Are We Consuming Too Much Groundwater?*

Brian Greaney
Univeristy of Washington

Joseph S. Shapiro
UC Berkeley
and NBER

Katherine R. H. Wagner
Univeristy of British Columbia

March 2025

Abstract

We study the optimality of groundwater extraction from most of the world's 3,400 aquifers. Many prominent aquifers have declining levels because human use exceeds recharge. For each aquifer, we use remote sensing and administrative data to estimate a dynamic model of water extraction, recover the discount factor that rationalizes observed groundwater extraction, and compare it against normative and market benchmarks. About two-thirds of the world's aquifers are extracted suboptimally; the rest have approximately optimal extraction rates. Suboptimal groundwater extraction creates $3.4 trillion in global present-value welfare costs, a large magnitude but less than 3% of aquifer present value. We calculate much larger losses from policies that subsidize water use or guarantee users fixed indefinite water quantities.

JEL Codes: E22, H23, H43, Q25
Keywords: Groundwater extraction, dynamic programming, renewable resource management, environmental taxes, tragedy of the commons, agricultural production

*We thank Eli Fenichel, Larry Karp, Matt Kotchen, and participants in seminars at the Paris School of Economics, Queen's, Simon Fraser University, and Temple University 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. Greaney: bg385@uw.edu; Shapiro: joseph.shapiro@berkeley.edu; Wagner: katherine.wagner@ubc.ca.

1 Introduction

Humans are rapidly depleting many of Earth's natural resources. People harvest timber, extract minerals, emit greenhouse gases, and drive species to extinction. Economists have long debated whether prevailing rates of natural resource use are optimal (Dasgupta and Heal 1979; Arrow et al. 2004; Dasgupta 2021). Hotelling (1931)'s description of the problem nearly a century ago is still relevant:

> 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 in many settings, some net depletion of a resource is socially optimal. Renewable resources like fisheries or timber have natural recharge rates, and optimal policy may involve depleting a resource to a lower steady-state, gradual depletion followed by gradual replacement, or variable extraction for consumption smoothing. Similarly, any consumption of nonrenewable resources like copper or other mineral deposits represents net depletion, but nonzero use of nonrenewable resources in many settings represents optimal policy.

We study optimal extraction rates and their relationship to actual extraction rates in the setting of groundwater extraction. Groundwater supplies 40% percent of the water used for global agriculture and 50% percent of the world's drinking water (Rodell et al. 2018; Ojha et al. 2018). Numerous aquifers around the world face rapid rates of extraction. In response, governments have begun debating and implementing laws to regulate groundwater extraction.

We combine remote sensing and administrative data on economic outputs and inputs, climate, and geology to compare actual and optimal extraction paths for each of the world's aquifers. We analyze data describing changes in groundwater stocks and flows over a decade for over 2,700 of the world's 3,400 aquifers that, combined, account for 93% of the Earth's surface.

We interpret these data through the lens of a classic dynamic model of renewable resource extraction. A representative agent in each aquifer produces output using a production technology with decreasing marginal product of water. The agent chooses water extraction to maximize the present discounted value of all future profits. Water extraction faces an intertemporal trade-off: consuming water today increases current profits at the expense of higher future extraction costs, since current extraction lowers the future water table. The representative agent's water consumption plan depends on the production technology, aquifer characteristics, and the discount

factor.

Connecting the model with data requires measuring three main terms. We infer groundwater extraction from remote sensing data. We calculate 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 estimate the willingness to pay for groundwater from production functions. Since groundwater use may be jointly determined with production decisions, our estimates of the production technology instrument for groundwater using lagged interactions of weather and geology.

We find that groundwater extraction is substantial, but that marginal costs of extraction are substantially below the willingness to pay for water from most aquifers. Overall, we estimate the highest willingness to pay in the same high-value agricultural areas. Estimated marginal extraction costs are a small fraction of median willingness to pay. The wedge between willingness to pay for groundwater and marginal costs helps explains why extraction exceeds optimal levels in many aquifers, though optimality of extraction depends both on current and future costs of today's extraction.

We use the calibrated model for three purposes. First, for each aquifer, we solve the dynamic model to recover the estimated discount factor that rationalizes observed groundwater extraction. We compare this estimated discount factor against two benchmarks—a normative discount factor (Ramsey 1928; Addicott et al. 2020); and long-term interest rates on government debt. Comparing the effective discount factor that rationalizes actual extraction against these benchmarks helps assess the optimality of current extraction. The estimated discount factor summarizes the impatience with which the aquifer is being extracted. If open access conditions, weak property rights, or other institutional failures make consumers fail to internalize effects of their water use on future extraction costs, the estimated discount factor may not equal the benchmark discount factor.

We find that two-thirds of the world's aquifers have extraction rates more rapid than is optimal, in that estimated discount factors are below benchmarks. This includes two US aquifers that have received substantial attention, the Central Valley Aquifer System in California and the Ogallala Aquifer in the Midwestern US. Other aquifers, such as the the North China Aquifer System in Asia, have approximately optimal extraction, despite experiencing net depletion. In the median aquifer, extraction rates exceed the optimum by 0.14% to 0.51% of aquifer capacity. Indeed, optimal extraction would allow the median global aquifer to refill completely in the long-run, above the steady state fill level of 94%.

Second, we study the welfare consequences of current extraction paths. For aquifers with excessively rapid extraction, this could be interpreted as the welfare cost of weak property rights or poor resource management. We also calculate the time-invariant tax rate on groundwater which

would make extraction in each aquifer close to the optimal path; we interpret this as the optimal tax rate.

The welfare analysis provides several informative results. Optimal extraction paths for many aquifers involve decades of net extraction, then a steady state with a high share of capacity filled. The global, present value welfare loss due to differences between current extraction paths and the optimal path we calculate is $3.4 trillion. While that is a large value, it represents under 3% of the present value that we estimate for all global aquifers. We also find that the aquifer-specific tax rates needed to implement the optimal extraction paths are moderate to large, with the median aquifer requiring an optimal tax equal to 18% of willingness to pay.

Third, we study the consequences of two other suboptimal but potentially realistic consumption paths—extracting groundwater at a constant rate equal to current consumption, and extracting groundwater myopically. Although neither scenario describes an optimal consumption path under a constant discount factor, both could be consistent with some existing water policies. Compared to current extraction paths, both of these scenarios lead to lower long-run fill, higher long-run costs, and, hence higher welfare losses. With constant extraction, long-run fill in the median aquifer falls by 3%. Long-run extraction costs double relative to the current extraction path with a constant discount factor. For the median aquifer, myopia increases extraction costs five-fold relative to the current path with a constant discount factor. The difference occurs because myopia decreases long-run fill much more. Both of these scenarios involve greater welfare losses than the benchmark discount factor discussed above because the benchmark discount factor scenario above assumes that extraction responds dynamically to changes in extraction costs, while these scenarios instead assume either constant consumption or use of a discount factor even lower than the benchmark. With myopia, the median welfare loss across all aquifers is nine times larger than the welfare loss under observed extraction. One broad qualitative takeaway from these comparisons is that existing weak property rights for groundwater generate modest to moderate welfare costs, on the order of several percent of aquifer value. Subsidies fixing the price of water or laws fixing water extraction in physical units generate many times greater social welfare costs. While we believe this comparison sounds intuitive, we do not believe prior research has shown it.

We build on several literatures. We provide the first empirical analysis of actual versus socially optimal extraction rates for most of the world's aquifers. This may also be the first analysis comparing actual versus optimal extraction paths separately for each global endowment of a non-market natural resource. Our general approach of interpreting remote sensing and administrative data for many global endowments of a natural resource through the lens of a dynamic programming model could in principle be applied to many other types of resources. A longstanding literature develops and extends Hotelling (1931)'s theory of natural resource extraction under many conditions. Estimating actual versus optimal resource extraction, however, poses challenging empirical

demands. Such analysis requires knowing information on resource use and recharge, demand for resource use, and a dynamic framework to compare these forces over time. Empirical analysis of optimal resource extraction often focuses on a single location with sufficient data (e.g., Fenichel et al. 2016). Research has developed alternative approaches to studying welfare and resource use, such as measuring net capital depletion for resources with market prices, but important resources suffer from missing markets or otherwise lack appropriate market prices (Carleton et al. 2024a,b). Timmins (2002) undertakes a version of this analysis for 13 California municipalities. Ryan and Sudarshan (2022) study groundwater extraction for villages in Northwest India, primarily sharing one aquifer. We build on Carleton et al. (2024a), who undertake a rich spatial analysis of global water flows, by incorporating intertemporal trade (dynamics) and analysis of optimal policy and welfare, but we also abstract from spatial (intranational and international) trade.

Our analysis of actual and optimal extraction also builds on a literature using remote sensing data to study resource extraction. This literature includes analysis of decentralization, trade, and deforestation (Burgess et al. 2012; Farrokhi et al. 2024), wetland regulation and depletion (Aronoff and Rafey 2024; Greenhill et al. 2024), and incentives for fallowing and land protection (Aspeland and Russo 2024). Scientists increasingly use remotely sensed data to measure groundwater trends but less often economics (Famiglietti et al. 2011; Richey et al. 2015a).

We also contribute to a more theoretical literature on externalities from natural resource extraction. Unlike some of this work, we interpret dynamic decision-making by estimating the discount factor that rationalizes actual extraction decisions, whereas many models instead study whether common pool or single-agent decision-making best characterizes resource extraction. We show a close relationship between the literature's and our interpretation, though we focus on the discount factor interpretation given its close relevance to macroeconomics and ease of providing a continuous measure of extraction practices. We also quantify the dynamic externality associated with current groundwater extraction and optimal time-invariant policy. Research dating back to Coase (1960) and Ostrom (1990) discusses how suboptimal extraction of resources with common pool management, including fisheries and groundwater, may arise because benefits from extraction are private while costs are spread across either current or future users. Several papers conduct empirical studies of how different management regimes can help internalize these externalities. Costello et al. (2016) show how competition for limited fish stocks results in the long-run decline of fisheries, while Burlig et al. (2021) studies how groundwater extraction responds to different management strategies, including price-based policies. Ferguson (2024) quantifies the welfare effects of policies to reduce distortions in surface water markets.

Before proceeding, a few clarifications may be useful. Calculating optimal extraction paths in a dynamic setting requires a view on the planner's true discount rate. While we consider both normative and market benchmarks for this rate, we can only quantify optimal outcomes up

to an assumption about the rate. Additionally, for tractability, we calculate the optimal time-invariant groundwater taxes. While we calculate that these are numerically close to optimal policy, computational obstacles prevent us from recovering the fully dynamic optimal policy. Finally, we abstract from spatial interactions and assume that each aquifer has independent management and factor prices. While existing work explores international trade, we focus on intertemporal trade (dynamics) and leave their combination for future work.

The rest of the paper is organized as follows. Section 2 provides background. Section 3 describes the model. Section 4 discusses the data. Section 5 explains model estimation and calibration. Section 6 presents results. Section 7 concludes.

2 Background on Water Resources

We discuss several types of water. Groundwater sits in underground aquifers formed over millennia. Groundwater represents one part of total terrestrial water storage, which includes water stored on or in land, moisture in surface soils, vegetation, snow, rivers, lakes, and wetlands. Surface waters like rivers lie above the earth's surface. Groundwater accounts for 20% of global freshwater. Agricultural irrigation accounts for 70% of global groundwater withdrawals; other uses like drinking water and manufacturing account for the remaining 30% (Rodell et al. 2018; Ojha et al. 2018). In recent years, groundwater use has grown overall and as a share of total water consumption because it increasingly substitutes for more variable surface water. Scientists worry that this trend is unsustainable, since groundwater is depletable (Siebert et al. 2010; Rodell et al. 2018).

Most of the earth's land surface stores groundwater. A consortium of global hydrological agencies identifies 3,400 distinct global aquifers (BGR 2025), ranging from major basins like the Central Valley Aquifer System in California to small and shallow aquifers. Aquifer systems cover nearby aquifers sharing geology and hydrology. Individual aquifers within some systems have interconnections through underground lateral water flow and may exchange groundwater slowly (Miller 2003).

Part of our analysis focuses on four case study aquifer systems: the Central Valley Aquifer System in California, the Indus Basin in South Asia, the North China Aquifer System, and the Ogallala Aquifer in the Midwestern US. We choose these four given public attention focused on them and their high levels of population density, irrigation demand, and groundwater use of all Earth's aquifer systems (Richey et al. 2015a; Borunda 2022). They are also among 37 major aquifer systems that hydrologists highlight. Appendix 2 briefly discusses the background of each case study.

Explaining aquifer terminology and function clarifies our analysis. Aquifer dynamics depend

on annual inflows and outflows, also called recharge and extraction, and on stocks. An aquifer's storage volume or reserves equal its current water, and its capacity is its maximum possible water storage. An aquifer's level represents the current top of its water table. When annual extraction from an aquifer exceeds its annual recharge, the aquifer's reserves fall and thus its level declines. Declining reserves increase the marginal cost of water extraction because users must lift groundwater further, requiring more energy. Aquifer recharge depends on rainfall, surface water runoff, and geology. A given amount of rainfall recharges an aquifer more in porous rock, while rainfall flows over non-porous rock as surface water runoff (e.g., streams). Water applied to agriculture may also recharge aquifers gradually as it flows through soil and rock.

Figure 1 illustrates an aquifer. Section 4.1 uses the relationship that an aquifer's capacity is the product of its vertical height, area, and specific yield. Groundwater fills only part of an aquifer's area; rock, soil, air, and other material fill the rest. The vertical extent of groundwater mixed with rock, soil, and air is the saturated thickness. Groundwater extraction ultimately depends on the specific yield—the fraction of the saturated thickness that users can drain. Specific yield varies within an aquifer depending on the type of rock, specifically the attraction of rock molecules to groundwater. Specific yield is lower further below the earth's surface (Richey et al. 2015a; Sutanudjaja et al. 2018). In calculating aquifer storage volume and capacity, we follow the hydrology literature and consider the portion of the aquifer that can be extracted (Richey et al. 2015a).

Consider a simple example. Imagine a cubic aquifer that begins 20 meters beneath the ground, extends up to 10 meters beneath the surface, and has width and length of 10 meters each. This aquifer has volume of 1000 m^3 ($= 10^3$). If the saturated thickness is 5 meters, then groundwater only occupies area from 20m to 15m below the ground; the top half of the aquifer has no water, so the aquifer has a storage volume of 500m^3 ($= 5 \cdot 10^2$). If the aquifer has specific yield of 0.5, then only half this value could ever be extracted, or 250 m^3.

Aquifer capacity could change over time. Groundwater depletion lowers water pressure upward on the Earth's surface, which can make the vertical distance above the water table decline ("subsidence"). Pollution, through saltwater intrusion or organic contamination, can reduce an aquifer's ability to store freshwater. We follow common practice in assuming that an aquifer's capacity is fixed over time because while subsidence or pollution can decrease aquifer capacity, these changes are typically small as a share of aquifer capacity (de Graaf et al. 2015; Ojha et al. 2018; Sutanudjaja et al. 2018).

Groundwater extraction requires pumps drilled into the aquifer's saturated thickness to draw water to the surface. Diesel or electricity typically power groundwater pumps. Energy costs for pumping groundwater can represent half of agricultural energy costs (Pradeleix et al. 2015). Pump costs depend on lift height, energy prices, and pump efficiency. Even an aquifer filled to capacity

has positive lift height since groundwater must be drawn through topsoil.

3 Model

3.1 Model Description

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

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

The $\min\{\cdot\}$ function ensures that available water cannot exceed aquifer capacity. Consumption is limited to available water plus recharge, since the agent cannot consume water when the aquifer is empty: $w \leq a + r_i$.

Agent i produces a numeraire output good using the following technology:

$$y_i = \mathcal{Z}_i k^{\alpha} l^{\gamma} w^{\phi} x_i^{1-\alpha-\gamma-\phi} \tag{2}$$

Here $\mathcal{Z}$ represents total factor productivity, k is capital, l is labor, and x is land. The agent chooses k, l, and w. We assume that the agent's choices do not affect capital or labor prices, that the agent owns the land, and that current technology, preferences, and non-resource factor prices persist in future years. Appendix C shows that these assumptions imply the following profit-maximizing output for a given water consumption choice w:

$$y_i(w) = z_i w^{\frac{\phi}{\theta}} \tag{3}$$

where $z_i = [\mathcal{Z}_i x_i^{1-\alpha-\gamma-\phi}(\alpha/p_i^k)^{\alpha}(\gamma/p_i^l)^{\gamma}]^{1/(1-\alpha-\gamma)}$ is a function of variables and parameters and $\theta \equiv 1 - \alpha - \gamma$. Appendix C also shows that profit, net of water costs, is

$$\tilde{\pi}_i(w) = \theta z_i w^{\frac{\phi}{\theta}} \tag{4}$$

The cost of water extraction reflects engineering forces. Let $\underline{h}_i$ and $\overline{h}_i$ denote the minimum and maximum lift height of aquifer i, 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 in

which initial water is a and water consumption is w, the mean lift height is

$$h_i(a, w) = \underline{h}_i + \left(1 - a + \frac{w - r_i}{2}\right)(\overline{h}_i - \underline{h}_i) \tag{5}$$

The cost of extracting w units of water when initial water is a is

$$\mathcal{C}_i(a, w) = c_i h_i(a, w) w \tag{6}$$

where c_i is an aquifer-specific cost per unit of water $\times$ lift height, reflecting pump efficiency and energy prices. Profit when available water is a and water consumption is w is then

$$\pi_i(a, w) = \tilde{\pi}_i(w) - \mathcal{C}_i(a, w) \tag{7}$$

The representative agent discounts the future at rate β, and solves the following Bellman equation:

$$V_i(a; \beta) = \max_{w \leq a + r_i} \pi_i(a, w) + \beta V_i(\min\{a + r_i - w, 1\}; \beta) \tag{8}$$

While the value function depends on all the model's parameters, our notation makes its dependence on the discount factor β explicit given the discount factor's importance in our analysis. The second term on the right-hand side of (8) represents the continuation value, which depends both on the agent's current choice of water consumption w and recharge r_i. The agent faces a tradeoff—increasing water consumption may raise both current profits and future extraction costs, thus decreasing profits in the future. This tradeoff is sometimes called the dynamic externality, since today's extraction raises the cost of future resource extraction. We let $w_i^*(a; \beta)$ denote the optimal policy function of (8).

The necessary first-order condition (FOC) of (8) is

$$\partial_w \pi_i(a, w) \geq \beta \partial_a V_i(\min\{a + r_i - w, 1\}; \beta) \tag{9}$$

which holds with equality for an interior solution where $w < a + r_i$. The FOC can also be written as

$$\partial_w \tilde{\pi}_i(w) \geq \partial_w \mathcal{C}_i(a, w) + \beta \partial_a V_i(\min\{a + r_i - w, 1\}; \beta). \tag{10}$$

Equation (10) 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 costs. The marginal cost on the right-hand side includes two terms—the current marginal cost $\partial_w \mathcal{C}_i(a, w)$, and the discounted marginal dynamic cost, $\beta \partial_a V_i(\min\{a - w + r_i, 1\}; \beta)$.

While equation (8) characterizes the value of the optimal consumption path, our framework

can also value arbitrary, sub-optimal consumption paths. Let $\mathbf{w} = \{w_{it}\}_{t=0}^{\infty}$ represent an arbitrary water consumption path that satisfies 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, equals the present discounted value of future profits:

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

where $\{a_{it}\}_{t=0}^{\infty}$ is determined by the law of motion (1). Let $\mathbf{w}_i^*(a_0; \beta)$ denote the optimal water consumption path when initial fill is a_0 and the discount factor is β. The value of the optimal water consumption path corresponds to the solution to (8):

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

3.2 Welfare Analysis

3.2.1 Welfare Loss from Suboptimal Discounting

We estimate the welfare loss from actual rather than optimal extraction paths by calculating the difference in the aquifer's profit between the two scenarios.

Section 4.4 discusses the benchmark discount factors β_i^*, calculated from Ramsey discounting or the long-run rate on government debt, which represent estimates of the planner's true discount factor. We solve the agent's problem (8) to obtain the maximum value of aquifer i with observed initial fill a_{i0}, $V_i(a_{i0}; \beta_i^*)$, using the benchmark discount factor, β_i^*. The estimated discount factor β_i^{e}, which rationalizes observed water use, will not generally equal β_i^*. Failures of property rights, common pool management, or political discount rates may lead to inefficiently rapid discounting, or $\beta^e < \beta^*$. Alternatively, distortions in other factor markets could inefficiently depress water use, giving $\beta^e > \beta^*$. Solving (8) with the estimated discount factor yields the estimated water consumption path $\mathbf{w}_i^*(a_{i0}; \beta_i^{\mathrm{e}})$. The value of the estimated water consumption path according to the benchmark discount factor β_i^* is $\mathcal{V}_i(\mathbf{w}_i^*(a_{i0}; \beta_i^{\mathrm{e}}), a_{i0}; \beta_i^*)$. Since we expect that estimated and benchmark discount factors will not equate ($\beta_i^{\mathrm{e}} \neq \beta_i^*$), the estimated value of the aquifer according to the benchmark β_i^* will be less than the maximum possible value: $\mathcal{V}_i(\mathbf{w}_i^*(a_{i0}; \beta_i^{\mathrm{e}}), a_{i0}; \beta_i^*) < V_i(a_{i0}; \beta_i^*)$.[1] The welfare loss due to water use being chosen using the discount factor β^{e} rather than β_i^*, expressed as a fraction of the maximum value, is

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

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

Section 6 reports estimates of ω_i for each aquifer.

3.2.2 Welfare Gains from Taxing Water

We also analyze how a tax or subsidy τ on groundwater extraction affects water use and welfare. We assume that the representative agent receives the water tax as a lump sum rebate, but that the agent takes this rebate as parametric when choosing water consumption, and thus it does not distort consumption choices. A representative agent facing water tax τ solves

$$V_i(a;\beta,\tau) = \left\{ \max_{w \leq a+r_i} \pi_i(a,w) - \tau w + \beta V_i(\min\{a-w+r_i, 1\};\beta,\tau) \right\} + \tau w \tag{13}$$

Let τ_i^* denote the welfare-maximizing water tax:[2]

$$\tau_i^* = \underset{\tau}{\text{argmax}}\ \mathcal{V}_i(\mathbf{w}^*(a_{i0};\beta_i^{\text{e}},\tau), a_{i0};\beta_i^*)$$

Optimally taxing water undoes a fraction

$$\frac{\omega_i - \omega_i(\tau_i^*)}{\omega_i}$$

of the welfare loss from non-social discounting.[3]

3.3 Recovering Discount Factors

Appendix D presents the full computational algorithm we use to solve the agent's problem and recover the estimated discount factor β_i^{e} for each aquifer i; we summarize here. Given an arbitrary

[2]Since water tax revenue is rebated, the value of an arbitrary water consumption path $\mathbf{w}$ when initial water level is a_0 and water is taxed at rate τ is

$$\sum_{t=0}^{\infty} \beta^t[\pi_i(a_{it},w_{it}) - \tau w_{it} + \tau w_{it}] = \sum_{t=0}^{\infty} \beta^t \pi_i(a_{it},w_{it}) = \mathcal{V}_i(\mathbf{w},a_0;\beta)$$

We define $\mathbf{w}^*(a;\beta,\tau)$ as the solution to (13) and

$$\omega_i(\tau) = \frac{V_i(a_{i0};\beta_i^*,\tau) - \mathcal{V}_i(\mathbf{w}^*(a_{i0};\beta_i^{\text{e}},\tau), a_{i0};\beta_i^*)}{V_i(a_{i0};\beta_i^*,\tau)}$$

as the welfare loss due to the representative agent choosing water use according to the discount factor β^{e} instead of β_i^*, given tax τ.

