbolic sine instruments (column (7)), 0.07 with gridded population as labor (column (8)), and 0.17 without temperature and

Table A.3: Heterogeneous Effects of Groundwater Extraction on GDP

	(1) Cropped %	(2) Irr %	(3) Any Crops	(4) Any Irr	(5) Dominant Ag	(6) High Ag, OLS	(7) High Ag, IV all
log(Groundwater) (ϕ_0)	-0.210*** (0.080)	-0.004 (0.066)	-0.204* (0.115)	-0.151 (0.098)	-0.138** (0.070)	0.014*** (0.003)	0.274*** (0.063)
log(Groundwater) $\times$ Value (ϕ_i)	1.743*** (0.447)	3.266 (2.011)	0.456*** (0.177)	0.451*** (0.171)	1.214*** (0.346)		
1 - log(Capital) - log(Labor) (θ)	0.439*** (0.153)	0.487*** (0.160)	0.405** (0.161)	0.366** (0.160)	0.456*** (0.155)	0.702*** (0.058)	-0.345 (0.323)
Implied log(Land) ($\theta - \phi$)	0.460*** (0.164)	0.544*** (0.163)	0.210 (0.224)	0.223 (0.202)	0.393** (0.164)	0.688*** (0.058)	-0.619* (0.356)
Aquifer FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Weather controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
F statistic	4.287	2.389	4.505	3.814	3.099	—	5.633
Value Mean	0.108	0.019	0.731	0.533	0.062	—	—
Observations	29,854	29,777	29,854	29,854	29,854	17,782	17,782

Notes: All models are estimated using instrumental variables, except column (6), which is estimated using OLS. Log(Groundwater), its interaction, log(Labor), and log(Capital) are instrumented in all models. Instruments and controls come from the two-step model selection process (Section 5.2). Log(Groundwater) is gross groundwater extraction in m^3 . Value, the variable in the column heading, measures the aquifer-specific value of groundwater and so gives heterogeneous water input shares (ϕ_i): the cropped share of the aquifer (column (1)), its irrigated share (column (2)), and indicators for any cropland (column (3)), any irrigation (column (4)), or agriculture as the dominant land use (column (5)) above the aquifer. Columns (6)–(7) exclude low-agriculture aquifers. Log(Capital) is the aquifer-specific value of capital in million 2023 U.S. dollars. Log(Labor) is aquifer population employed in agriculture. Implied log(Land) is the implied value of the land share under the assumption of constant returns to scale, evaluated at the mean of the value variable. The outcome variable is log(Night Lights Luminosity) by aquifer $\times$ year. F statistic is the Kleibergen–Paap rk Wald F statistic. Standard errors clustered by aquifer are in parentheses. Asterisks denote p-value less than 0.01 (***), 0.05 (**), 0.10 (*).

precipitation controls (column (9)). Elasticities are smaller, about 0.02, with gridded GDP, available for fewer years, as the outcome (column (3)) and with only the deepest aquifers (column (4)), though we cannot reject that all of these equal our main estimates. Estimates of θ are often close to our main estimates in sign, magnitude, and precision but are somewhat less stable across specifications than the estimates of ϕ .

Several features of these alternatives help explain the estimates. Column (1)’s larger input share suggests slightly higher returns to surface water than to groundwater in irrigation. Column (2) replaces night lights with country-level agricultural GDP and column (3) with

model-based gridded GDP adjusted for purchasing power (Kummu et al. 2018; Prabhakar et al. 2023), both re-weighted to the aquifer level; groundwater likely matters more for agricultural output than for total output, and the gridded GDP estimate may also reflect the purchasing power adjustment and fewer years of data. Column (4) excludes aquifers that hydrological agencies classify as shallow-water, which are more likely to be close to the surface, under thicker bedrock, and less likely to approximate the “bathtub” shape in Figure 1 (BGR 2025). Columns (5) and (6) split the sample by the satellite that measured night lights, 2013-2016 (the “VIIRS” years) and 2003-2012 (the “DMSP” years); the data are harmonized across the two, and the four VIIRS years give noisier estimates. Column (7)’s log instruments are robust to units but drop about 200 aquifers with zero or negative average temperatures, which the inverse hyperbolic sine keeps. Column (8) interpolates gridded population between five-year observations, and with this noisier labor measure no parameter is statistically significant. Column (9) shows that climate controls do not drive the main result, though without them the labor, capital, and land shares are less precise.

