



Federal Reserve
Bank of Dallas

Are We Consuming Too Much Groundwater?

Brian Greaney, Joseph S. Shapiro and Katherine R. H. Wagner

Working Paper 2626

Research Department

<https://doi.org/10.24149/wp2626>

August 2026

Working papers from the Federal Reserve Bank of Dallas are preliminary drafts circulated for professional comment. The views in this paper are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Dallas or the Federal Reserve System. Any errors or omissions are the responsibility of the authors.

Are We Consuming Too Much Groundwater?*

Brian Greaney[†], Joseph S. Shapiro[‡] and Katherine R. H. Wagner[‡]

July 15, 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, Kendra Marcoux, Zhehao (Colin) Lu, and Brant Walker for excellent research assistance, 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.

[†]Brian Greaney, Federal Reserve Bank of Dallas, brian.greaney@dal.frb.org.

[‡]Joseph S. Shapiro, UC Berkeley and NBER, joseph.shapiro@berkeley.edu.

[‡]Katherine R. H. Wagner, University of British Columbia, 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 today:

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 in the important setting of groundwater extraction. Groundwater 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 ([O’Neill et al. 2023](#)).

Priors vary on the efficiency of prevailing extraction. Because formal groundwater institutions are uncommon and weak, many observers have the prior that most aquifers have a common pool with many atomistic agents effectively acting myopically, that current groundwater levels are below the social optimum, and that current extraction exceeds the optimum ([Stavins 2011](#)). We are unaware of systematic tests of this prior. The existence of informal institutions that use many approaches to achieve non-myopic extraction may temper this prior ([Ostrom 1990](#); [Meinzen-Dick and Bruns 2024](#)).¹

We begin with a classic dynamic model of renewable resource extraction. Groundwater extraction produces output using a technology with decreasing marginal product of water. Water extraction is chosen to maximize the present discounted value of all future profits. Water extraction faces a tradeoff: consuming water today increases current profits at the expense of higher future extraction costs, since current extraction lowers the future water table.

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 marginal return

¹These 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](#)).

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 characterizes the impatience with which the representative agent extracts water from the aquifer, so it provides a useful 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 of decision and benchmark discount factors helps assess the optimality of current extraction. If institutional or market failures make the representative agent incompletely internalize effects of water extraction today on future extraction costs, the decision discount factor that the representative agent uses in choosing extraction may not equal the benchmark discount factor that appears in the representative agent’s utility function. We find that 61% of the world’s aquifers have decision discount factors below a Ramsey benchmark and 79% 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, actual extraction in most of the world’s aquifers instead has many individual farms, and even within an aquifer’s area, different regions may have heterogeneous local institutions. We essentially ask what (potentially non-optimal) extraction decisions by a representative agent would lead to the same outcomes that we observe from actual decentralized management. We also find 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 important 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. For example, we estimate that California’s Central Valley Aquifer System has a decision discount factor (i.e., revealed preference from observed extraction decisions) of 0.89. We can also rationalize observed water extraction by assuming that a continuum of atomistic water users use the Ramsey benchmark discount factor but act as if 41% of future extraction costs depend on others’ common pool extraction.

Our second set of model-based calculations compares the current groundwater extraction path

against the counterfactual path that would result from extraction under the benchmark discount factor. We find that the median aquifer is in a steady state, but not the optimal one. In the median aquifer, we find that actual current annual net extraction rates exceed the optimal rate by 0.1% to 0.5% of aquifer capacity; this is a large change from the baseline, where the median aquifer has 0.0% net extraction. Although zero net extraction characterizes the median aquifer, most individual aquifers have non-zero net extraction. More importantly, the current fill level of most aquifers is inefficiently low: although the median aquifer has zero net extraction (i.e., gross extraction equal to recharge), its extraction is nonetheless above the optimal level because its stock is below the optimal level. A transition path with diminished short- to medium-run extraction would obtain a long-term steady state at higher fill and greater aquifer present value.

Third, we study the difference in the present value of social welfare between the current extraction path and the extraction path reflecting the benchmark discount factor. We interpret the difference as the cost of non-optimal resource management, potentially due to institutional or market failures.

This welfare analysis shows that global aquifers are an enormously valuable asset: we estimate a present value for all global aquifers of \$103 trillion or more, comparable in magnitude to the value of all publicly traded firms on the planet, at \$117 trillion (OECD 2025), or global GDP, at \$111 trillion (World Bank 2023). Though we interpret this far-out-of-sample statistic cautiously, since it compares against an extreme scenario of completely extracting every aquifer, its order of magnitude illustrates the potentially large value of natural capital in this setting. The global welfare loss due to non-optimal extraction represents 2.6% of aquifer present value. We discuss reasons for the modest magnitude of this percentage welfare loss, including 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 the consequences of several other non-optimal but potentially realistic counterfactual extraction paths—maintaining current gross groundwater extraction levels forever (‘constant extraction’), extracting groundwater myopically, or maintaining current aquifer fill levels indefinitely by consuming an amount equal to recharge in each year (‘sustainability’). Constant extraction lightly echoes the Prior Appropriation doctrine in western U.S. surface water law (DWR 2019). Constant extraction could also result from government subsidies that fix real water prices over time. Myopic extraction is equivalent to common pool atomistic agents using a discount factor of zero. Zero net extraction resembles management policies that use quantity restrictions, such as California’s Sustainable Groundwater Management Act.

All three counterfactuals are costly relative to the optimal extraction path, and the first two are costlier than observed extraction. Under constant extraction, steady state extraction completely

drains the stock of 30.8% of aquifers. Compared to extraction under the benchmark discount factor, the constant extraction counterfactual generates \$4 to \$13 trillion in present value welfare costs. Myopic extraction also imposes large welfare costs—in steady state, myopia completely drains 10.1% of aquifers, and long-run marginal extraction costs under myopic extraction exceed the marginal costs under current extraction more than three-fold. Myopia and constant extraction generate welfare costs of 3.4% to 12.9% of aquifer present value. Under zero net extraction (sustainability), no aquifers are fully drained. The welfare costs are smaller than in the other two counterfactuals, though welfare remains below that of the optimal path because the sustainable path does not let aquifers reach a steady state with higher fill.

As one test of model validation, we compare results using separate estimates from the first and second half of our sample period. This includes comparing predicted extraction in the years 2010-2016 from a model calibrated to the period 2003-2009 against actual remotely-sensed extraction in 2010-2016. Across aquifers, model-based predictions of extraction from the second half of this period have extremely high correlation with remotely-sensed actual extraction during the second half of this period. Discount factors, estimated long-run fill, and estimated long-run marginal extraction costs also have high correlation when estimated separately across the two periods.

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](#); [Fenichel et al. 2016](#); [Ayres et al. 2021](#); [Burlig et al. 2021](#); [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](#); [Montginoul et al. 2016](#); [Edwards and Guilfoos 2021](#); [Rodella et al. 2023](#); [Zhao et al. 2024](#)). We provide a systematic test.

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](#)) and empirical studies of how resource stocks respond to management regimes for forests, wetlands, fish, biodiversity, and land ([Burgess et al. 2012](#); [Costello et al. 2016](#); [Greenhill et al. 2024](#); [Frank and Sudarshan 2024](#); [Aronoff and Rafey 2026](#)). Our normative analysis of inter-temporal water optimization builds on [Carleton et al. \(2025\)](#)’s positive analysis of inter-regional water optimization. Their analysis studies how agricultural and trade policies affect observed levels of water extraction, while ours focuses on a classic single-agent autarky model of resource extraction paths. Our aquifer valuations can inform natural capital accounts ([Prabhakar et al. 2023](#)).

We also relate a new empirical measure of resource management to a classic theoretical litera-

ture on extraction and the commons (Gordon 1954; Coase 1960; Dasgupta and Heal 1979; Ostrom 1990; Provencher and Burt 1993; Karp 2017; Dasgupta 2021). 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.²

Our findings recall Gisser and Sanchez (1980)’s classic argument that the welfare gap between common-pool and optimal groundwater extraction is small, though ours includes depletion, exhaustion, and lift constraints, features that 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 counterfactuals like myopia or constant extraction produce losses several times larger than the current extraction path.

Our paper has important limitations. First, we follow a longstanding resource literature since Hotelling (1931) of studying single-agent dynamic resource optimization; the conclusion discusses the scope for work linking inter-temporal and inter-regional optimization. Second, in line with many hydrological models, we study a single-cell geophysical model of an aquifer with vertical sides and known depth. While we allow recharge, lift height, specific yield, capacity, and most other empirical and model inputs to differ across aquifers, which capture some important components of differences across aquifers that models of hydrodynamics within an aquifer and time period generate, we abstract from spatial forces within an aquifer such as lateral flow and local cones of depression, which are more complex to address globally. Third, we focus on the classic stock externality of resource extraction, in which depleting a resource in the current period increases marginal extraction costs in future periods. Groundwater aquifers can generate other externalities—groundwater can interact with surface water availability, for example. We focus on the stock externality since it has been central to policy and academic discussions of resource management for decades, has close ties to theories of dynamic optimization, and has sufficient data to permit quantification 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.

²Quaas et al. (2012) estimate the shadow interest rate at which 13 European fisheries borrow from natural capital stocks, a related concept to our estimate of the discount factor that rationalizes observed extraction.

2 Background on Water Resources

Water resources come in several types. Groundwater fills underground aquifers, which can form over millennia. Groundwater accounts for one component of total terrestrial water storage, which includes moisture in soils and vegetation, snow, and surface waters. Surface waters (e.g., rivers and lakes) lie above the Earth’s surface. Groundwater accounts for 20% of global freshwater. Agricultural irrigation accounts for 70% of global groundwater extraction (Ojha et al. 2018). In recent years, groundwater extraction has grown because it increasingly substitutes for surface water. 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 aquifers (BGR 2025). Aquifer systems combine nearby aquifers sharing geology and hydrology, where lateral water flow may exchange groundwater slowly across aquifers (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. Throughout, we briefly discuss the Central Valley Aquifer System for a running illustration; additional details, including for the Ogallala, Indus, and North China systems, appear in Appendix F.

Explaining aquifer terminology clarifies our analysis (Figure 1). Aquifer dynamics depend on inflows and outflows, also called recharge and 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 represents the current top of its water table. When gross extraction exceeds recharge, reserves fall and water levels decline. Declining reserves increase the marginal cost of water extraction because users must lift groundwater further, requiring more energy. Recharge depends on rainfall, surface water runoff, and geology. Rainfall recharges an aquifer more in porous rock; rainfall flows over non-porous rock as surface water runoff (e.g., streams). Water applied to agriculture can also recharge aquifers as it flows through soil and rock.

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.³ Specific yield varies within an aquifer by rock type and is below 100%, partly since molecular forces attract soil and rock particles to water. Lower depths have lower specific yield (Sutanudjaja et al. 2018). Much of our analysis considers groundwater availability, which equals surface area times saturated thickness

³We use the terms groundwater extraction, consumption, use, and depletion somewhat interchangeably.

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 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.⁴

Consider a simple example: a cubic aquifer 10 m on each side whose top lies 10 m below the ground. The aquifer occupies 1000 m³ of volume below the surface of the Earth; 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 to pump water to the surface. Diesel or electricity typically power groundwater pumps. 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. Many regions of India ration electricity as an indirect tool for groundwater management (Ryan and Sudarshan 2022).

3 Model

This section explains the model and our use of it to measure welfare and estimate decision discount factors. Appendix Table A.1 collects notation.

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 law of motion

$$a' = \min\{a + r_i - w, 1\} \tag{1}$$

⁴Water in topsoil is typically not included in groundwater, though is part of total terrestrial water storage.

Gross extraction w is limited to available water reserves plus recharge $a + r_i$ since the agent cannot consume water from an empty aquifer.

Agent i produces a numeraire agricultural good using the following technology:⁵

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

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 capital and labor prices are fixed, that the agent owns the land, and that current technology and preferences persist in future years.⁶ Sensitivity analyses include versions of w accounting for surface water and precipitation.⁷ 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

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

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

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

⁵Much 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).

⁶Sufficient conditions for fixed factor prices would be that each aquifer is a small open economy with frictionless intra-national factor mobility. The assumption of perfectly elastic labor supply is more plausible for the foreseeable smooth changes in water extraction along a balanced growth path that we study than for a transition path in response to an unanticipated shock. Elastic labor could also reflect surplus workers in agricultural areas, particularly in developing countries.

⁷We abstract from backstop technologies like desalination, which is rarely used for irrigation due to its considerable cost (Zolghadr-Asli et al. 2023; Nargi 2024).