[3]The tax authority could weakly increase welfare by setting a time-varying tax rate. For tractability, our empirical application studies time-invariant optimal taxes, since allowing dynamic taxation makes the agent's dynamic problem time-varying and introduces additional choice variables. We calculate the welfare loss between the optimal extraction path and the extraction path that the time-invariant taxes introduce, which represents the potential welfare gain due to using the time-invariant rather than dynamic optimal taxation. We find that the unrealized welfare gain due to our restriction of time-invariant taxation is small.

candidate discount factor β, we solve the agent's dynamic water use problem (8). This solution reveals water use at the aquifer's current observed water level, $w_i^*(a_{i0};\beta)$. We use bisection to find an estimated discount factor $\beta_i^e \in [0,1)$ such that the model-implied water use $w_i(a_{i0};\beta_i^e)$ equals current observed water use w_i.[4] We interpret β_i^e as the discount factor which rationalizes observed water consumption.

The estimated policy function $w_i^*(a;\beta_i^e)$ describes how much water the representative agent of aquifer i consumes given beginning-of-period stock a and the estimated discount factor β_i^e. Given this policy function, we iterate forward on the law of motion (1), beginning with the observed stock in data a_{i0}, to obtain model-predicted paths of water use, 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 $T = 1000$ periods in the future. We verify ex-post that water levels in all aquifers have converged by this period.

3.4 Interpreting β

The discount factor β shapes water consumers' impatience. Lower values of β both value current profits and discount future extraction costs relatively more. Optimal water use at a given aquifer level a decreases in β: a more patient agent (higher β) consumes less today, reflecting greater sensitivity to future extraction costs. Our main analysis estimates the discount factor for each aquifer that rationalizes observed water consumption. For example, in an aquifer where marginal current profit $\partial_w \pi_i(a,w)$ is low and marginal dynamic cost $\partial_a V_i(a';\beta)$ is high, we infer a low discount factor (see the FOC equation 9). Comparing aquifer-specific discount factors versus benchmark values helps assess the optimality of observed water consumption behavior.

A common alternative approach for rationalizing observed resource extraction emphasizes that property rights shape whether consumers internalize the costs of their extraction (Karp 2017). Intuitively, a single agent with full present and future control over a resource will fully internalize the future effects of extraction decision. In contrast, atomistic "common pool" consumers ignore dynamic effects of extraction decisions. To explore the relationship between our "discount factor" approach and the alternative "property rights" approach, we extend our baseline model to incorporate property rights. Suppose each aquifer has a unit mass of infinitesimally small agents who each extract water from their own wells. The law of motion for a well's water level depends on both the well owner's own extraction w_{it} and average extraction $\overline{w}_{it}$, which atomistic agents take as given:

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

[4] If no $\beta_i^e \in [0,1)$ exists that matches observed water use, we exclude the aquifer from our analysis. This occurs in 20% of cases (14% of which model-implied water use is always too low, and 6% in which it is always too high).

The term $\lambda \in [0,1]$ measures the extent of a well owner's control over the water level. The value $\lambda = 0$ corresponds to a common pool, where individual agents have no control over the water level and so take it as given. The term $\lambda = 1$ corresponds to a single agent who fully internalizes the effects of their extraction. The agent's problem is then

$$V_i(a;\beta,\lambda) = \max_{w \leq a + r_i} \pi_i(a,w) + \beta V_i(\min\{a + r_i - \lambda w - (1-\lambda)\overline{w}_i(a), 1\}; \beta, \lambda) \tag{14}$$

which has optimal solution $w_i^*(a;\beta,\lambda)$. We focus on symmetric equilibria in which $\overline{w}_i(a) = w_i^*(a;\beta,\lambda)$. Solving (14) involves finding a fixed point in which, given $\overline{w}_i(a)$, the optimal policy function is $w_i^*(a) = \overline{w}_i(a)$. Appendix D describes our computational algorithm for solving this fixed point.

We analyze the relationship between the discount factor β and the property rights parameter λ as follows. First, we solve (8) (which corresponds to (14) with $\lambda = 1$) for a grid of discount factors β^n, $n \in \{1, ..., N\}$. Each discount factor predicts a different current water use $w_i^*(a^{\text{data}}; \beta^n)$. For each predicted water use, we find the λ for which predicted water use equals $w_i^*(a^{\text{data}}; \beta^n)$ when the discount factor is set to a benchmark value β^*.[5] Denote the property rights value that maps to β^n by λ^n:

$$w_i^*(a^{\text{data}}; \beta^*, \lambda^n) = w_i^*(a^{\text{data}}; \beta^n).$$

Figure 2 shows the relationship between β and λ for the Central Valley aquifer. This figure demonstrates that various observed water consumption levels can be rationalized either by property rights or the discount factor.

4 Data

This section describes geographic and administrative data we use to calibrate and apply this model, and Section 5 explains precisely how we use these data. Appendix A provides additional details.

We use a few conventions throughout. For gridded datasets, we take an unweighted average of gridpoints within each aquifer $\times$ year. For data reported by country rather than aquifer, we calculate the share of each aquifer in each country, then weight from country to aquifer level using these shares as weights. We deflate to 2023 dollars using the US GDP implicit price deflator (FRED 2024). Given aquifer capacity κ_i, for any model variable x_i, we use the notation $\tilde{x}_i \equiv x_i \kappa_i$. In other words, we let $(\tilde{a}, \tilde{r}, \tilde{w})$ denote variables from the model not scaled by aquifer capacity. For example, w in the model represents water withdrawals as a share of aquifer capacity, while $\tilde{w}$

[5] For this exercise, we use the Ramsey discount factor.

in the data represents water consumption in m^3.

Appendix A.1 describes additional steps we take to clean data. This includes imputing aggregate administrative data that are missing for a small number of countries where government records have low quality, excluding aquifers in Antarctica, and calculating recharge from ranges provided for each aquifer. The final analysis data set contains approximately 30,000 observations on 2,773 aquifers over the years 2006 to 2016.

4.1 Data: Groundwater

We use several data sources to calculate groundwater flows and stocks for each aquifer. The Gravity Recovery and Climate Experiment (GRACE) provides annual data on the change in the height water must be lifted, Δh_{it}. 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 2016, letting us observe the full calendar year. Due to using up to three lags in regressions, we calibrate the model and set up other datasets for the years 2006-2016. Appendix A.1 describes several steps we take to clean the GRACE data, including applying a scaling factor to reduce noise, calculating annual changes, separating groundwater from total water, and linking grid points to aquifers.[6]

We obtain additional aquifer attributes from the World-wide Hydrogeological Mapping and Assessment Programme (WHYMAP), a repository that harmonizes groundwater data from several hydrological research groups (BGR, 2025). WHYMAP provides data on the surface land area s_i above each aquifer and mean annual recharge $\tilde{r}_i$. In the absence of annual recharge data r_{it}, we assume that the annual recharge of an aquifer equals its mean $\tilde{r}_i$ over the years 2006-2016. This use of average recharge echoes similar practice in some global hydrological models (de Graaf et al. 2015).

We obtain other aquifer attributes from WaterGAP (de Graaf et al. 2015), a global hydrological model. WaterGAP reports gridded data on the saturated thickness of each aquifer in the year 2013, along with time-invariant measures of topsoil thickness $\mathcal{T}_i$ and specific yield $\mathcal{Y}_i$. We calculate each aquifer's volume in 2013, $\tilde{a}_{i,2013}$. Since one aquifer may have different types of rock with different specific yields, we follow Richey et al. (2015a) and use the minimum specific yield measured across gridpoints within an aquifer to account for the most binding constraint on groundwater flow.[7]

[6] GRACE directly 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 over many meters of porous rock. We divide by specific yield to convert the change in the height of water extracted from GRACE to the change in lift height.

[7] Richey 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%. Since specific yield declines with depth below the surface of

We use data on lift height $h_{i,2008}$ for each aquifer in the year 2008 from Fan et al. (2013). They calculate lift height data from 2,200 global observations of water table depth, supplemented with a groundwater model.

We obtain additional aquifer attributes by applying model and ancillary assumptions to these data. These data report change in aquifer lift height, recharge, surface area, topsoil thickness, specific yield, volume, and baseline lift height $(\Delta h_{it}, \tilde{r}_i, s_i, \mathcal{T}_i, \mathcal{Y}_i, \tilde{a}_{i,2013}, h_{i,2008})$. We calculate additional aquifer attributes as follows:

$$\Delta \tilde{a}_{it} = s_i \Delta h_{it} \tag{15}$$

$$\tilde{w}_{it} = \tilde{r}_i - \Delta \tilde{a}_{it} \tag{16}$$

$$\tilde{a}_{it+1} = \tilde{a}_{it} + \Delta \tilde{a}_{it} \tag{17}$$

$$\kappa_i = v_i s_i \mathcal{Y}_i \tag{18}$$

$$h_{it+1} = h_{it} + \Delta h_{it} \tag{19}$$

Equation (15) calculates the annual change in fill by implicitly assuming that aquifers have perfectly vertical sides, as in Rodell et al. (2007), Rodell et al. (2009), and Richey et al. (2015a). Equation (16) calculates annual groundwater withdrawals from each aquifer in m^3 as recharge minus change in fill, an assumption underlying the law of motion in (1). Equation (17) obtains the fill $\tilde{a}_{it}$ in each year 2006-2016 by applying the law of motion (1), where the change in fill $\Delta \tilde{a}_{it}$ is from (16) and the recursion begins from baseline fill $\tilde{a}_{i,2013}$. Equation (18) calculates each aquifer's capacity κ_i as the product of the aquifer's maximum vertical height (i.e., lift height plus saturated thickness), surface area s_i, and specific yield $\mathcal{Y}_i$. Equation (19) calculates lift height in each aquifer$\times$year using a recursion beginning from observed baseline values $h_{i,2008}$.

Richey et al. (2015a) provide one benchmark for these calculations. They summarize estimates of aquifer capacity under different assumptions. Their best available estimate of storage for the Central Valley Aquifer System in California, for example, is 176 km^3 of groundwater, while we calculate that the Central Valley has a storage capacity of 178 km^3.

Global Maps of Groundwater Extraction and Recharge

Figure 3a maps the percent of each global aquifer's capacity withdrawn annually, w_i, averaged over the years 2006–2016 (i.e., gross extraction). Most global land has some form of groundwater, even if not all aquifers are widely used. Gross extraction exceeds 2% of aquifer capacity for many aquifers, though not all, and a subset have extraction above 4%. Gross extraction varies by location. Aquifers underneath arid areas with inhospitable climates or soils for agriculture, like the

the Earth, Richey et al. (2015a) show that using mean specific yield at the surface of the Earth overstates aquifer capacity.

Australian Outback, the Arabian peninsula, or the Sahara desert, have limited withdrawals. More fertile areas like the California Central Valley or the breadbaskets of Eastern India have greater withdrawal rates. Figure 3b graphs the distribution of gross extraction across aquifers; median withdrawals are 2% annually. Without recharge, after half a century, 2% use would consume nearly two-thirds of an aquifer's capacity ($\approx 0.98^{50}$). Of course, recharge replaces much use. Figure A.2a maps net groundwater withdrawals across aquifers. Net water use equals 0 for the median aquifer, consistent with no overall net trend in fresh water availability in arable land (Carleton et al. 2024b). Of course, observing that the median aquifer has no net depletion does not mean its current level is optimal.

How do these patterns of withdrawal compare to recharge? Appendix Figure A.2c shows the global distribution of annual recharge r_i across aquifers, which generally varies between 0.5% and 1.5% of aquifer capacity. Recharge rates have substantial spatial heterogeneity, partly reflecting local precipitation. Spatial patterns of recharge have less obvious spatial correlation with patterns of agricultural production. Aquifers beneath the Amazon rain forest and the snow-covered Russian tundra, for example, have among the highest recharge levels.

4.2 Data: Production

We use several spatial variables related to production.

We use nighttime luminosity data from Li et al. (2020), measured as an index they define. We use the temporally harmonized version of their data spanning 2006-2016.. They harmonize records from the two satellites. We report sensitivity analyses that separate records from two satellites; Appendix A.3 discusses accuracy of these data.

We also use data on purchasing power-adjusted GDP data, which cover 2006-2015 (Kummu et al., 2018). These data rescale official subnational GDP statistics, so loosely approximate spatial variation in GDP within a region. We average nighttime luminosity and GDP across gridpoints within each aquifer$\times$year.

We obtain gridded data on unit production of 26 common crops from the Global Agro-Ecological Zones (GAEZ) project. Section 5.2 explains how we use water needs and value of the types of crops cultivated above each aquifer to estimate the willingness to pay for groundwater. GAEZ uses agronomic models and high-resolution data on hydrogeological characteristics to predict yields under different growing conditions, and much economic research uses it (e.g., Costinot et al. 2012; Alesina et al. 2013; Donaldson and Storeygard 2016). We measure actual unit output of each crop in tons.

We also use estimates of the dollar marginal return to an additional cubic meter of groundwater in agricultural production from D'Odorico et al. (2020). They obtain these estimates from informa-

tion on increased crop production in irrigated land relative to rain-fed conditions, for 16 different crops, in the year 2000.[8] Section 5.2 explains how we use these data to calculate the willingness to pay for groundwater for each aquifer as the weighted sum of the crop- and continent-specific values, weighted by each crop's share of total GAEZ production.

4.3 Data: Production Inputs and Costs

We use several measures of other agricultural inputs and their costs per country or country×year. We measure population and the employment share in agriculture from the World Bank (2023) Development Indicators. We measure capital stocks from the Penn World Tables (GGDC 2022). We measure wages from the International Labor Organization Statistics (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 World Bank Development Indicators, as in Eaton and Kortum (2002). We obtain data on interest rates used in monetary policy from the International Monetary Fund (IMF 2023).

We use data on electricity prices and pump efficiency to estimate groundwater extraction costs. We obtain estimates of annual industrial electricity prices from the International Energy Agency over 2006-2016 (IEA 2021); industrial is the closest pricing category to agriculture.[9] Engineering estimates of pump efficiency range from 35% to 85% (Martin-Candilejo et al. 2020). We use the estimate of 50% from Burlig et al. (2021) based on groundwater pump microdata from California (Pradeleix et al. 2015).[10]

We obtain gridded data on climate and geology, and sum or average each variable within each aquifer. We measure precipitation and temperature from the Climate Research Unit at the University of East Anglia (Harris and Jones 2019). We measure soil characteristics from the Global Hydrogeology Maps, a high-resolution map that distinguishes rock types, for the year 2014 (Gleeson et al. 2014). We assume that soil porosity and permeability are fixed over short timescales, following Gleeson et al. (2014) and other hydrology papers.

[8] For the 10 crops in GAEZ but not in D'Odorico et al. (2020), we average willingness to pay values for the 16 crops observed in both datasets.

[9] Some farms in developing countries use more expensive, diesel-powered groundwater pumps. For example, 25% of groundwater pumps in India use diesel rather than electricity for power (Rajan et al. 2020). In the absence of country-specific data on pump type , we measure energy costs from electricity prices.

[10] Larger electric pumps commonly operate in the range of 75-85% efficiency (Robinson 2002). Diesel groundwater pumps are about 10% less efficient than electric pumps (Rajan et al. 2020).

4.4 Data: Benchmark Discount Factors

Ramsey (1928) proposed the following normative discount factor:

$$\beta_i^* = d_i + \eta_i g_i$$

where d_i is the pure rate of time preference (i.e., 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). They calculate the utility discount rates implied by data on age-specific mortality rates and life expectancy for most countries. For each country, these provide a discount rate capturing a social planner's pure rate of time preference. They add these to estimates of the elasticity of the marginal utility of consumption and the rate of consumption growth. Figure A.3b shows that Ramsey discount factors are concentrated around 0.95, with most variation between 0.91 and 0.99.[11]

We also discuss a benchmark market-based interest rate of 2%, reflecting long-term rates of low-risk government debt. This is a central value in estimates of the social cost of carbon (Carleton and Greenstone 2022; Rennert et al. 2022).

5 Parameter Estimation and Model Calibration

This section explains how we estimate the production technology, productivity values, and cost function we use to take the model to the data.

5.1 Production Function Elasticities

Methods

To estimate production function parameters, we take a time-varying version of equation (2):

$$y_{it} = \mathcal{Z}_{it} k_{it}^{\alpha} l_{it}^{\gamma} w_{it}^{\phi}$$
$$\mathcal{Z}_{it} = e^{\Gamma X_{it}} \kappa_i \tau_t \varepsilon_{it}$$

Combining and taking logs gives

$$\log y_{it} = \alpha \log k_{it} + \gamma \log l_{it} + \phi \log w_{it} + X_{it}^{'}\Gamma + \kappa_i + \tau_t + \varepsilon_{it} \tag{20}$$

[11] As one comparison, using median demographic information from around the world yields a rate of 1.33% (Addicott et al. 2020); a recent survey estimates a value of 1.1% for the social utility discount factor on average in the economics research (Drupp et al. 2018).

As in the model, y_{it} represents output; k_{it}, l_{it}, and w_{it} are capital, labor, and water inputs, respectively. We assume that productivity $\mathcal{Z}_{it}$ depends on temperature and precipitation X_{it}; an aquifer fixed effect κ_i that comprises time-invariant productivity ζ_i and land inputs x_i; a year fixed effect τ_t; and an error term ε_{it} for aquifer i in year t. The main parameters we use from equation (20) are the elasticity of output with respect to groundwater extraction ϕ and $\theta = 1-\alpha-\gamma$, which affects profit maximizing output in equation (3). We cluster standard errors by aquifer..

Estimating equation (20) involves several challenges. First, water inputs w_{it} may include both groundwater and surface water. We explore use of groundwater, surface water, or combinations. In equation (20), to isolate the effect of groundwater use on output independently of other water sources, we control for aquifer-level rainfall and temperature X_{it}.

Second, since many sets of climate controls are plausible, we use Belloni et al. (2012)'s post-double selection to choose relevant controls from among total precipitation, average temperature, their squares, and their logs.

Third, no data report aquifer-level output, so y_{it} is unobserved. Our main results use remotely sensed (satellite) measures of nighttime luminosity $\bar{y}_{it}$ as a proxy for output y_{it}, though sensitivity analyses use gridded GDP for a subset of years. We assume that output is an aquifer-specific multiple χ_i of nighttime luminosity $y_{it} = \chi_i \bar{y}_{it}$. Many studies find that nighttime luminosity is correlated with output (Donaldson and Storeygard 2016). In older remote sensing data products, the relationship was stronger in urban than rural areas (Henderson et al. 2012). Newer night lights data more strongly correlate with both rural and urban economic activity (Chen and Nordhaus 2015; Perez-Sindin et al. 2021).

The fourth econometric challenge is a combination of possible measurement error and omitted variables bias in the inputs. Measurement error is a common concern in remotely sensed data such as GRACE, which are typically the smoothed average of several images (Donaldson and Storeygard 2016; Proctor et al. 2023). Measurement error may also arise in our aquifer level estimates of labor and capital inputs, which are convex combinations of country-level data. Omitted variables bias is also a classic concern in production function estimates since here the representative agent may observe productivity shocks which we do not observe.

We use instrumental variables to help address this potential bias. Our instruments for groundwater extraction include two- and three-period lags of total precipitation and average temperature, their interactions with soil permeability and porosity, and their logs. Similarly, we use up to third-lagged prices for capital and labor (i.e., interest rates and wages) as instruments for capital and labor inputs.

The intuition underlying these instruments is that precipitation and average temperature in previous years determine groundwater availability today through effects on available aquifer storage. The law of motion for groundwater in equation (1) shows this relationship. Rainfall has a

larger effect on groundwater stocks for more porous aquifers. We use climate variables beginning at least two years prior to limit potential confounding effects through asset storage or other channels. The use of lagged input prices as instruments for capital and labor is common in production function estimation (Ackerberg et al. 2015; Doraszelski and Jaumandreu 2018; Pindyck and Rotemberg 1983); this approach assumes that lagged prices are uncorrelated with current shocks to agricultural productivity that might affect current demand and thus production.[12] Since many variants of these instruments are possible (e.g., multiple lags), we select relevant instruments using a variant of the least absolute shrinkage and selection operator (Tibshirani 1996) that identifies instruments with the highest explanatory power from the set of possible variables (Belloni et al., 2012); this post-double selection procedure parallels our selection of controls from among contemporaneous temperature and precipitation variables.

The corresponding first stage equation is:

$$log(w_{it}) = \beta Z_{it} + \Gamma X_{it} + \zeta_i + \xi_t + \nu_{it} \tag{21}$$

where w_{it} is groundwater use, labor, or capital, Z_{it} are the potential instruments; and ν_{it} is the error.

Results

Appendix Table A.1 shows the first stage regressions of groundwater use, capital, and labor on the instruments. The candidate instruments selected by the Belloni et al. (2012) algorithm comprise a combination of the lagged precipitation and temperature variables, as well as wages and interest rates. We find that these instruments predict input use and the first stage F-statistic from estimating equation (21) for groundwater, capital, and labor inputs is 10.4, which marginally exceeds the critical value for a reduction in bias of at least 90% in the IV estimates with our three endogenous variables and the twelve selected instruments (Stock and Yogo 2003).

Table 1 shows estimates of two model parameter inputs: the return to groundwater ϕ and $\theta = 1 - \alpha - \gamma$, where ϕ, α, and γ are the water, capital, and labor shares from equation (20). We also report the implied return to land $\theta - \phi$, assuming constant returns. Columns (2)-(4) use instrumental variables. Column (2) includes groundwater only as a regressor. Column (3) includes capital and labor inputs but instruments for groundwater only. Column (4) instruments for labor and capital inputs in addition to water inputs.

Table 1 suggests moderate importance for water in output. OLS estimates in column (1)

[12]Several of these production function estimation papers use current, rather than lagged, input prices as instruments, which would be valid in our setting if aquifer-specific productivity were uncorrelated with aggregate demand. Since many aquifers produce significant output of specific crops, we think this assumption is less valid in our setting and prefer to use lagged prices that do not require this stronger assumption.

find a small and statistically insignificant elasticity of output with respect to groundwater use. Simultaneity and omitted variables bias due to measurement error may attenuate this estimate towards zero. The sign of the bias is consistent with unobserved shocks to productivity, such as introduction of drought-resistant seeds, that boost output and reduce water use. Instrumenting for groundwater use only (column 2) and controlling for other inputs (column 3) increase the groundwater elasticity to somewhat over 0.2, with a standard error of 0.08. Our preferred estimate in column (4) also instruments for capital and labor, which reduces the groundwater elasticity to just over 0.1. The column (4) estimate of the capital and labor input share is $\alpha + \gamma = 0.49$. The implied land share is 0.41, implying decreasing returns due to this fixed factor.

Most estimates of production technology, even for agriculture, do not account for water. Ryan and Sudarshan (2022), one exception, estimate water input shares of 0.04 to 0.18, and labor and capital shares of 0.42 to 0.54. While their analysis of farmers in the Indian state of Rajasthan has a substantially different setting than our global analysis, these are in the same ballpark.[13]

Appendix Table A.2 finds that regions with greater potential benefits from groundwater have a greater estimated elasticity of output with respect to groundwater; Appendix E discusses details. Appendix Table A.3 reports estimates of this production function using alternative measures of GDP. Most are statistically indistinguishable from the main estimates of the groundwater share, though point estimates vary somewhat across columns

5.2 Productivity

We calculate aquifer-specific productivity as follows. Since we do not observe aquifer-level output in dollars, we cannot directly use the production function estimates of z_i as an approximation of total factor productivity. Differentiating the production technology (3) with respect to water then solving for productivity gives

$$z_i = \frac{\theta \kappa_i^{\frac{\phi}{\theta}} y_i'(\tilde{w}_i)}{\phi(\tilde{w}_i)^{\frac{\phi}{\theta}-1}} \tag{22}$$

where, as above, $\tilde{w}_i = w_i \kappa_i$ is cubic meters of water.