F Additional Discussion of Results

F.1 Model Validation

The model has good out-of-sample fit: estimated on data from 2003-2009, it predicts extraction in 2010-2016 with a correlation above 0.95 against remotely sensed actual extraction (Figure A.8, Panel A). Most points in the binned scatterplot lie near the 45-degree line, so the predictions are essentially unbiased as well as accurate.

A second, somewhat weaker test asks whether the model recovers similar results for an aquifer when fit separately to each half of the sample. It does (Figure A.8, Panels B through D). Across aquifers, the correlation between the two periods’ estimates exceeds 0.95 for decision discount factors and welfare losses, and is 0.77 for long-run fill.


These tests have an important limitation: available remote sensing data permit validation only seven years ahead, while our counterfactuals extend much further, so strong fit here does not guarantee accuracy over the horizons that matter for welfare. Still, a misspecified model would likely produce biased or inaccurate seven-year forecasts, and this model does not.

F.2 Long-Term Management of the Ogallala Aquifer


Under its decision discount factor, the Ogallala Aquifer has relatively rapid gross extraction over the century (Figure A.9, Panel A). Extraction is slower under the benchmark Ramsey

Figure A.8: Comparison of 2003-2009 Versus 2010-2016 Model Predictions and Values


(a) Annual Net Extraction in 2010-2016 (% of Capacity)




(b) Decision Discount Factor



(c) Long-Run Fill (% of Capacity)




(d) Percentage Loss in Present Value Due to Non-Optimal Extraction, Ramsey Benchmark Discount Factor



Notes: In Panel A, the x-axis shows annual net extraction in 2010-2016, as a percent of aquifer capacity, predicted by a model fit to 2003-2009, and the y-axis shows remotely sensed actual net extraction in 2010-2016. In Panels B-D, the axes show estimates from models fit to 2003-2009 (x-axis) and 2010-2016 (y-axis). Each graph shows a binned scatterplot and the 45-degree line.

Figure A.9: Resource Management Paths by Discount Factor, Ogallala Aquifer



Notes: This figure shows predicted gross extraction, fill, extraction costs, and annual profit for the Ogallala Aquifer under different discount factors; the solid blue line is the current extraction path.


and market discount factors. As a result, the benchmark paths steadily raise fill and lower extraction costs, increasing long-run profits (Panels B-D). Unlike the Central Valley, where extraction in all three scenarios reaches steady state within a few decades, the Ogallala is still converging at the end of the century.

F.3 Model Sensitivity Analyses

Appendix Table A.5 reports alternative model-based estimates. Column (1) repeats the main estimates, and columns (2)-(7) use different parameters.

Columns (2)-(4) use different production function parameters. Column (2) calibrates the

Figure A.10: Long-Run Fill, by Discount Factor and Counterfactual



Notes: This figure shows histograms of steady state aquifer fill, as a percent of capacity, under the decision discount factor (i.e., the current extraction path; Panel A) and under counterfactuals in which the planner uses the Ramsey (Panel B) or market (Panel C) discount factor or is myopic (Panel D), users keep gross extraction constant unless the aquifer is drained (Panel E), or users consume recharge every year (Panel F).