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 , which plays a central role in our analysis, describes the representative agent's patience in making consumption decisions, given market or institutional failures like weak property rights. Our notation therefore emphasizes the value function's dependence on β_i , which we assume is time-invariant for each aquifer, though we do empirically test for its changes over time within our analysis period. The second term on the right-hand side of (9) represents the continuation value, which depends both on the agent's current choice of water extraction w and on recharge r_i . The agent faces a tradeoff—increasing water consumption may raise both current profits and future extraction costs, decreasing future profits, since current extraction raises the cost of future extraction. 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 $w < a + r_i$, and has an intuitive interpretation—the left-hand side represents the marginal revenue from extracting w units of water, while the right-hand side represents 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 $\beta_i \partial V_i(\cdot)/\partial a$. One could think of our quantitative approach later in the paper as using calibration and estimation to find values for all components of equation (10) besides the discount factor β_i , and then inferring the decision discount factor as the value which makes this first-order condition (10) 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 profits between the extraction path under the decision discount factor β_i^d and the extraction path under the benchmark discount factor β_i^* that the representative agent would use in the absence of institutional or market failures.⁸ 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

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

Later we estimate ω_i for each aquifer. 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 $\mathcal{V}_i(\mathbf{w}_i^*(a_{i0}; \beta_i^d), a_{i0}; \beta_i^*)$ and the aquifer's estimated value under the current extraction path

⁸We call this discount factor the benchmark rather than the optimal discount factor partly to avoid confusion because given any discount factor, we solve for the optimal consumption path. Additionally, a preference parameter like the discount factor is not usually interpreted as an object that is optimized. Behavioral analysis of decision versus experienced utility (Kahneman and Sugden 2005) has a parallel to our comparison of decision versus benchmark discount factors.

will be less than the maximum possible value: $\mathcal{V}_i(\mathbf{w}_i^*(a_{i0}; \beta_i^d), a_{i0}; \beta_i^*) < V_i(a_{i0}; \beta_i^*)$.⁹

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 presents the algorithm, and we summarize here. Given an arbitrary candidate discount factor β_i , we solve the agent’s dynamic water extraction problem (9). This solution reveals the agent’s choice of water extraction at the aquifer’s current 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 decision discount factor $\beta_i^d \in [0, 1)$ exists that matches observed water extraction, we exclude the aquifer from analysis, though a sensitivity analysis includes these aquifers with the closest feasible discount factor.

The estimated policy function $w_i^*(a; \beta_i^d)$ describes how much water the representative agent of aquifer i consumes given beginning-of-period stock a and the decision discount factor β_i^d . We then iterate forward on the law of motion (1), beginning with the observed stock in data a_{i0} , to obtain model-predicted paths of water extraction, aquifer level, profit, and extraction costs T periods into the future: $\{w_{it}, a_{it}, \pi_{it}, ec_{it}\}_{t=0}^T$. Our numerical analysis iterates out to $T = 1,000$ periods in the future. We verify ex-post that water levels in all aquifers converge by this period.

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.2 summarizes data sources.

Several measurement conventions apply across our datasets. For gridded datasets, we take an unweighted average of gridpoints within each aquifer $\times$ year. We deflate all values to 2023 dollars using the U.S. GDP implicit price deflator (FRED 2024). Given aquifer capacity κ_i , for any water quantity, we let q_i denote its value as a share of capacity and $\tilde{q} \equiv q_i \kappa_i$ its absolute level in m^3 . For example, w represents gross water extraction as a share of aquifer capacity, while $\tilde{w}$ represents gross water extraction in m^3 . Similarly, we measure aquifer recharge r_i and fill a_i as

⁹The 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 $V_i(a_{i0}; \beta_i^*)$.

a share of capacity.¹⁰ We analyze all data at the annual level, since most of our data lack sub-annual frequency. For data reported by country, we calculate the value for each aquifer that spans international borders as the weighted mean across countries, with weights equal to the share of the aquifer’s area in each country. For the case study aquifer systems, we aggregate the aquifer-level data as the weighted average across component aquifers, using the share of each aquifer in the total area of the system as weights.

The final analysis dataset has 38,822 observations on 2,773 aquifers over the years 2003 to 2016.¹¹ Many analyses exclude aquifers with limited agricultural activity (i.e., less than 1% cropped area). To reduce noise from individual years, the main results use the 2003-2016 mean of all variables. Some results distinguish the 37 major aquifer systems, which have large surface area that increases the precision of remotely sensed inputs (Richey et al. 2015a).

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. To apply the law of motion of groundwater in the absence of annual recharge data, we assume that an aquifer’s annual recharge equals its mean over previous years, echoing some hydrological models (Siebert et al. 2010; de Graaf et al. 2015). A sensitivity analysis obtains an estimate of time-varying recharge by applying a proportionality assumption between observed time-varying precipitation and unobserved time-varying recharge, an idea related to Rodell et al. (2018).¹²

To measure the saturated thickness $\mathcal{T}_{i,2013}$ of each aquifer in the year 2013, we use a widely-used global hydrological model (WaterGAP, from de Graaf et al. (2015)). To calculate the saturated thickness of each global aquifer, the authors use a global hydrogeological model parameterized using terrain geometry and lithology. Comparisons of model-derived estimates with observational data in selected areas where they are feasible have a correlation of 0.85 – 0.87 for the preferred model parameterization. Estimates of saturated thickness and, by extension, aquifer capacity that are based on observed aquifer properties improve upon capacity estimates from previous literature that assume constant values for inputs like specific yield across space (de Graaf et al. 2015; Richey

¹⁰We use this notation only for variables measured in terms of m^3 . For variables in levels or area, such as height h or surface area s , we do not scale by capacity, so h represents height in m and s represents area in m^2 .

¹¹We therefore refer to the 2003-2016 data as either the current or baseline level.

¹²For example, this sensitivity analysis seeks to reflect the extent to which floods, droughts, and other weather drive variation in groundwater recharge and depth.

et al. 2015a). Specific yield varies widely across aquifers (Appendix Figure A.1, Panels A and B). We assume 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: net extraction as a share of aquifer capacity equals 0.0% and 0.1% for the median and mean aquifer (Table 2), consistent with the near-zero mean overall net trend in fresh water availability in arable land (Carleton et al. 2024). The distribution is somewhat symmetric around zero (Appendix Figure A.2, Panel B). Since any long-run extraction path ends with a steady state that has zero net extraction, one summary of our results is that the average aquifer is in a steady state (zero net extraction), 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 crop production in irrigated land relative to rain-fed conditions, for 16 different 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 Tables ([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](#)). While we observe these variables at the country rather than aquifer level, as controls in some specifications of the production function estimates, they play a less direct role in our welfare and counterfactual analyses than aquifer management data like annual extraction.

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 years 2003-2016 (Appendix Figure [A.3](#)). Industrial energy prices provide the closest metric to agricultural prices that is available for all countries.¹³ Engineering estimates of pump efficiency range from 35% to 85%, with estimates centered around 60% ([Pradeleix et al. 2015](#); [Martin-Candilejo et al. 2020](#)). Our main results use the estimate of 54% from [Burlig et al. \(2021\)](#) based on groundwater pump microdata from California. Our main results measure energy costs from electricity prices given the absence of data from most countries on pump energy source, though a sensitivity analysis uses estimates of country-specific energy sources for groundwater pumping as well as energy-source specific estimates of pump efficiency ([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 Climate Research Unit at the University of East Anglia ([Harris and Jones 2019](#)). We measure rock porosity and 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. We compute the share of each aquifer that is covered by cropland by taking aquifer-averages over the share of each pixel that is cropland. We classify an aquifer as low agriculture if less than 1% of its area is cropped, and exclude low agriculture aquifers from parts of the analysis.

¹³In the U.S. at least, some farmers pay industrial electricity rates ([EIA 2014](#)). In other cases, electricity prices for agriculture can vary substantially from electricity prices for other sectors, depending on agricultural subsidies, time of use or seasonal pricing, or nonlinear pricing ([Badiani-Magnusson and Jessoe 2018](#)). Industrial prices to some extent proxy for marginal cost, which is ultimately the relevant variable to the planner rather than subsidized market prices.

4.4 Data: Benchmark Discount Factors and Others

The benchmark discount factor β_i^* represents the value the representative agent would use for resource management decisions in the absence of institutional or market failures. It is also the discount factor that the representative agent 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.3, 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 only show the Ramsey benchmark since it varies by aquifer while the market rate does not.¹⁴ 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.

Finally, as correlates for the decision discount factors that we recover for each aquifer, we use data on GDP per capita, regions with formal or informal water markets, the extent of democracy, and other variables (see Appendix D.5).

5 Parameter Estimation and Model Calibration

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

¹⁴While long-term interest rates on government debt vary by country due to risk, we adopt the standard assumption that different countries have the same risk-free long-term interest rate. For example, although Zimbabwe’s debt has higher interest rates than U.S. debt, both countries have the same relevant risk-free rate.

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:¹⁵

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

The annual change in available water is then the product of aquifer area, specific yield, and the change in aquifer height, where a rising lift height translates into decreasing 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, and Richey et al. 2015a). Gross annual ground-water 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 the available water (i.e., fill) $\tilde{a}_{it}$ in each year, with the annual change in available water $\Delta \tilde{a}_{it}$ taken from (13). The recursion begins from baseline fill $\tilde{a}_{i,2013}$, the year in which we can recover fill directly from the observed saturated thickness data:

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

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:

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

¹⁵Richey et al. (2015a), who also estimate aquifer capacity, use a stronger assumption of a uniform specific yield of 1% across all aquifers, based on prior estimates of the specific yield of subsurface rock. We allow specific yield to vary across aquifers and to differ from 1% in some cases, though the median estimate is 1%.

Log capacity varies considerably across aquifers (Appendix Figure A.1, Panels C and D).¹⁶

5.2 Estimation: Production Functions

Methods

We use versions of the following equations to estimate production function parameters:¹⁷

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

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

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), an annual 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 model-based analysis of welfare and counterfactuals in Sections 6 and 7 primarily uses 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.¹⁸ 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 therefore instrument for water, capital, and labor. The first stage equation for (20) is

¹⁶Capacity has more uncertainty than our other data inputs because it depends on saturated thickness, which is not directly observed. Older studies assume a uniform specific yield and maximum depth for all aquifers, producing a wide range of capacity estimates (Richey et al. 2015a). We instead follow recent work that combines hydrological models with observational data on well depths and rock characteristics (de Graaf et al. 2015).

¹⁷Market prices of water provide an alternative possible estimate of the marginal valuation of water resource, but water markets are rare globally and markets that do exist suffer from substantial market frictions.

¹⁸In our model of final and no intermediate goods, GDP equals output. We assume that the log of true output equals an aquifer-specific multiplier χ_i of nighttime luminosity plus normally-distributed measurement error, $\log y_{it} = \chi_i \log \bar{y}_{it} + v_{it}$. Richer machine-learning estimates of gridded GDP (e.g., Rossi-Hansberg and Zhang 2025) are only available beginning in 2013, limiting overlap with our groundwater data.

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

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 (Pindyck and Rotemberg 1983; Akerberg et al. 2015; Doraszelski and Jaumandreu 2018). Identification requires that lagged input prices and lagged weather be uncorrelated with current shocks to aquifer productivity.¹⁹

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

In first-stage regressions of groundwater extraction, capital, and labor on the instruments, lasso selects lagged precipitation, temperature, wages, and interest rates as instruments, including some interactions (Appendix Table A.3). 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 (10.01) (Stock and Yogo 2005). Given the observed instrument strength, we also report model results calibrated to match production function values from existing literature.

The estimates find moderate importance of water for output (Table 3). 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.²⁰ Instrumenting for groundwater raises the elasticity (columns 2 and 3); our preferred spec-

¹⁹Some 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.

²⁰For example, regional adoption of drought-resistant seeds would raise output while reducing water extraction, biasing OLS downward.

ification in column (4) also instruments for capital and labor. Column (5) shows that calibrating the land share from the literature (Gollin and Udry 2021; Ryan and Sudarshan 2022), rather than estimating it, yields a capital + labor share of 0.40 rather than 0.49. 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).

Comparison benchmarks are scarce because few production estimates account for water. Ryan and Sudarshan (2022)’s analysis of farmers in the Indian state of Rajasthan, one exception, estimates water input shares of 0.04 to 0.18, and labor + capital shares of 0.42 to 0.54. While our preferred instrumental variables estimates of 0.10 for the water share and 0.49 for the labor + capital share are in the middle of this range, the settings of course differ.

Appendix E discusses alternative production function estimates. Higher cropland, irrigation, and agricultural dependence predict higher returns to water (Appendix Table A.4). 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.5).