We calibrate the baseline productivity of each aquifer z_i using empirical analogues to each term in (22). Data report capacity κ_i and water use $\tilde{w}_i$. Table 1 estimates the parameters ϕ and θ. To calculate aquifer-level willingness to pay for groundwater, we proceed in several steps. We first obtain crop- and continent-specific estimates of the marginal revenue of groundwater ("willingness to pay") from D'Odorico et al. (2020). The authors calculate the crop- and continent-specific value

[13] Ryan and Sudarshan (2022) also include land as a production factor in their (cross-sectional) analysis; this input is absorbed by the aquifer fixed effects in our analysis. The estimated land share in this other paper is approximately 0.5, which combined with our estimated input shares implies approximately constant returns to scale. Gollin and Udry (2021) similarly find a land share of approximately one-half.

of water as the ratio of the increase in crop production from irrigation and the associated increase in irrigation water.[14] We then use crop-level raster data on irrigated crop production from GAEZ to calculate the share of agricultural production attributable to each crop on the grid. We calculate willingness to pay for groundwater at the pixel level by summing the crop-specific willingness to pay estimates using the crop shares as weights. We assign pixels without crop production a willingness to pay for groundwater of 0 when averaging across the aquifer's area. Finally, we calculate the average willingness to pay for groundwater across the aquifer.

Figure 3c maps the global distribution of willingness to pay for groundwater by aquifer. Areas with limited and unproductive agriculture, like the Sahara desert, have low willingness to pay. Regions with productive staple crops like wheat in the midwestern US have moderate willingness to pay. Areas with high value crops, like fruits and nuts in California's Central Valley, have high willingness to pay.

5.3 Cost Function

We measure groundwater extraction by constructing an empirical analogue to equation (6). Following engineering practice (Martin-Candilejo et al. 2020; Pradeleix et al. 2015; Robinson 2002), we measure the cost c_i per unit of water×lift height as

$$c_i = \frac{\text{electricity price}_i \times \text{constant}_i}{\text{pump efficiency}_i} \times \kappa_i$$

Data report electricity prices; Appendix Figure A.3e maps their distribution. Pumping efficiency and the physical constant are calibrated as discussed above.[15] We scale the resulting marginal cost of extraction in m^3 by observed aquifer capacity κ_i since w_{it} in the model is measured as a share of aquifer capacity, rather than m^3, for ease of computation.

The dispersion in lift heights and electricity prices across aquifers translates into heterogeneity in extraction costs. Figure 3e maps our estimates of the marginal cost of lifting one cubic meter of groundwater to the surface, $c_i \times h_{it},$, averaged over 2006-2016, and the histogram Figure 3f shows their distribution. Lift costs are typically between \$0.01 and $\$0.04/m^3$ with most mass around $\$0.02/m^3$, and never above $\$0.15/m^3$.[16] As points of reference, we estimate an extraction cost

[14] Specifically, they estimate

$$\text{wtp}_i = \frac{P_i \Delta Y_i}{IWR_i}$$

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. wtp_i is the willingness to pay for water in aquifer i.

[15] Pumping efficiency numbers in the engineering literature are typically in the range of 0.60-0.75 (Martin-Candilejo et al., 2020; Pradeleix et al., 2015). Burlig et al. (2021)'s estimate of 0.54 is based on measured efficiency of the pumps in their dataset.

[16] The kernel density estimate in Appendix Figure A.1b shows that the distribution of marginal costs in major aquifers is approximately $\$0.01/m^3$ more costly.

of $\$0.02/m^3$ for California's Central Valley Aquifer System and $\$0.02/m^3$ for the Indus Basin.[17] Appendix Figure A.1c shows the density of the difference between willingness to pay and marginal cost; the mean difference is around $\$0.16/m^3$, or several times larger than the marginal cost of extraction itself.[18] The long right tail of costs mirrors the distribution of lift heights across all aquifers, which is concentrated around 20 m and with a long right tail (Figure A.3d); by contrast, the electricity price distribution, which is the other main component of variation in marginal costs across aquifers, is bi-modal, with prices concentrated around \$0.025/kWh and \$0.075/kWh (Figure A.3f).

5.4 Model Inputs Summary

Table 2 summarizes values of several model inputs for the four case study aquifers and for the mean and median of all global aquifers. Appendix Table A.5 shows results weighted by crop production, GDP, and population for all global aquifers and the 37 major aquifer systems. We show our estimates of aquifer capacity, groundwater use, willingness to pay, 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 2 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 Ogallala Aquifer has the largest capacity, which contributes to its comparatively low depletion as a share of its capacity. All four of these aquifers are relatively high-value, with willingness to pay estimates that range from $\$0.16/m^3$ for the Ogallala to $\$0.22/m^3$ for the Central Valley, in the upper half of the distribution (Figure 3d). Marginal costs are between 10% and 30% of willingness to pay for these aquifers, with differences largely driven by variation in lift height. This table also shows that the Ramsey discount factor is between 0.94 and 0.96 for all four aquifers; we will interpret a finding of β less than these Ramsey discount factors as suboptimally high extraction rates.

The four case studies are largely representative of all global aquifers, with a few main differences. Mean and median groundwater usage as a share of capacity is approximately 2-3% for all aquifers, which is similar in magnitude to the case studies that have comparable capacity (i.e., excluding the largest Ogallala aquifer). The mean and median global aquifers are characterized by willingness to pay estimates that are weakly below the Central Valley Aquifer System, Indus Basin, and North China Aquifer System; weighting by aquifer crop value, GDP, or population increases willingness to pay by approximately 10%. Median and mean extraction costs are on the lower end of the case

[17] The literature finds comparable estimates: Burlig et al. (2021) calculate an extraction cost of $\$0.05/m^3$ for California's Central Valley Aquifer and Ryan and Sudarshan (2022) report extraction costs of $\$0.04/m^3$ from surveying farmers in India's state of Rajastan, which is situated above the Indus Basin.

[18] Surface water prices often exceed groundwater extraction costs due to higher treatment costs. Surface water prices in California, for example, are approximately \$0.18 per cubic meter during the time period of our data (Hagerty, 2023).

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 ranging from 0.94 to 0.96 across the different examples.

6 Results: Efficiency of Groundwater Extraction

6.1 Visualizing How We Estimate Discount Factors

To further clarify how we estimate discount factors, we first graphically describe the representative agent's problem within a period. Visualization is clearest one aquifer at a time. In Figure 4, each panel shows one case study aquifer. Within each graph, the horizontal axis describes groundwater consumption as a share of aquifer capacity (w) and the vertical axis describes current-period marginal revenues or costs in $\$/m^3$ of groundwater. In each panel, the thick downward-sloping marginal revenue curve shows how additional groundwater extraction today affects current profits. The positive value of this line reflects the value of water extraction, while the downward slope reflects decreasing returns to water. The three dashed lines show the marginal cost of groundwater extraction under different discount factors. The asterisks show the myopic $\beta = 0$; the yellow squares describe our estimated β; and the benchmark β^* in purple triangles represents Ramsey discounting. The vertical dashed line shows actual water extraction observed in data. We infer the estimated discount factor as the value which makes the marginal revenue and cost curves intersect at observed water consumption w.

Figure 4 shows several substantive results for these case study aquifers. In each panel, the benchmark factor exceeds the estimated discount factor, in that the triangles lie above the squares, and implying that all these aquifers are being extracted more rapidly than is optimal.

The agent extracts groundwater until marginal revenue equals marginal cost in each case, and water use decreases in the discount factor. When $\beta = 0$, the myopic agent chooses water use so that marginal revenue equals current marginal cost only, which is a function of current, but not future, lift height. When $\beta > 0$, an additional dynamic cost captures the effect of current extraction on future lift height and, by extension, future cost. Since marginal cost is higher when the agent considers the increase in future costs associated with current extraction, water use falls in these two scenarios. It falls further in the benchmark Ramsey discount factor case because this Ramsey discount factor exceeds the estimated discount factor.

Consider, for example, the Central Valley aquifer. GRACE and capacity data indicate that in the mean year in 2006-2016, people consumed water equal to 2.4% of the aquifer's capacity (Table 3). The production function estimates in Table 1, column (4), imply that the marginal revenue per unit of water consumed at this quantity is about \$0.12 per m^3 of water. At this

quantity of water consumption, complete myopia would imply a marginal cost of water extraction around \$0.04/m^3, which is lower than the marginal revenue we estimate. Ramsey discounting would imply a marginal cost of water extraction around \$0.18/m^3, which is above the marginal revenue we estimate. The reason is that Ramsey discounting assigns a fairly high weight to the additional future extraction costs due to having to lift the water further while myopia assigns zero weight to those future costs. To recover the estimated discount factor we ask: what discount factor implies a marginal cost of water extraction equal to 0.11, the marginal cost of extraction at the observed water use? The answer is a discount factor of 0.87 (Table 3).

6.2 Estimated Discount Factors

Panel A of Figure 5 maps the discount factor that we estimate for each aquifer, by applying the algorithm from Appendix D to solve equation (8). Dark green shows aquifers with a discount factor above 0.95; moderate green represents aquifers with a discount factor between 0.90 and 0.95; and light green represents aquifers with a discount factor below 0.90. The histogram in Panel B shows the distribution of estimated discount factors across aquifers.

The map in Figure 5, panel A, shows a wide geographic distribution of discount factors across aquifers. Most North American aquifers have estimated discount factors below 0.90, although the Ogallala does not. Aquifers in Southeast Asia and Australia also have low estimated discount factors, although parts of Northeast and Western Asia do not. While the classification boundaries of 0.90 and 0.95 between groups in this chloropleth map are somewhat arbitrary, they provide one initial set of information about which aquifers are likely to have water extraction paths in line with benchmark discount factors.

The histogram in Figure 5, panel A, shows that a majority of aquifers have high estimated discount factors, though not all. Close to half of all aquifers have an estimated discount factor above 0.95 percent. This would loosely suggest that in these aquifers, the extraction path for groundwater is in the ballpark of what might occur under a discount factor in line with long-term rates on government debt.

The first row of Table 3 presents summary statistics on estimated discount factors. Column (1) describes the median aquifer, column (2) describes the mean aquifer, and columns (3)-(6) describe the four case study aquifers.

The first row of Table 3 shows that the median aquifer has a discount factor of 0.92. This is below the long-term rate on US government debt, which could suggest these aquifers are extracted at an inefficiently rapid rate. At the same time, this discount factor is in line with some countries' Ramsey discount rates or higher government bond rates, so extraction in the median aquifer is not necessarily inefficient. Column (3) shows that a discount factor of 0.87 best accounts for extraction

patterns in the Central Valley in California. This is a fairly low discount factor, and is in line with the prevailing informal view that this aquifer is suffering from inefficiently rapid extraction. Column (4) shows an estimated discount factor of 0.94 for the US Ogallala aquifer. The Ogallala has experienced less recent net extraction than many aquifers, explaining this moderate discount factor. Columns (5) and (6) report lower estimated discount factors of 0.82 and 0.91 for the Indus Basin and North China Aquifer System, suggesting that extraction here may be more rapid than is optimal, since these values are likely to be below benchmarks. Weighting aquifers by crop value or population gives a lower median global discount factor of 0.86 or 0.88. Weighting aquifers by GDP gives a higher median global estimated discount factor of 0.93.

Appendix Table A.4 reports correlates of these estimated discount factors. Columns (1) through (9) regress each aquifer's estimated discount factor on one aquifer characteristic, while column (1) includes the characteristics all at once. Columns (1) through (4) consider model inputs, columns (5) through (7) consider Ramsey discount rates and their components, and columns (8) and (9) consider other country characteristics. Column (10) considers them all at once. We include continent fixed effects to account for basic geographic heterogeneity.

Many patterns in Appendix Table A.4 are sensible though some are more surprising. Column (1) shows that aquifers with more benefits from extraction are more patient, while column (2) shows that aquifers with more costs from extraction or costs due to electricity prices or lift height are less patient. Point estimates imply that aquifers with higher Ramsey discount factors also have higher estimated discount factors, though they are statistically insignificant. Richer and more democratic countries are more impatient, perhaps because they empower more parties to compete in rapidly extracting resources, echoing open access.

6.3 Comparing Estimated Versus Benchmark Discount Factors

Estimated discount factors may differ across aquifers for efficient or inefficient reasons. Differences in the benchmark discount factors provide an efficient reason. For example, some aquifers have rapid economic growth or higher interest rates, giving them lower benchmark discount factors, and leading to rapid but optimal groundwater extraction. Alternatively, estimated discount factors may differ across aquifers for inefficient reasons. For example, some aquifers may suffer from common pool open access, where atomistic farmers each extract water rapidly, myopically, and unconstrained. While each farmer acts unilaterally and optimally in such aquifers, the extraction path is inefficiently rapid from the planner's perspective.

Panel A of Figure 6 maps the difference between estimated and Ramsey benchmark discount factors (Addicott et al. 2020) across aquifers, while the histogram in Panel B describes their distribution. In Panel A, aquifers in light green have an estimated discount factor 2.5 percentage

points or more below the benchmark, suggesting inefficiently rapid extraction, or impatience. Aquifers in dark green have an estimated discount factor more than 2.5 percentage points above the benchmark, implying excessive patience. Aquifers shown in moderate green have somewhat similar estimated and Ramsey rates.

Figure 6 shows that most aquifers have estimated discount factors 2.5 percentage points or more below the benchmark. This provides another piece of evidence that extraction in most aquifers is inefficiently rapid, and that farmers less than fully account for the dynamic externality of current extraction on future extraction costs. Panel A shows that California's Ogallala aquifer and much of Oceania and Southeast Asia have inefficiently rapid extraction, while some parts of Northeast and West Asia do not. The similar geographic patterns in Panel A of Figures 5 and 6 suggests that most of the variation in estimated discount factors across aquifers does not reflect differences in benchmark discount factors, i.e., we estimate that most of the geographic variation in extraction choices driving estimated discount factors is inefficient. This echoes the finding discussed above of no significant correlation between estimated discount factors and Ramsey benchmark discount factors.

The histogram in Panel B of Figure 5 shows large gaps between estimated and benchmark discount factors. While the modal difference is close to zero, many aquifers have estimated discount factors more than 20 percent below the benchmark, representing extremely rapid depletion relative to benchmarks.

Figure 6 shows a modest but small set of aquifers with estimated discount factors above the benchmark. We see a few possible interpretations. Distortions in the market for labor, capital, land, intermediates, and other goods due to policy or other failures may lead to extraction rates slower than benchmark. Such market failures are commonplace, especially in developing countries, although we do not model them directly. More broadly, every model abstracts from some features of reality, data can contain measurement error, some model parameters are estimated with sampling variation. The discount factor most consistent with extraction paths in an aquifer may therefore not numerically equal the discount factor we estimate for the aquifer. Ongoing work undertakes a bounding exercise to assess how accounting for these sources of variation in our estimates of discount factors may account for estimates of discount factors that exceed benchmarks, and how patterns of discount factors would change in the absence of these sources of variation.

Table 3, row 2, shows that the differences between estimated and benchmark discount factors is 3 percentage points for the median aquifer but 10 for the mean aquifer, due to the skewness of estimated discount factors. The benchmark exceeds the estimated discount factor for all four case studies.

Appendix Table A.6, row 2, summarizes other moments of the difference between estimated and benchmark discount factors. Weighted by crop value or population, the difference between

these discount factors is 6 percentage points, while weighted by GDP, the difference is 1 percentage point.

6.4 Long-Run Outcomes

The last three panels of Tables 3 summarize the implications of inefficiently low discount factors for long-run extraction, fill, and cost using the model, for the median and mean global aquifer and for the case studies. First, in both tables, observed water use exceeds the optimal Ramsey rate for all aquifers, including the Ogallala Aquifer, which is already refilling. Optimal extraction according to the Ramsey discount factor would reduce net water use and imply that extraction should fall enough to allow the Central Valley Aquifer System, the Indus Basin, the median and mean global aquifers to refill. Optimal extraction on net is negative for $\beta = 0.98$, indicating that the aquifers should refill gradually in all of these cases. This is the opposite pattern of extraction currently observed in the data for all of our examples except for the (refilling) Ogallala Aquifer. Second, and as a result, long-run steady state aquifer fill is higher under the Ramsey discount factor relative to the observed long-run fill under current extraction rates in all examples; using $\beta = 0.98$ implies that it is optimal to allow all aquifers to refill fully and consume only recharge. Finally, long-run extraction costs fall relative to observed costs in all examples because long-run fill rises and, by extension, long-run lift height falls. Appendix Table A.6 shows that these same patterns persist for the weighted global aquifer medians and the median major aquifer systems.

For each case study aquifer, Figure 7 shows fill levels over the years 2016-2100. The solid line describes fill when extraction is chosen using our estimated discount factor, while the dashed line describes fill when extraction is chosen using the benchmark Ramsey rate. For the Central Valley, Indus, and North China aquifers, observed extraction depletes 2% to 10% of aquifer capacity over time. Near the end of the 21st century, marginal profit would slow extraction and increase fill. Extraction declines as marginal extraction costs rise, given the lower water table, and fill converges to a steady state. In contrast, for each of these aquifers, extraction using the Ramsey discount factor would slow extraction and increase fill. For the Ogallala aquifer in Panel B, which is already filling under the estimated discount factor, extracting under the Ramsey discount factor accelerates fill, though by the year 2100 the aquifer is still gradually filling.

Figure 8 shows the corresponding long-run water extraction costs for the same examples. Predicted costs in each scenario decrease in predicted fill levels: as aquifer fill increases, lift height and, by extension, extraction costs fall. As Figure 7 shows, extracting at the socially optimal rate increases long-run fill relative to the observed business-as-usual scenario, and so the corresponding costs (in long-dashed red) are lower for all aquifers. The Ogallala aquifer again is an example of an aquifer that is refilling: observed increases in fill cause costs to decline in the business-as-usual

scenario, and costs decline further when extraction is slowed in the optimal extraction scenario.

Figure 9 summarizes the long-run distributions of future fill for all global aquifers under our estimated discount factors (Panel A), under Ramsey discounting (Panel B), and under market discounting (Panel C). The main take-away from this figure is that current inefficiently high rates of extraction result in much lower and more dispersed aquifer fill levels in the future. Under our estimated discount factors, approximately 40% of global aquifers will fully recharge by the end of the century; this fraction increases to approximately 60% under Ramsey discounting and further to over 80% under the higher discount factors associated with market discounting.

Overall, the deviations of the discount factors that rationalize observed water use from the socially optimal discount factors have important implications for agricultural profits. Figure 10 summarizes the differences between annual estimated and socially optimal profits under Ramsey discounting for our case study aquifers. Since the estimated discount factors for these aquifers are below the Ramsey discount factors, estimated profits are initially higher than socially optimal profits. This is due to higher initial water use in the estimated scenario, which leads to higher future extraction costs. After a few years, annual profits in the socially optimal case exceed that of the estimated scenario.

6.5 Welfare Losses Due to Estimated Versus Benchmark Discount Factors

We calculate the economic consequences of this gap between estimated and benchmark discount factors from equation (12). For each aquifer, Panel A of Figure 11 maps the welfare loss due to extraction paths inconsistent with benchmark discount factors. The map expresses losses as a percent of aquifer value. For areas in dark red, non-optimal extraction paths create a cost that exceeds 4.5 percent of aquifer value. In contrast, for areas in beige, the cost of non-optimal extraction is less than 1.5 percent of the aquifer's present value. Panel B shows a histogram of these losses across aquifers.

Panel A of Figure 11 shows that the costs of non-optimal groundwater extraction are large in the same areas where earlier results show that estimated discount factors are especially low—parts of the Western US including California; much of Mexico; and much of Oceania.

Figure 11, Panel B, shows that while the modal aquifer has a welfare loss less than 1 percent of aquifer present value. At the same time, the distribution has a right tail. Several percent of aquifers have losses exceeding 5 percent of aquifer value, and a few above 10 percent of aquifer value.

Table 4, Panel C, summarizes the welfare losses due to non-optimal depletion. For the median aquifer, the present-value welfare loss due to non-optimal extraction is 1.4% of the aquifer's value

and almost double this value for the mean aquifer (2.66%). In the case study aquifers, this proportion is high for the Central Valley (3.4%), low for the Ogallala (0.2%), moderate for the Indus Basin (1.3%), and practically nonexistent for the North China Aquifer System (0.1%). The second row of Panel C expresses these losses in billions. For the median aquifer, we estimate losses of $60 million in present value due to non-optimal extraction. We estimate that California's Central Valley Aquifer System has losses of $1.8 billion in present value due to non-optimal extraction. We estimate that losses in the Ogallala are smaller, at $0.3 billion, and losses in the North China Aquifer System are even smaller, at $20 million.

Do these statistics represent large or small losses? Of course, subjective magnitudes depend on the comparison and judgment of the reader. Global present value losses in the trillions due to non-optimal extraction are certainly not small. At the same, these losses aggregate over all future years and aquifers. Current-year damages, or damages as a proportion of current profits, are modest.

Our sense is that losses of 1 to 2 percent of aquifer value due to non-optimal groundwater extraction are an order of magnitude smaller than many people's priors. This paper's introductory quote from Hotelling (1931), Ostrom (1990)'s Nobel-winning institutional analysis of groundwater management, or the tone of the Arrow et al. (2004) or the Dasgupta (2021) syntheses of sustainability, all give a sense of urgency and concern, that humanity may be rapidly squandering the limited natural resources available. If these costs are 1 or 2 percent of natural resource value, should we really be so concerned?

These statistics assume that the representative agent chooses the optimal extraction path, conditional on a discount factor. The inefficiency comes from actions reflecting a different discount factor than the benchmark. But in the absence of policy which distorts intertemporal optimization, the representative agent would consume water until the marginal benefit of doing so equals the marginal cost. The estimated costs of 1 to 2 percent primarily reflect the additional energy required to extract water once the groundwater table gradually declines to increasingly low levels. Optimization also imposes a moderating force—as the water table declines, the increasing extraction cost leads the representative agent to consume less water.

Our conclusion from this analysis is that if intertemporal extraction is optimal conditional on the discount rate we estimate, then the costs of inefficient extraction are on the order of 1 to 2 percent of aquifer value. Below we consider scenarios where the representative agent departs from the optimal extraction path, even conditional on a discount factor which differs from the benchmark, and we find larger losses from such choices. Those extraction paths could represent inefficient policy or incomplete optimization.

6.6 Taxes to Address Differences Between Estimated and Benchmark Discount Factors

While the previous subsection estimates losses from inefficient extraction, it says nothing about how policy could change these decisions. We therefore turn to study corrective taxes which shift the extraction path towards the program consistent with benchmark interest rates, as in equation (13).

Figure 12 describes the distribution of these tax rates. Panel A shows a map of the tax rate for each aquifer. Aquifers in white either lack sufficient data to calculate the tax rate or the calculation fails to converge, though we can sign the direction of the welfare loss. Panel B plots a histogram of the distribution of the tax rate across aquifers. We report all tax rates in $\$/m^3$, though below we relate them to users' valuation of water.

Panel A of Figure 12 shows wide variation in tax rates. Many aquifers have optimal tax rates below 1.5 cents per cubic meter, while others have rates above 4.5 cents per cubic meter. As a reference to help interpret these magnitudes, recall Table 2 reports that willingness to pay for water in the median aquifer is 0.20, while extraction costs are only 0.02 $\$/m^3$. Panel B shows that the majority of aquifers have a time-invariant optimal tax of several cents per m^3. We calculate that water subsidies are optimal in a minority of aquifers, given our estimate that this minority of aquifers consume less than is optimal.

Table 3, Panel B, summarizes the time-invariant taxes which maximize welfare. The median welfare has a time-invariant tax of 2 cents per m^2. This optimal tax is much higher in the Central Valley at 10 cents per m^3, reflecting its rapid extraction. The optimal tax is near zero in the North China Aquifer System, since its estimated discount factor is near the benchmark, so it has nearly optimal extraction.