Table A.4: Model Outputs: Discount Factor and Intermediate Outcomes, Medians

	Crop Weighted (1)	GDP Weighted (2)	Population Weighted (3)
<i>Panel A: Discount Factor</i>			
Estimated	0.840	0.943	0.881
Estimated - Ramsey discount	-0.107	0.003	-0.057
<i>Panel B: Baseline Net Water Use (% of aquifer capacity)</i>			
Observed	0.00%	0.03%	0.08%
Ramsey discounting	-0.98%	0.07%	-0.50%
Market discounting	-1.37%	-0.40%	-1.16%
<i>Panel C: Long-Run Fill (% of aquifer capacity)</i>			
Observed extraction path	92.49%	93.50%	92.49%
Ramsey discounting	100.00%	92.19%	100.00%
Market discounting	100.00%	100.00%	100.00%
<i>Panel D: Long-Run Marginal Extraction Cost (\$/m³)</i>			
Observed extraction path	0.021	0.018	0.028
Ramsey discounting	0.000	0.028	0.000
Market discounting	0.000	0.000	0.000

Notes: Columns (1)-(3) report medians across global aquifers weighted by crop value, GDP, and population. Results use the decision discount factor, the Ramsey benchmark (from [Addicott et al. 2020](#)), and the market benchmark (0.98). Panels B and C are in percentage points (e.g., 0.33 is a third of a percentage point). All estimates are based on the model in Section 3. All dollar values are in 2023 U.S. dollars.

land share from [Ryan and Sudarshan \(2022\)](#). With our estimated groundwater input share ϕ , this pins down the resource share θ without estimating the capital and labor shares in equation (20). Welfare losses from the current extraction path are similar with our estimated or these calibrated production function parameters. In this alternative calibration, the welfare loss in dollars rises by 1-2%, but as a percent of aquifer value it falls, because global aquifers have higher value. The share of aquifers over-extracted falls by 2-4 percentage points.

Column (3) estimates equation (20) with all water (surface and groundwater) rather than groundwater only, which assumes the two are perfect substitutes in production. The water input share ϕ rises slightly (0.121 versus 0.103) and the resource share θ falls slightly (0.479 versus 0.514; Table A.7). Including surface water slightly increases the dollar welfare loss from non-optimal extraction; the percentage loss is higher because surface water substitutes lower the value of groundwater.

Table A.5: Consequences of Non-Optimal Discounting and Other Counterfactuals, Sensitivity Analyses

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
<i>Panel A: Ramsey benchmark discount factor</i>							
Welfare loss from:							
Observed extraction path (%)	2.55%	2.10%	3.39%	4.71%	2.29%	2.54%	2.63%
Observed extraction path (trillion \$)	2.57	2.63	2.61	2.83	2.50	2.61	2.87
Myopia (%)	9.16%	7.48%	12.02%	16.88%	9.13%	9.19%	9.34%
Constant water (%)	3.40%	2.71%	4.56%	6.55%	3.05%	3.39%	3.57%
Zero net extraction (%)	2.87%	2.38%	3.79%	5.34%	2.61%	2.86%	2.97%
Share of aquifers over-extracted:							
Unweighted	60.85%	56.45%	62.65%	56.43%	62.41%	61.27%	63.76%
Capacity weighted	71.68%	67.63%	72.01%	68.04%	70.39%	72.06%	72.76%
Population weighted	77.39%	64.52%	78.90%	64.51%	77.82%	77.90%	81.40%
GDP weighted	48.98%	46.17%	50.24%	46.17%	50.57%	49.75%	50.72%
Global value of all aquifers (trillion \$)	103.26	143.59	74.72	49.52	105.17	104.54	102.20
<i>Panel B: Market benchmark discount factor</i>							
Welfare loss from:							
Observed extraction path (%)	2.90%	2.24%	3.92%	4.95%	2.83%	2.91%	3.12%
Observed extraction path (trillion \$)	8.36	8.41	8.37	8.51	8.93	8.47	9.29
Myopia (%)	12.88%	10.67%	16.68%	22.89%	12.90%	12.91%	13.05%
Constant water (%)	5.44%	4.09%	7.37%	10.37%	5.34%	5.46%	5.88%
Zero net extraction (%)	2.91%	2.25%	3.94%	5.01%	2.83%	2.91%	3.15%
Share of aquifers over-extracted:							
Unweighted	81.52%	80.03%	82.87%	79.80%	84.18%	82.58%	83.51%
Capacity weighted	89.00%	87.90%	88.98%	87.32%	89.17%	89.40%	88.60%
Population weighted	95.99%	94.60%	96.38%	94.29%	96.35%	95.69%	96.65%
GDP weighted	70.35%	69.73%	70.21%	69.72%	90.75%	70.22%	71.01%
Global value of all aquifers (trillion \$)	244.91	340.39	177.05	115.61	248.30	248.32	242.44
Baseline	Yes						
Calibrated land share		Yes					
Surface water added			Yes				
Heterogeneous water returns				Yes			
Heterogeneous fuel sources					Yes		
Variable recharge						Yes	
Drilling costs added							Yes
Observations	1,369	1,387	1,360	1,391	1,410	1,366	1,352

