Groundwater in a changing climate

R. Aaron Hrozencik ⓡ Matt Woerman ⓡ Fiona Burlig ⓡ Louis Preonas*

March 27, 2026

Preliminary and incomplete
Please do not cite or circulate without permission

Abstract

To what extent is climate adaptation privately valuable, and how costly is this adaptation for society? We study these questions in agriculture, a highly vulnerable sector where irrigation is a key method of adaptation. Using rich data on groundwater pumping from California, we estimate that farmers pump substantially more water in response to extreme heat. However, this adaptation comes at a social cost, increasing groundwater depths, which increases future pumping costs. We develop a dynamic simulation model in which groundwater pumping responds to both temperature realizations and aquifer depth per our reduced-form estimates, and where aggregate extraction feeds back into future aquifer depth. Under climate change, aquifer depth increases by an additional 68–162% by 2100, compared to a world without temperature-induced adaptive irrigation. Aggregate welfare losses across the aquifer will exceed $1–3 billion annually by end-of-century. In ongoing work, we quantify the impact of climate change on the open access externality, which a simple model shows is theoretically ambiguous.

Keywords: groundwater, agriculture, climate adaptation
JEL Codes: Q15, Q25, Q41

*Hrozencik: University of Georgia. Conner Hall, 147 Cedar St, Athens, GA 30602. Email: Aaron.Hrozencik@uga.edu. Woerman: Department of Agricultural and Resource Economics, Colorado State University. Nutrien Building, 301 University Avenue, Fort Collins, CO 80523. Email: matt.woerman@colostate.edu Burlig: University of Chicago Harris School of Public Policy and NBER. Keller Center, 1307 E. 60th Street, Chicago, IL 60637. Email: burlig@uchicago.edu. Preonas: Department of Agricultural and Resource Economics, University of Maryland. 2200 Symons Hall, 7998 Regents Drive, College Park, MD 20742. Email: lpreonas@umd.edu. The order in which the authors' names appear has been randomized using the AEA Author Randomization Tool (mw0-7IwOF3ft), denoted by ⓡ. We thank Nick Hagerty for graciously sharing data with us. For helpful comments and suggestions, we thank Chris Costello, Koichiro Ito, Erin Kelley, Ashley Langer, and seminar participants at UC Berkeley, the University of Chicago, the Coase Project, and Montana State. Camellia Ye provided truly excellent research assistance. Karen Notsund provided invaluable support in obtaining data. We received generous funding in support of this project from the USDA Economic Research Service under Cooperative Agreement 58-6000-0-0052 and NIFA Grant 2023-09508. All remaining errors are our own.

1 Introduction

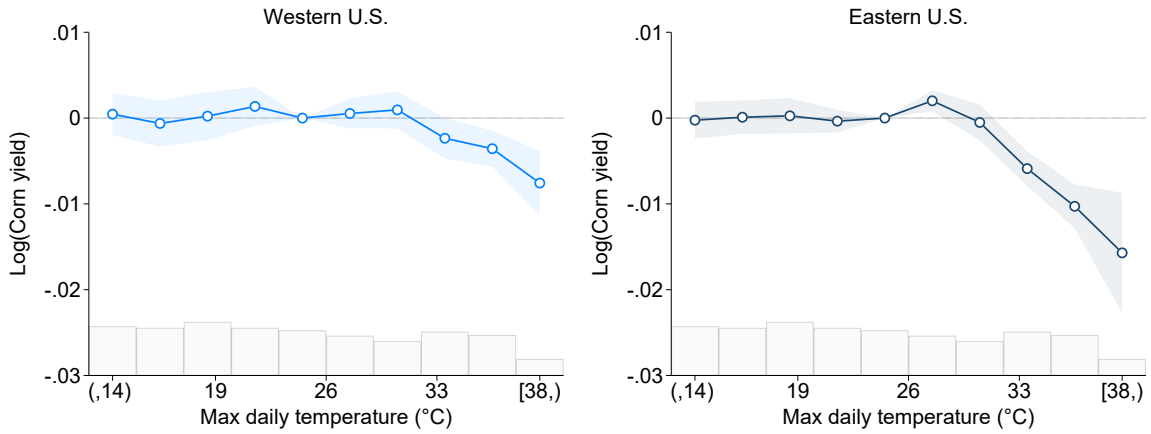
Anthropogenic climate change will impose enormous economic costs (Hsiang et al. (2017); Nordhaus (2019)). Current trajectories of greenhouse gas (GHG) emissions indicate that mitigation alone will not be sufficient to avoid these costs (Lee et al. (2023)). As such, adaptation to an evolving climate will be vital in averting future costs and ensuring continued human well being (Carleton et al. (2024)). Despite its growing importance, there exists relatively little empirical research focused on climate change adaptation.

Climate change adaptation is the set of actions taken in direct response to current or projected changes in the climate (Boser et al. (2024)). Though it may provide important private protective benefits, adaptation comes at a cost. Prior work has focused on the private costs of adaptation (e.g., Carleton et al. (2022) and Hultgren et al. (2025)), but many important forms of adaptation also impose costs on society. For example, when households cope with hot temperatures by using air conditioning, this increases local pollution and speeds climate change (Auffhammer (2022) and Abajian et al. (2025)). We focus on the interaction between climate adaptation and non-renewable resources. As resource stocks become depleted, marginal extraction costs rise. Thus, adaptation via resource use can exacerbate well-known open access externality problems (Ostrom (1990), Provencher and Burt (1993)).

We study this interaction in the context of agriculture, a particularly climate-exposed sector (Mendelsohn, Nordhaus, and Shaw (1994) and Hultgren et al. (2025)). Agriculture has access to a powerful adaptation technology: irrigation, which protects plants against hot temperatures by keeping soil moist (Tack, Barkley, and Hendricks (2017); Li et al. (2019); Li et al. (2020); Luan et al. (2021)). Figure 1 uses standard climate change econometric techniques to show how the prevalence of irrigation can mitigate yield losses caused by hot temperatures. In the western U.S., where the majority of cropland is irrigated, the impacts of hot temperatures on yields is meaningfully attenuated relative to the eastern U.S., where agriculture is mostly rainfed. At the same time, groundwater is a key source of irrigation water, but aquifers are rapidly being depleted (Jasechko et al. (2024)), potentially threatening

both high-value crop production and climate change adaptation (Schlenker, Hanemann, and Fisher (2007) and Carleton, Crews, and Nath (2023)).

Figure 1: Effect of temperature on yields differs in the Western vs. Eastern U.S.



Notes: This figure plots the results of a regression of log corn yield against a flexible function of temperature. The left panel uses data from counties west of the 100th Meridian, and is restricted to counties where more than 50% of cropland is irrigated, while the right panel uses data from east of the 100th Meridian, where cropland is largely rain-fed. Data come from the USDA's National Agricultural Statistics Service county-level yield estimates from 1980-2020. Regressions include a growing season precipitation control, county fixed effects, and year fixed effects. Standard errors are clustered by county.

In this paper, we take the importance of irrigation, its role as an adaptation tool, and the dynamics defining groundwater extraction costs as a starting point for answering two research questions. First, to what extent do farmers adapt to high temperatures through irrigation, and how privately valuable is this adaptation? Second, how costly is climate change for society, accounting for the open-access "stock externality" (Provencher and Burt (1993)) whereby private farmers fail to internalize the effect of current groundwater use on future aquifer depletion?