Because most water use is unpriced, we cannot calculate ad valorem equivalents of these tax rates by comparing them to market prices. Additionally, these rates are per unit of water, so as a specific rather than ad valorem tax, the ad valorem equivalent would change over time. Regardless, to provide a useful benchmark, we can calculate these tax rates to the representative agent's willingness to pay for water from an aquifer. Table 3 shows that the tax is around 18 percent of willingness to pay in the median aquifer. This ratio is quite high in the Central Valley, where the time-invariant optimal tax is nearly half of willingness to pay for groundwater. The ratio of taxes to willingness to pay is lower in the other case study aquifers.

6.7 Counterfactual Extraction Patterns

Table 5 shows the long-run model results under three different counterfactual scenarios, for the median and mean global aquifers and the case studies. Appendix Table A.8 shows the same

results for weighted medians and the median of the major aquifer systems. In both tables, Panel A summarizes the welfare loss resulting from observed extraction, from a counterfactual scenario where extraction is fully myopic (i.e., $\beta = 0$), and a counterfactual constant water scenario in which current water use is maintained each period until either the aquifer empties or fills up, at which point water use is equal to recharge. Panel B shows net water use for these same scenarios, while Panels C and D show long-run fill and cost, respectively. All of the results are measured relative to the socially optimal results using the benchmark Ramsey discount factor.

There are four main takeaways from these tables. First, the loss in social welfare is generally significantly higher for both fully myopic extraction and constant extraction counterfactuals, relative to observed extraction patterns: the constant water scenario welfare loss is at least 20% higher for most case studies, the mean, and the medians, while welfare loss in the common pool scenario varies between approximately 4 times (Central Valley) to almost 700 times (North China) as large as the welfare loss under observed extraction; an exception is the Ogallala aquifer, where extraction at a constant rate decreases the welfare loss. The intuition here is that current water use is less than what the optimizing agent will choose in the future when it optimizes according to the estimated discount factor, and this future water use is closer to the optimal plan under the Ramsey discount factor (which is greater than the estimated β) than the observed plan. Second, and intuitively, the myopic common pool scenario results in an increase in water use relative to current extraction since all estimated discount factors are well above 0. Third, and as a result, long-run fill is lower in both counterfactual scenarios and both tables; again, an exception is the Ogallala aquifer, which completely refills under the constant water scenario since the aquifer is currently refilling. Conversely, if extraction continues at its current rate, the Central Valley, Indus, and North China aquifers will be completely depleted. Finally, the pattern of long-run costs mirror long-run fill: the myopic common pool scenario reduces fill and therefore extraction costs increase, while costs also decrease in the constant water scenario except if the aquifer fully refills (i.e., the Ogallala).

7 Conclusions

This paper studies the extent to which current groundwater extraction is economically efficient. To do so, we compile global remote sensing and administrative data to estimate a dynamic model of natural resource extraction for almost all of the world's aquifers. We recover aquifer-specific discount factors rationalizing observed extraction, the welfare loss resulting from deviating from socially optimal extraction paths, and the Pigouvian taxes needed to implement them. We find that two-thirds of global aquifers are suboptimally extracted: the discount factor that rationalizes observed groundwater extraction patterns is lower than what a social planner would choose. We

estimate that the present discounted value of welfare costs of this suboptimal extraction exceeds $3.43 trillion USD, or roughly 3.4% of annual global GDP. This estimate is comparable to the mortality cost of climate change, which Carleton et al. (2022) estimate to be 3.2% of global GDP in 2100.

This paper's approach to studying economic efficiency and sustainability of extraction is general enough to apply to any natural resource. Administrative and remotely sensed data provide some of the necessary information to calculate willingness to pay and harvesting costs of fisheries and timber, for example. Recent U.S. governmental and international agency initiatives consider valuing domestic and international natural resources to compile more complete measures of national wealth (Prabhakar et al. 2023; UN 2024). Given appropriate data, this paper provides a way for future research to do so.

References

Ackerberg, D. A., K. Caves, and G. Frazer (2015): "Identification properties of recent production function estimators," *Econometrica*, 83, 2411–2451.

Addicott, E., E. Fenichel, and M. Kotchen (2020): "Even the Representative Agent Must Die," *Journal of the Association of Environmental and Resource Economists*, 7, 379–415.

Addicott, E. T., E. P. Fenichel, and M. J. Kotchen (2020): "Even the representative agent must die: using demographics to inform long-term social discount rates," *Journal of the Association of Environmental and Resource Economists*, 7, 379–415.

Alesina, A., P. Giuliano, and N. Nunn (2013): "On the Origins of Gender Roles: Women and the Plough," *Quarterly Journal of Economics*, 128, 469–530.

Aronoff, D., and W. Rafey (2024): "Conservation Priorities and Environmental Offsets: Markets for Florida Wetlands," NBER Working Paper no. 31495.

Arrow, K., P. Dasgupta, L. Goulder et al. (2004): "Are We Consuming Too Much?" *Journal of Economic Perspectives*, 18, 147–172.

Aspeland, K. M., and A. Russo (2024): "Additionality and Asymmetric Information in Environmental Markets: Evidence from Conservation Auctions," Working Paper.

Belloni, A., D. Chen, V. Chernozhukov, and C. Hansen (2012): "Sparse Models and Methods for Optimal Instruments With an Application to Eminent Domain," *Econometrica*, 80, 2369–2429.

BGR (2025): "World-wide Hydrogeological Mapping and Assessment Programme (WHYMAP)," Jan, https://www.whymap.org/whymap/EN/Home/whymap_node.html.

Borunda, A. (2022): "The deceptively simple plan to replenish California's groundwater," *National Geographic*.

Burgess, R., M. Hansen, B. A. Olken, P. Potapov, and S. Sieber (2012): "The Political Economy of Deforestation in the Tropics," *Quarterly Journal of Economics*, 127, 1707–1754.

Burlig, F., L. Preonas, and M. Woerman (2021): "Energy, Groundwater and Crop Choice," Working Paper.

Carleton, T., L. Crews, and I. Nath (2024a): "Agriculture, Trade, and the Spatial Efficiency nof Global Water Use," Mimeo, UCSB.

——— (2024b): "Is the World Running Out of Fresh Water?" *AEA Papers and Proceedings*, 114, 31–35.

Carleton, T., and M. Greenstone (2022): "A Guide to Updating the US Governmentâs Social Cost of Carbon," *Review of Environmental Economics and Policy*, 196–218.

Carleton, T., A. Jina, M. Delgado et al. (2022): "Valuing the Global Mortality Consequences of Climate Change Accounting for Adaptation Costs and Benefits," *Quarterly Journal of Economics*, 137, 2037–2105.

Chen, X., and W. Nordhaus (2011): "Using Luminosity Data as a Proxy for Economic Statistics," *Proceedings of the National Academy of Sciences*, 108, 8589–8594.

——— (2015): "A Test of the New VIIRS Lights Data Set: Population and Economic Output in Africa," *Remote sensing*, 7, 4937–4947.

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

Costello, C., D. Ovando, T. Clavelle et al. (2016): "Global fishery prospects under contrasting management regimes," *Proceedings of the National Academy of Sciences*, 113, 5125–5129.

Costinot, A., D. Donaldson, and C. Smith (2012): "Evolving Comparative Advantage and the Impact of Climate Change in Agricultural Markets: Evidence from a 9 Million-Field Partition of the Earth," Mimeograph, MIT.

Dasgupta, P. (2021): *The Economics of Biodiversity: The Dasgupta Review*: HM Treasury.

Dasgupta, P. S., and G. Heal (1979): *Economic Theory and Exhaustible Resources*: Cambridge University Press.

D'Odorico, P., D. D. Chiarelli, L. Rosa, A. Bini, D. Zilbermand, and M. C. Rulli (2020): "The Global Value of Water in Agriculture," *PNAS*, 117, 21985–21993.

Donaldson, D., and A. Storeygard (2016): "The View from Above: Aplications of Satellite Data in Economics," *Journal of Economic Perspectives*, 30, 171–198.

Doraszelski, U., and J. Jaumandreu (2018): "Measuring the bias of technological change," *Journal of Political Economy*, 126, 1027–1084.

Drupp, M. A., M. C. Freeman, B. Groom, and F. Nesje (2018): "Discounting disentangled," *American Economic Journal: Economic Policy*, 10, 109–134.

Eaton, J., and S. Kortum (2002): "Technology, Geography, and Trade," *Econometrica*, 70(5), 1741–1779.

Famiglietti, J. S., M. Lo, S. L. Ho et al. (2011): "Satellites Measure Recent Rates of Groundwater Depletion in California's Central Valley," *Geophysical Research Letters*, 38, 1–4.

Fan, Y., H. Li, and G. Miguez-Macho (2013): "Global patterns of groundwater table depth," *Science*, 339, 940–943.

Farrokhi, F., E. Kang, H. S. Pellegrina, and S. Sotelo (2024): "Deforestation: A Global and Dynamic Perspective," Working Paper.

Feng, W., M. Zhong, J.-M. Lemoine, R. Biancale, H.-T. Hsu, and J. Xia (2013): "Evaluation of groundwater depletion in North China using the Gravity Recovery and Climate Experiment (GRACE) data and ground-based measurements," *Water Resources Research*, 49, 2110–2118.

Fenichel, E., J. K. Abbott, J. Bayham, W. Boone, E. M. K. Haacker, and L. Pfeiffer (2016): "Measuring the value of groundwater and other forms of natural capital," *Proceedings of the National Academy of Sciences*, 113, 2382–2387.

Ferguson, B. (2024): "Trade Frictions in Surface Water Markets," Working Paper.

FRED (2024): "Gross Domestic Product: Implicit Price Deflator," Oct, https://fred.stlouisfed.org/series/GDPDEF.

GGDC (2022): "Penn World Table version 10," August, https://www.rug.nl/ggdc/productivity/pwt/.

Gibson, J., S. Olivia, and G. Boe-Gibson (2020): "Night Lights in Economics: Sources and Uses," *Journal of Economic Surveys*, 34, 955–980.

Gleeson, T., N. Moosdorf, J. Hartmann, and L. van Beek (2014): "A glimpse beneath earth's surface: GLobal HYdrogeology MaPS (GLHYMPS) of permeability and porosity," *Geophyscial Research Letters*, 41, 3891–3898.

Goldewijk, C. K. (2017): "Anthropogenic land-use estimates for the Holocene; HYDE 3.2," Sept, https://doi.org/10.17026/dans-25g-gez3.

Gollin, D., and C. Udry (2021): "Heterogeneity, measurement error, and misallocation: Evidence from African agriculture," *Journal of Political Economy*, 129, 1–80.

de Graaf, I. E. M., E. H. Sutanudjaja, L. P. H. van Beek, , and M. F. P. Bierkens (2015): "A high-resolution global-scale groundwater model," *Hydrological Earth System Science*, 19, 823–837.

Greenhill, S., H. Druckenmiller, S. Wang et al. (2024): "Machine Learning Predicts Which Rivers, Streams, and Wetlands the Clean Water Act Regulates," *Science*, 383, 406–412.

Hagerty, N. (2023): "What Holds Back Water Markets? Transaction Costs and the Gains from Trade," Working Paper.

Harris, I., and P. Jones (2019): "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, doi:/10.5285/7ad889f2cc1647efba7e6a356098e4f3.

Henderson, J. V., A. Storeygard, and D. N. Weil (2012): "Measuring Economic Growth from Outer Space," *American Economic Review*, 102, 994–1028.

Hotelling, H. (1931): "The Economics of Exhaustible Resources," *Journal of Political Economy*, 39, 137–75.

IEA (2021): "IEA World Energy Prices," Sept, https://www.iea.org/data-and-statistics/data-product/energy-prices.

ILO (2023): "ILOSTAT Database," June, https://ilostat.ilo.org/.

IMF (2023): "Interest Rate Data: Selected Indicators," June, https://data.imf.org/regular.aspx?key=63087881.

Karp, L. (2017): *Natural Resources as Capital*: MIT Press.

Kummu, M., M. Taka, and J. H. A. Guillaume (2018): "Gridded global datasets for Gross Domestic Product and Human Development Index over 1990-2015," *Scientific Data*, 5, 1–15.

Li, X., Y. Zhou, M. Zhao, and X. Zhao (2020): "A harmonized global nighttime light dataset 1992-2018," *Scientific Data*, 7, 3891–3898.

Martin-Candilejo, A., D. Santillan, and L. Garrote (2020): "Pump Effciency Analysis for Proper Energy Assessment in Optimization of Water Supply Systems," *Water*, 132, 575–605.

Miller, J. A. (2003): *The National Atlas of the United States of America*, Chap. Principal Aquifers.

Muzammil, M., A. Zahid, and L. Breuer (2021): "Economic and environmental impact assessment of sustainable future irrigation practices in the Indus Basin of Pakistan," *Scientific Reports*, 11, 23466.

Ojha, C., M. Shirzaei, S. Werth, D. Argus, and T. Farr (2018): "Sustained groundwater loss in California's Central Valley exacerbated by intense drought periods," *Water Resources Research*, 54, 4449–4460.

Ostrom, E. (1990): *Governing the Commons*: Cambridge University Press.

Perez-Sindin, X. S., T.-H. K. Chen, and A. V. Prishchepov (2021): "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*, 24, 100647.

Pindyck, R. R., and J. J. Rotemberg (1983): "Dynamic factor demands and the effects of energy price shocks," *American Economic Review*, 73, 1066–1079.

Prabhakar, A., G. M. Raimondo, and S. D. Young (2023): "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.

Pradeleix, L., P. Roux, S. Bouarfa, B. Jaouani, Z. Lili-Chabaane, and V. Bellon-Maurel (2015): "Environmental Impacts of the Contrasted Groundwater Pumping Systems Assessed by Life Cycle Assessment Methodology: Contribution to the Water-Energy Nexus Study," *Irrigation and Drainage*, 64, 124–138.

Proctor, J., T. Carleton, and S. Sum (2023): "Parameter recovery using remotely sensed variables," NBER Working Paper 30861.

Rajan, A., K. Ghosh, and A. Shah (2020): "Carbon footprint of India's groundwater irrigation," *Carbon Management*, 3, 265–280.

Ramsey, F. P. (1928): "A Mathematical Theory of Saving," *Economic Journal*, 38, 543–59.

Rennert, K., F. Errickson, B. C. Prest et al. (2022): "Comprehensive evidence implies a higher social cost of CO_2," *Nature*, 610, 687–692.

Rhodes, E. C., H. L. Perotto-Baldivieso, E. P. Tanner, J. P. Angerer, and W. E. Fox (2023): "The Declining Ogallala Aquifer and the Future Role of Rangeland Science on the North American High Plains," *Rangeland Ecology and Management*, 87, 83–96.

Richey, A. S., B. F. Thomas, M.-H. Lo, J. S. Famiglietti, S. Swenson, and M. Rodell (2015a): "Uncertainty in global groundwater storage estimates in a Total Groundwater Stress framework," *Water Resources Research*, 51, 5198–5216.

Richey, A. S., B. F. Thomas, M.-H. Lo, J. T. Reager, J. S. Famiglietti, K. Voss, S. Swenson, and M. Rodell (2015b): "Quantifying renewable groundwater stress with GRACE," *Water Resources Research*, 51, 5217–5238.

Robinson, D. (2002): "Construction and Operating Costs of Groundwater Pumps for Irrigation in the Riverine Plain,"Technical report, CSIRO Land and Water.

Rodell, M., J. Famiglietti, D. Wiese, J. Reager, H. Beaudoing, F. Landerer, and M. Lo (2018): "Emerging Trends in Global Freshwater Availability," *Nature*, 557, 651–673.

Rodell, M., J. Chen, H. Kato, J. S. Famiglietti, J. Nigro, and C. R. Wilson (2007): "Estimating Groundwater Storage Changes in the Mississippi River Basin (USA) using GRACE," *Hydrology Journal*, 15, 159–166.

Rodell, M., I. Velicogna, and J. S. Famiglietti (2009): "Satellite-based estimates of groundwater depletion in India," *Nature*, 460, 999–1003.

Ryan, N., and A. Sudarshan (2022): "Rationing the Commons," *Journal of Political Economy*, 130, 210–257.

Scanlon, B. R., C. C. Faunt, L. Longuevergne, R. C. Reedy, W. M. Alley, V. L. McGuire, and P. B. McMahone (2012): "Groundwater depletion and sustainability of irrigation in the US High Plains and Central Valley," *Proceedings of the National Academy of Sciences*, 109, 9320–9325.

Schmied, H. M., D. Caceres, S. Eisner et al. (2021): "The global water resources and use model WaterGAP v2.2d: model description and evaluation," *Geoscientific Model Development*, 14, 1037–1079.

Siebert, S., J. Burke, J. Faures, K. Frenken, J. Hoovegeen, P. Doell, and F. T. Portmann (2010): "Groundwater for Use in Irrigation: A Global Inventory," *Hydrological Earth System Sciences*, 14, 1863–1880.

Siebert, S., M. Kummu, M. Porkka, P. Doll, N. Ramankutty, and B. R. Scanlon (2015): "Historical Irrigation Dataset (HID)," Mar, doi:/10.13019/M20599.

Stock, J. H., and M. Yogo (2003): "Testing for Weak Instruments in Linear IV Regression," *Unpublished Mimeograph, Department of Economics, Harvard University.*

Sutanudjaja, E. H., R. van Beek, N. Wanders et al. (2018): "PCR-GLOBWB 2: A 5 arcmin global hydrological and water resources model," *Geoscience Model Development*, 11, 2429–2453.

Tibshirani, R. (1996): "Regression Shrinkage and Selection via the Lasso," *Journal of the Royal Statistical Society. Series B*, 58, 267–288.

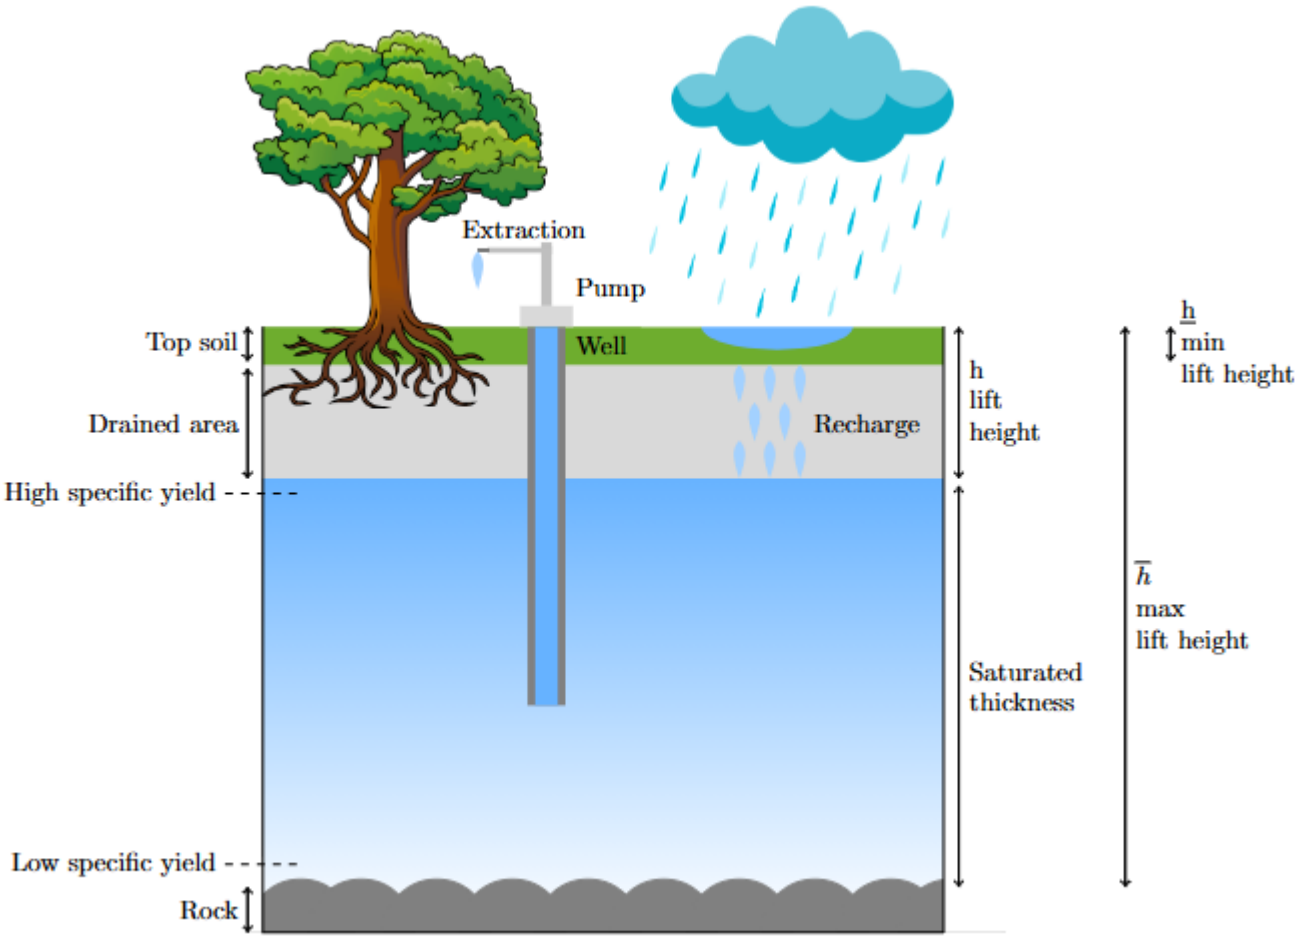
Timmins, C. (2002): "Measuring the Dynamic Efficiency Costs of Regulators' Preferences: Municipal Water Utilities in the Arid West," *Econometrica*, 70, 602–629.

UN (2024): "System of Environmental-Economic Accounting: Ecosystem Accounting,"Technical report, United Nations.

World Bank (2023): "World Bank Development Indicators," Dec, https://www.worldbank.org/en/home.

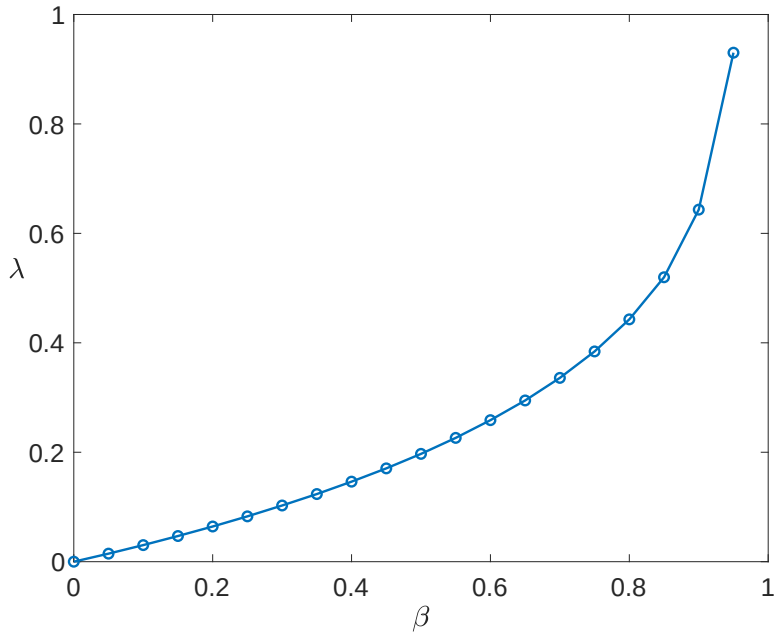
Figures and Tables

Figure 1: Aquifer Terminology



Notes: This figure shows a stylized example of an aquifer to illustrate hydrological terminology used in the text.