Notes: This table displays model results for all aquifers. Column (1) repeats the main estimates. Column (2) calibrates the land share from [Ryan and Sudarshan \(2022\)](#) rather than estimating the capital and labor shares in equation (20). Column (3) estimates equation (20) with all water (surface and groundwater) rather than groundwater only (Table A.7). Column (4) lets returns to water vary across aquifers (Table A.3). Column (5) lets marginal extraction cost vary with pump fuel source (electricity versus diesel) and pump efficiency. Column (6) uses recharge estimates that vary by year. Column (7) includes drilling costs. The Ramsey discount factors are based on [Addicott et al. \(2020\)](#), re-weighted to the aquifer level. The market discount factor is based on a market discount rate of 2%. All estimates are based on the model in Section 3 and are unweighted unless otherwise noted. All dollar values are in 2023 U.S. dollars.

Column (4) lets the water input share vary across aquifers, using the parameters in Table A.3. As in column (3), the dollar welfare loss barely changes, again suggesting that aquifers are similarly far from the optimal path. The welfare loss as a percent of aquifer value increases because the global value of aquifers falls.

Column (5) lets marginal extraction costs vary across aquifers with pump fuel source (electricity versus diesel) and pump efficiency. Welfare losses and global aquifer values are all within 10% of the main estimates, for both Ramsey and market discounting. Most of the over-extracted aquifer shares are similar or higher.

Column (6) calibrates recharge to reflect variation in precipitation in the years of GRACE extraction data. Since unusually high or low rainfall in 2003-2016 would make recharge calibrated to the long-run average understate or overstate extraction, this column scales recharge, and so extraction, by the percentage deviation of rainfall from its average over the 100 prior years. Results are again quantitatively similar.

Column (7) adds annualized well drilling costs to marginal extraction costs. The U.S. Department of Agriculture (USDA) reports that annual drilling expenditure is about 12% of annual pumping costs (USDA 2024); raising extraction costs by 12% yields quantitatively similar results.

Appendix Table A.6 reports further sensitivity analyses using different subsets of aquifers. Column (2) uses only major aquifer systems. Column (3) includes aquifers with decision discount factors outside 0 to 1. Column (4) adds aquifers with less than 1% agricultural land. Column (5) combines the samples from columns (3) and (4). In each, the share of aquifers with inefficiently rapid extraction is similar to the main estimates, at 61% to 67% for Ramsey discounting and 79% to 85% for market discounting, and non-optimal extraction generates a welfare cost of several trillion dollars.

F.4 Policy Functions

Appendix Figure A.11 shows optimal water consumption as a function of aquifer fill (the “policy function”) for each of the four case study aquifer systems. Each panel shows four discount factors: “myopic” ($\beta = 0$), the decision discount factor, the Ramsey discount factor, and the “market” discount factor ($\beta = 0.98$). The policy function shifts down as β increases: a more patient agent uses less water for a given aquifer fill. The dashed horizontal line is recharge; where a policy function crosses it, extraction equals recharge, which defines long-run fill. More patient agents reach higher long-run fill.