We focus our analysis on California's Central Valley, one of the world's most valuable crop-producing regions, which is also on the front lines of global groundwater depletion due to over-extraction under an open-access regime. We start by developing a simple theoretical model of optimal groundwater management, which shows that the impact of climate change on the welfare losses due to open access externalities are ambiguous. Next, we leverage rich panel data on well-level groundwater extraction (first described in Burlig, Preonas, and Woerman (2025)) to estimate how California farmers respond to both temperature shocks and variation in groundwater depth. These reduced-form estimates identify farmer responses

using year-over-year variation in temperature and depth within location. Combining these estimates with basic theory, we can bound farmers' willingness to pay to avoid hot days and further aquifer drawdown.

Next, we use our reduced-form estimates to parameterize a simple dynamic model of agricultural groundwater pumping as a function of temperature and groundwater depth. Our model explicitly accounts for the effect of current-year groundwater extraction on groundwater depths in the next year. This lets us forward-simulate adaptive irrigation as a function of future climate projections and dynamically updated groundwater depths. This reveals the extent to which climate change will increase aquifer depletion. Finally, we use our simulation results to bound the welfare consequences of climate change, and to conduct optimal policy counterfactuals—allowing us to quantify the size of the stock externality with vs. without climate change.

Our reduced-form estimates reveal several insights. First, we find evidence that farmers adapt to climate change by pumping more groundwater: farmers increase annual pumping by 15.5 acre-feet (AF), or 2.8%, in response to an additional day above 38°C. We find no evidence that farmers respond to hot temperatures by changing the use of other variable inputs (fertilizers or pesticides), implying that the amount farmers spend on this additional pumping is a lower bound on their willingness to pay to avoid high heat: $775 per additional hot day per year. Adapting to climate change through irrigation is thus privately beneficial to farmers.

Second, we find that farmers reduce their groundwater pumping as the aquifer becomes deeper, and pumping becomes more costly. This aligns with existing evidence that farmers exhibit downward-sloping demand for groundwater (Bruno, Jessoe, and Hanemann (2024); Burlig, Preonas, and Woerman (2025)). Farmers reduce their pumping by 67 AF per year (12%) for every 10% increase in groundwater depth, and are willing to pay $2,014 per year to avoid such a drop in the aquifer level. Private adaptation to climate change thus imposes an important social cost.

Our dynamic simulation results predict that the aggregate cumulative effect of adaptive irrigation in response to climate change will increase the depth of the San Joaquin Valley aquifer from 153 feet in 2020 to 289 feet in 2100 under RCP4.5, and to 364 feet in 2100

under RCP8.5. Projected aquifer depth is 68% (RCP4.5) to 162% (RCP8.5) greater at end-of-century than what would have occurred absent climate change, with 41% (62%) of the increase in depth from 2020 to 2100 resulting from climate change-induce adaptive irrigation.

Our simulations also reveal the importance of the "natural scarcity" effect, whereby as groundwater depths rise, pumping costs increase, causing farmers to pump less. If we shut down this effect, aquifer depth increases to 617 feet (705) in 2100 under RCP4.5 (8.5). Thus, the natural scarcity mechanism mitigates 71% (62%) of the aquifer depletion likely to occur over this 80-year horizon, and 85% (72%) of the depletion resulting from climate change.

Finally, applying our theory model to these simulation results, we find that climate change creates annual welfare losses for farmers in the San Joaquin Valley of over $1.1 billion per year by end-of-century under RCP4.5. At least half of this loss is the direct result of hotter temperatures, while the remainder results from the indirect effect of hotter temperatures on aquifer depletion. Under RCP8.5, this annual loss is much larger: over $3.5 billion per year by end-of-century, 72% of which comes from aquifer depletion. For comparison, these future losses represent 4.5% and 14.5% of current aggregate annual crop revenues in the San Joaquin Valley. Both the direct costs of climate change and the social costs resulting from climate adaptation are thus economically meaningful.

This paper contributes to several literatures. First, we build on the nascent literature empirically measuring the extent of climate adaptation. Much of the climate change economics literature has focused on mitigation e.g., Nordhaus (1993) and Pindyck (2013), or the costs of climate change e.g., Deschênes and Greenstone (2007) and Hsiang et al. (2017). However, current climate science indicates that mitigation alone is likely not sufficient to avoid the costs of climate change (Lee et al. 2023). More recent research highlights the potentially important role of adaptation in determining the future costs of climate change (Carleton et al. 2022; Hultgren et al. 2025), but do not measure specific adaptation strategies. We build on this literature by documenting that adaptation has important private benefits, but can come at a large social cost, in the context of an important and generally overlooked margin of adaptation: irrigation (Burke and Emerick 2016). We find that farmers *are* adapting to climate change by irrigating their crops, and that this is privately beneficial, with farmers revealing a substantial willingness to pay to avoid high heat by spending more

on irrigation. However, we also find that farmers have a large willingness to pay to avoid aquifer depth increases, demonstrating that this private adaptation is socially costly.

Second, we build on a rich literature on the dynamics of non-renewable (Hotelling 1931) and renewable (Gordon 1954) natural resources. We build on Provencher and Burt (1993) and Timmins (2002), who consider the dynamics of the groundwater resource, by adding the costs of climate change. In doing so, we contribute to a literature of hydro-economic models of the interaction between water resource dynamics and water use (Gisser and Sanchez 1980; Allen and Gisser 1984; Hrozencik et al. 2017; Rad et al. 2020). This literature has primarily focused on modeling feedback loops between aquifer levels and pumping costs, where groundwater depletion increases aquifer depths and the marginal cost of groundwater pumping. We build on this literature by incorporating both temperature and depth responsiveness into a hydro-economic modeling framework and demonstrating how these margins counteract each other in future projections of groundwater use and depletion. Our work is also related to that of Carleton, Crews, and Nath (2023), who focus on the impacts of trade on global water use in a similar modeling framework. Our estimates demonstrate that climate change imposes substantial costs on the California aquifer under open access: under RCP4.5 (8.5), we estimate annual welfare losses approaching 1 (3) billion USD by end of century.

Finally, we build on the long strand of literature evaluating the impacts of climate change on the agricultural sector. Our results contribute to long-standing debates on the effect of climate change on agricultural production (Mendelsohn, Nordhaus, and Shaw 1994; Schlenker, Hanemann, and Fisher 2005; Schlenker and Roberts 2009; Deschênes and Greenstone 2007; Burke and Emerick 2016; Hultgren et al. 2025; Burke et al. 2024). Specifically, we demonstrate how irrigation acts as a form of climate change adaptation, potentially mitigating some of the damages associated with high temperature exposure in the short-run. However, our dynamic model suggests that the irrigation margin of adaptation may be less effective in the long-run as depth responsiveness begins to constrain temperature induced increases in pumping. This feedback between temperature and depth responsiveness highlights an important tension between natural resource dynamics and adaptation previously unexplored in the literature.

2 Background

California is the most economically important agricultural state in the U.S. in terms of farm-level sales of agricultural goods (Johnson and Cody 2015). The unique climate of the state facilitates the production of many high value crops (e.g., perennial fruit and tree nuts). About half of the fruits and nuts consumed in the U.S. are grown in California. Agricultural production in the state is concentrated in the relatively arid Central Valley. Given the arid climate of the region, irrigation is necessary for the production of most crops. Large-scale state and federal investments in surface water irrigation infrastructure facilitated the state's agricultural development during the first half of the 20^{th} century. The Central Valley also overlies several aquifers with historically abundant stocks of groundwater. Advancements in groundwater pumping technology in the 1940s and 1950s opened these groundwater stocks to widespread use by the agricultural sector. Since then, groundwater has become an increasingly important source of water for the state's agricultural sector. As of 2023, about 50% of irrigated land in the state relied on groundwater as a primary or secondary water source (United States Department of Agriculture 2023).