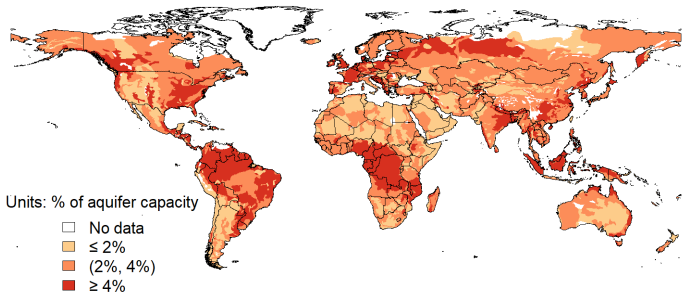
Figure 2: Estimated Relationship Between Discount Factor and Property Rights Interpretations



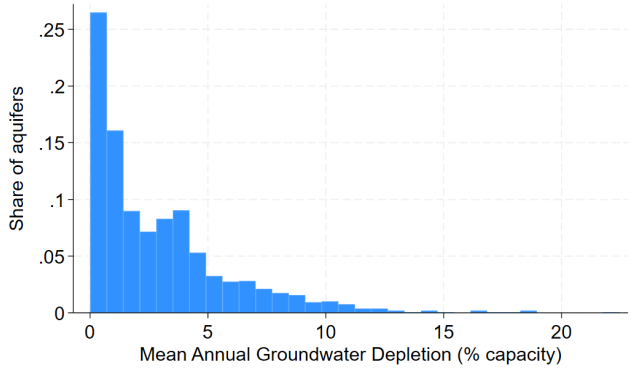
Notes: This figure displays the relationship between the property rights parameter λ and the discount factor β for the Central Valley aquifer. Section 3.4 provides additional details on these two interpretations of our results.

Figure 3: Aquifer Characteristics

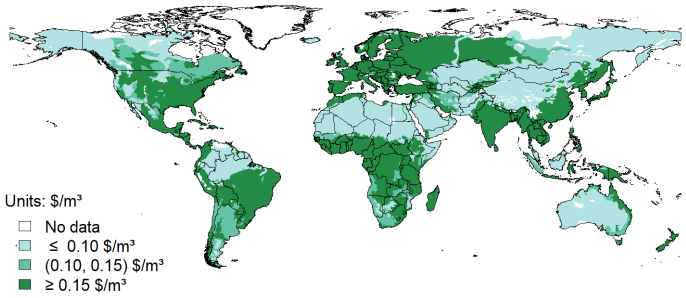
(a) Mean Groundwater Extraction, by Aquifer



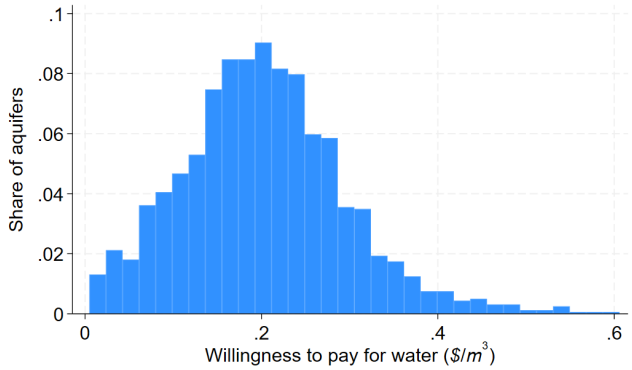
(b) Distribution of Mean Groundwater Use



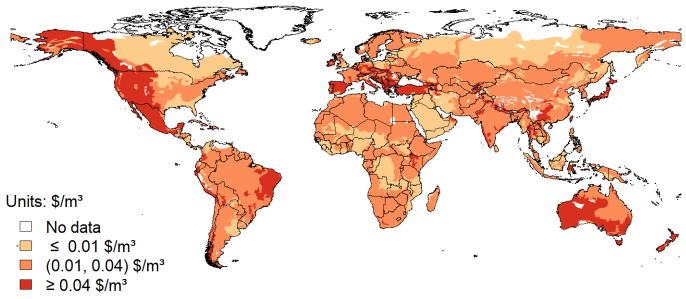
(c) Willingness to Pay for Groundwater, by Aquifer



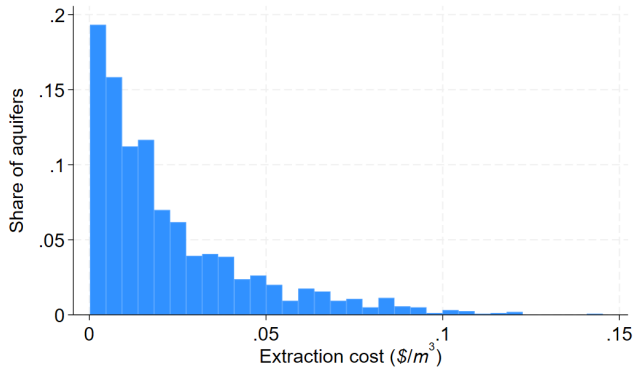
(d) Distribution of Willingness to Pay



(e) Marginal Cost of Groundwater, by Aquifer

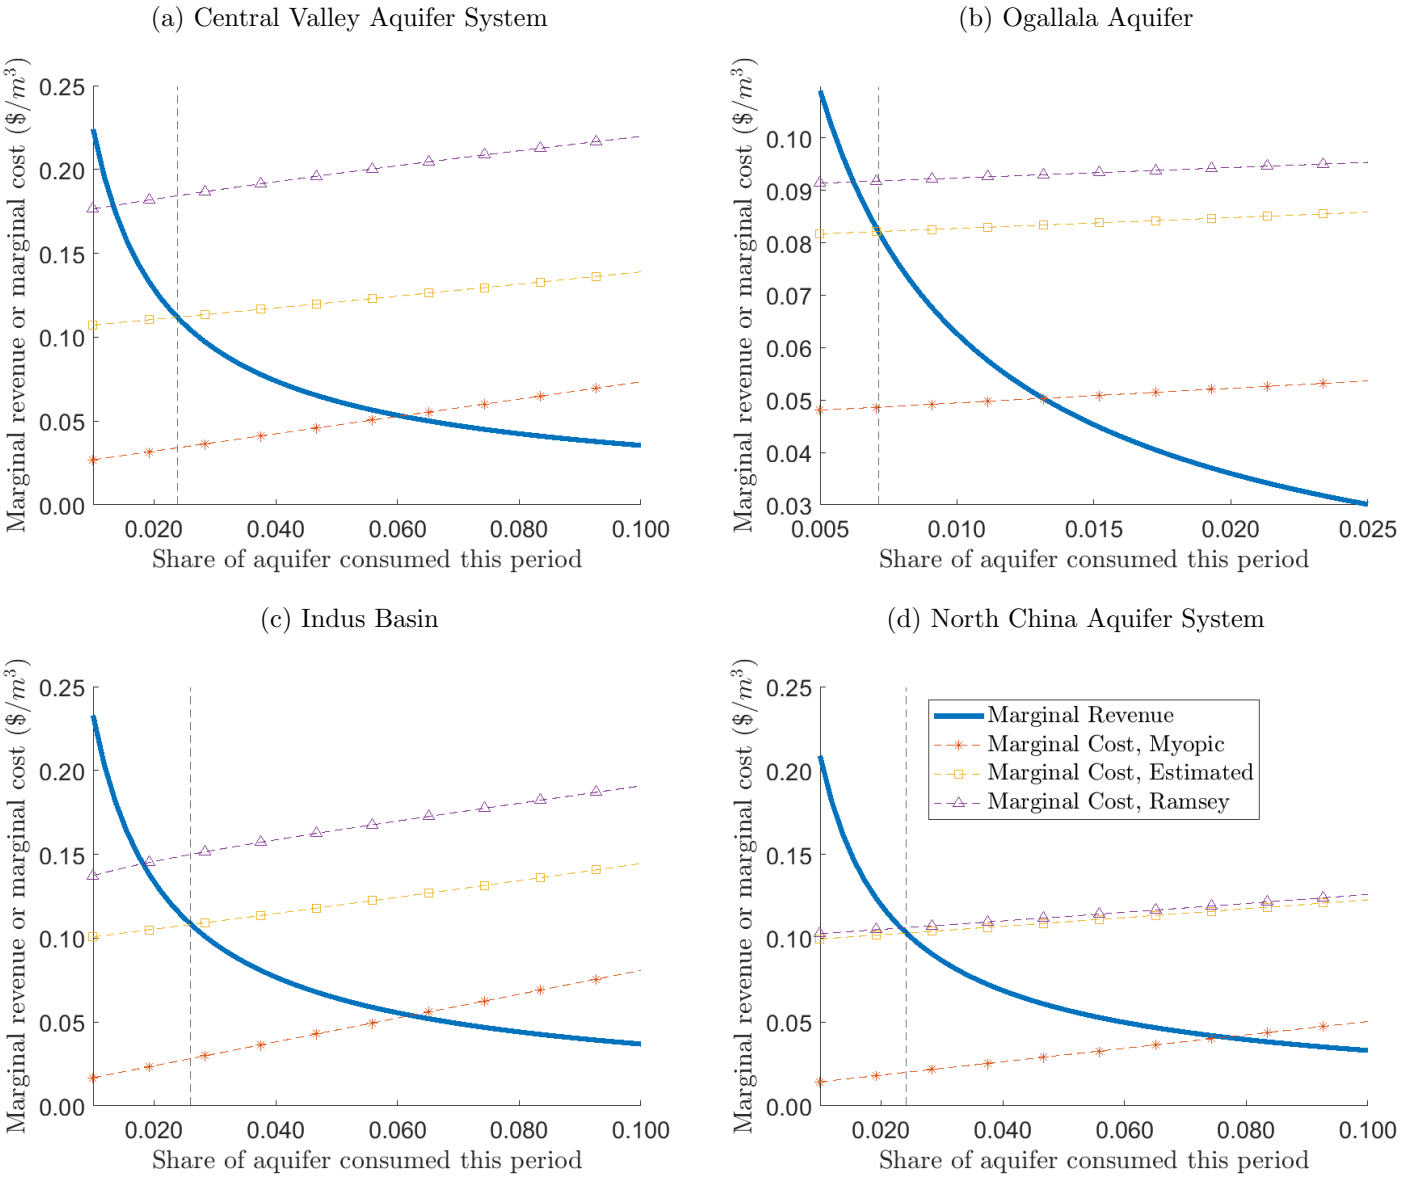


(f) Distribution of Marginal Cost of Groundwater



Notes: These figures show a map and a histogram of mean annual gross groundwater extraction as a share of the aquifer's capacity from GRACE and WHYMAP (Panels A and B), a map and a histogram of willingness to pay per m^3 of groundwater calibrated using data from GAEZ and D'Odorico et al. (2020)(Panels C and D), and a map and a histogram of marginal cost of extracting one m^3 of groundwater calibrated using data from Fan et al. (2013) and the International Energy Agency (Panels E and F). These are aquifer-level estimates averaged over the years 2006–2016. We calculate willingness to pay and marginal cost following the approaches in Sections 5.2 and 5.3 respectively and these are measured in USD/m^3 of groundwater.

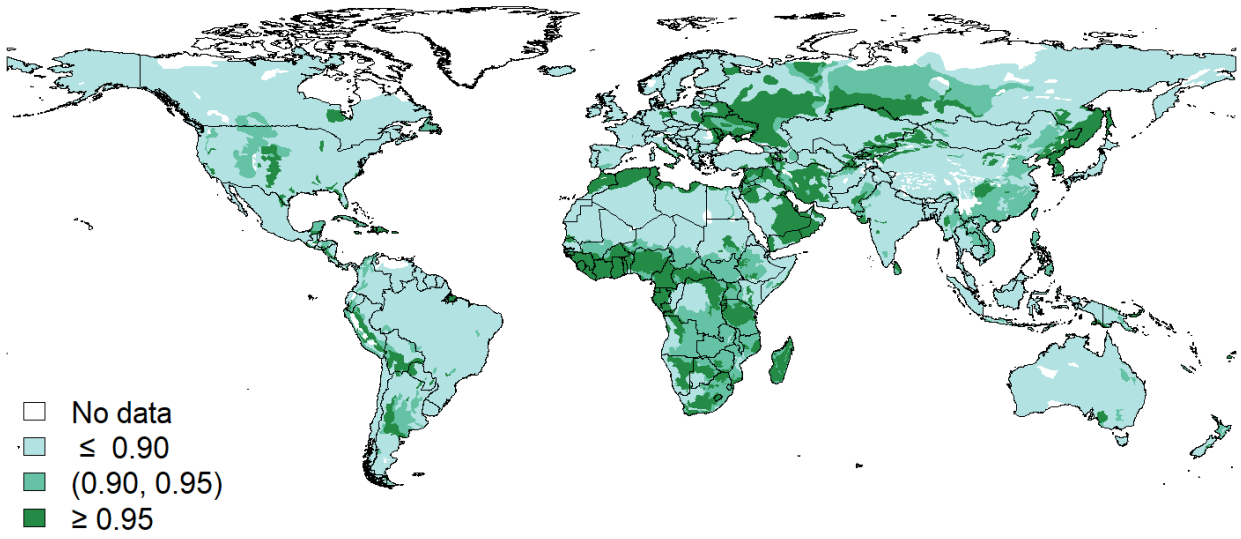
Figure 4: Marginal Cost and Marginal Revenue Curves for Case Study Aquifers



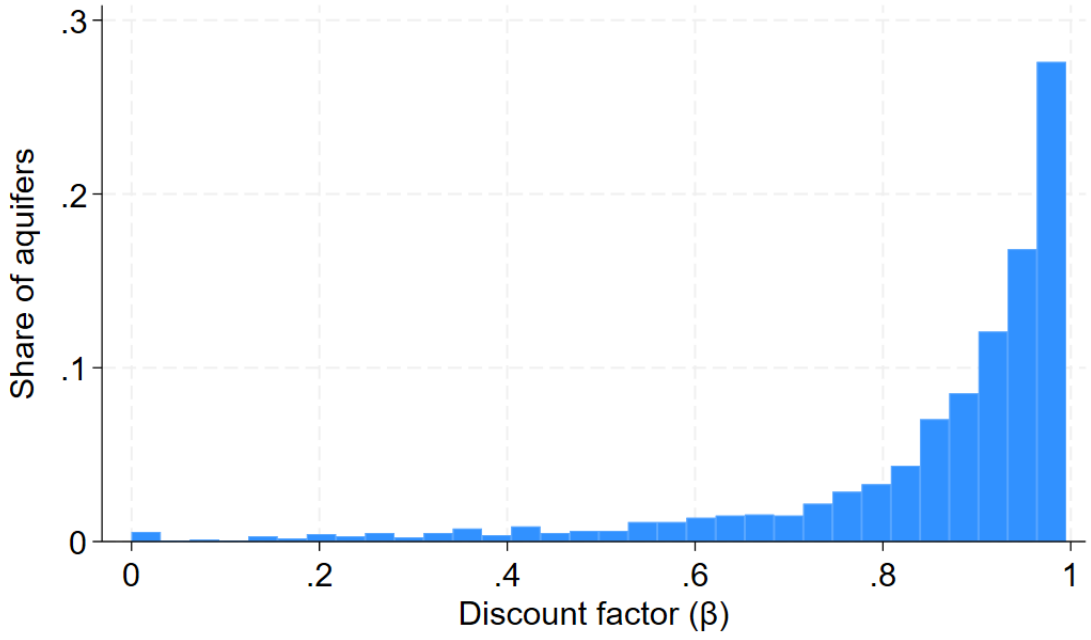
Notes: These figures show the marginal revenue and cost curves for the four case study aquifers as a function of groundwater consumption (w) as a share of aquifer capacity. The dashed vertical line shows observed water use from GRACE. The estimated discount factor β represents the value where marginal revenue equals marginal cost at observed water use (i.e., the β for which observed water use is optimal). Graphs show marginal costs for three scenarios: complete myopia, estimated discount factors, and the socially optimal Ramsey benchmark. With complete myopia, $\beta = 0$ in all four graphs. In the Central Valley aquifer, estimated $\beta = 0.87$ and Ramsey $\beta = 0.96$. In the Ogallala aquifer, estimated $\beta = 0.94$ and Ramsey $\beta = 0.96$. In the Indus aquifer, estimated $\beta = 0.82$ and Ramsey $\beta = 0.90$. In the North China aquifer, estimated $\beta = 0.91$ and Ramsey $\beta = 0.91$.

Figure 5: Estimated Discount Factors

(a) Estimated Discount Factors, by Aquifer



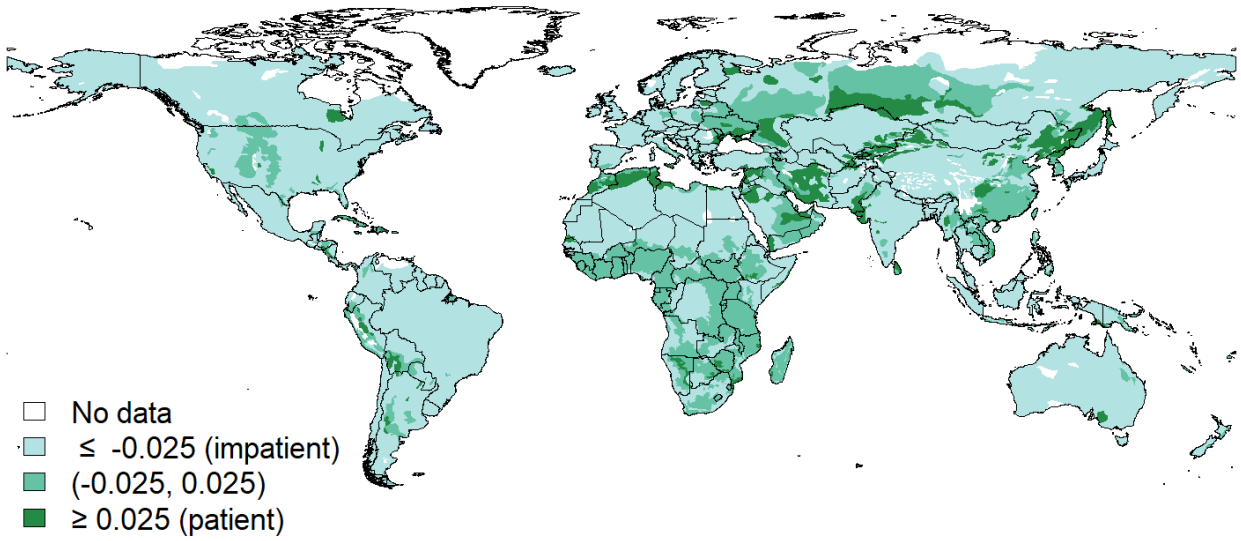
(b) Distribution of Estimated Discount Factors Across Aquifers



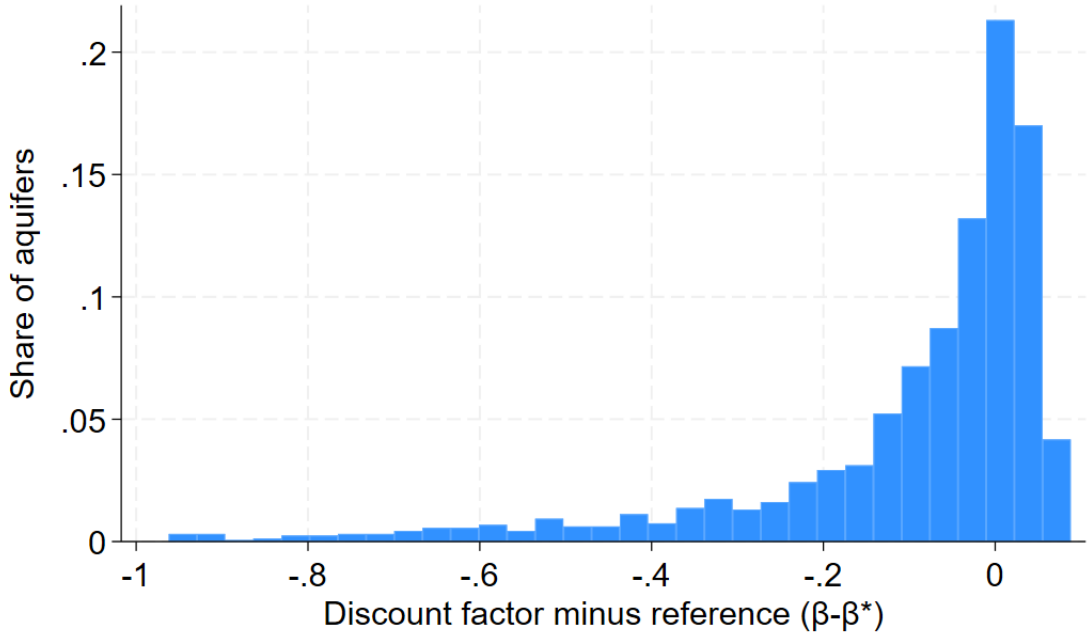
Notes: Panel A shows the estimated discount factor (β) for each aquifer. Panel B shows the distribution of estimated discount factors across aquifers. Section 3 explains the methodology for estimating the discount factor.