Table A.6: Implications of Observed and Counterfactual Consumption Paths for Social Welfare, Sensitivity Analyses


	(1)	(2)	(3)	(4)	(5)
<i>Panel A: Ramsey benchmark discount factor</i>					
Welfare loss from:					
Observed extraction path (%)	2.55%	2.30%	3.01%	2.66%	3.40%
Observed extraction path (trillion \$)	2.57	1.65	3.02	2.88	3.43
Myopia (%)	9.16%	9.99%	8.69%	9.35%	8.77%
Constant water (%)	3.40%	3.57%	7.82%	3.58%	39.14%
Zero net extraction (%)	2.87%	2.84%	8.45%	2.97%	21.56%
Share of aquifers over-extracted:					
Unweighted	60.85%	66.67%	61.01%	62.57%	64.97%
Capacity weighted	71.68%	79.49%	71.25%	70.28%	72.65%
Population weighted	77.39%	82.25%	77.43%	77.26%	77.37%
GDP weighted	48.98%	64.59%	49.14%	51.14%	49.52%
Global value of all aquifers (trillion \$)	103.26	68.99	107.91	115.55	121.25
<i>Panel B: Market benchmark discount factor</i>					
Welfare loss from:					
Observed extraction path (%)	2.90%	2.89%	3.34%	3.07%	3.88%
Observed extraction path (trillion \$)	8.36	5.31	9.47	9.02	10.40
Myopia (%)	12.88%	13.28%	12.28%	12.93%	12.11%
Constant water (%)	5.44%	6.22%	10.45%	5.74%	52.94%
Zero net extraction (%)	2.91%	3.15%	8.20%	3.09%	20.44%
Share of aquifers over-extracted:					
Unweighted	81.52%	84.93%	79.01%	82.06%	80.64%
Capacity weighted	89.00%	92.25%	85.65%	88.37%	86.42%
Population weighted	95.99%	97.26%	95.20%	95.94%	95.15%
GDP weighted	70.35%	83.81%	67.91%	71.79%	67.15%
Global value of all aquifers (trillion \$)	244.91	160.82	255.51	273.02	285.93
Baseline	Yes				
Major aquifers	Yes				
Extremal discount factors included	Yes				
Low agriculture aquifers	Yes				
Observations	1,369	219	1,572	1,611	2,004

Notes: This table displays model results for all aquifers and for different aquifer samples. Column (1) repeats the main estimates. Column (2) uses only major aquifers. Column (3) adds aquifers with discount factors outside 0 to 1, column (4) adds low-agriculture aquifers, and column (5) adds both. The Ramsey discount factors are based on [Addicott et al. \(2020\)](#), re-weighted to the aquifer level. The market discount factor is based on a market discount rate of 2%. All estimates are based on the model in Section 3 and are unweighted unless otherwise noted. All dollar values are in 2023 U.S. dollars.