5.3 Calibration: Productivity

Applying the model quantitatively requires an estimate of transformed productivity z_i from equation (4), not only the production function parameters. We assume nighttime luminosity is proportional to GDP but does not numerically equal output in each aquifer, so we cannot directly use the production function estimates alone to 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 3 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.²¹

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

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

²¹D’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.

groundwater revenue for crop c .

The marginal returns to groundwater vary by aquifer (Figure 2, Panel C). Areas with little 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 estimating the cost function (7). Following engineering practice (Robinson 2002; Pradeleix et al. 2015; Martin-Candilejo et al. 2020), we measure the total extraction cost of lifting w_i cubic meters of groundwater to the surface from (7) as the height h_{it} in meters that the water is lifted times an aquifer-specific engineering cost c_i and 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:²²

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

As discussed above, we calibrate pump efficiency and the physical constant to Burlig et al. (2021), though a sensitivity analysis allows fuel type, price, and pump efficiency to vary by region.²³

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 has a far more skewed distribution than electricity prices, and therefore drives more of the dispersion in marginal extraction costs (Appendix Figure A.3, Panels D and F).

5.5 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 7, 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

²²Engineering 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.8, column 5).

²³Since in the model $C_i(a, w)$ 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.

when fit separately to each half of the sample. It does (Figure 7, 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 permits 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.

6 Results: Discount Factors

This section presents the decision discount factors that rationalize observed extraction in each aquifer, and compares them against the Ramsey and market benchmarks.

6.1 Graphical Explanation of Decision Discount Factors

A graph clarifies how the model estimates the decision discount factor that rationalizes an aquifer’s extraction path, using the Central Valley Aquifer System as an example (Figure 3). The horizontal axis describes gross groundwater extraction as a share of aquifer capacity (w). The vertical axis describes marginal revenues or marginal costs in $\$/m^3$. The thick downward-sloping marginal revenue curve, $\partial\tilde{\pi}_i(a, w)/\partial w$, shows how current groundwater extraction affects current profits. The positive value of this line reflects the return to water extraction, while the downward slope reflects decreasing returns to water. The dashed upward-sloping marginal cost curves show the marginal cost of extraction under different discount factors, including the stock externality that extracting groundwater today increases extraction costs in future periods. The benchmark β_i^* in orange crosses describes market discounting. The benchmark β_i^* in yellow triangles describes Ramsey discounting. The purple squares describe the estimated decision discount factor β_i^d . The green asterisks show the myopic scenario, where $\beta_i = 0$. The vertical dashed line shows observed gross water extraction w .

Interpreted graphically, our estimation algorithm essentially 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 upwards, by accounting more for future costs due to the greater future lift heights that current extraction causes. Policy functions then map aquifer fill to water consumption (see Appendix F.5 and Appendix Figure A.5).

In the Central Valley, the Ramsey and market benchmark factors exceed the decision discount factor, since the triangles and crosses lie above the squares (Figure 3). This implies that 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 Aquifer in the U.S. Great Plains does not. Aquifers in South Asia and Australia also have low decision discount factors; parts of the Middle East and West Africa do not. The major aquifer systems have a similar range of decision discount factors (Appendix Figure A.6).

The mean and median aquifers have decision discount factors of 0.86 and 0.92 (Table 2, Panel A). Because the benchmark market discount factor is 0.98 and most Ramsey discount factors are between 0.93 and 0.97, aquifers are split between having discount factors well below these benchmarks and having a discount factor somewhat near these benchmarks. Few aquifers have a decision discount factor below 0.50. The aquifers with decision discount factors in the ballpark of these benchmarks will contribute to our subsequent finding that the welfare losses from excess extraction are not enormous as a percent of aquifer present value for the mean aquifer. Weighting across aquifers by crop value, GDP, or population gives a similar range (Appendix Table A.6). Our estimate that decision discount factors for most aquifers substantially exceed zero implies that most aquifers are far from myopia.

These decision discount factors have sensible correlations with other variables (Appendix Table A.7). 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 geographic patterns in Panel A of Figures 4 and 5 suggest that most of these patterns reflect spatial patterns of decision discount factors rather than spatial patterns of Ramsey benchmark discount factors. This finding echoes the low correlation between decision and Ramsey discount factors discussed earlier.

We do estimate large gaps between some decision and benchmark discount factors (Figure 5, Panel B). While the modal difference is close to zero, many aquifers have decision discount factors far below the benchmark. A few aquifers also have decision discount factors slightly above the benchmark, potentially due to distortions in markets for factors or intermediates. The share of

aquifers with a discount factor below the benchmark is 61.0% with the Ramsey benchmark and 79.0% with a market benchmark (Table 4). These shares are moderately higher when weighted by aquifer capacity.

Decision discount factors are 2 versus 9 percentage points below the Ramsey discount factor for the median versus mean aquifer (Table 2, Row 2). This difference reflects 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.6).

6.4 Discount Factor Versus Property Rights Interpretation of β_i

We primarily interpret resource management in terms of the gap between the decision discount factor and the benchmark discount factor. A more standard approach to analyzing efficiency, 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 explores the relationship between our discount factor interpretation and the more standard property rights interpretation of observed extraction. Some research equates the existence of single-agent extraction with optimal resource management; we instead allow that a single representative agent for an aquifer may engage in non-optimal resource management, by extracting under a decision discount factor that differs from the benchmark discount factor.

Suppose each aquifer has a unit mass of atomistic agents who each extract water from private individual wells on the 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 $\lambda_i = 0$ corresponds to a common pool, where agents have no control over the water level and so take it as given. The value $\lambda_i = 1$ corresponds to a single agent who fully internalizes the effects of the agent's 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 $\bar{w}_i(a) = w_i^*(a; \beta_i, \lambda_i)$. Solving (24) involves finding a fixed point where, given the atomistic agent's water consumption $\bar{w}_i(a)$, the optimal policy function is $w_i^*(a) =$

$\bar{w}_i(a)$. Appendix C.2 describes our computational algorithm to solve for this fixed point.²⁴

Our discount factor interpretation is analytically equivalent to this property rights interpretation for the polar cases. The common pool atomistic agent case corresponds to $\lambda_i = 0$, and equivalently, $\beta_i = 0$ (myopia). Similarly, the single agent optimal extraction path corresponds to $\lambda_i = 1$ and $\beta_i = \beta_i^*$ (discounting at the benchmark rate). Many resource models describe these two cases, and our use of a discount factor rather than property rights is isomorphic to such models for these cases.²⁵

To further clarify how our discount factor interpretation of resource extraction compares to the literature’s property rights interpretation, we compute (β_i^d, λ_i) pairs such that the main model, which uses decision discount factor β_i^d , induces the same extraction as the property rights model, which relies on the parameter λ_i . For the property rights model, we set β_i equal to either the Ramsey or market discount factor.

We estimate a simple relationship between the decision discount factor β_i^d and the property rights regime λ_i that we recover for the Central Valley Aquifer System across a wide range of possible water consumption levels (Figure 6). Either the property rights or discount factor interpretation can rationalize observed levels of water extraction. The figure shows a strictly monotonic, increasing, nonlinear relationship between the discount factor interpretation (β_i^d) and the property rights interpretation (λ_i) for this aquifer. Its 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 aquifers (Appendix Figure A.7).²⁶ All these patterns hold whether using the Ramsey or market benchmark. These one-to-one functions suggest that while most of this paper uses the discount factor interpretation of resource management, one could think of the discount factor as analogous to the more long-standing property rights interpretation. The standard comparison of open access to single agent extraction provides only a binary characterization of property rights, while this analysis provides a continuous measure.

7 Results: Resource Management and Welfare

This section turns to the model’s remaining outputs: the extraction paths implied by decision versus benchmark discount factors, the welfare cost of the gap between them, and the consequences

²⁴Simpler 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).

²⁵In 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.

²⁶We have calculated these relationships only for the case study aquifers since this calculation is computationally intensive, including relative to other results in the paper.

of counterfactual extraction rules.

7.1 Results: Resource Management

Under either benchmark discount factor, the Central Valley Aquifer System would cut extraction nearly in half in order to refill the aquifer (Figure 8, 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 increases to a steady state of 1.6% by 2040, or 25 years after the decrease in extraction.²⁷ Decreasing extraction by half is an enormous change that would have large short-run consequences for farmers.²⁸ Steady state extraction under the decision and benchmark discount factors is equal because in a steady state for any aquifer, extraction equals recharge.²⁹

The same benchmark paths that refill the aquifer also reduce extraction costs and eventually raise annual profit (Figure 8). Under the decision discount factor, the Central Valley Aquifer System fill declines from 95% of capacity today to 90% by the year 2100. Under either benchmark discount factor, in contrast, the aquifer completely fills by 2040. Under the decision discount factor, extraction costs rise by about 40% over the 21st century. Under either benchmark discount factor, by contrast, the rising water table makes extraction costs plummet over a few decades to nearly zero. The current extraction path increases short-run profit. Within a decade, however, annual profit under the extraction path for either discount factor exceeds annual profit under the current extraction path. In steady state, annual profit under both benchmarks moderately exceeds annual profit under the current extraction path.

As discussed earlier, under the decision discount factor, net water extraction in the mean or median aquifer nearly equals zero (Table 2, Panel B). 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 decrease short-term extraction in order to increase long-term fill. The Ogallala Aquifer is refilling, but the optimal path would decrease extraction further.³⁰ In all cases either Ramsey or market discounting

²⁷This rapid refill is easily explained—the Central Valley has annual recharge of about 1.6% of capacity and has high baseline fill. Considerably decreasing gross extraction therefore can refill the aquifer within a limited time period.

²⁸As with other analysis of optimal policy (e.g., optimal taxation in public finance), we analyze the policy decisions that maximize social welfare, abstracting from political considerations. Political economy considerations would be important to practically pursuing the optimal extraction path.

²⁹At the capacity boundary $a = 1$, an aquifer can remain at a steady state level of fill of 100% as long as extraction is weakly less than recharge.

³⁰This 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).

would initially let recharge moderately exceed gross water extraction, producing negative net water extraction (i.e., refilling the aquifer).

These water extraction plans affect long-run fill (Table 2, Panel C). Under the current extraction path, the median aquifer will only be 94.1% full in steady state. In contrast, in steady state, Ramsey or market discounting leads the median aquifer to have 100% fill. Compared to the current extraction path, the benchmark discount factor substantially decreases steady state extraction costs as aquifers reach higher fill (Panel D). We find similar patterns for the median global aquifer weighted by crop value, GDP, or population (Appendix Table A.6).

These summary statistics mask differences in outcomes across aquifers. Under the decision discount factor, 46% of aquifers achieve more than 95% fill (Figure 9). Under the Ramsey and market benchmark discount factors, 61% to 84% of aquifers fill completely.

7.2 Results: Social Welfare

Aquifer Present Value

All global aquifers together have global value of \$103 trillion under Ramsey discounting and \$245 trillion under market discounting (Appendix Table A.8). Sensitivity analyses vary but all exceed \$49 trillion. Valuing an aquifer in full is equivalent to a counterfactual that completely drains it, which lies far outside the moderate changes in extraction paths that most of our scenarios consider. We therefore interpret the order of magnitude as informative — particularly since, to our knowledge, no existing work has estimated it — but nonetheless 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 earlier results show that decision discount factors are low, including in parts of the Western U.S. and Mexico (Figure 10, Panel A). The modal aquifer has a welfare cost of non-optimal extraction below 1.0% of aquifer present value, though the distribution has a long right tail (Panel B). Several percent of aquifers have losses due to non-optimal extraction paths that exceed 5.0% of aquifer present value, and a few have losses exceeding ten percent of aquifer value. No aquifers have a welfare cost of observed relative to optimal extraction which exceeds 16% of the aquifer’s value. We obtain comparable results for the market benchmark discount factor (Appendix Figure A.9).

These welfare costs of non-optimal aquifer depletion have large dollar though small percentage magnitudes (Table 4, Panel A). 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. In California’s Central Valley Aquifer System, for example, the loss is \$2.1 to \$9.0 billion in present value. Sensitivity analyses obtain qualitatively similar results (Appendix F.3).

A few reasons explain why non-optimal extraction may create modest percentage welfare costs. First, non-optimal extraction produces a natural dampening force. As the representative agent extracts water more quickly than is optimal, the water table declines and extraction costs rise. Thus, excessive extraction today slows extraction in the future through increasing future extraction costs. Second, extraction under the benchmark discount factor typically decreases short-run profits as the aquifer refills. Although this extraction path increases annual profit in future years, discounting diminishes the present value of those future 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. We consider relatively high benchmark discount factors, but this general force is still relevant to our setting. Third, a majority of aquifers have a decision discount factor above 0.90, so the typical gap between decision and benchmark discount factors is not enormous.