Figure 6: Estimated Minus Benchmark Discount Factors, by Aquifer

(a) Estimated Minus Benchmark Discount Factors ($\beta - \beta^*$), by Aquifer

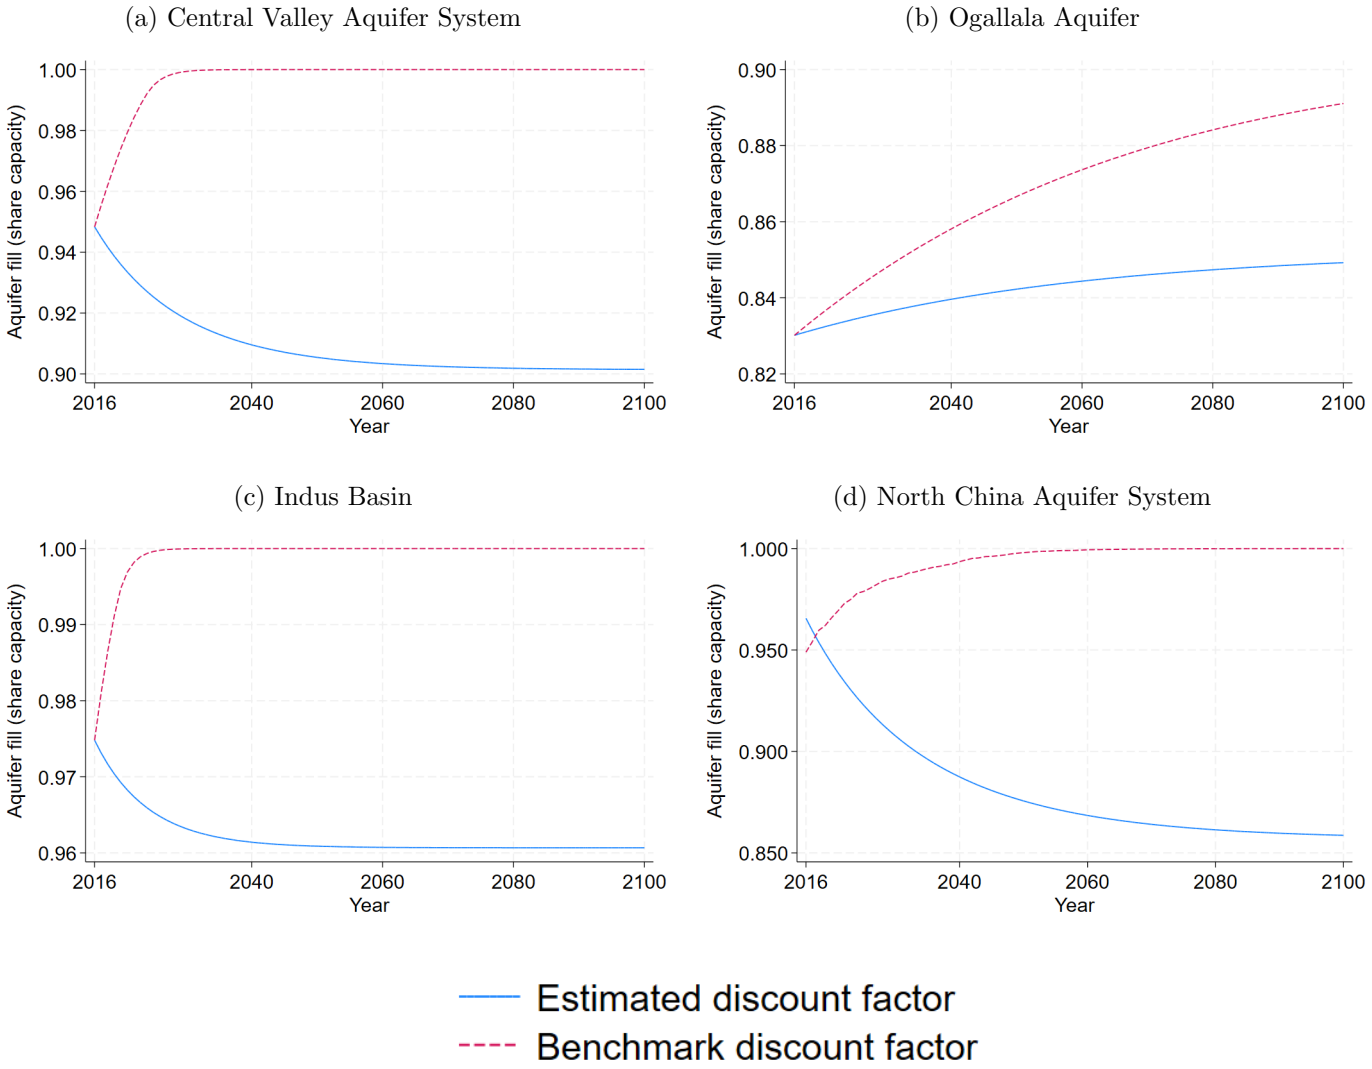


(b) Distribution of Estimated Minus Benchmark Discount Factors



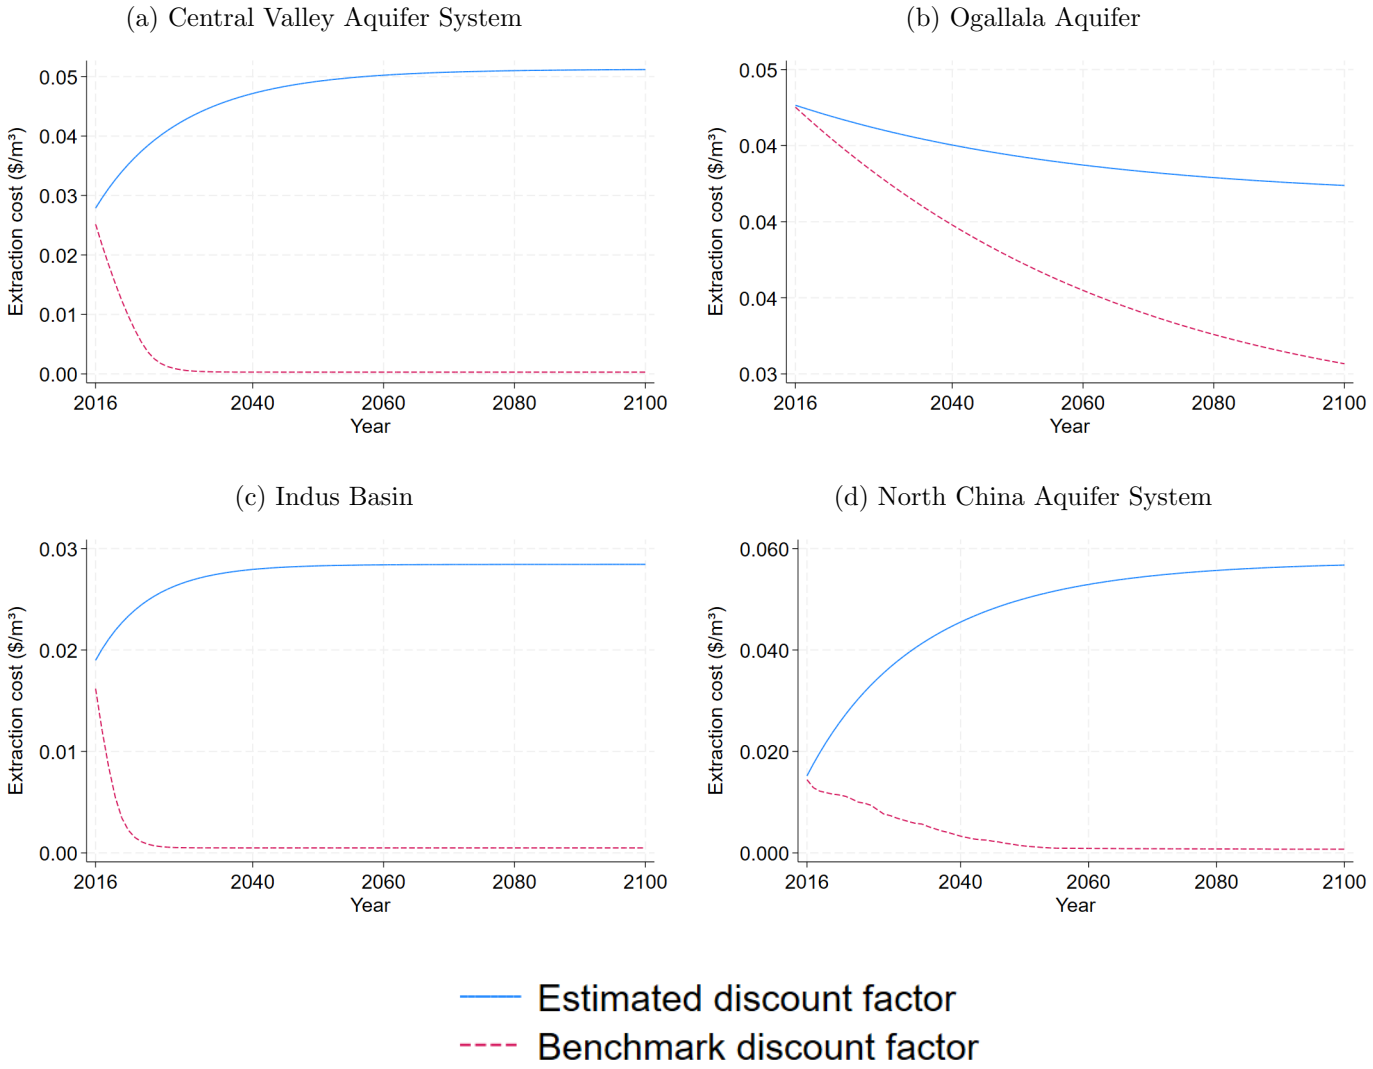
Notes: Panel A maps, for each aquifer, the difference between the discount factor that we estimate using the methodology from Section 3 and the Ramsey benchmark discount factors from Addicott et al. (2020). Panel B graphs the distribution of this difference across aquifers.

Figure 7: Fill as Share of Capacity, by Aquifer and Year



Notes: This figure shows the predicted fill levels (as a share of each aquifer capacity) for each of the four case study aquifers. "Estimated" refers to an agent who discounts the future at the estimated discount factor. "Benchmark" refers an agent who discounts the future at the Ramsey benchmark discount factor.

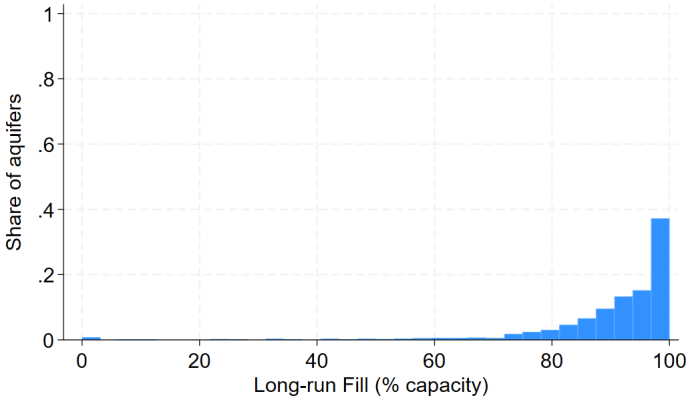
Figure 8: Water Extraction Cost, by Aquifer and Year



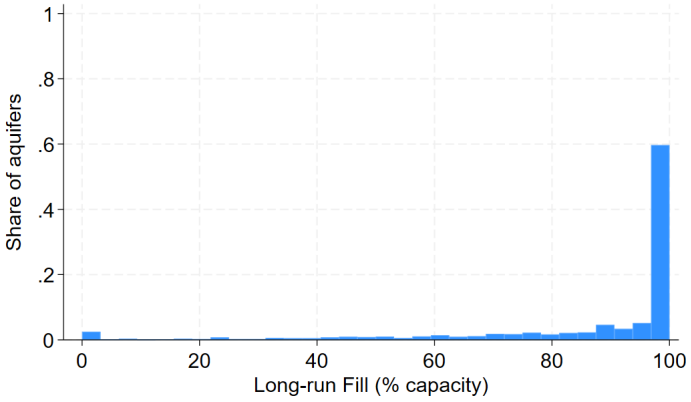
Notes: This figure shows the predicted future water extraction costs ($/m^3) for our four case study aquifers. "Estimated" refers to an agent who discounts the future at the estimated discount factor. "Benchmark" refers an agent who discounts the future at the Ramsey benchmark discount factor.

Figure 9: Long-Run Fill Under Different Discount Factors

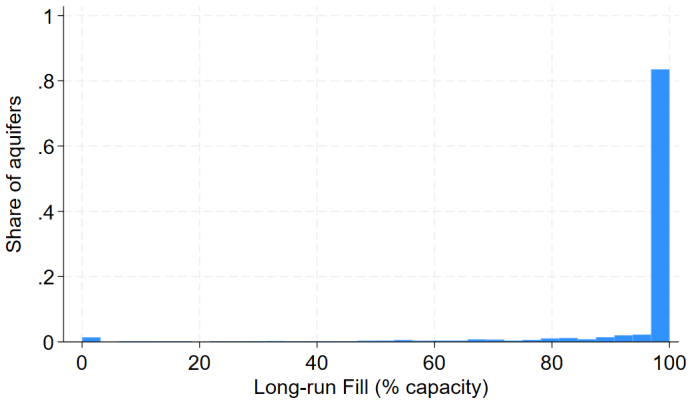
(a) Under Estimated Discounting



(b) Under Ramsey Discounting

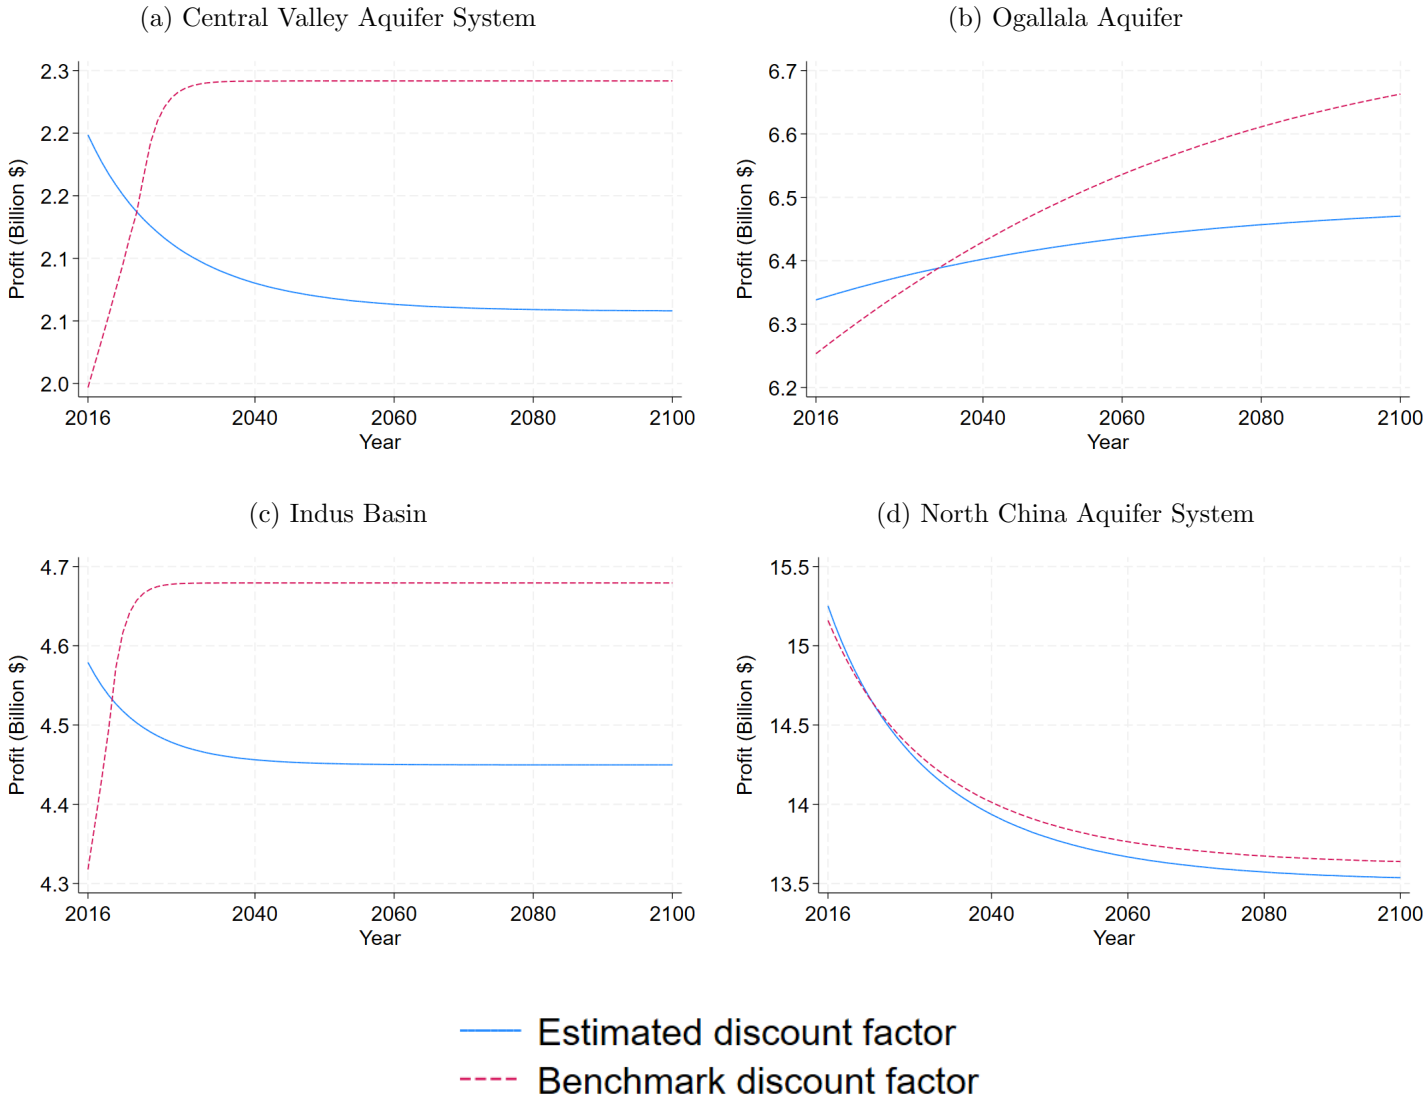


(c) Under Market Discounting



Notes: This figure shows shows histograms of long-run fill (measured as a percent of aquifer capacity) under three different discount factors: estimated discounting in Panel (A), Ramsey discounting in Panel (B) and market discounting in Panel (C).

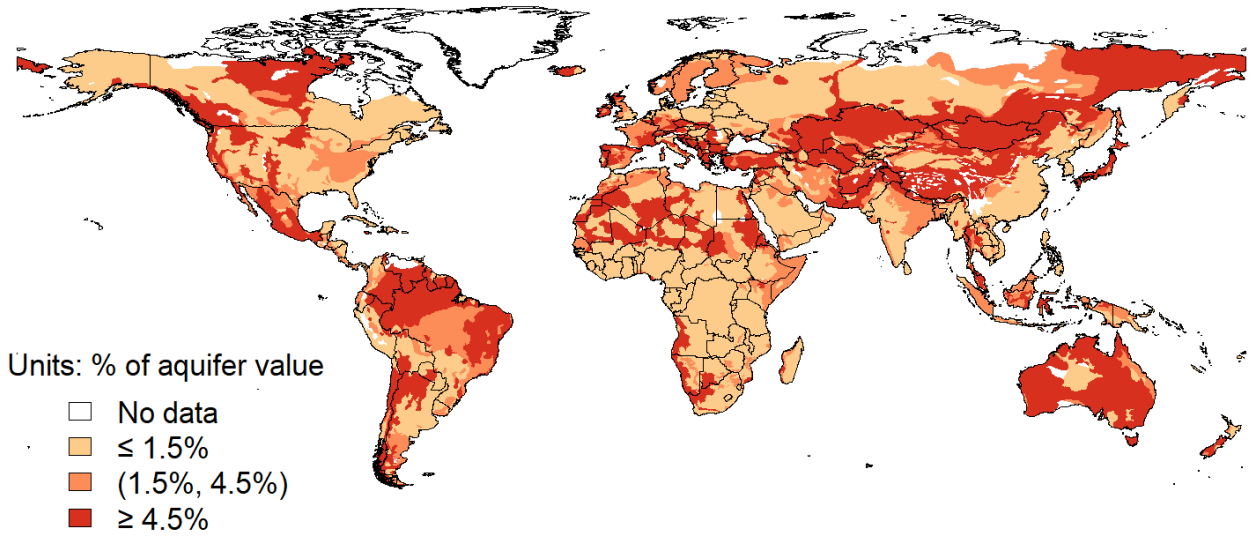
Figure 10: Profit, by Aquifer and Year



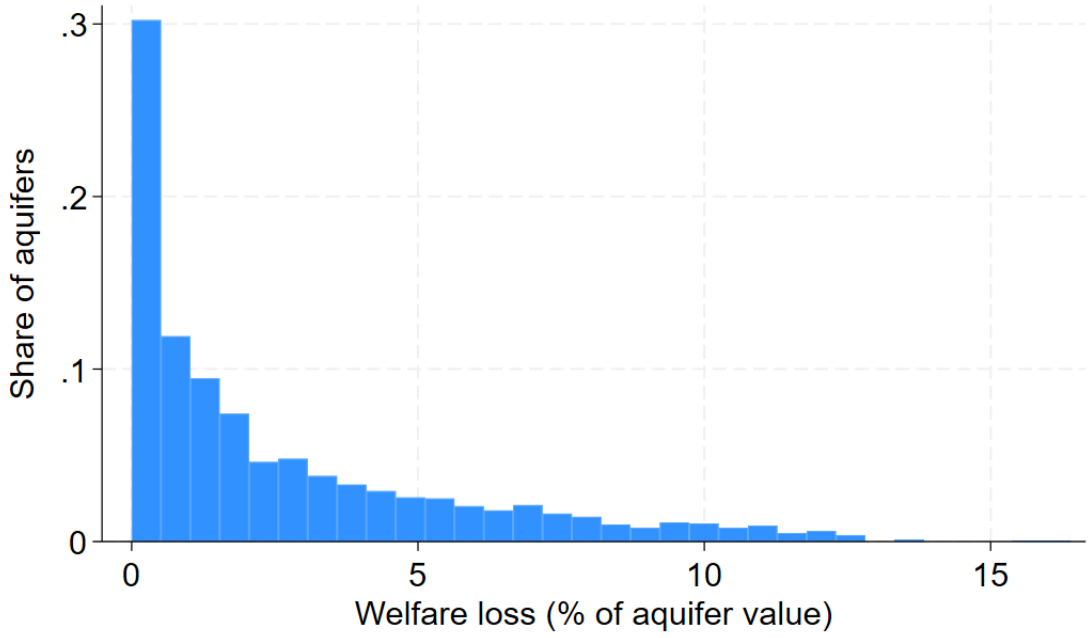
Notes: These figures show predicted annual profit for the four case study aquifers. "Estimated" refers to the case in which water is extracted by an agent who discounts the future at the estimated discount factor. "Benchmark" refers to the case in which water is extracted by an agent who discounts the future at the socially optimal rate using the Ramsey discount factor. Units are billions of 2023 USD.

Figure 11: Welfare Losses Due to Difference Between Estimated and Benchmark Discount Factors

(a) Welfare Losses Due to Deviation from Benchmark Discount Factor, by Aquifer



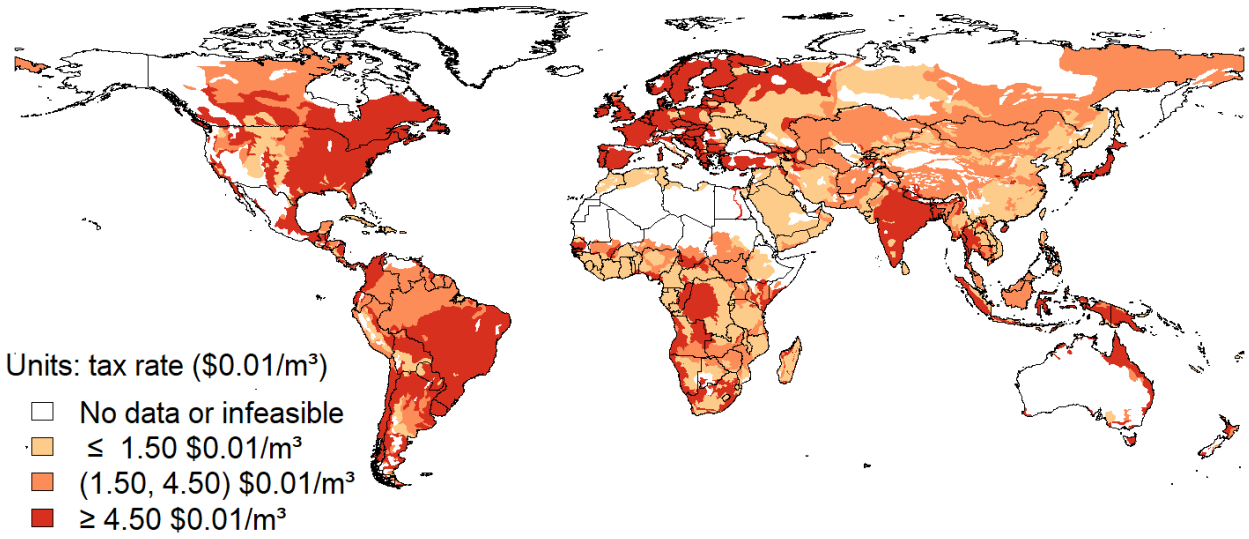
(b) Distribution of Losses Due to Deviation from Benchmark Discount Factor



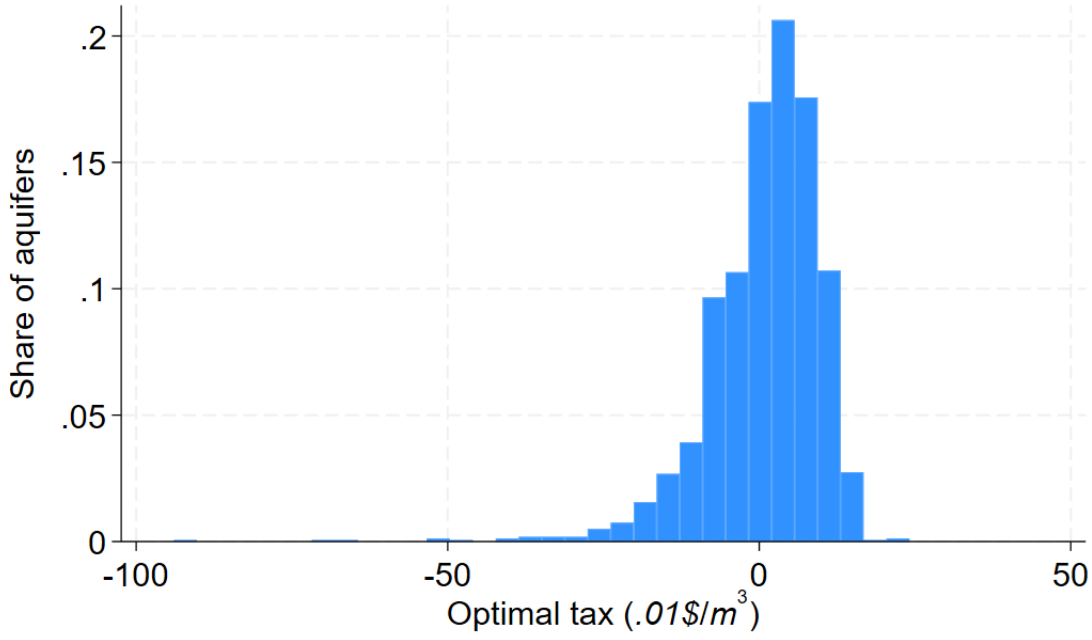
Notes: Panel A maps the welfare loss in each aquifer due to the estimated extraction path deviating from the Ramsey benchmark, measured as a percent of each aquifer's present value. Panel B maps the distribution of these losses across aquifers.