The agricultural sector's increasing reliance on groundwater has severely taxed the state's aquifers. Pumping rates in excess of recharge have caused significant depletion of the aquifers supplying groundwater to farms, households, and municipalities of the Central Valley (Famiglietti et al. 2011; Liu et al. 2022). To address groundwater depletion concerns, the state passed the landmark Sustainable Groundwater Management Act in 2014, intended to bring aquifers into "sustainability" (i.e., ensure that recharge is greater than or equal to extraction). Our sample period occurs prior to the enforcement of SGMA, and therefore our estimates currently do not consider this policy.[1]

3 Theory

We develop a stylized model of agricultural production, groundwater use, and groundwater stocks. We briefly overview the model here, and present it in full in Appendix C.

1. In future work, we plan to compare our findings to SGMA targets, in order to inform how stringent SGMA will be required to be with and without climate change.

We consider N identical farmers indexed by i. Farmer i's profits in year t are:

$$\pi(w_t; D_t, T) \;=\; p\,Q(w_t; T) \;-\; c(D_t)\,w_t$$

The single input w_t is water used for irrigation in year t by farmer i (we suppress the i subscript, since $w_{it} = w_t$ for all N identical farmers). D_t is groundwater depth at the beginning of year t's growing season. Marginal groundwater pumping costs $c(D_t)$ increase linearly in depth D_t, consistent with physics (Hurr and Litke (1989)):

$$\frac{\partial c(D_t)}{\partial D_t} > 0 \quad , \quad \frac{\partial^2 c(D_t)}{\partial D_t^2} = 0$$

T is a stylized temperature parameter, which we use to model a shift in the temperature distribution caused by climate change; this parameter abstracts from uncertainty of temperature realizations, nonstationarity of the temperature distribution, and the timing of climate change. $Q(\cdot)$ is the crop production function, which is increasing and concave in water input w_t, decreasing in temperature T, and has a positive cross-partial (such that the marginal product of water increases with temperature):

$$\frac{\partial Q(w_t; T)}{\partial w_t} > 0 \quad , \quad \frac{\partial^2 Q(w_t; T)}{\partial w_t^2} < 0 \quad , \quad \frac{\partial Q(w_t; T)}{\partial T} < 0 \quad , \quad \frac{\partial^2 Q(w_t; T)}{\partial w_t\, \partial T} > 0$$

Finally, the crop output price p is fixed and exogenous.

We assume that all farmers have an identical profit function in order to simplify the presentation of this stylized model.[2] Our empirical application relaxes this homogeneity assumption to simulate changes to farmer welfare and quantify the stock externality.

Aquifer dynamics We model a stylized equation of motion to capture how aggregate groundwater pumping in year t (i.e., Nw_t) impacts depth in year $t+1$ (i.e., D_{t+1}):

$$D_{t+1} \;=\; D_t \;+\; \Phi N w_t$$

2. This stylized model also abstracts from farmers' crop choices (the focus of Burlig, Preonas, and Woerman (2025)) for simplicity, which appears reasonable in this setting because we find no evidence that farmers' temperature and depth responses vary with crop choice (see Appendix Figure B9).

The constant $\Phi > 0$ subsumes specific yield of the soil and the effect of evapotranspiration (i.e., ρ and $(1-\psi)$ in Equation (9), respectively), and we assume it is homogeneous across the full aquifer.[3]

3.1 Social planner's problem

The Planner dynamically optimizes groundwater extraction by choosing $\{w_t\}$ to maximize total discounted profits of all farmers over an infinite horizon, subject to the equation of motion:

$$\max_{\{w_t\}} \sum_{i=0}^{N} \sum_{t=0}^{\infty} \beta^t \, \pi_{it} \;=\; N \sum_{t=0}^{\infty} \beta^t \Big[p \, Q(w_t; T) \;-\; c(D_t) \, w_t \Big]$$

$$\text{s.t.} \quad D_{t+1} \;=\; D_t \;+\; \Phi w_t N$$

Combining the Planner's FOC with respect to control variable w_t and state variable D_t yields:

$$p \frac{\partial Q(w_t; T)}{\partial w_t} \;=\; c(D_t) \;+\; N\Phi \sum_{s=t+1}^{\infty} \beta^{s-t} w_s \frac{\partial c(D_s)}{\partial D_s} \tag{1}$$

The Planner sets each farmer's marginal value of water equal to their marginal pumping cost *plus* the marginal increase in discounted pumping costs imposed on all N farmers over all future time periods.

3.2 Farmer's problem

We assume that farmer i is small relative to the aquifer, maximizes her own profits, and is forward-looking with respect to future groundwater pumping costs. The assumption of large N is reasonable in our setting—as of 2023, there were 34,500 distinct farms in the San Joaquin Valley, which operate over 38,000 groundwater pumps (Joaquín Morales and Ayres

3. This equation of motion also ignores recharge through rainfall (i.e., $-\rho\overline{R}_t$ in Equation (9)), which is not expected to change markedly in our geographic setting due to climate change (see Appendix Figure A2). We make these simplifications only for ease of presentation.

(2023))—and means that unlike in Timmins (2002), farmers do not adjust their pumping behavior in order to optimize the depth trajectory of the aquifer.[4]

It follows that farmer i solves a dynamic problem similar to that of the Planner, but only internalizes $1/N$ of the value of future groundwater stocks:

$$\max_{\{w_{it}\}} \sum_{t=0}^{\infty} \beta^t \, \pi_{it} \;=\; \sum_{t=0}^{\infty} \beta^t \Big[p \, Q(w_{it}; T) \;-\; c(D_t) \, w_{it} \Big]$$
$$\text{s.t.} \quad D_{t+1} \;=\; D_t \;+\; \Phi[w_{it} \;+\; (N-1) w_{-it}]$$

Here, the equation of motion decomposes the change in depth due to farmer i's extraction vs. due to extraction by all other farmers (i.e., $(N-1)w_{-it}$).

Combining the farmer's FOCs similarly yields:

$$p \frac{\partial Q(w_{it}; T)}{\partial w_{it}} \;=\; c(D_t) \;+\; \underbrace{\Phi \sum_{s=t+1}^{\infty} \beta^{s-t} w_{is} \frac{\partial c(D_s)}{\partial D_s}}_{\approx 0} \tag{2}$$

This final term is $1/N$ the size of the analogous term in the Planner's FOC because farmer i only internalizes her impact on *her own* future pumping costs. This term becomes trivially small with large N, such that we can rewrite Equation (2) as:

$$p \frac{\partial Q(w_t; T)}{\partial w_t} \;=\; c(D_t) \tag{3}$$

Farmer i sets the marginal value of water equal to her own (private) pumping cost. This is equivalent to static profit maximization ignoring future impacts on aquifer depths.

4. Even in light of recent corporate farm consolidation in the Central Valley, ownership remains almost certainly sufficiently diffuse to yield collective action to jointly conserve groundwater. For example, per recent reports, the top 10% of farmer owners in Kings County (a large producing county in the San Joaquin Valley) now own 86% of the county's farmland (https://www.sfchronicle.com/opinion/openforum/article/farms-agriculture-corporations-21015386.php). However, collective action is still unlikely with more than 5–15 large owners across the full San Joaquin groundwater basin (Olson (1965)).