7.3 Counterfactual Extraction Patterns

This subsection analyzes three counterfactuals. In the first (myopia), the representative agent statically optimizes and ignores any impacts on future extraction costs (so $\beta_i^d = 0$). In the second (constant extraction), the current level of extraction persists in physical units for all future years. If an aquifer is fully drained, then extraction equals recharge. In the third (sustainability), the representative agent consumes recharge in every period.

We choose these counterfactuals for a few reasons. Common pool extraction can induce myopia. The Prior Appropriation doctrine from the Western U.S. or growing water subsidies could induce constant extraction. Sustainable extraction reflects the spirit of California’s Sustainable Groundwater Management Act. Additionally, our main results find modest welfare costs of the current extraction path. By weakening the complete dynamic optimization assumption of the main results, the first two counterfactuals explore whether alternative interpretations of the current extraction path find it to have larger welfare costs. Finally, the third counterfactual highlights classic trade-offs between extraction that is sustainable and extraction that maximizes the aquifer’s present value (Arrow et al. 2004).

Table 5 compares each counterfactual to the current extraction path along four dimensions: baseline net extraction, long-run fill, long-run marginal extraction costs, and welfare losses relative to the Ramsey and market benchmark paths. Figure 9 shows the corresponding distributions of

long-run fill across aquifers. We discuss each counterfactual in turn.

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 89% 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 the losses under the current extraction 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, its baseline net extraction equals that of the current path (a tenth of a percent of capacity in the mean aquifer). Long-run fill in the mean aquifer is 61%, similar to myopia, but the distribution across aquifers is bimodal rather than compressed: constant extraction completely drains 30.8% of aquifers and completely fills 45.2%.³¹ 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, and the fill distribution is compressed—no aquifer

³¹Steady state fill under the constant extraction counterfactual is somewhat mechanistic. In aquifers where observed extraction exceeds recharge, constant extraction fully drains the aquifer in steady state. In aquifers where recharge exceeds observed extraction, constant extraction completely fills an aquifer in steady state.

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 from reaching the fill level that maximizes the aquifer’s present value. This counterfactual therefore shows that a policy that is sustainable can still have large welfare costs relative to extraction that is optimal.

8 Conclusions

Economists have long worried that inefficiently rapid extraction of natural resources generates large social costs. To put numbers on these concerns, we modernize a classic model of resource extraction, applied with 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, groundwater extraction equals recharge today, just as it would in any counterfactual steady state, although most individual aquifers have non-zero net extraction. In a long term sense, therefore, median extraction levels reflect a steady state. At the same time, we calculate that most current levels are inefficient, since most aquifers have inefficiently low stocks. The optimal extraction plan in most aquifers would therefore decrease short-run extraction in order to refill most aquifers, and then would reach a steady state plan in which extraction equals recharge. While decreasing short-run water extraction decreases short-run profits, these costs are outweighed by the long-term benefit of a higher water table, which decreases extraction costs and increases profits.

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. Two counterfactual scenarios—myopia or repeating current extraction levels indefinitely—increase the costs of inefficient extraction, though under either scenario, costs remain under 13% of aquifer 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 which best characterizes the current extraction path. We show a monotonic numerical relationship between the discount factor and an

interpretation in terms of property rights.

We leave several avenues for future work. Work on spatial analysis of the environment and natural resources has burgeoned recently (Balboni and Shapiro 2025). Future work could also combine our analysis of intertemporal optimization (dynamics) with analysis of inter-regional optimization (trade and spatial). While this combination has been informative for analysis of climate change, its potential for groundwater is unknown.

Additionally, future work could study the importance of stochastic fluctuations around the long-run dynamic policy we analyze. One important constraint is the limited availability of evidence on stochastic patterns of recharge for each global aquifer. The ability to use aquifers as a buffer stock would tend to enhance their value beyond our nonstochastic analysis, amplifying our finding that a majority of aquifers are extracted more rapidly than is efficient, and that aquifers' aggregate global present value is enormous. At the same time, stochastic recharge is most important when an aquifer is empty, which is rare in scenarios we consider, particularly in the short-run. Future work could test the idea that the welfare importance of long-run extraction trajectories exceeds the importance of short-run fluctuations. Growth versus business cycle analysis in macroeconomics lightly echoes this comparison (Lucas 1988). Our nonstochastic dynamic analysis also slightly echoes the focus of climate change research on long-term dynamics (Nordhaus 2008) rather than on short-run fluctuations (Felkner et al. 2009). Some case study evidence finds modest importance of stochastic refill relative to the importance of long-run extraction (e.g., Merrill and Guilfoos 2018), though the extent to which this generalizes to other aquifers is unknown.

Aquifers also interact with other natural resources and ecosystem services, such as through saltwater intrusion or interactions of ground and surface waters. While data are limited to account for these in all global aquifers, analysis to clarify the broad importance of other ecosystem interactions 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 and exciting question is how to adapt the analysis from this paper for statistical agencies to measure dynamic values of natural resources.

References

- Akerberg, Daniel A, Kevin Caves, and Garth Frazer, “Identification properties of recent production function estimators,” *Econometrica*, 2015, 83 (6), 2411–2451.
- 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.
- Aronoff, Daniel and Will Rafey, “Conservation priorities and environmental offsets: Markets for Florida wetlands,” *American Economic Review*, 2026, 116 (5), 1723–1764.
- Arrow, Kenneth, Partha Dasgupta, Lawrence Goulder, Gretchen Daily, Paul Ehrlich, Geoffrey Heal, Simon Levin, Karl-Goeran Maeler, Stephen Schneider, 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.
- 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.
- 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.
- Borunda, Alejandra, “The deceptively simple plan to replenish California’s groundwater,” *National Geographic*, 2022, March 23.
- 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.
- Burgess, Robin, Matthew Hansen, Benjamin A. Olken, Peter Potapov, and Stefanie Sieber, “The Political Economy of Deforestation in the Tropics,” *Quarterly Journal of Economics*, 2012, 127 (4), 1707–1754.
- 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, “Energy, Groundwater and Crop Choice,” 2021. 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, pp. 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.
- 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.
- Cline, William R.**, “The economics of global warming,” *Institute for International Economics, Washington, DC*, 1992, p. 399.
- 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.
- Costello, Christopher, Daniel Ovando, Tyler Clavelle, C. Kent Strauss, Ray Hilborn, Michael C. Melnychuk, Trevor A. Branch, Steven D. Gaines, Cody S. Szuwalski, Reniel B. Cabral, Douglas N. Rader, and Amanda Leland**, “Global fishery prospects under contrasting management regimes,” *Proceedings of the National Academy of Sciences*, 2016, 113 (18), 5125–5129.
- 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**, *The Economics of Biodiversity: The Dasgupta Review*, HM Treasury, 2021.
- 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,” *The 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.
- Felkner, John, Kamilya Tazhibayeva, and Robert Townsend**, “Impact of climate change on rice production in Thailand,” *American Economic Review*, 2009, 99 (2), 205–210.
- 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.
- Fenichel, Eli, Joshua K. Abbott, Jud Bayham, Whitney Boone, Erin M. K. Haacker, and Lisa Pfeiffer**, “Measuring the value of groundwater and other forms of natural capital,” *Proceedings of the National Academy of Sciences*, 2016, 113 (9), 2382–2387.
- Fowlie, Meredith, Michael Greenstone, and Catherine Wolfram**, “Do energy efficiency investments deliver? Evidence from the weatherization assistance program,” *The Quarterly Journal of Economics*, 2018, 133 (3), 1597–1644.
- Frank, Eyal and Anant Sudarshan**, “The social costs of keystone species collapse: Evidence from the decline of vultures in India,” *American Economic Review*, 2024, 114 (10), 3007–3040.
- FRED**, “Gross Domestic Product: Implicit Price Deflator,” Oct 2024. <https://fred.stlouisfed.org/series/GDPDEF>, accessed 10/28/2024.
- Getirana, Augusto, Sujay Kumar, Manuela Giroto, 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.
- GGDC**, “Penn World Table Version 10,” August 2022. <https://www.rug.nl/ggdc/productivity/pwt/>, accessed 8/31/2022.
- Gibson, John, Susan Olivia, and Geua Boe-Gibson**, “Night Lights in Economics: Sources and Uses,” *Journal of Economic Surveys*, 2020, 34 (5), 955–980.
- 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.
- Goldewijk, C.G.M. Klein**, “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.
- 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,” *The 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.
- Greenhill, Simon, Hannah Druckenmiller, Sherrie Wang, David A. Keiser, Manuela Giroto, Jason K. Moore, Nobuhiro Yamaguchi, Alberto Todeschini, and Joseph S. Shapiro**, “Machine Learning Predicts Which Rivers, Streams, and Wetlands the Clean Water Act Regulates,” *Science*, 2024, 383 (6681), 406–412.

- 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.
- Hasan, M. F., R. Smith, Vajedian S., Pommerenke R., and Majumdar S.**, “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.
- 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, 137–75.
- Houthakker, HS**, “The Pareto distribution and the Cobb-Douglas production function in activity analysis,” *The 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.
- INSCR**, “Polity5: Regime Authority Characteristics and Transitions Datasets,” 2022. <https://www.systemicpeace.org/inscrdata.html>, accessed 10/19/22.
- 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), 3891–3898.

- 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, 132 (12), 575–605.
- 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.**, *The National Atlas of the United States of America* 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.
- Montginoul, Marielle, Jean-Daniel Rinaudo, Nicholas Brozović, and Guillermo Donoso**, “Controlling groundwater exploitation through economic instruments: Current practices, challenges and innovative approaches,” in “Integrated groundwater management: Concepts, approaches and challenges,” Springer International Publishing Cham, 2016, pp. 551–581.
- 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,” *The American Economic Review*, 1977, 67 (1), 341–346.
- , *A Question of Balance: Weighing the Options of Global Warming Policies*, Yale University Press, 2008.
- 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.
- O’Neill, Claire, Matt McCann, and Umi Syam**, “America Is Using Up Its Groundwater Like There’s No Tomorrow,” *New York Times*, 2023. August 28.
- 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,” Technical Report, Earthscan 1989.
- 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.
- Pindyck, Robert R. and Julio J. Rotemberg**, “Dynamic factor demands and the effects of energy price shocks,” *American Economic Review*, 1983, 73 (5), 1066–1079.
- 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), 543–59.
- 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.
- 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.
- 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*, 999–1003.
- , **Jay S. Famiglietti, David N. Wiese, John T. Reager, Henri 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 (7), 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.
- Rossi-Hansberg, Esteban and Jialing Zhang, “Local GDP Estimates Around the World,” 2025. BFI Working Paper No. 2025-17.
- Ryan, Nicholas and Anant Sudarshan, “Rationing the Commons,” *Journal of Political Economy*, 2022, 130 (1), 210–257.
- Sarah, S., Ahmed S., Violette S., and de Marsily S., “Groundwater sustainability challenges revealed by quantification of contaminated groundwater volume and aquifer depletion in hard rock aquifer systems,” *Journal of Hydrology*, 2021, 597, 126286.
- 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.
- Schmied, Hannes Mueller, Denise Caceres, Stephanie Eisner, Martina Floerke, 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 Doll, “The global water resources and use model WaterGAP v2.2d: model description and evaluation,” *Geoscientific Model Development*, 2021, 14 (2), 1037–1079.
- Siebert, S., J. Burke, J.M. Faures, K. Frenken, J. Hoogeveen, P. Doell, and F. T. Portmann, “Groundwater for Use in Irrigation: A Global Inventory,” *Hydrology and Earth System Sciences*, 2010, 14, 1863–1880.
- Siebert, Stefan, Matti Kummu, Miina Porkka, Petra Doll, Navin Ramankutty, and Bridget R. Scanlon, “Historical Irrigation Dataset (HID),” Mar 2015. <https://mygeohub.org/publications/8/2>, accessed 11/8/2021.
- 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, *Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg*, Cambridge: Cambridge University Press, 2005.
- 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, Jannis 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,” *Geoscience 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), 602–629.
- 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.
- USDA, “2023 Irrigation and Water Management Survey,” Technical Report, United States Department of Agriculture 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 Bank Development Indicators,” Dec 2023. <https://www.worldbank.org/en/home>.

Zhao, Jinhua, Nathan P Hendricks, and Haoyang Li, “Groundwater institutions in the face of global climate change,” *Annual Review of Resource Economics*, 2024, 16.

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.