Figure 12: Optimal Time-Invariant Tax Rates on Groundwater Extraction

(a) Optimal Time-Invariant Tax Rates to Maximize Present Value of Aquifer, by Aquifer



(b) Distribution of Optimal Time-Invariant Tax Rates to Maximize Present Value of Aquifer



Notes: Panel A maps the time-invariant tax rate τ which maximizes the present value of each aquifer, calculated following the model in Section 3. Negative values indicate taxes while positive values indicate subsidies. Panel B graphs the distribution of these tax rates across aquifers.

Table 1: Model Parameters and Estimated Effect of Groundwater Extraction on GDP

	(1) OLS	(2) IV, water only	(3) IV, water only	(4) IV, all
ϕ: log(Groundwater)	0.004	0.227***	0.253***	0.103**
	(0.003)	(0.077)	(0.079)	(0.044)
θ: 1 - log(Capital) - log(Labor)	0.863***		0.893***	0.514***
	(0.049)		(0.051)	(0.155)
θ - ϕ: Implied log(Land)	0.858***		0.639***	0.411**
	(0.049)		(0.085)	(0.173)
Aquifer FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Climate Controls	Yes	Yes	Yes	Yes
F stat		21.608	21.105	10.401
Observations	29854	29854	29854	29854

Notes: The outcome variable is log(Night Lights Luminosity) by aquifer-year. Log(Groundwater) is instrumented in models 2-4 and log(Labor) and log(Capital) are instrumented in model 4; the instruments and precipitation and temperature controls are chosen as part of the two-step model selection process described in Section 5.1. Log(Groundwater) is groundwater use in m3. Log(Capital) is the aquifer-specific value of capital in million 2023 USD. 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. Standard errors clustered by aquifer are in parentheses.

Table 2: Model Inputs: Average Aquifers and Case Studies

	Median (1)	Mean (2)	Central Valley (3)	Ogallala (4)	Indus (5)	North China (6)
Capacity (billion m^3)	43.01	275.32	177.60	2449.66	336.62	1262.77
Water: annual use (% capacity)	2.01%	2.93%	2.38%	0.71%	2.60%	2.42%
Willingness to pay for water ($/$m^3$)	0.20	0.20	0.22	0.16	0.21	0.20
Extraction cost ($/$m^3$)	0.02	0.02	0.03	0.05	0.02	0.02
Lift height (m)	17.86	20.96	41.23	41.27	12.64	37.72
Electricity price ($/kWh)	0.07	0.07	0.07	0.07	0.10	0.04
Ramsey discount factor	0.95	0.95	0.96	0.96	0.90	0.91

Notes: This table displays statistics for the model inputs for all aquifers' raw median (Column 1) and mean (Column 2) and for the four case study aquifers. The Ramsey discount factors are from Addicott et al. (2020), re-weighted at the aquifer-level. All dollar values are in 2023 USD.

Table 3: Model Outputs: Discount Factor and Intermediate Outcomes

	Median (1)	Mean (2)	Central Valley (3)	Ogallala (4)	Indus (5)	North China (6)
Panel A: Discount Factor						
Estimated	0.92	0.85	0.87	0.94	0.82	0.91
Estimated - Ramsey discount	-0.03	-0.10	-0.09	-0.02	-0.09	-0.01
Panel B: Net Water Use (% of aquifer capacity)						
Observed	0.00%	0.08%	0.35%	-0.05%	0.16%	0.61%
Ramsey d.f	-0.14%	-0.47%	-0.70%	-0.14%	-0.62%	0.53%
Market d.f	-0.51%	-0.94%	-0.85%	-0.33%	-0.94%	-0.54%
Panel C: Long-Run Fill (share of capacity)						
Estimated	0.94	0.90	0.90	0.85	0.96	0.86
Ramsey d.f	1.00	0.87	1.00	0.90	1.00	0.87
Market d.f	1.00	0.95	1.00	1.00	1.00	1.00
Panel D: Long-Run Cost ($\$/m^3$)						
Estimated	0.02	0.03	0.05	0.04	0.03	0.06
Ramsey d.f	0.00	0.03	0.00	0.03	0.00	0.05
Market d.f	0.00	0.01	0.00	0.00	0.00	0.00

Notes: This table displays model results for the raw median (Column 1) and mean (Column 2) of all global aquifers and the case study aquifer results. We report results using our estimated discount factor, the socially optimal Ramsey discount factor, and the socially optimal market discount factor. d.f. is discount factor. 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%. Many table entries show $share * 100$, e.g. 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 USD.

Table 4: Model Outputs: Welfare and Optimal Taxation

	Median (1)	Mean (2)	Central Valley (3)	Ogallala (4)	Indus (5)	North China (6)
Panel A: Welfare gain from following...						
Ramsey d.f (% of aquifer value)	1.44%	2.66%	3.36%	0.18%	1.34%	0.01%
Market d.f (%)	1.79%	3.08%	5.48%	2.51%	3.99%	4.87%
Ramsey d.f (billion $)	0.06	1.79	1.81	0.29	0.64	0.02
Market d.f (billion $)	0.16	5.62	6.20	8.26	9.28	35.77
Panel B: Market Tax						
Ramsey d.f: ($/m^3)	0.02	0.01	0.10	0.01	0.06	0.00
Market d.f: ($/m^3)	0.06	0.06	0.12	0.05	0.10	0.12
Share of willingness to pay: Ramsey	0.18	0.08	0.47	0.08	0.30	0.02
Share of willingness to pay: market	0.42	0.29	0.56	0.34	0.50	0.60

Notes: This table displays model results for all aquifers. For each statistic, we report the raw median (Column 1) and mean (Column 2) of all global aquifers and the case study aquifer results. We report results using our estimated discount factor, the socially optimal Ramsey discount factor, and the socially optimal market discount factor. d.f. is discount factor. 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%. Many table entries show $share * 100$, e.g. 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 USD.

Table 5: Counterfactual Results: Average Aquifers and Case Studies

	Median (1)	Mean (2)	Central Valley (3)	Ogallala (4)	Indus (5)	North China (6)
Panel A: Net Water Use (%)						
Estimated	0.00%	0.08%	0.35%	-0.05%	0.16%	0.61%
Myopic	3.71%	6.11%	4.03%	0.55%	3.79%	5.85%
Constant Water	0.00%	0.08%	0.35%	-0.05%	0.16%	0.61%
Panel B: Long-Run Fill (%)						
Estimated	94.29%	90.18%	90.14%	85.22%	96.07%	85.69%
Myopic	73.61%	63.07%	76.47%	72.63%	85.23%	68.38%
Constant Water	91.68%	61.70%	0.00%	100.00%	0.00%	0.00%
Panel C: Long-Run Cost						
Estimated	0.02	0.03	0.05	0.04	0.03	0.06
Myopic	0.10	0.10	0.12	0.08	0.11	0.13
Constant Water	0.04	0.14	0.52	0.00	0.71	0.40
Panel D: Welfare Loss (Ramsey) (%)						
Estimated	1.44%	2.66%	3.36%	0.18%	1.34%	0.01%
Myopic	9.37%	9.37%	12.47%	3.96%	10.13%	6.67%
Constant Water	1.77%	3.59%	7.25%	0.11%	1.94%	0.65%

Notes: This table displays model results for the raw median (Column 1) and mean (Column 2) of all global aquifers and for the case study aquifers. We report results using our estimated discount factor, using a discount factor of $\beta = 0$ ("Myopic"), and using constant groundwater extraction equal to observed extraction today ("Constant Water"). Many table entries show $share * 100$, e.g. 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 USD.

Appendix

A Data Details

A.1 General Data Details

The groundwater and recharge data sets both merit additional discussion. We take several steps to obtain annual changes in groundwater storage for each aquifer from the raw gridded GRACE data; these steps closely follow recommendations in the scientific literature to isolate groundwater storage changes from changes in other terrestrial water storage types (e.g., moisture in vegetation and snow) (Richey et al. 2015a; Rodell et al. 2007; Rodell et al. 2009). 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 density 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, available from the Global Land Data Assimilation System (GLDAS); we transform the GLDAS data to the same units as GRACE (i.e., annual deviations from the 2004-2009 mean). Fourth, we intersect the resulting gridded data set with global maps of aquifer extent from the World-wide Hydrological Mapping and Assessment Programme (WHYMAP) 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 exclude them.

The recharge dataset is built 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).[19] These recharge data account for differences in local precipitation, soil porosity, and in-flow from returning surface water runoff.

The recharge data are modeled using the WaterGAP version 2 hydrological model and available through the WHYMAP data set, which assigns each aquifer a binned estimate of annual recharge that takes into account local hydrogeological factors such as precipitation, rock bed porosity, and surface water runoff in the model (Schmied et al., 2021).[20] 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.

Groundwater height is measured differently in the groundwater and the recharge data sets than in the lift height data set. The groundwater and recharge data report the change in "equivalent water thickness"—that is, the distance by which lift height would rise if the rock bed were comprised of 100% water. In practice, the porosity of the rock is well below 100%, and so we divide these data by specific yield to obtain the change in lift height resulting from the change in equivalent water thickness.

[19] To 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 correlation coefficient of 0.86-0.97 between the model-predicted and observational data sets, depending on the parameters used in the model, which suggests that it lines up reasonably well with observational data.

[20] While 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).

A.2 Willingness to Pay and Costs Data

Our proxy for aquifer GDP is gridded remotely sensed (i.e., satellite) data of global luminosity from Li et al. (2020). The harmonized version of the data set that we use, which spans 1992-2018, are recorded by two different satellites: the Defense Meteorological Satellite Program (DMSP) for the years 1992-2013 and the Visible Infrared Imaging Radiometer Suite (VIIRS) for the years 2012-2018. We assess the robustness of our results to analyzing the data separately for each of the two time periods.

We also draw on crop production and the value of 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 the following: banana, barley, cassava, cotton, groundnut, fodder, fruit, maize, millet, oil palm, olives, other cereals, potato, pulses, rapeseed, rice, stimulants, sorghum, soy, sugarbeet, sugarcane, sunflower, tobacco, vegetables, wheat, yams, and one "rest of crops" category. The value of groundwater in irrigation 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.[21] The crops are 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).

The country-level production input, cost, and Ramsey discount factor data from disparate sources differ in their spatial and temporal coverage. For years between 2006 and 2016 that are missing data, we linearly interpolate between the existing observations for countries with missing data. We impute missing country-level data using the annual continent mean.

In addition to the labor data from the World Bank 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 allows us to 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.

A.3 Production Proxies

How well does nighttime luminosity proxy economic activity? These data have 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: a survey of the literature by Gibson et al. (2020) approximates 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 between 0.28 and 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 (Chen and Nordhaus 2011; 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

[21] Specifically, they estimate

$$wtp_i = \frac{P_i \Delta Y_i}{IWR_i}$$

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.

agricultural share of GDP is less than 20%; in our data set, approximately 80% of aquifers are less than 20% cropped.

We also draw on two alternate measures of GDP to use as outcome variables to assess the robustness of our preferred estimates. The first comprises standard-use country-level measures of GDP from the World Bank Development Indicators for the years 2006 to 2016, which we weight by the share of the population employed in agriculture from the same source to proxy for agricultural GDP. We reweight these data to construct aquifer-level estimates using the fraction of the aquifer located in each country, as with the other country-level data sets. The second is gridded, purchasing-power adjusted GDP data from (Kummu et al., 2018), available for the years 1990 to 2015. These gridded data are constructed based by scaling official national statistics, and hence provide only an approximation of the spatial distribution of economic activity at the sub-national level. In general, the availability of global sub-national GDP is limited.

A.4 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 using 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.

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 cropland and irrigated land coverage are from the History Database of the Global Environment (HYDE) and the Aquaknow Historical Irrigation Data set, respectively, 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.

Finally, we note that the raw temperature data from the Climate Research Unit at the University of East Anglia (Harris and Jones 2019) are measured in Fahrenheit. We convert average temperature from degrees Celsius to Fahrenheit to eliminate almost all negative values, allowing us to take logs of temperature rather than relying on inverse hyperbolic sine transformations. The other variables that are measured in logs (wages, interest rates, labor, and capital) are always positive when aggregated at the aquifer-level.

B Additional Background on Water Resources

This appendix provides additional background on the four aquifer systems that we use as case studies. The Central Valley Aquifer System has received particular media and policy attention; it produces 40% of American fresh produce, despite recently experiencing a 20-year drought (Borunda 2022). Groundwater in California has historically been unmetered, which means that farmers can pump as much groundwater as they need with limited to no oversight. Recent drought conditions have contributed to increasing extraction rates over time, with groundwater levels falling by approximately 10% per year since 2008. During these drought years, groundwater use increased from 30-40% of irrigation needs to 60%. To

counter this trend, California passed the Sustainable Groundwater Management Act in 2014, with the goal of achieving sustainable extraction of the Central Valley Aquifer System by 2042. Proposed policies include price mechanisms similar to the taxes that we consider in this paper, as well as non-price policies (Burlig et al. 2021).

The Central Valley Aquifer System and the Ogallala Aquifer together account for approximately 50% of total groundwater withdrawals in the US since 1900. The Ogallala and the Central Valley are ranked first and second nationally for total withdrawals. Despite its larger surface area, the Ogallala is less intensively cultivated than the Central Valley and produces approximately 30% of the country's agricultural output (Rhodes et al. 2023). Groundwater levels in the Ogallala Aquifer have risen in the last 20 years on average, unlike in the Central Valley and the other case studies that we consider (Richey et al. 2015b). However, there is significant variation in depletion rates within the Ogallala Aquifer, with approximately 4% of its land area accounting for approximately one-third of total withdrawals (Scanlon et al. 2012).

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 the major aquifer systems, a trend that is driven primarily by a combination of high demand from irrigated agriculture and low precipitation (Richey et al. 2015b). Depletion is accelerating with limited groundwater regulation, with an increase in withdrawals of 50% predicted by 2050 (Muzammil et al. 2021).

The North China Aquifer System is also among the aquifers with the highest depletion rates in the world. The arid plain above the North China Aquifer System is the largest wheat and maize production zone in China and, as with the other case studies, groundwater extracted from this aquifer is primarily used for irrigated agriculture. Between 2003 and 2010, the aquifer lost approximately 50 km^3 of groundwater, an amount with exceeds the capacity of the country's Three Gorges Dam, the largest power station in the world. The Chinese government is aware of this trend; they are implementing South-North Water Transfer Project, a massive series of canals that, when completed, will transport water from Southern China to the arid region of North China with the goal of reducing reliance on groundwater (Feng et al. 2013).

C Derivation of Equations 3 and 4

Given w, the agent faces the static optimization problem

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

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

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

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

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

which can be written as

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

where $z_i = [\mathcal{Z}_i x_i^{1-\alpha-\gamma-\phi}(\alpha/p_i^k)^\alpha(\gamma/p_i^l)^\gamma]^{1/(1-\alpha-\gamma)}$ and $\theta = 1-\alpha-\gamma$ are functions of parameters. Substituting the first-order conditions into the objective function gives profit net of water costs:

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

D Computational Appendix

D.1 Solving the Water-use Problem (8)

We solve the water extraction problem (8) using value function iteration. We discretize aquifer fill level $a \in [0,1]$ using an evenly spaced grid with 100 nodes. Throughout this Appendix, we suppress aquifer indicators i and denote discretized aquifer fill levels by a_j. Our numerical algorithm is as follows:

1. Guess an initial discretized value function $V^0(a_j)$. 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. The unconstrained version of the first-order condition (9), evaluated at the discretized fill levels a_j, is

$$\phi z w_j^{\phi/\theta-1} - c[\underline{h} + (1 - a_j - r/2)(\overline{h} - \underline{h})] - c(\overline{h} - \underline{h})w_j + \beta\partial_a V^n(\min\{a_j - w_j + r, 1\}) = 0 \quad (23)$$

   Compute the unconstrained water uses w_j^u that satisfy (23).[22] 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 (8) 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

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

   for a numerical tolerance parameter $\varepsilon > 0$, stop.[23] Otherwise, set $n = n + 1$ and return to step (2).

D.2 Welfare Analysis

We compute the welfare loss ω defined in equation (12) as follows. First, we compute the discretized policy function w_j implied by the estimated discount factor β^e using the algorithm described in Appendix D.1. 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

[22] We use Matlab's fzero function to perform this task.
[23] We use $\varepsilon = 10^{-4}$.

sufficiently large that a has (approximately) converged by time T, then we can approximate the value of the estimated water consumption path according to the benchmark discount factor (defined in equation 11) using[24]

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

Once this is done, compute the maximum value of the aquifer according to the benchmark discount factor β^*, $V(a_0;\beta^*)$, using the algorithm described in Appendix D.1. The welfare loss ω can then be computed using equation (12).

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 measure heterogeneity in the importance of groundwater using the percentile of the distribution of groundwater irrigation effect on crop yields (column 1), the share of the aquifer that is cropland (column 2), the share of the aquifer that is irrigated (column 3), and an indicator variable equal to 1 if any agriculture occurs above the aquifer (column 4), if any irrigation takes place (column 5), or if agriculture is the dominant land use above the aquifer (column 6). We then estimate ϕ_i for each aquifer i using the following regression:

$$\log y_{it} = \log z_i + \alpha \log k_{it} + \gamma \log l_{it} + \phi \log w_{it} + \psi \log w_{it} \times v_i + \varepsilon_{it} \tag{24}$$

where v_i measures differences in the value of groundwater across aquifers and the aquifer-specific water input share is $\phi_i = \phi + \psi \times v_i$. We use the same instrumental variables approach as in equation (20) to estimate this equation.

The heterogeneous effects estimates, in Appendix Table A.2, support that groundwater is more valuable in areas with more agriculture. For example, column 1 includes the interaction between groundwater use and the aquifer's percentile in the distribution of groundwater irrigation's effect on crop yields, between 0 and 1; an aquifer in the highest percentile of the distribution has a groundwater elasticity of 0.35, while the corresponding elasticity for an aquifer in the middle of the distribution is 0.11. The results are qualitatively similar for the other measures of aquifer agricultural activity as well: aquifers at the 90th percentile of the distributions of cropland and irrigation have elasticities of 0.41 and 0.19 (Appendix Table A.2, columns 2 and 3, respectively), above the mean of 0.10 across all aquifers (Table 1).

Appendix Table A.3 shows that the input shares are 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 find estimates of the input shares that are similar in sign, magnitude, and precision when we use water inputs from all sources instead of groundwater only, when we use administrative data on agricultural GDP as an outcome variable, when we use data from time periods when the night lights satellite data were measured differently, and when we use 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 2), 0.17 when using re-weighted agricultural GDP as an outcome in the place of night lights (column 3), 0.15 when using the more recent night lights satellite data and 0.09 when using the older satellite data (columns 6 and 7), respectively), 0.10 when we exclude negative temperature instruments, to use logs instead of inverse hyperbolic sine (column 8), and 0.07 when we measure labor using gridded population (column

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

9). We estimate smaller elasticities of approximately 0.03 when we measure groundwater inputs based on net extraction, rather than gross extraction (column 1), when we use gridded GDP available for a more limited set of years as an outcome instead of night lights (column 4), and when we restrict the data set to the deepest aquifers only (column 5), though cannot reject that all of these are equal to our main estimates. Appendix F discusses these alternative specifications in greater detail.

F Robustness

F.1 Technology Parameters

Appendix Table A.3 assesses the robustness of the estimated model parameters ϕ and θ to different definitions of water use 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 $\theta - \phi$.

Columns 1 and 2 estimate the input shares using different definitions of water use. Column 1 defines groundwater use based on net extraction, rather than gross extraction as in Table 1; since net extraction can be negative if recharge exceeds extraction, we use an inverse hyperbolic sine transformation, rather than a log, to obtain the water input share. We find that the point estimate of the water input share is smaller than our main estimate, though statistically indistinguishable from it. Column 2 combines surface water and groundwater use in the measurement of water inputs. 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 1, which is suggestive of slightly higher returns to the surface water component of irrigation. In both alternate 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 3 and 4 define the profit outcome variable differently. Rather than using observational night lights luminosity, column 3 uses administrative data on country-level agricultural GDP and column 4 uses model-based gridded estimates of total GDP adjusted for purchasing power (Kummu et al., 2018), 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 1 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 nightlights 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 4 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 5, 6, and 7 use the same variables as our main estimates, but different data subsamples. Column 5 excludes shallow-water aquifers from the data, while columns 6 and 7 split the data into two time periods corresponding to the two different satellites that measured night lights luminosity, our outcome variable. Column 6 includes 2013-2016 (i.e., the "Visible Infrared Imaging Radiometer Suite (VIIRS)" years) and column 7 includes 2006-2012 (i.e., the "Defense Meteorological Satellite Program (DMSP)" years); the full night lights data set is harmonized across the two time periods, so we expect, and find, similar results across the two time periods. Eight out of the nine estimates of the input shares are statistically indistinguishable from the main estimates, with water input shares that vary between 0.03 and 0.15.

Column 8 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.

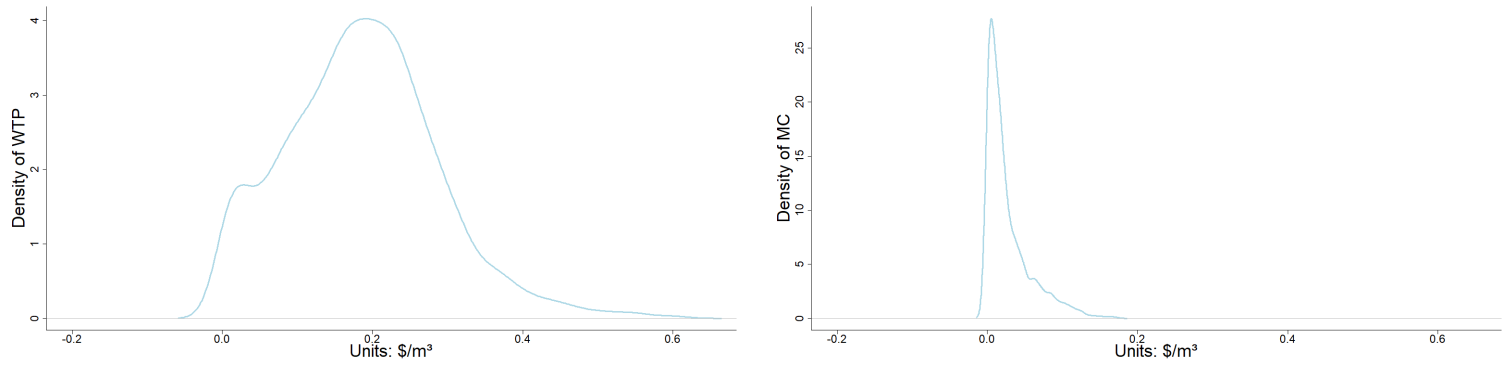
Column 9 uses the same groundwater inputs, capital inputs, and instruments as our main analysis, but replaces our measure of labor with gridded population estimates. Gridded population data are available every five years; we linearly interpolate intervening years. Using this alternate measure of labor, 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.