3.3 The stock externality

Comparing the first-order conditions for the Planner (Equation (1)) vs. farmer i (Equation (3)), the Planner internalizes an additional term:

$$\tau_t \;\equiv\; N\Phi \sum_{s=t+1}^{\infty} \beta^{s-t} w_s \frac{\partial c(D_s)}{\partial D_s} \;>\; 0 \tag{4}$$

This term τ_t is the stock externality, equivalent to the marginal effect of farmer i's pumping in period t on the present value of aggregate future profits.

We can approximate the total welfare loss from open-access over-extraction as the product of τ_t, the marginal stock externality per unit of water, and the difference in aggregate extraction Nw_t between a no-policy (NP) vs. an optimal policy (OP) scenario. This present discounted sum is:

$$LOSS \;\approx\; N \sum_{t=0}^{T} \beta^t \tau_t \Big(w_t^{NP} - w_t^{OP} \Big) \tag{5}$$

The approximation results from the fact that τ_t is a marginal value along the optimal path $\{w_t^{OP}\}$, whereas $(w_t^{NP} - w_t^{OP})$ is a non-marginal change in extraction levels. Comparing Equation (1) vs. Equation (3), it is clear that $w_t^{NP} > w_t^{OP} > 0$. Since $\tau_t > 0$, this means that $LOSS > 0$.

3.4 Climate change

Next, we consider the impacts of climate change by taking comparative statics with respect to T. For simplicity, we consider a one-time positive shift in temperatures (i.e., $dT > 0$) that occurs in year $t = 0$. This shock is permanent and (for the purposes of this stylized model) unanticipated. The key empirical objects of interest are $\frac{d\tau_t}{dT}$, how the stock externality varies with climate, and $\frac{dLOSS}{dT}$, how the total welfare loss from open-access over-extraction varies with climate. As we show in Appendix C, these are theoretically ambiguous because of two counterveiling forces: as temperatures rise, the marginal benefit of groundwater pumping increases, because irrigation is increasingly needed to protect crops from the damaging impacts of heat; but as pumping increases, so too does its marginal cost, as the aquifer is driven

closer to depletion. Which of these effects dominates determines the net effect of climate change on the stock externality and the loss from open access, and is an empirical question.

4 Data

Electricity and groundwater data We use confidential customer-level data from Pacific Gas and Electric (PG&E) to measure electricity and groundwater usage.[5]We observe monthly electricity bills for the universe of 108,172 agricultural consumers in PG&E's service territory from 2008 to 2019. We also observe the universe of agricultural pump tests—21,851 tests at 17,107 locations—from PG&E's Advanced Pumping Efficiency Program (APEP). These tests provide detailed measurements of pump characteristics and performance, including operating efficiency, lift, and flow. Matching electricity bills to APEP pump tests yields a sample of 10,277 confirmed agricultural groundwater pumpers.[6] To compute the quantity of groundwater extracted by each pump in acre-feet (AF), we apply basic physics (Hurr and Litke (1989)):

$$\frac{\text{kWh}}{\text{AF}} = \frac{[\text{Lift (feet)}] \times 1.0241}{\text{Operating pump efficiency (\%)}} \tag{6}$$

where kilowatt-hours (kWh) of billed electricity consumption is scaled by pump efficiency and the inverse vertical distance that groundwater is lifted to the surface.[7] Greater groundwater depths increase lift, requiring more electricity to pull groundwater to the surface, thereby increasing pumping costs per AF.

Tax parcels We identify farms using California county assessor tax parcels (following Bruno, Jessoe, and Hanemann (2024); Burlig, Preonas, and Woerman (2025)). Using parcel polygons, we aggregate annual groundwater use to the farm level by summing across pumps contained in each parcel. This yields a yearly panel of 7,817 unique parcels, for which we present summary statistics in Table 1. The average parcel contains 1.91 observed pumps,

5. PG&E is California's largest investor-owned electric utility, which services the majority of California's cropland. Burlig, Preonas, and Woerman (2025) provides a detailed overview of these data.

6. Appendix Figure A1 plots the locations of these 10,277 pumpers.

7. Lift is largely governed by groundwater depths, along with other pump-specific adjustments (such as gauge pressure, and the height of the pump above the surface).

Table 1: Summary statistics

Unique parcels	7,817
Count of observed pumps per parcel	1.91
	(1.70)
"Croppable" area of parcel (acres)	121.77
	(187.23)
Average electricity use (kWh/year)	286,161
	(462,100)
Average electricity bill ($/year)	43,706
	(58,386)
Average groundwater use (AF/year)	1,004
	(1,648)
Average groundwater use (AF/acre per year)	4.83
	(24.66)
Average count of days/year above 32°C	87.94
	(34.05)
Average annual precipitation (mm)	289.70
	(173.93)
Average groundwater depth in sub-basin (feet)	119.51
	(55.38)
Parcel within surface water district (1/0)	0.61
	(0.49)
Parcel within San Joaquin Valley groundwater basin (1/0)	0.72
	(0.45)

Notes: We report means and standard deviations (in parentheses) across parcel-specific averages, weighted by (time-invariant) croppable acreage. The number of pumps, electricity use, and electricity bill are derived from PG&E electricity data. The groundwater use statistics use PG&E data in conjunction with groundwater depth data from CASGEM, following Equation (6). Temperature and precipitation data come from PRISM. Surface water data come from Hagerty (2022).

122 "croppable" acres, and consumes 1,000 AF of groundwater per year (or 4.83 AF per acre per year) at a cost of $43,706 in electricity bills.[8]

Weather data Our weather data come from PRISM, and consist of 4 kilometer gridded daily maximum temperatures and precipitation totals. This is a standard dataset in the climate impacts literature (see e.g., Schlenker and Roberts (2009); Burke and Emerick (2016)). We create daily rasters of maximum temperature and precipitation, and extract pointwise daily weather realizations at the centroid of each parcel in our sample. California farmers

8. As in Burlig, Preonas, and Woerman (2025), we define croppable acres using USDA's Cropland Data Layer, subtracting land not suited to crop production (e.g., forests, buildings, bodies of water) from each parcel's total acreage.

face a hot and dry climate: from 2008–2019, the average parcel in our sample experienced 88 days per year above 32°C (90°F), and only 289 mm (0.94 feet) of precipitation.[9]

Groundwater depth data We use groundwater data from the California Department of Water Resources' California Statewide Groundwater Elevation Monitoring (CASGEM) Program. CASGEM reports thousands of depth measurements across locations for each month in our sample. We leverage these depth measurements in two ways. First, we create monthly rasters of depth that let us assign a monthly depth at the location of all 10,277 pumps in our sample. We use these pump-specific depths to parameterize Equation (6), incorporating month-to-month variation in lift to generate more accurate conversions from kWh to AF.[10]

Second, we average all monthly depth measurements within each groundwater sub-basin.[11] Then, we assign each sub-basin an initial annual depth based on its January average.[12] This lets us control for year t's depth *before* weather during year t's growing season has realized, and with a high enough spatial aggregation to preclude being influenced by any individual parcel's pumping behavior. The average parcel in our sample faces an initial depth of 120 feet in its sub-basin. Our forward simulation analysis focuses on the San Joaquin Valley basin, which contains 72% of parcels in our sample.