Figures and Tables

Figure 1: Aquifer Structure and Terminology




Notes: This figure depicts an example aquifer to illustrate hydrological terminology.

Figure 2: Aquifer Characteristics



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 $\$/m^3$ 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 $\$/m^3$, 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.


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 $\beta^* = 0.96$.

Figure 4: Decision Discount Factors, by Aquifer

(a) Spatial Patterns of Decision Discount Factors




(b) Distribution of Decision Discount Factors Across Aquifers




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

(a) Spatial Patterns of Decision Minus Ramsey Discount Factors ($\beta^d - \beta^*$)




(b) Distribution of Decision Minus Ramsey Discount Factors Across Aquifers



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.


Figure 6: Estimated Relationship Between the Discount Factor and Property Rights Interpretations for the Central Valley Aquifer System




Notes: This figure describes the relationship between the property rights parameter λ and the decision discount factor β^d for the Central Valley Aquifer System, using the methodology described in Appendix C.2. We estimate λ assuming that β^* equals the Ramsey or the market discount factor. Appendix Figure A.7 shows estimates for the other case study aquifers.

Figure 7: 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



● Quantile bin mean — 45-degree line


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

Figure 8: 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.


Figure 9: Long-Run Fill, by Discount Factor and Counterfactual




Notes: This figure shows histograms of steady state aquifer fill, expressed as a percent of aquifer capacity, under alternate extraction paths. Panel A shows steady state fill under the decision discount factor (i.e., the current extraction path). Panels B through F show steady state fill under counterfactuals. In Panel B, the planner uses a Ramsey discount factor. In Panel C, the planner uses a market discount factor. In Panel D, the planner is myopic. In Panel E, users have the same gross extraction in each future year, unless the aquifer is drained. In Panel F, users consume recharge in every year.

Figure 10: Welfare Loss Due to Deviation from Ramsey Benchmark Discount Factor, by Aquifer

(a) Spatial Patterns of Welfare Losses



(b) Distribution of Welfare Losses



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.

Table 1: Summary Statistics

	Median aquifer (1)	Mean aquifer (2)	Central Valley (3)	Ogallala (4)	Indus (5)	North China (6)
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

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. \(2020\)](#). Currency values are in 2023 U.S. dollars.

Table 2: Model Outputs: Discount Factor and Intermediate Outcomes

	Median aquifer (1)	Mean aquifer (2)	Central Valley (3)	Ogallala (4)	Indus (5)	North China (6)
<i>Panel A: Discount Factor</i>						
Estimated	0.924	0.860	0.890	0.935	0.714	0.890
Estimated - Ramsey discount	-0.022	-0.086	-0.069	-0.023	-0.206	-0.022
<i>Panel B: Baseline Net Water Use (% of aquifer capacity)</i>						
Observed	0.00%	0.08%	0.23%	-0.30%	0.43%	0.36%
Ramsey discounting	-0.12%	-0.47%	-0.55%	-0.62%	-0.25%	-0.09%
Market discounting	-0.53%	-0.98%	-0.74%	-0.83%	-0.64%	-0.94%
<i>Panel C: Long-Run Fill (% of aquifer capacity)</i>						
Observed extraction path	94.09%	89.84%	88.51%	98.38%	86.99%	91.54%
Ramsey discounting	100.00%	86.52%	100.00%	100.00%	97.48%	98.07%
Market discounting	100.00%	94.89%	100.00%	100.00%	100.00%	100.00%
<i>Panel D: Long-Run Marginal Extraction Cost (\$/m³)</i>						
Observed extraction path	0.018	0.031	0.057	0.005	0.089	0.032
Ramsey discounting	0.001	0.032	0.000	0.000	0.018	0.007
Market discounting	0.000	0.013	0.000	0.000	0.000	0.000

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. Panels B and C are in percentage points (i.e., 0.33% represents a third of a percentage point). Ramsey 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. Currency values are in 2023 U.S. dollars.

Table 3: Estimated Production Function Parameters

	OLS (1)	IV, water only (2)	IV, water only (3)	IV, all (4)	IV, calibrated (5)	IV, GAEZ (6)
log(Groundwater) (ϕ)	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(Groundwater) $\times$ Value (ϕ_i)						0.479** (0.194)
1 - log(Capital) - log(Labor) (θ)	0.863*** (0.049)	—	0.893*** (0.051)	0.514*** (0.155)	0.603*** (0.044)	0.405** (0.162)
Implied log(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(Groundwater), where groundwater represents gross extraction in m³. Columns (4) and (6) instrument for log(Labor) and log(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(Land) is the implied value of the land share under the assumption of constant returns to scale. The dependent variable is log(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 (*).

Table 4: 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 from following...</i>							
Ramsey discounting (% of aquifer value)	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%
Ramsey discounting (billion \$)	0.07	1.88	2,582	2.12	0.54	5.13	0.54
Market discounting (billion \$)	0.20	6.11	8,386	9.01	2.93	91.87	44.88
<i>Panel B: Fraction of aquifers over-extracted</i>							
Ramsey discounting	—	—	0.610	—	—	—	—
Market discounting	—	—	0.790	—	—	—	—
Ramsey discounting (capacity weighted)	—	—	0.713	—	—	—	—
Market discounting (capacity weighted)	—	—	0.857	—	—	—	—

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 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. Currency values are in 2023 U.S. dollars.

Table 5: 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 (%)</i>						
Observed extraction path	0.00%	0.08%	0.23%	-0.30%	0.43%	0.36%
Myopic extraction	4.12%	6.44%	2.85%	2.64%	1.08%	7.45%
Constant water extraction	0.00%	0.08%	0.23%	-0.30%	0.43%	0.36%
Zero net extraction	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
<i>Panel B: Long-Run Fill (%)</i>						
Observed extraction path	94.09%	89.84%	88.51%	98.38%	86.99%	91.54%
Myopic extraction	71.65%	61.41%	75.46%	73.42%	84.06%	67.65%
Constant water extraction	91.09%	60.74%	0.00%	100.00%	0.00%	0.00%
Zero net extraction	94.82%	93.19%	92.44%	90.49%	91.87%	96.58%
<i>Panel C: Long-Run Marginal Extraction Cost (\$/m³)</i>						
Observed extraction path	0.02	0.03	0.06	0.00	0.09	0.03
Myopic extraction	0.10	0.10	0.12	0.07	0.11	0.12
Constant water extraction	0.04	0.14	0.50	-0.00	0.68	0.37
Zero net extraction	0.02	0.02	0.04	0.03	0.06	0.01
<i>Panel D: Welfare Loss (Ramsey) (%)</i>						
Observed extraction path	1.35%	2.55%	2.75%	0.48%	2.58%	0.18%
Myopic extraction	8.93%	9.16%	10.66%	10.23%	5.42%	8.43%
Constant water extraction	1.68%	3.40%	4.85%	1.23%	6.79%	0.52%
Zero net extraction	1.49%	2.87%	1.73%	1.75%	0.84%	0.12%
<i>Panel E: Welfare Loss (Market) (%)</i>						
Observed extraction path	1.69%	2.90%	5.54%	1.24%	10.83%	3.37%
Myopic extraction	14.29%	12.88%	14.64%	14.38%	14.28%	16.01%
Constant water extraction	2.24%	5.44%	12.27%	2.44%	38.88%	11.25%
Zero net extraction	1.64%	2.91%	3.74%	4.07%	6.61%	1.57%

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. Myopia describes extraction under a discount factor of $\beta = 0$. Constant water extraction describes a scenario where in each future year, the representative agent extracts the same amount as under the baseline period of the current extraction path. Zero net extraction describes a scenario where the representative agent extracts recharge in each year, so that aquifer fill is constant in all years. All estimates are based on the model in Section 3.3 and Appendix C.1. Dollar values are in 2023 U.S. dollars.

Appendix

A Additional Background on Case Study Aquifer Systems

This appendix provides additional background on the four case study aquifer systems.³² As the main text mentions, we choose these four given their capacity, population, GDP, and common public discussion. These four also have the modest research benefit that one aquifer almost completely contains the aquifer system in the hydrological map data (BGR 2025).

The Central Valley Aquifer System supports 40% of American fresh produce (Borunda 2022). California has historically not metered groundwater extraction, so farmers can pump groundwater with little oversight. In response to a declining water table, California in 2014 passed the Sustainable Groundwater Management Act, which aims for sustainable extraction of the Central Valley Aquifer System by 2042, using price and non-price policies (Burlig et al. 2021; Bruno et al. 2022).

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

The Indus Basin straddles the arid regions of India, Pakistan, China, and Afghanistan. Some research considers the Indus Basin to have the highest net depletion of all major aquifer systems, 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 this aquifer system is the largest wheat and maize production area in China. As with the other case studies, agriculture accounts for most groundwater extracted from this aquifer. Between 2003 and 2010, the aquifer lost 50 km³ of groundwater, exceeding the capacity of the country’s Three Gorges Dam, the world’s largest power station (Feng et al. 2013).

B Derivation of Model Results

Given w , the agent faces the static optimization problem

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

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

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

³²Naming conventions for these water resources vary slightly across researchers; we follow the naming conventions from 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 follows the convention of referring to the whole aquifer system by the name of this main unit. The boundaries of some major aquifer systems, such as the Indus Basin, draw their names from surface water hydrological basins with which they overlap almost perfectly.

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

$$y_i(w) = \mathcal{Z}_i \left(\frac{\alpha y_i(w)}{p_i^k} \right)^\alpha \left(\frac{\gamma y_i(w)}{p_i^l} \right)^\gamma w^{\phi_i} x_i^{1-\alpha-\gamma-\phi_i}$$

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

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) using value function iteration. We discretize aquifer fill level $a \in [0, 1]$ using an evenly spaced grid with 100 nodes. Throughout this subsection, we suppress aquifer indicators i and use a_j to denote discretized aquifer fill levels. Our numerical algorithm is as follows:

1. Guess an initial discretized value function $V^0(a_j)$.³³ Set the iteration counter to $n = 0$.
2. Use cubic spline interpolation to obtain the interpolating function $V^n(a)$ and its derivative $\partial_a V^n(a)$.
3. Compute the unconstrained water extraction w_j^u that satisfies the unconstrained version of the first-order condition (10), evaluated at the discretized fill levels 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. Substitute the discretized policy function into the Bellman equation (9) to update the discretized value function:

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

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 mentioned in the main text, if no $\beta_i^d \in [0, 1)$ exists that matches observed water extraction, we exclude the aquifer from the main analysis, though a sensitivity analysis adds these aquifers back into the analysis. In 14% of aquifers, model-implied water extraction is at the lower bound of the range of β_i^d values for which the model converges, and in 6% of aquifers, model-implied water extraction is at the upper bound.

³³We implement this with an initial guess of $V^0(a_j)$ everywhere.

C.2 Computing the Property Rights Parameter λ_i

This subsection explains how we calculate the property rights parameter λ_i that Section 6.4 describes. For a given $\lambda_i \in [0, 1]$, we solve the agent’s problem (24) using the following algorithm:

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) using the algorithm described 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

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

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

To estimate the property rights parameter, we search for a $\lambda_i \in [0, 1]$ such that the model-predicted extraction equals observed extraction: $w(a^{\text{data}}; \lambda_i) = w^{\text{data}}$.

C.3 Welfare Analysis

We compute the welfare loss ω from excess extraction from equation (12) as follows. First, we use the algorithm from Appendix C.1 to compute the discretized policy function w_j implied by the decision discount factor β_i^d . We then use cubic spline interpolation to approximate the policy function $w(a)$, $a \in [0, 1]$. Given the fill level at time t , we can use the law of motion (1) and policy function to predict the fill level at $t + 1$:

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

Starting with the current observed fill level a_0 , iterate forward in this manner to obtain the predicted time series of fill levels $\{a_t\}_{t=0}^T$ and water consumptions $\{w_t\}_{t=0}^T$ from time $t = 0$ to $t = T$. If T is sufficiently large that a has (approximately) converged by time T , then we can approximate the value of the current water consumption path according to the benchmark discount factor (defined in equation (11)) using³⁴

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

We then compute the maximum value of the aquifer according to the benchmark discount factor β_i^* , $V(a_0; \beta_i^*)$, using the algorithm described in Appendix C.1. Finally, we compute the welfare loss ω from equation (12).

D Data Details

D.1 Groundwater Storage and Recharge

GRACE uses a pair of satellites to measure changes in Earth’s gravity fields in each aquifer. GRACE reports 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 the annual change in storage for the years 2003 to

³⁴We set $T = 1000$ and verify ex-post that the fill level of every aquifer converges over this horizon.