G Figures and Tables

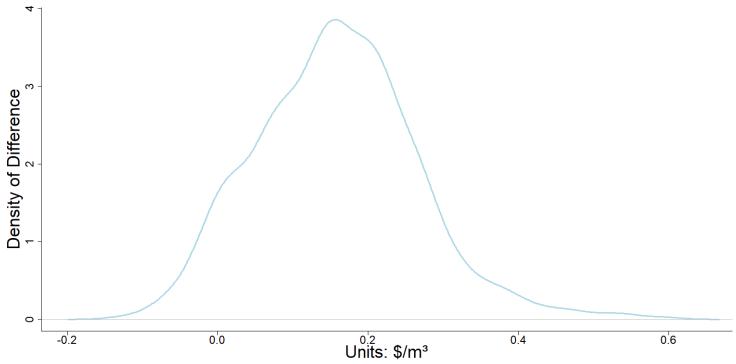
Figure A.1: Kernel Density of Willingness to Pay for and Marginal Cost of Groundwater

(a) Willingness to Pay

(b) Marginal Cost



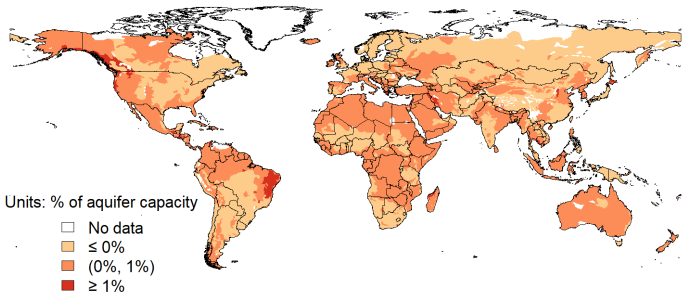
(c) Willingness to Pay Minus Marginal Cost



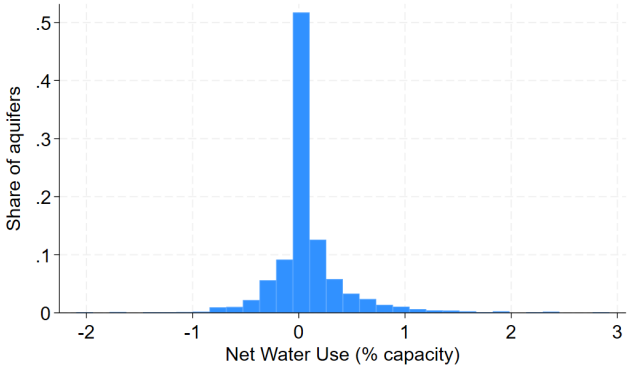
Notes: This figure shows the density of willingness to pay (WTP) for groundwater (Panel A), of marginal cost (MC) of groundwater (Panel B), and of the difference between WTP and MC (Panel C). All estimates are in USD per cubic meter of groundwater extracted.

Figure A.2: Maps of Aquifer Recharge and Extraction

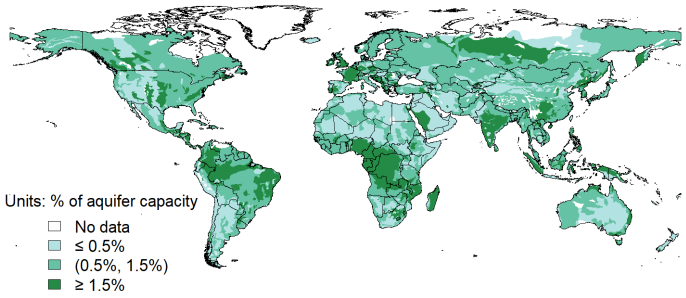
(a) Mean Annual Net Groundwater Use, by Aquifer



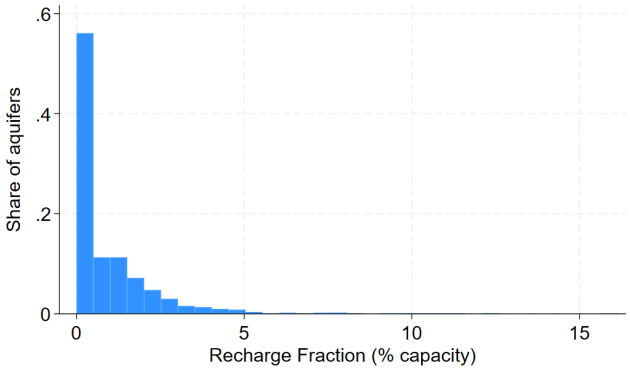
(b) Distribution of Annual Net Groundwater Use



(c) Mean Annual Groundwater Recharge, by Aquifer



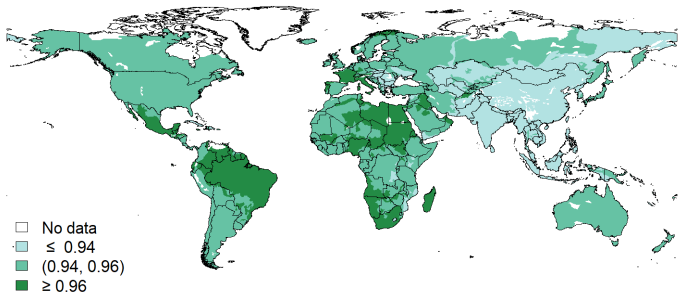
(d) Distribution of Annual Groundwater Recharge



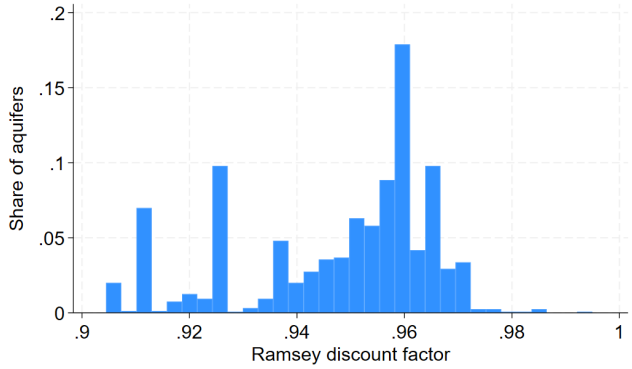
Notes: These figures show maps and histograms of aquifer-level mean net groundwater use as a share of aquifer capacity (Panel A and B) and mean recharge as a share of the aquifer's capacity from the WaterGAP model (de Graaf et al. 2015) (Panel C and D) between 2006 and 2016.

Figure A.3: Additional Aquifer Characteristics

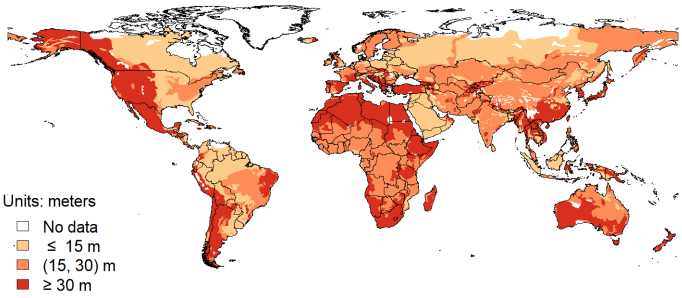
(a) Ramsey Discount Factor, by Aquifer



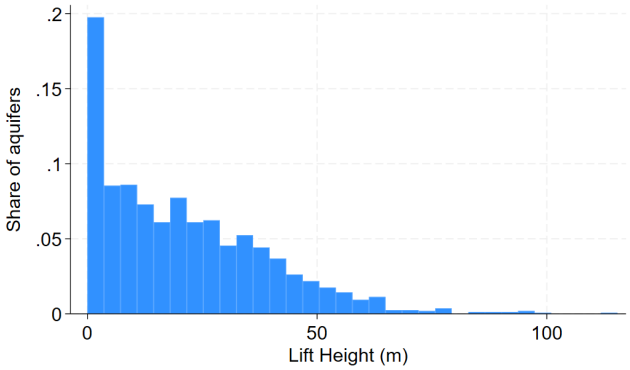
(b) Distribution of Ramsey Discount Factor



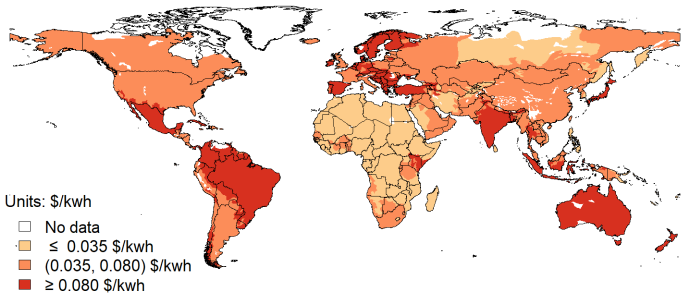
(c) Mean Lift Height, by Aquifer



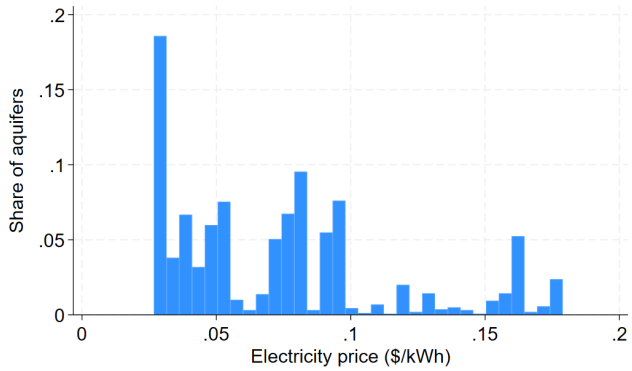
(d) Distribution of Mean Lift Height



(e) Mean Electricity Price, by Aquifer



(f) Distribution of Mean Electricity Price



Notes: These figures show maps and histograms of Ramsey discount factors from Addicott et al. (2020), re-weighted to the aquifer level (Panels A and B), the mean lift height from Fan et al. (2013) (Panels C and D) and mean electricity prices from the International Energy Agency (Panels E and F). The means are aquifer-level averages over 2006 and 2016. The changes in lift height are measured in m and electricity prices are measured in USD per kWh.

Table A.1: First Stage for Instrumental Variables Regressions

	(1) log(Groundwater)	(2) log(Capital)	(3) log(Labor)
Average Temperature, Lag 2	-0.008	-0.001	-0.002**
	(0.006)	(0.001)	(0.001)
Average Temperature, Lag 2 X Porosity	0.001	-0.017*	-0.010*
	(0.042)	(0.009)	(0.005)
Average Temperature, Lag 3 X Porosity	-0.043*	-0.028***	-0.006*
	(0.024)	(0.005)	(0.003)
asinh(Average Temperature, Lag 2)	-0.052***	-0.049***	-0.020***
	(0.019)	(0.004)	(0.002)
log(Total Precipitation, Lag 2) X Porosity	-0.683***	0.276***	0.030
	(0.139)	(0.046)	(0.018)
asinh(Average Temperature, Lag 3)	0.062***	-0.035***	-0.030***
	(0.016)	(0.004)	(0.002)
log(Wage, Lag 1)	-0.002	-0.001	-0.069***
	(0.016)	(0.007)	(0.006)
log(Wage, Lag 2)	0.048***	0.011***	-0.028***
	(0.017)	(0.004)	(0.003)
log(Interest Rate, Lag 2)	0.094***	0.045***	-0.031***
	(0.014)	(0.002)	(0.001)
log(Interest Rate, Lag 1)	-0.073***	0.096***	-0.032***
	(0.012)	(0.004)	(0.002)
log(Wage, Lag 1)	-0.042***	0.025***	-0.014***
	(0.014)	(0.005)	(0.004)
log(Interest Rate, Lag 3)	0.021*	-0.001	0.004**
	(0.012)	(0.004)	(0.002)
Aquifer FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Climate Controls	Yes	Yes	Yes
F stat	10.401	10.401	10.401
Observations	29854	29854	29854

Notes: This table shows the first stage regressions for Table 1. The outcome variables are the log of groundwater use, capital, and labor by aquifer-year, respectively. The instruments and precipitation and temperature controls are chosen as part of the two-step model selection process (see text). Log(Groundwater) is groundwater use in m3. Log(Capital) is the aquifer-specific value of capital in million 2023 USD. Log(Labor) is aquifer population employed in agriculture. Standard errors clustered by aquifer are in parentheses.

Table A.2: Heterogeneous Effects of Groundwater Extraction on GDP

	(1) GAEZ	(2) Cropped %	(3) Irr %	(4) Any Crops	(5) Any Irr	(6) Dominant Ag
ϕ_0: log(Groundwater)	-0.130	-0.210***	-0.004	-0.204*	-0.151	-0.138**
	(0.095)	(0.080)	(0.066)	(0.115)	(0.098)	(0.070)
ϕ_i: log(Groundwater) × Value	0.479**	1.743***	3.266	0.456***	0.451***	1.214***
	(0.194)	(0.447)	(2.011)	(0.177)	(0.171)	(0.346)
θ: 1 - log(Capital) - log(Labor)	0.405**	0.439***	0.487***	0.405**	0.366**	0.456***
	(0.162)	(0.153)	(0.160)	(0.161)	(0.160)	(0.155)
θ - ϕ: Implied log(Land)	0.470***	0.462***	0.487***	0.409**	0.368**	0.464***
	(0.158)	(0.152)	(0.160)	(0.161)	(0.159)	(0.155)
Aquifer FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Climate Controls	Yes	Yes	Yes	Yes	Yes	Yes
F stat	5.082	4.287	2.389	4.505	3.814	3.099
Value Mean		0.108	0.019	0.731	0.533	0.062
Observations	29854	29854	29777	29854	29854	29854

Notes: All models are estimated using IV. The outcome variable is log(Night Lights Luminosity) by aquifer-year. 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 text). Log(Groundwater) is groundwater use in m3. Value is the aquifer-specific value of groundwater, measured by heterogeneity in the variable specified in the column heading. Column 1 uses the aquifer's percentile in the distribution of irrigation's effect on crop yield from GAEZ, column 2 uses the share of the aquifer that is cropped, column 3 uses the share of the aquifer that is irrigated, column 4 uses a binary indicator for any cropped land above the aquifer, column 5 uses a binary indicator for any irrigation above the aquifer, column 6 uses a binary indicator for agriculture as the dominant land use above the aquifer. Log(Capital) is the aquifer-specific value of capital in million 2023 USD. 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. Standard errors clustered by aquifer are in parentheses.

Table A.3: Robustness of Estimated Effects of Groundwater Extraction on GDP

	(1) Net Extraction	(2) Total Water	(3) Ag GDP	(4) Grid GDP	(5) Deep Aq	(6) VIIRS	(7) DMSP	(8) log IVs	(9) log(Pop)
ϕ: asinh(Net Groundwater)	0.021**								
	(0.011)								
ϕ: log(All Water)		0.121**							
		(0.049)							
ϕ: log(Groundwater)			0.169***	0.024	0.024	0.151*	0.098***	0.100**	0.073
			(0.024)	(0.063)	(0.044)	(0.088)	(0.018)	(0.042)	(0.045)
θ: 1 - log(Capital) - log(Labor)	0.507***	0.479***	0.551***	0.440*	0.645***	2.410***	0.612***	0.592***	0.138
	(0.146)	(0.151)	(0.134)	(0.239)	(0.140)	(0.360)	(0.079)	(0.162)	(0.173)
θ - ϕ: Implied log(Land)	0.486***	0.357**	0.383***	0.416*	0.620***	2.259***	0.513***	0.492***	0.066
	(0.147)	(0.173)	(0.134)	(0.242)	(0.162)	(0.350)	(0.084)	(0.178)	(0.191)
Aquifer FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Climate Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
F stat	12.872	9.138	10.401	9.103	9.267	4.237	7.293	10.830	12.741
N	29854	29854	29854	27248	23244	10717	19137	29772	29854

Notes: All models are estimated using IV. The outcome variable is log(Night Lights Luminosity) by aquifer-year. 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 text). Log(Groundwater) is groundwater use in m3. Log(Capital) is the aquifer-specific value of capital in million 2023 USD. 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 uses net groundwater extraction instead of gross groundwater extraction, column 2 combines groundwater and surface water withdrawals, column 3 uses agricultural GDP (re-weighted to the aquifer-level from the country-level) as an outcome variable, column 4 uses gridded purchasing power parity GDP as an outcome variable, column 5 excludes shallow water aquifers from the analysis sample, columns 6 and 7 split the sample into the two time periods where two different satellite operations collected night lights data, column 8 uses logs of the IVs instead of inverse hyperbolic sine and excludes aquifers with negative average temperatures, and column 9 uses the log of gridded population instead of agricultural labor. Standard errors clustered by aquifer are in parentheses.

Table A.4: Correlations between the Estimated Discount Factor (β) and Aquifer Characteristics

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Model Inputs										
Log(WTP)	0.580***									0.766***
	0.128									0.118
Log(Marginal Cost)		-3.085***								-3.181***
		0.440								0.633
Log(Elec. Price)			-1.594***							-0.629**
			0.274							0.259
Log(Lift Height)				-0.030***						0.002
				0.008						0.010
Ramsey d.f. and Components										
Ramsey Discount Factor					0.098					0.581
					0.891					0.509
Pure Rate of Time Pref.						0.028				0.048**
						0.040				0.024
Growth Rate, GDP/Pop							-0.001			0.001
							0.003			0.002
Other Country Characteristics										
Log(GDP/Pop)								-0.009***		-0.010***
								0.003		0.003
Log(Democracy)									-0.124***	-0.063**
									0.044	0.030
Continent FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1600	1600	1600	1600	1600	1600	1600	1600	1600	1600

Notes: The outcome variable is the aquifer-level estimate of the discount factor (β). The Ramsey discount factor and the pure rate of time preference are from Addicott et al. (2020), re-weighted to the aquifer-level. d.f. is discount factor. All models are estimated using OLS. Willingness to pay and marginal cost are in $\$/m^3$, electricity prices are in $\$/kWh$, and lift height is in m. Robust standard errors are in parentheses.

Table A.5: Model Inputs: Medians

	Crop-Weighted (1)	GDP-Weighted (2)	Population-Weighted (3)	Major Unweighted (4)
Capacity (billion m^3)	1044.42	17.58	1044.42	369.97
Water: annual use (% capacity)	3.85%	2.03%	3.60%	3.15%
Willingness to pay for water ($/$m^3$)	0.23	0.22	0.23	0.19
Extraction cost ($/$m^3$)	0.02	0.01	0.02	0.02
Lift height (m)	19.83	10.23	20.33	19.16
Electricity price ($/kWh)	0.07	0.05	0.06	0.05
Ramsey discount factor	0.95	0.96	0.94	0.96

Notes: This table displays statistics for the model inputs for all aquifers and the subset of major aquifer systems. For each statistic, we report the global aquifer crop value-weighted median (Column 1), the global aquifer GDP-weighted median (Column 2), the global aquifer population-weighted median (Column 3), and the major aquifer (unweighted) median (Column 4). The Ramsey discount factors are from Addicott et al. (2020), re-weighted to the aquifer-level. All dollar values are in 2023 USD.

Table A.6: Model Outputs: Discount Factor and Intermediate Outcomes, Medians

	Crop-Weighted (1)	GDP-Weighted (2)	Population-Weighted (3)	Major Unweighted (4)
Panel A: Discount Factor				
Estimated	0.86	0.93	0.88	0.92
Estimated - Ramsey discount	-0.06	-0.01	-0.06	-0.04
Panel B: Net Water Use (% of aquifer capacity)				
Observed	0.01%	0.02%	0.08%	0.02%
Ramsey d.f	-0.86%	-0.00%	-0.50%	-0.49%
Market d.f	-1.27%	-0.41%	-1.13%	-0.98%
Panel C: Long-Run Fill (share of capacity)				
Estimated	0.92	0.94	0.92	0.93
Ramsey d.f	1.00	0.96	1.00	1.00
Market d.f	1.00	1.00	1.00	1.00
Panel D: Long-Run Cost ($\$/m^3$)				
Estimated	0.02	0.02	0.03	0.02
Ramsey d.f	0.00	0.02	0.00	0.00
Market d.f	0.00	0.00	0.00	0.00

Notes: This table displays model results for all aquifers and the subset of major aquifer systems. For each statistic, we report the global aquifer crop value-weighted median (Column 1), the global aquifer GDP-weighted median (Column 2), the global aquifer population-weighted median (Column 3), and the major aquifer (unweighted) median (Column 4). We report results using our estimated discount factor, the socially optimal Ramsey discount factor, and the socially optimal market discount factor. d.f. is discount factor. 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%. Many table entries show $share * 100$, e.g. 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 USD.

Table A.7: Model Outputs: Welfare and Optimal Taxation, Medians

	Crop-Weighted (1)	GDP-Weighted (2)	Population-Weighted (3)	Major Unweighted (4)
Panel A: Welfare gain from following...				
Ramsey d.f (% of aquifer value)	0.91%	1.68%	0.91%	1.25%
Market d.f (%)	2.84%	0.99%	3.39%	1.61%
Ramsey d.f (billion $)	3.12	0.02	1.95	0.56
Market d.f (billion $)	25.16	0.05	22.82	1.60
Panel B: Market Tax				
Ramsey d.f: ($/m^3)	0.06	0.00	0.04	0.03
Market d.f: ($/m^3)	0.10	0.07	0.11	0.07
Share of willingness to pay: Ramsey	0.30	0.03	0.23	0.27
Share of willingness to pay: market	0.46	0.39	0.47	0.43

Notes: This table displays model results for all aquifers and the subset of major aquifer systems. For each statistic, we report the global aquifer crop value-weighted median (Column 1), the global aquifer GDP-weighted median (Column 2), the global aquifer population-weighted median (Column 3), and the major aquifer (unweighted) median (Column 4). We report results using our estimated discount factor, the socially optimal Ramsey discount factor, and the socially optimal market discount factor. d.f. is discount factor. 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%. Many table entries show $share * 100$, e.g. 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 USD.

Table A.8: Counterfactual Results: Medians

	Crop-Weighted (1)	GDP-Weighted (2)	Population-Weighted (3)	Major Unweighted (4)
Panel A: Net Water Use (%)				
Estimated	0.01%	0.02%	0.08%	0.02%
Myopic	5.55%	6.88%	5.34%	5.03%
Constant Water	0.01%	0.02%	0.08%	0.02%
Panel B: Long-Run Fill (%)				
Estimated	92.49%	94.04%	92.49%	93.37%
Myopic	71.53%	59.56%	71.20%	67.46%
Constant Water	78.65%	79.17%	12.60%	79.20%
Panel C: Long-Run Cost				
Estimated	0.02	0.02	0.03	0.02
Myopic	0.12	0.11	0.11	0.09
Constant Water	0.05	0.08	0.15	0.06
Panel D: Welfare Loss (Ramsey) (%)				
Estimated	0.91%	1.68%	0.91%	1.25%
Myopic	10.43%	7.95%	9.98%	10.69%
Constant Water	1.44%	1.70%	1.26%	1.73%

Notes: This table displays model results for all aquifers and the subset of major aquifer systems. For each statistic, we report the global aquifer crop value-weighted median (Column 1), the global aquifer GDP-weighted median (Column 2), the global aquifer population-weighted median (Column 3), and the major aquifer (unweighted) median (Column 4). We report results using our estimated discount factor, using a discount factor of $\beta = 0$ ("Myopic"), and using constant groundwater extraction equal to observed extraction today ("Constant Water"). Many table entries show $share * 100$, e.g. 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 USD.