References for the Online Appendix

- Beaudoing, Hiroko and M Rodell, “GLDAS Noah Land Surface Model L4 monthly 1.0 x 1.0 degree,” 2016. <https://doi.org/10.5067/LWTYSMP3VM5Z>, accessed 3/4/2019.
- Bruno, Ellen M, Nick Hagerty, and Arthur R Wardle, “The political economy of groundwater management: Descriptive evidence from California,” in “American agriculture, water resources, and climate change,” University of Chicago Press, 2022.
- Chen, Xi and William Nordhaus, “Using Luminosity Data as a Proxy for Economic Statistics,” *Proceedings of the National Academy of Sciences*, 2011, 108 (21), 8589–8594.
- Feng, Wei, Min Zhong, Jean-Michel Lemoine, Richard Biancale, Hou-Tse Hsu, and Jun Xia, “Evaluation of groundwater depletion in North China using the Gravity Recovery and Climate Experiment (GRACE) data and ground-based measurements,” *Water Resources Research*, 2013, 49 (4), 2110–2118.
- Getirana, Augusto, Sujay Kumar, Manuela Girotto, and Matthew Rodell, “Rivers and floodplains as key components of global terrestrial water storage variability,” *Geophysical Research Letters*, 2017, 44 (20), 10,359–10,368.
- Gibson, John, Susan Olivia, and Geua Boe-Gibson, “Night Lights in Economics: Sources and Uses,” *Journal of Economic Surveys*, 2020, 34 (5), 955–980.
- Hasan, M. F., R. Smith, S. Vajedian, R. Pommerenke, and S. Majumdar, “Global land subsidence mapping reveals widespread loss of aquifer storage capacity,” *Nature Communications*, 2023, 14, 6180.
- Henderson, J. Vernon, Adam Storeygard, and David N. Weil, “Measuring Economic Growth from Outer Space,” *American Economic Review*, 2012, 102 (2), 994–1028.
- INSCR, “Polity5: Regime Authority Characteristics and Transitions Datasets,” 2022. <https://www.systemicpeace.org/inscrdata.html>, accessed 10/19/2022.
- Klein Goldewijk, C. G. M., “Anthropogenic land-use estimates for the Holocene; HYDE 3.2,” Sept 2017. <https://archaeology.datastations.nl/dataset.xhtml?persistentId=doi:10.17026/dans-25g-gez3>, accessed 2/25/2022.
- Perez-Sindin, Xaquín S., Tzu-Hsin Karen Chen, and Alexander V. Prishchepov, “Are night-time lights a good proxy of economic activity in rural areas in middle and low-income countries? Examining the empirical evidence from Colombia,” *Remote Sensing Applications: Society and Environment*, 2021, 24, 100647.
- Rhodes, Edward C., Humberto L. Perotto-Baldivieso, Evan P. Tanner, Jay P. Angerer, and William E. Fox, “The Declining Ogallala Aquifer and the Future Role of Rangeland Science on the North American High Plains,” *Rangeland Ecology and Management*, 2023, 87, 83–96.
- Sarah, S., S. Ahmed, S. Violette, and G. de Marsily, “Groundwater sustainability challenges revealed by quantification of contaminated groundwater volume and aquifer depletion in hard rock aquifer systems,” *Journal of Hydrology*, 2021, 597, 126286.
- Siebert, Stefan, Matti Kummu, Miina Porkka, Petra Döll, Navin Ramankutty, and Bridget R. Scanlon, “Historical Irrigation Dataset (HID),” Mar 2015. <https://mygeohub.org/publications/8/2>, accessed 11/8/2021.
- USDA, “2023 Irrigation and Water Management Survey,” Technical Report, United States Department of Agriculture 2024.

Figure A.11: Policy Functions



Notes: This figure shows optimal water consumption as a function of aquifer fill (the "policy function") for the four case study aquifer systems. Each figure shows the policy function for four discount factors: "myopic" ($\beta = 0$), the decision discount factor, the Ramsey discount factor, and the "market" discount factor ($\beta = 0.98$). The dashed horizontal line is the aquifer's recharge.

Table A.7: Alternative Production Function Parameter Estimates

	(1) Total Water	(2) Ag GDP	(3) Grid GDP	(4) Deep Aq	(5) VIIRS	(6) DMSP	(7) log IVs	(8) log(Pop)	(9) No Xs
log(All Water) (ϕ)	0.121** (0.049)	— —	— —	— —	— —	— —	— —	— —	— —
log(Groundwater) (ϕ)	— —	0.169*** (0.024)	0.024 (0.063)	0.024 (0.044)	0.151* (0.088)	0.098*** (0.018)	0.100** (0.042)	0.073 (0.045)	0.170*** (0.050)
1 - log(Capital) - log(Labor) (θ)	0.479*** (0.151)	0.551*** (0.134)	0.440* (0.239)	0.645*** (0.140)	2.410*** (0.360)	0.612*** (0.079)	0.592*** (0.162)	0.138 (0.173)	0.072 (0.170)
Implied log(Land) ($\theta - \phi$)	0.357** (0.173)	0.383*** (0.134)	0.416* (0.242)	0.620*** (0.162)	2.259*** (0.350)	0.513*** (0.084)	0.492*** (0.178)	0.066 (0.191)	-0.098 (0.190)
Aquifer FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Weather controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No
F statistic	9.138	10.401	9.103	9.267	4.237	7.293	10.830	12.741	9.974
Observations	29,854	29,854	27,248	23,244	10,717	19,137	29,772	29,854	29,854