Surface water districts We assign parcels to water districts using polygons assembled by Hagerty (2022). 61% of parcels are assigned to a water district, and therefore are highly likely to receive surface water allocations.[13] We use these assignments both to account for

9. In California's Central Valley, 90% of precipitation occurs between November and April, when temperatures are lower and many crops are off-cycle.

10. Depth tends to vary cyclically within each calendar year, and also exhibits substantial year-to-year variation. We construct AF at the pump-month level to capture that variation before aggregating to the parcel-year (as in Burlig, Preonas, and Woerman (2025)).

11. California defines groundwater basins (and sub-basins) based on geological and hydrological conditions. Our sample spans 41 unique basins and 85 unique sub-basins.

12. January is both the start of the calendar year and the month with the least groundwater pumping in our sample. For sub-basin years without a January measurement, we use the above rasters to assign an interpolated initial depth.

13. Farmers pay lower marginal cost for these district water allocations than for self-pumped groundwater (Hagerty (2022)). Hence, we can treat district water as inframarginal to any observed groundwater use (as in Burlig, Preonas, and Woerman (2025)). Farmers also have the option to purchase surface water on the open market. However, such transactions comprise a very small share of total water, at prices far higher than marginal groundwater pumping costs (Hagerty (2023)).

spatial dependencies related to unobserved variation in district water allocations, and to test whether our estimates are robust to surface water use as a potential confounder.

Census of Agriculture We use data on other farm inputs from the USDA's Census of Agriculture. This Census is conducted every five years, our use incorporates 8 Census waves from 1982–2017. These data report farms' spending on fertilizer and chemical inputs per cropland acre, at the county-year level.

Climate projections To forward simulate groundwater use and aquifer depletion, we use annual county-level temperature and precipitation projections from 2020–2100. These data come from Hsiang et al. (2017); and Rasmussen, Meinshausen, and Kopp (2016). They derive climate projections using an ensemble approach: probability-weighting combinations of climate models to generate projections of daily temperature and precipitation data under alternate climate scenarios. We consider RCP4.5 (a moderate-emissions scenario) and RCP8.5 (a high-emissions scenario), which use 43 and 44 separate climate models, respectively.[14] Appendix Figure A2 plots how temperature and precipitation evolve across climate models under each scenario.

5 Reduced form estimation and results

5.1 Empirical strategy

We employ a standard panel fixed effects approach to identify the causal impact of temperature and groundwater depth on outcomes related to agricultural groundwater pumping. Our preferred econometric specification is:

$$Y_{it} \;=\; f(\mathbf{T}_{it}) \;+\; \beta\, \log(D_{st}) \;+\; g(P_{imt}) \;+\; \gamma_i \;+\; \delta_t \;+\; \varepsilon_{it} \tag{7}$$

14. These Representative Concentration Pathway (RCP) scenarios are the standard projections used in the economics literature on climate change impacts (e.g., Hultgren et al. (2025)). RCP4.5 is consistent with a stabilization of global fossil fuel emissions, while RCP8.5 is consistent with intensive growth of global fossil fuel emissions (Van Vuuren et al. (2011)).

where the dependent variable Y_{it} is an outcome for parcel (i.e., farm) i in year t. Our main outcome variables are (i) annual groundwater pumping and (ii) annual groundwater pumping expenditures (i.e. electricity bills), which we use to estimate the cost of adaptation.

$f(\mathbf{T}_{it})$ is a flexible function of temperature realizations $\mathbf{T}_{it}$ experienced at the location of parcel i in year t. We represent $f(\mathbf{T}_{it})$ as a set of 10 bins, where each bin sums the number of days in year t where the maximum daily temperature falls within that bin's defined range (Deschênes and Greenstone (2007); Hsiang et al. (2017)).[15] Each temperature bin coefficient represents the marginal effect of having one additional day in that bin, relative to the omitted bin of 23–26°C. These coefficients flexibly trace out farmers' adaptive response to temperatures.

$\log(D_{st})$ is the natural logarithm of groundwater depth. We spatially average depth D_{st} across the full sub-basin s that contains parcel i, in order to shut down any endogeneity between parcel i's own pumping behavior and its localized depth. We also use depth measured in January of year t, in order to minimize the confounding effects of contemporaneous year-t groundwater extraction on this regressor.[16] The coefficient β captures how groundwater pumping (and related outcomes) responds to changes in costs induced by the additional energy required to lift water to the surface from greater depths (Burlig, Preonas, and Woerman (2025)). Since the cost of groundwater pumping increases linearly in depth (i.e., differentiating Equation (6) with respect to lift), the logarithmic functional form attenuates groundwater demand response as depth increases.[17]

We control for precipitation $g(P_{imt})$, where $g(\cdot)$ is a linear function in the total precipitation at parcel i for each month m of the calendar year t. Estimating separate monthly precipitation coefficients accounts for the timing of rainfall, which predominantly occurs during winter months and not during the hot summer growing season. γ_i are parcel fixed effects and δ_t are year fixed effects, which flexibly control for unobservable cross-sectional factors and common time-varying shocks, respectively. Finally, ε_{it} is an idiosyncratic error term

15. We define eight bins of 3°C increments spanning 14–38°C, as well as one bin for all temperatures below 14°C and one bin for all temperatures greater than 38°C. Days that fall exactly at a bin cutoff between bins are included in the lower bin.

16. As discussed above, January is conveniently both the start of the calendar year and the calendar month with the least groundwater pumping in our sample.

17. Appendix Figure B3 shows that our results are robust to alternative functional forms, with the log functional form aligning with a (less parametric) piecewise linear function of depth.

that we cluster by water district to account for arbitrary correlation within water districts over time (following Burlig, Preonas, and Woerman (2025)).[18] We weight parcels by their cropable acreage, such that our estimated coefficients are representative of the average unit of cropland in our sample.

Identification For our temperature effect estimates, the identifying variation comes from comparing relatively warmer vs. cooler years at the same location. This requires the assumption of no time-varying unobserved factors that are correlated with year-to-year fluctuations in temperature, and are also correlated with the outcome variable. Given that weather realizations are as-good-as-random conditional on cross-sectional and temporal averages, this assumption is likely to be satisfied (Deschênes and Greenstone (2007); Schlenker and Roberts (2009); Hsiang et al. (2017)).

For the depth coefficient β, the identifying variation comes from comparing years where the sub-basin-wide groundwater level is deeper vs. shallower. We therefore require the assumption that farmers are "depth-takers" and cannot influence the average depth of their sub-basin, which is plausible given sub-basins in our sample contain 3,593 agricultural groundwater pumps on average. We must also assume that the outcome variable does not simultaneously impact depth, which is plausible given that D_{st} is measured in January, before the vast majority of year t's groundwater pumping has realized.[19] To address the potential for differential counterfactual trends in depth across sub-basins—if, say, certain sub-basins have relatively faster secular rates of aquifer depletion, perhaps due to differences in geology—, we add sub-basin-specific linear trends to Equation (7) as a robustness test, identifying β off of deviations from trends in D_{st}.