2016, letting us observe the full calendar year, and calibrate the model to these years. We drop the 2002 and 2017 months to exclude data for partial years; many of our other data sets are at the annual level and do not permit analysis of month-to-month seasonality. Due to using up to three lags in regressions, the estimation of the production function input shares uses the years 2006-2016.³⁵

GRACE reports the height of the change in water storage, rather than the change in the height that water must be lifted. The distinction arises because one meter of extractable groundwater may be spread vertically across many meters of porous rock. We divide the change in water thickness by specific yield to convert the change in the height of water extracted from GRACE to the change in lift height.³⁶

The raw gridded GRACE data provide measures of total terrestrial water storage, which includes groundwater, surface water, snow, and vegetation and soil moisture. We use gridded measures of plant and soil moisture and snow water-equivalent from the Global Land Data Assimilation System (GLDAS) to measure sources of terrestrial water storage other than groundwater (Beaudoing and Rodell 2016). Surface water, the remaining component of total terrestrial water storage (together with groundwater, plant and soil moisture, and snow), accounts for a small share of variation in total water storage; surface water accounts for only about 8% of changes in global total 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). This approach follows Rodell et al. (2007), Rodell et al. (2009), and Richey et al. (2015a).

We take several steps to obtain annual changes in groundwater storage for each aquifer from the raw gridded GRACE data. These steps also follow Rodell et al. (2007), Rodell et al. (2009), and Richey et al. (2015a) to isolate groundwater storage changes from changes in other terrestrial water storage types at the aquifer-level. First, we multiply the GRACE data by the land-grid-scaling factor provided by NASA to reduce post-processing noise. Second, we calculate the annual change in terrestrial water storage from the month of January of one year to the January of the following year using the average of the data sets generated by the three different agencies that pre-process the data into raster format, in units of deviations in groundwater storage relative to the 2004-2009 mean. Third, we isolate groundwater storage from total terrestrial water storage by subtracting gridded measures of plant and soil moisture and snow water-equivalent from the GLDAS (Beaudoing and Rodell 2016); we transform the GLDAS data to the same units as GRACE (i.e., annual deviations from the 2004-2009 mean and January-January changes). Fourth, we intersect the resulting gridded data set with global maps of aquifer extent from the World-wide Hydrological Mapping and Assessment Programme (WHYMAP) (BGR 2025) to obtain aquifer-level averages of groundwater storage changes. Fifth, since some resulting observations are clear outliers (e.g., increases or decreases in groundwater height exceeding hundreds of meters), we exclude the top and bottom 1% of observations. Finally, we exclude aquifers in Antarctica since several of our main data sets lack data on them.

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; Schmied et al. 2021).³⁷ These gridded recharge data account for differences in local precipitation, soil porosity, and in-flow from returning surface water runoff. Schmied et al. (2021) provide details: recharge is a function

³⁵A separate GRACE Follow-On (GRACE-FO) mission was launched in May 2018 to continue tracking terrestrial water storage. Our data set ends after the first mission due to the two-year January-January gap in the panel, the onset of the pandemic years and associated economic fluctuations, and uncertainty about continuity of data from the two missions.

³⁶For 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.

³⁷To assess accuracy of this model, de Graaf et al. (2015) compare the model-predicted lift heights in the WaterGAP model and observational well depth data from the global map of groundwater lift heights compiled by Fan et al. (2013). They find a reasonably high correlation coefficient of 0.85-0.87 between the model-predicted and observational data for their preferred model parameterization.

of daily total effective precipitation (including snow and rain), moisture stored in the soil, and soil-specific porosity and saturation characteristics. Effective precipitation is precipitation after deducting amounts absorbed by the plant canopy. Effective precipitation is adjusted by a soil-specific runoff coefficient that adjusts for differences in the extent to which water can permeate the soil. If the total of daily effective precipitation that does not run off the topsoil plus existing soil moisture exceeds the soil-specific saturation threshold, the excess becomes recharge in the underlying aquifer. The model adds estimates of return flow from irrigation and inflow from surface water bodies to obtain total recharge.³⁸ WHYMAP reports binned estimates of annual recharge based on this methodology (BGR 2025). We assign the midpoint of each bin to each aquifer, except for the highest bin, which includes recharge above 300 mm per year and to which we assign the minimum of the bin; and the lowest bin, which includes recharge of less than 2 mm per year and to which we assign the maximum of the bin. In sensitivity analyses, we allow aquifer-level recharge to vary from the 30-year average reported in WHYMAP by scaling their estimate up or down by the percentage deviation in rainfall during our sample time period from the 100-year average rainfall.

Calculating aquifer capacity in equation (17) uses data on specific yield, saturated thickness, topsoil thickness, lift height, and surface area. Since one aquifer may have rock with different specific yields, we follow Richey et al. (2015a) and use the minimum specific yield across gridpoints within an aquifer to account for the most binding constraint on groundwater flow. We assume aquifer capacity is fixed over time because capacity changes due to pollution or subsidence typically have small magnitude relative to aquifer capacity (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).³⁹ Quantifications of the effects of pollution on groundwater storage are limited; one study on one aquifer in India finds that dilution from recharge may fully offset pollution incursion (Sarah et al. 2021).

D.2 Measurement Error

Satellite- and model-derived data may have measurement error. Our averaging across years of datasets provides one response. Apart from our discussion of four case studies, we also emphasize moments and other aggregates of data across many aquifers, rather than emphasizing the precise value for each individual aquifer. In regression settings, some papers use Empirical Bayes to adjust value added estimates for teachers or similar results for sampling variation. We do not pursue this approach here for a few reasons—the use of dynamic programming rather than regressions implies that we do not have estimates of uncertainty for each aquifer to use for adjusting decision discount factors; we also lack estimates of uncertainty for most inputs; and even with these inputs, the computational burden of resampling the dynamic programming algorithm for each global aquifer would be large.

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

³⁸While this model takes into account flow between surface water and groundwater, this and other global hydrological models typically abstract from lateral groundwater flow from one aquifer to neighboring aquifers due to lack of data and model recharge as a function of precipitation and local land cover and geological characteristics (e.g., de Graaf et al. 2015; Sutanudjaja et al. 2018).

³⁹Subsidence that has small implications for storage capacity can have very large above-ground infrastructure costs (Hasan et al. 2023).

from 1992-2013. We assess the robustness of our results to analyzing the data separately for each of the two time periods.

How well does nighttime luminosity proxy economic activity? Nighttime luminosity has the advantages of gridded resolution and global scope, as GDP data varies in quality across countries and is not typically available sub-nationally (Chen and Nordhaus 2011; Donaldson and Storeygard 2016). Night lights data are commonly used in economic analysis for these reasons: Gibson et al. (2020) approximate that, since 2012, over 150 economics papers have used night lights data as a proxy for 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 across the two decades prior to 2012, with a higher correlation in urban areas than in rural ones. Since then, changes in the measurement and spatial resolution of the data have improved the approximation of rural economic activity after 2012 (Perez-Sindin et al. 2021). For the earlier, pre-2012 years of the data, Gibson et al. (2020) report a statistically significant, positive relationship between night lights and GDP in areas where the agricultural share of GDP is less than 20%; while aquifer-specific agricultural GDP data are not available, we note that, in our data, 80% of aquifers are less than 20% cropped.

We also draw on two alternate measures of GDP to use as outcome variables in production function estimates. The first comprises country GDP from the World Bank’s World Development Indicators for the years 2003 to 2016. We multiply these by the share of the population employed in agriculture, also from the World Development Indicators, to proxy for agricultural GDP. The second is gridded, purchasing power adjusted GDP data from Kummur et al. (2018), available for the years 1990 to 2015. These gridded data scale official national statistics, so only approximate the spatial distribution of sub-national economic activity.

D.4 Marginal Return to Groundwater and Costs Data

We also draw on crop production from irrigated and rainfed sources and the marginal return to groundwater in the irrigation of these crops from the GAEZ data set published by the Food and Agriculture Organization and the International Institute for Applied Systems Analysis. The crops included in GAEZ production under actual conditions data include banana, barley, cassava, cotton, groundnut, fodder, fruit, maize, millet, oil palm, olives, other cereals, potato, pulses, rapeseed, rice, rye, stimulants, sorghum, soy, sugarbeet, sugarcane, sunflower, tobacco, vegetables, wheat, yams, and one “rest of crops” category. The marginal return to groundwater of these crops varies by crop and, for some crops, by continent and is from D’Odorico et al. (2020). The authors calculate the crop- and continent-specific value of water as the ratio of the increase in crop production from irrigation and the associated increase in irrigation water.⁴⁰ The crops included in the D’Odorico et al. (2020) data are barley, cassava, groundnut, maize, millet, potatoes, oil palm, rapeseed, rice, rye, sorghum, soybean, sugarbeet, sugarcane, sunflower, and wheat. We use the continent-specific prices where available (i.e., for maize, rice, soybean, and wheat). For the 10 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 5.2 calculates the marginal return to groundwater for each aquifer. To do so, we calculate the marginal return to groundwater $y'_i(\tilde{w}_i) = \sum_{c=1}^{26} \sigma_{ic} \times v_c$ for each pixel within the aquifer, 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

⁴⁰They 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.

groundwater in irrigation for crop c in that pixel. We assign pixels without crop production a marginal return to groundwater of 0. To calculate the marginal return to groundwater at the aquifer-level, we average across pixels within the aquifer.

The various country-level data differ somewhat in spatial and temporal coverage. For years between 2003 and 2016 that are missing data, we linearly interpolate between the existing observations for each country. We impute missing country-level data using the annual continent mean.

In addition to the labor data from the World Bank’s World Development Indicators discussed in Section 4, we also draw on alternate gridded measures of population. A benefit of this data set is that its finer resolution lets us more precisely approximate aquifer labor inputs than through the use of the country-level statistics. A disadvantage is that this data set is only available every five years; we linearly interpolate the intervening years.

The extraction costs data analyzed in the main text use electricity prices and groundwater pump efficiency for electric pumps in the calculation of extraction costs. We also consider an alternative measure of extraction costs that incorporates differences in pump fuel type and efficiency across aquifers. We obtain electricity and diesel groundwater pump efficiency and electricity and diesel shares in total energy used for groundwater extraction from [Qin et al. \(2024\)](#) and electricity and diesel industrial input prices from [IEA \(2021\)](#). We re-weight these country-level variables to the aquifer-level. We calculate the aquifer-level extraction cost as the weighted sum of the cost of extraction using an electric pump and a diesel pump using the shares of each fuel in total groundwater extraction energy as weights.

D.5 Other Data

Our analysis of heterogeneous returns to groundwater relies on a variety of alternative measures of the value of water in irrigated agriculture. Our preferred estimate is based on a measure of the relative contribution of groundwater to agricultural production, using data from GAEZ. We measure groundwater irrigation’s effect on crop yields as the weighted sum across crops of the current actual yield under irrigation conditions less the current actual yield under rainfall conditions at the aquifer level. The weights are the share of each crop’s contribution to total yields. We sort the resulting irrigation effects by percentiles since the yield returns in physical quantities themselves are highly dispersed.

Many very small aquifers have missing cropland shares; if an aquifer is missing cropland share and has nonzero crop value in 2010 from GAEZ, we do not define it as low agriculture.

Additional alternate measures of the local importance of groundwater include gridded data on the spatial extent of cropped land, irrigated land, and dominant land use. The data on irrigated land coverage are from the History Database of the Global Environment (HYDE) and the Aquaknow Historical Irrigation Data set for the year 2005 ([Goldewijk 2017](#); [Siebert et al. 2015](#)). We use these data to construct measures of the fraction of the aquifer that is cropped and the fraction that is irrigated as the area dedicated to each of these uses divided by the aquifer land area, as well as binary measures indicating whether or not any cropland or irrigation activity occurs above the aquifer. We also measure whether the dominant use of land above the aquifer is agriculture using a separate gridded data set from GAEZ on dominant land use class.

Other, non-agricultural aquifer characteristics include GDP per capita, democracy, and water markets. We use these data to analyze determinants of the discount factors that we estimate. GDP per capita and democracy data are from [Kummu et al. \(2018\)](#) and [INSCR \(2022\)](#) respectively, re-weighted from country-level to aquifer level. 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 summarizes values of several model inputs for the four case study aquifers and for the mean and median of all global aquifers. We show our estimates of aquifer capacity, groundwater extraction, marginal returns, marginal extraction cost and some of its components, lift height and electricity prices, and the benchmark Ramsey discount factor from Addicott et al. (2020).