Notes: All models are estimated using instrumental variables. Log(Groundwater), log(Labor), and log(Capital) are instrumented in all models; the instruments and precipitation and temperature controls are chosen as part of the two-step model selection process (see Section 5.2). Log(Groundwater) is groundwater extraction in m³. Log(Capital) is the aquifer-specific value of capital in million 2023 U.S. dollars. Log(Labor) is aquifer population employed in agriculture. Implied log(Land) is the implied value of the land share under the assumption of constant returns to scale. Column (1) combines groundwater and surface water extraction. Column (2) uses agricultural GDP (Ag GDP) (re-weighted to the aquifer level from the country level) as an outcome variable. Column (3) uses gridded purchasing power parity GDP as an outcome variable. Column (4) excludes shallow-water aquifers from the analysis sample. Columns (5) and (6) split the sample into the two time periods where two different satellite operations collected night lights data: column (5) includes 2013-2016 (i.e., the “Visible Infrared Imaging Radiometer Suite (VIIRS)” years) and column (6) includes 2003-2012 (i.e., the “Defense Meteorological Satellite Program (DMSP)” years). Column (7) uses logs of the IVs instead of inverse hyperbolic sine and excludes aquifers with negative average temperatures. Column (8) uses the log of gridded population instead of agricultural labor. Column (9) excludes precipitation and temperature controls. The outcome variable is log(Night Lights Luminosity) by aquifer-year, except in columns (2) and (3). F statistic is the first-stage Kleibergen–Paap rk Wald F statistic. Standard errors clustered by aquifer are in parentheses. Asterisks denote p-value less than 0.01 (***), 0.05 (**), 0.10 (*).

Table A.8: Correlations Between the Decision Discount Factor (β_i^d) and Aquifer Characteristics

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)
Model Inputs													
Log(Groundwater Returns)	0.522*** (0.156)	—	—	—	—	—	—	—	—	—	—	—	0.745*** (0.153)
Log(Marginal Cost)	—	-3.062*** (0.447)	—	—	—	—	—	—	—	—	—	—	-2.991*** (0.564)
Log(Elec. Price)	—	—	-1.862*** (0.327)	—	—	—	—	—	—	—	—	—	-0.913*** (0.243)
Log(Lift Height)	—	—	—	-0.027*** (0.008)	—	—	—	—	—	—	—	—	0.003 (0.008)
Ramsey Discount Factor and Components													
Ramsey Discount Factor	—	—	—	—	0.122 (0.948)	—	—	—	—	—	—	—	0.751 (0.517)
Pure Rate of Time Pref.	—	—	—	—	—	0.019 (0.044)	—	—	—	—	—	—	0.061*** (0.021)
Growth Rate, GDP/Pop	—	—	—	—	—	—	-0.001 (0.004)	—	—	—	—	—	0.000 (0.002)
Other Country Characteristics													
Formal water market exists	—	—	—	—	—	—	—	-0.020 (0.028)	—	—	—	—	0.055** (0.027)
Informal water market exists	—	—	—	—	—	—	—	—	0.046 (0.046)	—	—	—	0.023 (0.030)
International Aquifer	—	—	—	—	—	—	—	—	—	0.016 (0.020)	—	—	0.024* (0.015)
Log(GDP/Pop)	—	—	—	—	—	—	—	—	—	—	-0.010*** (0.003)	—	-0.010*** (0.002)
Log(Democracy)	—	—	—	—	—	—	—	—	—	—	—	-0.128*** (0.046)	-0.058** (0.028)
Continent FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,369	1,369	1,369	1,369	1,369	1,369	1,369	1,369	1,369	1,369	1,369	1,369	1,369

Notes: The outcome variable is the aquifer-level estimate of the decision discount factor (β_i^d). The Ramsey discount factor and the pure rate of time preference are from [Addicott et al. \(2020\)](#). Country characteristics are discussed in Appendix D.5. All models are estimated using OLS. Marginal returns and marginal cost are in \$/m³, electricity prices are in \$/kWh, and lift height is in m. Robust standard errors are in parentheses. Asterisks denote p-value less than 0.01 (***), 0.05 (**), 0.10 (*).