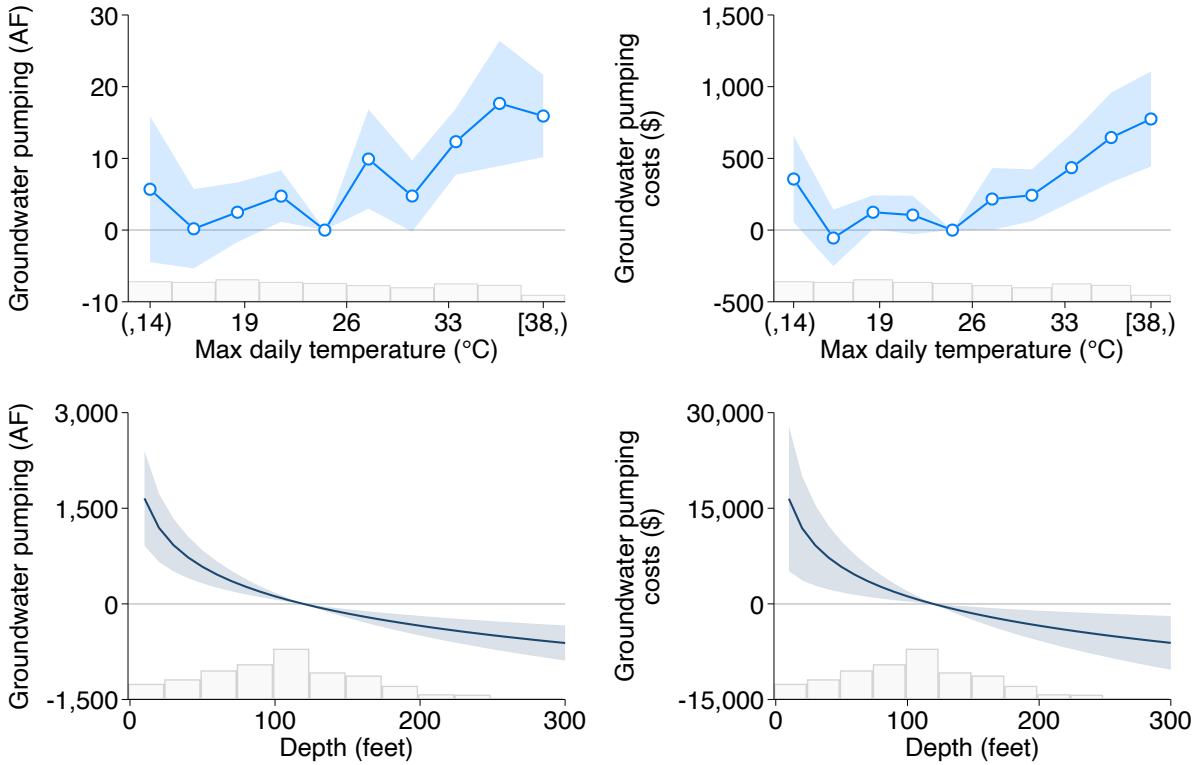
18. This accounts for unobserved shocks to surface water allocations common to all farms with a water district, serial correlation within parcels (or between neighboring parcels), and other unobserved spatial correlations. For the 39% of parcels not assigned to a water district, we construct analogous spatial clusters at the county level.

19. A related concern is that aggregate groundwater pumping in year t could have a dynamic feedback effect on aquifer levels that affects $D_{s(t+1)}$. Appendix F discusses how this effect could bias $\hat{\beta}$ towards zero in the presence of serial correlation in ε_{it}. However, Monte Carlo simulations reveal that the strength of the dynamic feedback and serial correlation in our sample are not strong enough to meaningfully bias $\hat{\beta}$.

5.2 Results

Figure 2 presents results from estimating Equation (7) using our two main outcome variables: the quantity of groundwater consumption (left panels) and groundwater pumping expenditures (right panels). Both outcome variables are at the parcel-year level. We plot causal estimates of how temperature (top panels) and groundwater depth (bottom panels) impact adaptive pumping behavior. We also report these results numerically in Appendix Table B1.

Figure 2: Temperature and depth meaningfully impact groundwater pumping



Notes: This figure presents the results of two regressions estimating Equation (7). Our two outcome variables are: annual groundwater consumption (in AF, left panels), and annual groundwater pumping expenditures (in $, right panels). The top panels plot the coefficients on ten temperature bins (omitting the 23–26°C bin). The bottom panels display the effects of depth, estimated using a log functional form, and plotted relative to the mean depth of 119.5 feet. 95% confidence intervals are shaded. Bars report histograms of temperature and depth within our sample. Regressions include month-specific linear precipitation controls, parcel fixed effects, and year fixed effects. Regressions weight parcels by croppable acres. Standard errors are clustered by water district.

We find that temperatures below 20°C have at most a modest impact on groundwater pumping or spending on groundwater extraction. However, hotter temperature cause farmers to meaningfully increase groundwater pumping. At the extreme, each day with a maximum temperature above 38°C increases annual groundwater consumption at the farm

level (i.e., aggregated across all pumps in a parcel) by 15.9 AF (1.6% of the mean) and annual groundwater expenditures by $775 (1.8% of the mean), compared to a day with a maximum temperature in the reference bin of 23–26°C. This is consistent with farmers applying groundwater to adapt to high temperatures, which numerous studies have shown damage crop yields (e.g., Schlenker and Roberts (2009)).

We also find that greater groundwater depths cause farmers to pump less groundwater. For every 10% increase in the average January depth of the sub-basin, farmers reduce their annual groundwater consumption by an average of 67 AF per year (6.6% of the mean). Since a 10% increase in depth increases groundwater pumping costs by (roughly) 10%, this estimate implies a groundwater cost elasticity of -0.66, aligning with groundwater cost elasticity estimates of -0.46 to -0.79 from the same sample in Burlig, Preonas, and Woerman (2025). If farmers pump less as the costs of pumping rise, it follows that greater depths due to aquifer depletion serve as a natural scarcity mechanism to stem over-extraction. Finally, we estimate that annual groundwater expenditures fall by $665 (1.5% of the mean)—a much smaller effect in percentage terms, since greater depths simultaneously reduce pumping quantities and increase pumping costs.[20]

These results suggest that climate change has the potential to accelerate groundwater depletion, as hotter temperatures increase farmers' demand for adaptive irrigation. However, they also suggest that groundwater cost response will attenuate the impacts of climate-induced depletion. The relative magnitude of these two effects is an empirical question. In Section 6, we develop a simple dynamic model that embeds both responses, and use forward simulations to quantity the overall impact of climate change on groundwater depths.

Robustness One potential source of bias in Figure 2 is the time-varying availability of surface water. Annual surface water allocations are governed by a complicated mix of riparian and appropriative rights (see Hagerty (2023) for an excellent overview). While these allocations are typically determined by weather realizations prior to the growing season (e.g., winter snowpack), they still could bias our temperature estimates—if, say, hotter tempera-

20. Per Equation (6), the relationship between electricity and groundwater changes as depth increases: the same quantity (and thus cost) of electricity yields less groundwater. The temperature-response functions in the top panels of Figure 2 hold this relationship fixed, yielding quantity and expenditure effects that are similar in percentage terms (and implying an average groundwater cost of $\$755 \div 15.9$ AF $= \$47$ per AF).

tures induce more groundwater pumping not as adaptive irrigation, but as a substitute for scarce surface water. Moreover, if surface water allocations are negatively correlated with January depths (e.g., due to drought), then substitution away from scarce surface water could attenuate our estimates of groundwater cost response. To test for these potential biases, we estimate Equation (7) separately for parcels with vs. without surface water access (as measured by being locating inside vs. outside a surface water district). We recover nearly identical temperature and depths estimates for both groups (see Appendix Figure B1), indicating that surface water allocations are unlikely to be biasing our pooled estimates.