Table 1 shows results for our four case study aquifers: the Central Valley Aquifer System, the Ogallala Aquifer, the Indus Basin, and the North China Aquifer System. Of these, the Indus Basin has the largest capacity, which contributes to its relatively low depletion as a share of its capacity. All four of these aquifers are relatively high-value, with marginal returns estimates that range 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 2d). Marginal costs are between 5% and 35% of marginal return to groundwater for these aquifers, with differences largely driven by variation in lift height. This table also shows that the Ramsey discount factor is between 0.91 and 0.96 for all four aquifers.

The four case studies are somewhat similar to the mean or median global aquifer, with a few main differences. Mean and median gross groundwater extraction as a share of capacity are approximately 2-3% for all aquifers, which is similar in magnitude to the case studies. The mean and median global aquifers are characterized by marginal returns estimates that are weakly below the Central Valley Aquifer System and the North China Aquifer System; weighting by aquifer crop value, GDP, or population increases marginal returns by approximately 10%. Median and mean extraction costs are on the lower end of the case studies, in the range of $\$0.01/m^3$ to $\$0.02/m^3$. Differences in lift height, rather than electricity prices, appear to drive this distinction. Ramsey discount factors are comparable, also within the range of 0.91 to 0.96 across the different examples.

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 where groundwater is difficult to access or with 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 has gross annual extraction below 1.0%, though a large share have extraction up to 5.0% of capacity, and a smaller share up to 10.0% (Panel B). The mean aquifer has gross annual extraction equal to 3.1% of capacity (Table 1).

Panel C of Appendix Figure A.2 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. Spatial patterns of recharge have less clear correlation with spatial patterns of agricultural production. Annual recharge in most aquifers ranges from 0.5% to 1.5% of aquifer capacity (Panel D).⁴¹ Because recharge has a long right tail, however, the mean across aquifers is larger, at 2.9% (Table 1).

E Heterogeneous Production Function Estimates

We examine heterogeneity in the return to groundwater across aquifers by extending equation (20) to include an interaction term between groundwater extraction and a measure of the importance of groundwater in each aquifer. We then estimate ϕ_i for each aquifer i using the following regression:

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

⁴¹Most aquifers have some level of recharge, and only 6% of aquifers with data have annual recharge below 0.1% of aquifer capacity.

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 use the same instrumental variables approach as in equation (20) to estimate this equation. We also estimate equation (26) without the interaction term but exclude low agriculture aquifers from the sample, i.e., aquifers with less than 1% share of cropland. We report results in Appendix Table A.4. We measure heterogeneity in the importance of groundwater using the share of the aquifer that is cropland (column 1), the share of the aquifer that is irrigated (column 2), and an indicator variable equal to 1 if any agriculture occurs above the aquifer (column 3), if any irrigation takes place (column 4), or if agriculture is the dominant land use above the aquifer (column 5). Columns 6 and 7 report OLS and IV results on the high agriculture data subsample.

The heterogeneous effects estimates, in Appendix Table A.4, support that groundwater is more valuable in areas with more agriculture. For example, aquifers at the 90th percentile of the distributions of cropland and irrigation have elasticities of 0.41 and 0.19 (Appendix Table A.4, columns 1 and 2, respectively), above the mean of 0.10 across all aquifers (Table 3). Restricting the sample to high agricultural aquifers only yields an IV estimate of 0.274 (column 7), again above the mean across all aquifers.

Appendix Table A.5 shows that the input shares are generally robust to different GDP outcomes, different groundwater extraction definitions, and different data subsamples. They are also robust to the inclusion of surface water in w_{it} in addition to groundwater. We estimate the input shares using water inputs from all sources instead of groundwater only, using economic data on agricultural GDP as an outcome variable, using data from time periods when the night lights satellite data were measured differently, and using alternate measures of labor based on gridded population estimates. Relative to our main estimate of $\phi = 0.10$, we estimate a groundwater input share of 0.12 when including all water inputs in w_{it} (column 1), 0.17 when using re-weighted agricultural GDP as an outcome in the place of night lights (column 2), 0.15 when using the more recent night lights satellite data and 0.10 when using the older satellite data (columns 5 and 6, respectively), 0.10 when we use logs instead of inverse hyperbolic sine (column 7), 0.07 when we measure labor using gridded population (column 8), and 0.17 when we exclude temperature and precipitation controls (column 9). We estimate smaller elasticities of approximately 0.03 when we use gridded GDP available for a more limited set of years as an outcome instead of night lights (column 3), and when we restrict the data set to the deepest aquifers only (column 4), though we cannot reject that all of these are equal to 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 .

F Additional Discussion of Results

F.1 Graphical Analysis of Decision Discount Factors

This appendix subsection explains the graphical discussion of decision discount factors in detail for the Central Valley Aquifer System. The dashed vertical line in Figure 3 shows that in the mean year for our data, gross water extraction equaled 1.8% of aquifer capacity. At this quantity, we calculate marginal revenue from water of \$0.12/m³, which is the value on the y-axis where the thick downward-sloping marginal revenue line intersects the dashed vertical consumption line. At this quantity, myopia implies a marginal extraction cost of \$0.04/m³, shown in the dashed line with asterisks. At this quantity, this myopic marginal cost is below marginal revenue, so myopia has too low a discount factor to rationalize observed water extraction. The Ramsey discount factor of 0.96 (Table 1) implies the marginal cost of water extraction equals \$0.18/m³ at the observed quantity, shown in the dashed line with triangles. Because this exceeds marginal revenue, Ramsey has too high a discount factor to rationalize observed extraction. The market discount factor of 0.98 shown with ×'s is even higher. Ramsey and market

discounting consider future extraction costs due to lifting the water further; myopia does not. A discount factor of 0.89 makes the discounted marginal cost of water extraction of $\$0.12/\text{m}^3$ equal marginal revenue at the observed quantity, and therefore best describes observed extraction.

F.2 Long-Term Management of Ogallala Aquifer

This appendix subsection explains long-term resource management outcomes for one case study, the Ogallala Aquifer. Figure A.8 shows that the decision discount factor for the Ogallala Aquifer leads to relatively rapid gross extraction over the century (Panel A). Extraction is slower under the benchmark Ramsey and market discount factors. Panels B and C show that, as a consequence, the benchmark extraction paths steadily increase fill and decrease extraction costs, leading to higher long-run profits. The main difference between the Central Valley and Ogallala case studies is that the Central Valley rapidly reaches steady state, while the Ogallala is still converging by the end of the century. Similarly, Central Valley extraction in the three scenarios converges within a few decades.

F.3 Model Sensitivity Analyses

This section summarizes the sensitivity of the model results to changes in the estimated and calibrated model parameters. Table A.8 reports these results. Column 1 summarizes our model estimates from the main text and Columns 2-7 report model results for different model parameters.

Columns 2-4 use different production function parameters. Column 2 uses a calibrated value of the land share from Ryan and Sudarshan (2022) in the production function. Together with our estimated groundwater input share ϕ , calibrating the land share pins down the resource share θ without requiring estimation of the capital and labor shares in equation (20). We find similar welfare losses using our production function estimates or calibrating the production function parameters based on these estimates from the literature, which suggests that aquifers in this scenario are similarly far from the optimal extraction path on average. The dollar value of the welfare loss increases by 1-2% using the alternative production function, but the welfare loss as a percent of aquifer value is lower because the total value of global aquifers is higher. The share of aquifers that are over-extracted falls by 2-4 percentage points in the Ramsey and market benchmarks.

Column 3 uses production function parameters resulting from estimating equation (20) using all water (surface+groundwater) in production, rather than groundwater only. This sensitivity analysis assumes that surface water and groundwater are perfectly substitutable in production. Table A.5 shows that the water input share ϕ rises slightly relative to the groundwater only input share (0.121 versus 0.103) and the resource share θ falls slightly (0.479 versus 0.514). Including surface water slightly lowers the welfare loss in dollars from non-optimal extraction. The welfare loss in percentage terms is higher because the value of groundwater resources is lower, reflecting the availability of surface water substitutes.

Column 4 uses aquifer-specific production function parameters resulting from allowing the water input share to vary across aquifers; Table A.4 reports these production function parameters. As in column 3, we find a very small effect on the welfare loss in dollars from non-optimal extraction, again suggesting that aquifers are similarly far from the optimal extraction path in this scenario. The effect on the welfare loss as a percent of aquifer value increases because the global value of aquifers falls.

Column 5 reports results that use an alternate calibration of marginal extraction costs. This sensitivity analysis allows marginal extraction costs to vary across aquifers based on differences in aquifer-specific pump fuel source (electricity versus diesel) and pump efficiency. We find that using this alternative marginal cost estimate has very little effect on the model outcomes. We find welfare loss estimates and global aquifer valuations that are all within 10% of the outcomes that we report in the main text, for both Ramsey and market discounting. Most of the over-extracted aquifer shares are similar or higher.

Column 6 calibrates recharge to the distribution of precipitation during the years where we observe extraction data from GRACE. If rainfall is unusually high or low during this 2003-2016 period, calibrating recharge to the long-run average will overstate or understate extraction. In this sensitivity analysis, we allow recharge and, by extension, extraction to be higher or lower than the long-run average by the same percent that rainfall is higher or lower than the average of the 100 prior years. Again, we find quantitatively similar results.

Column 7 uses an estimate of marginal extraction costs that includes annualized well drilling costs. The U.S. Department of Agriculture (USDA) reports that aggregate annual expenditure on drilling is approximately 12% of aggregate annual pumping costs (USDA 2024); in this sensitivity analysis, we increase annual extraction costs by 12% and continue to find quantitatively similar results.

Appendix Table A.9 reports further sensitivity analyses using different subsets of aquifers. Column (2) uses only major aquifer systems. Column (3) includes aquifers with decision discount factor outside 0 to 1. Column (4) includes estimates for aquifers with less than 1% of land devoted to agriculture, and Column (5) includes aquifers with low agriculture and decision discount factors outside of 0 to 1. In each of these alternatives, the share of aquifers where extraction is inefficiently rapid is similar to the main estimates, at 61% to 65% for Ramsey discounting and 79% to 81% for market discounting, and non-optimal extraction generates a welfare cost of several trillion dollars.

F.4 Technology Parameters

Appendix Table A.5 assesses the robustness of the estimated model parameters ϕ and θ to different definitions of water extraction and profits, different data subsamples, and different specifications of the instruments. We report our estimates of the groundwater input share ϕ , $\theta = 1 - \alpha - \gamma$, and the implied return to land $\theta - \phi$.

Column 1 combines surface water and groundwater extraction in the measurement of water inputs. The estimate of the total water input share is statistically indistinguishable from our estimate using groundwater only. Using this alternative definition results in a water input share point estimate that is approximately 20% larger than our estimate based on groundwater only in Table 3, which is suggestive of slightly higher returns to the surface water component of irrigation. In both models, the capital and labor input shares are similar in sign, magnitude, and precision to our estimates that include only groundwater, leading to quantitatively similar estimates of θ .

Columns 2 and 3 define the profit outcome variable differently. Rather than using observational night lights luminosity, column 2 uses economic data on country-level agricultural GDP and column 3 uses model-based gridded estimates of total GDP adjusted for purchasing power (Kummu et al. 2018; Prabhakar et al. 2023) both of which we re-weight to the aquifer-level. We find that the point estimates of the input shares are slightly larger than our main estimates in Table 3 when we use agricultural GDP as the outcome variable and slightly smaller when we use gridded GDP as the outcome variable, though in both cases the point estimates are statistically indistinguishable from the main estimates. In particular, the larger groundwater effect on agricultural GDP relative to our benchmark estimate of 0.10 using night lights and 0.03 using gridded GDP seems to accurately capture the importance of groundwater in agriculture, relative to all sectors. The smaller gridded GDP estimate in column 3 may also reflect the purchasing power adjustment and the smaller sample size due to the fewer number of years in which this data set is available.

Columns 4, 5, and 6 use the same variables as our main estimates, but different data subsamples. Column 4 excludes shallow-water aquifers from the data; aquifers classified as shallow-water by hydrological agencies 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). The results on this sample are statistically indistinguishable from our main estimates, though the smaller sample size reduces precision. Columns 5 and 6 split

the data into two time periods corresponding to the two different satellites that measured night lights luminosity, our outcome variable. Column 5 includes 2013-2016 (i.e., the “VIIRS” years) and column 6 includes 2003-2012 (i.e., the “DMSP” years); the full night lights data set is harmonized across the two time periods, so we expect, and find, similar results across the subsample that includes much of our data (i.e., the DMSP years) and noisier estimates for the three-year VIIRS sample.