We might also worry about non-parallel counterfactual trends, especially since certain sub-basins exhibit strong secular trends in depth during our sample period. As a robustness test, we estimate Equation (7) adding sub-basin-specific linear trends, such that the depth coefficient β is identified off of year-on-year deviations from these trends (see Appendix Figure B2). This yields slightly attenuated depth estimates: a 10% increase in depth reduces annual groundwater consumption by 49 AF (4.9% of the mean), and reduces annual groundwater expenditures by \$516 (1.2% of the mean).[21] We use this attenuated estimate for groundwater use as an alternative parameterization in our forward simulations below.

Our results in Figure 2 are also robust to: alternative functional forms for both the left- and right-hand side variables (Appendix Figures B3–B6); two-way clustering by water district and county-year (Appendix Figure B7); alternative weights and aggregation (Appendix Figure B8); and controlling for crop choice (Appendix Figures B9–B10).

Comparison to engineering estimates The engineering and crop science literature addresses how water use responds to temperature shocks and other factors using process-based models simulating crop growth (Foster, Brozović, and Butler (2014); Maucieri et al. (2025)). We use Aquacrop, a state-of-the-art crop model (Vanuytrecht et al. (2014)), to simulate how much additional water California crops require when exposed to high heat (See Appendix G for more details). We estimate that an additional day with maximum temperature above 38°C requires a 0.54% increase in annual water use, or about a third of our estimated treatment effect. This difference highlights the importance of empirically measuring the impact

21. This attenuation may reflect the removal of bias, or simply less identifying variation. The corresponding temperature estimates are not sensitive to the inclusion of sub-basin-specific linear trends.

of temperature on farmer behavior; engineering models such as Aquacrop make the strong assumptions that farmers have perfect knowledge of their field's soil moisture content and never irrigate past their field's water holding capacity. Our estimates demonstrate that understanding farmer responses to climate change requires real-world estimates, as simply using the engineering model would substantially understate the effect of temperature on groundwater extraction.

Substitution to other agricultural inputs While we find that groundwater expenditures respond positively to hotter temperatures and negatively to greater depths, it is possible that farmers are reoptimizing to compensate for these expenditure changes by adjusting spending on non-water inputs. To test for this this substitution mechanism, we estimate corresponding temperature and depth effects using expenditures on two key inputs: fertilizer and chemicals (i.e., pesticides, herbicides, fungicides, etc.).[22] The USDA's Census of Agriculture reports these expenditures as county-level averages, allowing us to estimate a modified version of Equation (7) at the county-year level.

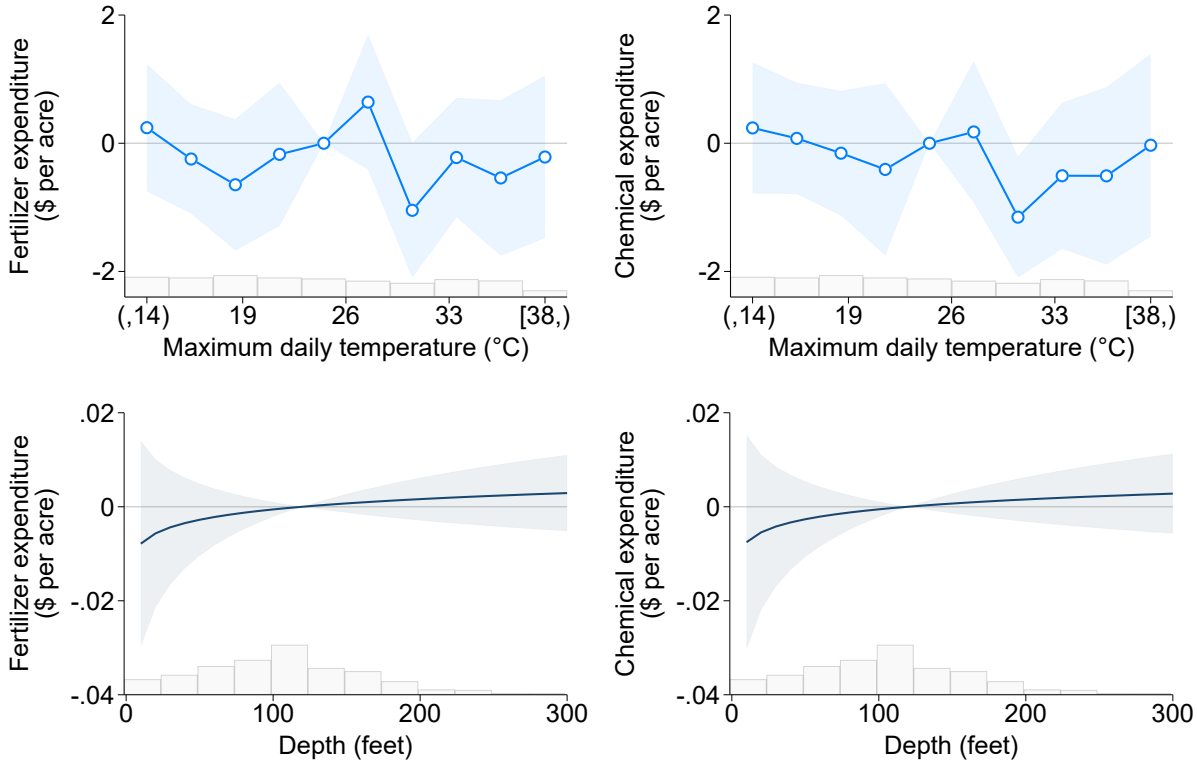
Figure 3 presents these results, for the natural log transformations of fertilizer expenditures in the left panels and chemicals expenditures in the right panels.[23] We find no systematic evidence that higher temperatures or greater depths cause farmers to change their expenditures on either input. These results suggest that increased water use is the primary short-run temperature adaptation mechanism for California farmers, and that farmers do not reoptimize other inputs to offset the costs of adaptive irrigation.

Interacting temperature and depth Equation (7) treats farmers' temperature and depth responses as additively separable. However, farmers' tradeoff—balancing the benefits of adaptive irrigation vs. the cost of increased pumping—could shift toward less adaptive irrigation at greater depths or toward less cost response at higher temperatures. We estimate a more flexible version of Equation (7) that adds an interaction between each of the 10 temperature bins and log depth. This regression yields heterogeneous temperature response

22. We choose these specific outcomes because they are plausibly impacted by growing conditions.

23. We log-transform these expenditure variables to account for a strongly right-skewed distributions across counties. We also winsorize both expenditure variables at their 95th percentiles.

Figure 3: Temperature and depth do not impact other farm inputs



Notes: This figure presents the results of two regressions estimating a version of Equation (7) at the county-year level. Our two outcome variables are the natural logarithms of average fertilizer expenditures (in $ per acre; left panels), and chemical expenditure (in $ per acre; right panels). Outcome data come from the USDA Census of Agriculture five-yearly waves from 1982 through 2017, and are reported at the county level. The top panels plot the coefficients on ten temperature bins (omitting the 23–26°C bin). The bottom panels display the effects of depth, estimated using a log functional form, and plotted relative to the mean depth of 119.5 feet. 95% confidence intervals are shaded. Bars report histograms of temperature and depth within our sample. Regressions include month-specific linear precipitation controls, county fixed effects, and year fixed effects. Temperature, precipitation, and depth variables are aggregated up to the county-year level. Regressions are unweighted. Standard errors are clustered by county.