Column 7 replicates our main estimates using instruments in logs, rather than inverse hyperbolic sine. The inverse hyperbolic sine transformation of average temperature and total precipitation in the instruments used to generate our main estimates of the input shares allows us to include approximately 200 aquifers with negative or zero average temperatures, but is sensitive to the choice of units. Log-transforming temperature and precipitation generates instruments that are robust to different measurement choices, at the cost of excluding approximately 200 observations. We find estimates of the water, capital, and labor input shares that are virtually identical using either set of instruments.

Column 8 uses the same groundwater inputs, capital inputs, and instruments as our main analysis, but replaces our measure of labor with gridded population estimates. Gridded population data are available every five years; we linearly interpolate intervening years. Using this alternate measure of labor, our estimate of the groundwater input share is statistically indistinguishable from our main estimate, though none of the model parameters are statistically significant using this noisier labor measure.


Column 9 again uses the same production function inputs and instruments as our main analysis, but excludes the temperature and precipitation controls. The groundwater input share is almost identical to our main estimates, suggesting that climate is not driving our main result. The labor and capital shares and, by extension, the land share are less precisely estimated.

F.5 Policy Functions

Appendix Figure A.5 shows optimal water consumption as a function of aquifer fill (the “policy function”) for each of the four case study aquifers. 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 policy function shifts down as β increases: a more patient agent uses less water for a given aquifer fill. In each figure, the dashed horizontal line is the aquifer recharge. Water use equals recharge where a policy function intercepts this line. The long-run fill of an aquifer is the fill level at which water consumption equals recharge. The more patient an agent is, the higher will be the long-run fill level.

G Figures and Tables


Figure A.1: Patterns of Specific Yield and Capacity, 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).

Figure A.2: 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.

Figure A.3: 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.

Figure A.4: Marginal Cost and Marginal Revenue Curves for Other Case Study Aquifers



Notes: The horizontal axis shows gross extraction as a share of aquifer capacity. In each graph, the dashed vertical line shows observed gross groundwater extraction based on GRACE (Rodell et al. 2019). The decision discount factor β represents the value where marginal revenue equals marginal cost at observed water extraction (i.e., the β for which observed water extraction is optimal). Graphs show marginal costs for four discount factors: complete myopia, the decision discount factor, the Ramsey benchmark discount factor, and the market benchmark discount factor. With complete myopia, $\beta = 0$ in all four graphs. The market discount factor is $\beta = 0.98$ in all four graphs. In the Ogallala Aquifer, decision $\beta^d = 0.94$ and Ramsey $\beta^* = 0.96$. In the Indus Basin, decision $\beta^d = 0.71$ and Ramsey $\beta^* = 0.92$. In the North China Aquifer System, decision $\beta^d = 0.89$ and Ramsey $\beta^* = 0.91$.


Figure A.5: Policy Functions



Notes: This figure shows optimal water consumption as a function of aquifer fill (the "policy function") for the four case study aquifers. 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.


Figure A.6: Decision Discount Factors, Major Aquifers

(a) Spatial Patterns of Decision Discount Factors




Notes: The decision discount factor represents the value consistent with observed extraction. This figure maps the decision discount factor (β^d) for each component aquifer of the major aquifer systems. Case study aquifers are outlined in red.

Figure A.7: Estimated Relationship Between Discount Factor and Property Rights Interpretations: Other Case Study Aquifers



Notes: This figure displays the relationship between the property rights parameter λ and the discount factor β for the Ogallala, Indus, and North China Aquifer Systems, using the methodology from Section 6.4 and Appendix C.2. We estimate λ assuming β^* is given by the Ramsey or market discount factor.


Figure A.8: Paths of Resource Management Outcomes Under Different Discount Factors, Ogallala Aquifer




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 Ogallala Aquifer. Each line represents a prediction under a different discount factor. The solid blue line represents the current extraction path.

Figure A.9: Welfare Losses Due to Difference Between Estimated and Market Discount Factors

(a) Welfare Losses, by Aquifer



(b) Distribution of Welfare Losses



Notes: Panel A maps the welfare loss in each aquifer, measured as a percent of each aquifer's present value, due to the current extraction path deviating from the market benchmark. Panel B maps the distribution of these losses across aquifers.

Table A.1: Model Notation

Variable:	Meaning
i	Aquifer
t	Period
$a, a_i, a_{it} \in [0, 1]$	Amount of water available at beginning of a period
c_i	Aquifer-specific cost per unit of water $\times$ lift height
d_i	Pure rate of time preference (discount rate on utility)
e_i	Industrial electricity price
g_i	Growth rate of consumption per capita
$h_{it}, \underline{h}_i, \bar{h}_i$	Lift height, min lift height, max lift height of aquifer
k	Capital, chosen by agent
l	Labor, chosen by agent
p_i^k, p_i^l	Rental price of capital; price of labor
$r, r_i \in [0, 1]$	Recharge quantity of water for aquifer i
s_i	Surface land area above aquifer
v_c	Marginal groundwater revenue for crop c
$w, w_i, w_{it} \in [0, 1]$	Water consumption
$\mathbf{w}, \mathbf{w}^*$	Water consumption path, optimal water consumption path
$\bar{w}_{it}$	Average extraction
$\tilde{a}, \tilde{w}, \tilde{r}$	Not scaled by aquifer capacity: amount of water available, water consumption, recharge quantity
x	Land
y_i	Numeraire output good produced by agent i
$\bar{y}_{it}$	Remotely sensed nighttime lights; proxy for output y_{it}
z_i	Transformed productivity
H_c	Mean years of schooling in country c
V_i	Value of aquifer i
X_{it}	Temperature and precipitation
Z_{it}	Vector of instruments (lagged weather, input prices)
α	Share capital
$\beta, \beta_i^*, \beta_i^d$	Discount factor, benchmark discount factor, decision discount factor
γ	Share labor
ζ_t	Annual global productivity disturbance
η_i	Elasticity of marginal utility of consumption
θ	Cost share of water plus land
κ_i	Capacity
$\lambda_i \in [0, 1]$	Extent of well owners' control over water level
ν_i	Measure of the value of groundwater in aquifer i (e.g., GAEZ return to irrigation)
ξ	Pump efficiency
π_i	Profit when available water is a and water consumption is w
$\tilde{\pi}_i$	Profit, net of water costs
ρ	Physical constant reflecting flow rate, gravitational acceleration, and water density
σ_{ic}	Share of crop c in total agricultural production for aquifer i

(Continued on next page)

Table A.1: Model Notation (Continued from previous page)

Variable:	Meaning
ϕ_i, ϕ	Elasticity of output w.r.t. groundwater extraction (aquifer-specific; common component)
φ_i	Time-invariant aquifer productivity characteristics (soil quality; subsumes land)
χ_i	Aquifer-specific multiple of nighttime luminosity
ω_i	Welfare loss due to using non-benchmark discount factor
$\mathcal{C}_i$	Cost of extracting w units of water when initial water is a
$\mathcal{H}_i$	Topsoil thickness of aquifer
$\mathcal{T}_i$	Saturated thickness of aquifer
$\mathcal{V}$	Value of a consumption path on aquifer, discounted at β
$\mathcal{Y}_i$	Specific yield of aquifer
$\mathcal{Z}_i, \mathcal{Z}_{it}$	Total factor productivity

Table A.2: Data Sources

(1) Data Source	(2) Variables	(3) Unit	(4) Description	(5) Citation
GRACE	Δh_{it}	cm/year	Height of the change in water storage; used to compute change in lift height (m)	Rodell et al. (2019)
WHYMAP	$\tilde{r}_i$	mm/year	Mean annual aquifer recharge; input to aquifer dynamics model	BGR (2025)
	s_i	m^2	Surface land area; used to derive aquifer characteristics such as capacity	
WaterGAP	$\mathcal{H}_i, \mathcal{Y}_i$	m, unitless	Topsoil thickness and specific yield; used to compute aquifer capacity $\kappa_i = (\tilde{h}_i - \mathcal{H}_i)s_i Y_i$	de Graaf et al. (2015)
Fan et al. (2013)	$h_{i,2008}$	meters	Initial lift height; benchmark for water table depth in 2008	Fan et al. (2013)
Li et al. (2020)	$\bar{y}_{it}$	Index	Nighttime light intensity; proxy for aquifer-level economic output	Li et al. (2020)
Kummu et al. (2018)		USD	Gridded purchasing power-adjusted GDP; alternative proxy for output	Kummu et al. (2018)
GAEZ		tons, %	Yields of 26 crops under various conditions; combined with marginal returns to estimate $y'_i(\tilde{w}_i)$	FAO and IIASA (2024)
D’Odorico et al. (2020)		\$/m ³	Marginal value of irrigation water by crop; used to estimate marginal returns to groundwater	D’Odorico et al. (2020)
World Bank WDI		people	Population, employment share in agriculture; used to compute effective labor input	World Bank (2023)
	H_c	years	Mean years of schooling; adjusts for differences in human capital	
Penn World Table	K_{it}	USD	Capital stock; capital input in the production function	GGDC (2022)
ILO		USD	Wages; used in productivity estimation and calibration	ILO (2023)
IMF	r_c	%	Interest rates; used as proxy for cost of capital in production function	IMF (2023)
IEA	$p_{e,c}$	\$/kWh	Electricity prices; used to calculate marginal cost of groundwater extraction	IEA (2021)
Climate Research Unit	$precip_{it}, temp_{it}$	mm, °C	Climate data; used as instruments in production function and to estimate recharge	Harris and Jones (2019)
Global Hydrogeology Maps			Soil characteristics; informs recharge and aquifer structure	Gleeson et al. (2014)
Addicott et al. (2020)	β_i^*		Benchmark discount factor; used to compute welfare-optimal extraction path	Addicott et al. (2020)
HYDE / AquaKnow	IrrigArea _i	km ² , %	Irrigated area; used to estimate groundwater dependence of agriculture	Goldewijk (2017)

Table A.3: 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 X Porosity	0.001 (0.042)	-0.017* (0.009)	-0.010* (0.005)
Average Temperature, Lag 3 X 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) X 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 3. The outcome variables are the log of groundwater extraction, capital, and labor by aquifer, respectively. The instruments and precipitation and temperature controls are chosen as part of 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 (*).

Table A.4: 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	.108	.019	.731	.533	.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. 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 gross groundwater extraction in m^3 . Value is the aquifer-specific value of groundwater, measured by heterogeneity in the variable specified in the column heading; this gives us heterogeneous water input shares (ϕ_i). Column (1) uses the share of the aquifer that is cropped. Column (2) uses the share of the aquifer that is irrigated. Column (3) uses a binary indicator for any cropped land above the aquifer. Column (4) uses a binary indicator for any irrigation above the aquifer. Column (5) uses a binary indicator for agriculture as the dominant land use 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 (*).

Table A.5: Alternative Estimated Effects of Production Function Parameters

	(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), its interaction, 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 (Shallow Aq) 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.6: 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: For each statistic, we report the global aquifer crop value-weighted median (Column 1), the global aquifer GDP-weighted median (Column 2), and the global aquifer population-weighted median (Column 3). We report results using our decision discount factor, the benchmark Ramsey discount factor, and the benchmark market discount factor 0.98. The Ramsey discount factors are from [Addicott et al. \(2020\)](#). Panels B and C are in percentage points, i.e., 0.33, which represents 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.

Table A.7: 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. \(2020\)](#). d.f. is discount factor. Country characteristics are discussed in Appendix D.5. All models are estimated using OLS. Marginal returns and marginal cost are in $\$/m^3$, 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 (*).

Table A.8: 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%
Fraction 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%
Fraction 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					
Add surface water			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) summarizes our model estimates from the main text. Columns (2) through (7) report model results for different sets of model parameters. Column (2) uses a calibrated value of the land share from [Ryan and Sudarshan \(2022\)](#) in the production function instead of the value implied by estimating the capital and labor shares in equation (20). Column (3) uses production function parameters resulting from estimating equation (20) using all water (surface+groundwater) in production, rather than groundwater only; Table A.5 reports these production function parameters. Column (4) uses aquifer-specific production function parameters resulting from allowing returns to water to vary across aquifers; Table A.4 reports these production function parameters. Column (5) uses estimates of marginal extraction cost that incorporate heterogeneity across aquifers in groundwater 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. All dollar values are in 2023 U.S. dollars.

Table A.9: 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%
Fraction 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%
Fraction 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 factor included			Yes		Yes
Low agriculture aquifers				Yes	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) summarizes our model estimates from the main text. Column (2) uses only major aquifers. Column (3) uses all aquifers, not just those with discount factor between 0 and 1. 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%. Column (4) includes low agriculture aquifers. Column (5) includes low agriculture aquifers, and aquifers regardless of discount factor. All estimates are based on the model in Section 3. All dollar values are in 2023 U.S. dollars.