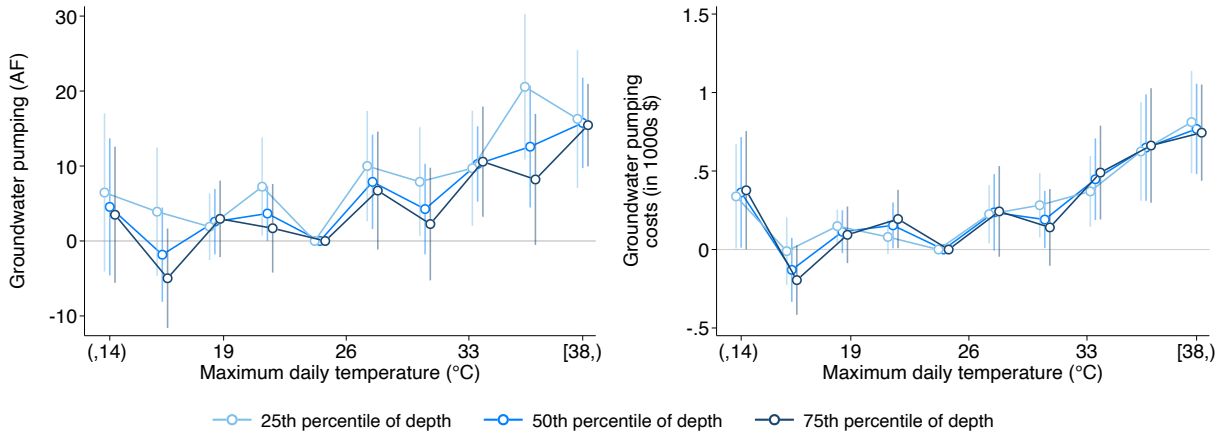
estimates that vary continuously in sub-basin depth and heterogeneous depth response estimates that vary by temperature bin.

Figure 4 summarizes these heterogeneous responses. The top panels plot our estimated temperature responses at the 25th, 50th, and 75th percentiles of groundwater depth in our sample (78 feet, 119 feet, and 150 feet respectively). This indicates a weaker temperature response at greater depths, consistent with farmers being less willing to pay for adaptive irrigation at higher pumping costs. While not pictured, we also recover nearly identical depth responses for the warmest vs. coolest years in our sample (2017 and 2019, respectively)—

albeit with slightly less depth response in the warmer year, consistent with farmers being more willing to incur pumping costs when the benefits of adaptive irrigation are greater.

While the regression coefficients on the interaction terms underlying these coefficients are statistically different from zero—confirming the above intuition on the directions of these interaction effects—the heterogeneous responses in Figure 4 are not statistically different from each other. This indicates that our preferred model with additively separable temperature and depth is sufficient to characterize these two opposing effects *within our sample*. However, as climate change intensifies and groundwater depletion continues over the coming decades, the importance of temperature-depth interactions will grow. For this reason, we use the interacted model as our preferred specification in our forward simulations.

Figure 4: Temperature × depth interaction effects are not statistically different in sample



Notes: This figure plots heterogeneous temperature response estimates using a modified version of Equation (7) that fully interacts 10 temperature bins with log depth. We summarize these interactions by plotting the (linear combinations of) temperature response estimates at the 25th, 50th, and 75th percentiles of groundwater depth in our sample (78 feet, 119 feet, and 150 feet, respectively). Plots are analogous the top panels of Figure 2. Regressions include 10 temperature bins (omitting the 23–26°C bin), log depth, the interaction of these temperature bins with log depth, month-specific linear precipitation controls, and year fixed effects. Regressions are weighted by croppable acres. Standard errors are clustered by water district. Vertical lines denote 95% confidence intervals.

5.3 Willingness to pay

We use our regression estimates, combined with simple theory, to derive (bounds on) farmers' marginal willingness to pay to avoid increases in temperature and increases in aquifer depth. Appendix C.6 provides more details on this approach.

Temperature increases The effect of shifting a day from 23–26°C to above 38°C on groundwater pumping expenditures provides a lower bound on the willingness to pay to avoid high heat (i.e., the coefficients plotted in the top-right panel of Figure 2). To see this, note that this coefficient represents the additional amount farmers are willing to spend to pumping on a hot day. Because we find that farmers do not adjust other inputs in response to high heat, we treat this as a lower bound. The full willingness to pay is weakly greater because our regression does not account for the yield losses due to extreme heat, which are highly likely even net of adaptive irrigation (see Figure 1).

Using this approach, our results in Figure 2 imply that the average farmer is willing to pay at least $775 (per parcel-year) to avoid shifting an additional day from the omitted 23–26°C to above 38°C. If we perform the same calculation using our model with temperature-depth interactions (i.e., the top-right panel of Figure 4), we find a similar lower bound of $768 (per parcel-year) to avoid an additional day above 38°C—conditional on the average groundwater depth of 119.5 feet. Since both models are linear in temperature bins, we can linearly extrapolate these per-day bound to bound the willingness to pay to avoid non-marginal increases in temperature.

Depth increases Our theory model implies that the marginal decrease in profit from an increase in depth is equal to the optimized quantity of groundwater use scaled by the marginal change in pumping cost due to the depth increase (this follows from the Envelope Theorem):

$$\frac{d\pi^*}{dD_t} = -w_t^* \frac{\partial c(D_t)}{\partial D_t}$$

We can quantify this expression empirically for each parcel-year using (i) observed w_{it}, the dependent variable from our regressions, and (ii) microdata on electricity prices and operating pump efficiencies, which let us use Equation (6) to parameterize $\left(\frac{\partial c(D_t)}{\partial D_t}\right)_{it}$. Averaging their product across parcel-years and weighting by acreage, the average marginal willingness to pay to avoid a 1-foot depth increase is $103 per parcel-year.

To assess non-marginal increases in depth, we can integrate the above expression substituting the fitted regression model (i.e. $\hat{w}_{it}$ as a function of depth) for w_t^*. For the uninteracted model, the average willingness to pay to avoid a 10% increase in depth is $2,014 per parcel-

year. For the model that interacts temperature with depth, this average willingness to pay increases slightly to $2,023 per parcel-year.

6 Effects of climate change on the California aquifer

In this section, we move beyond our reduced-form effects to measure the impacts of climate change on California's aquifer and on welfare. Doing so requires a dynamic model, because groundwater extraction in one period will raise aquifer depth (i.e., lower water levels) in the subsequent period.

6.1 A model of aquifer depth under climate change

We specify a stylized dynamic model of adaptive irrigation and groundwater depletion for a unified groundwater basin. The basin contains N heterogeneous farmers (indexed by i) who extract groundwater for irrigation. Farmer i's total AF of groundwater use in year t is:

$$w_{it} \;=\; F\Big(\mathbf{T}_{it}, R_{it}, D_{st}; \boldsymbol{\beta}\Big) \;+\; \gamma_i \;+\; \delta_t \;+\; \varepsilon_{it} \tag{8}$$

Hotter temperatures ($\mathbf{T}_{it}$) increase w_{it} through adaptive irrigation, while rainfall (R_{it}) reduces farmers' water use (i.e., $\frac{\partial F}{\partial \mathbf{T}} > 0$, $\frac{\partial F}{\partial R} < 0$). Groundwater depth ($D_{st}$) for sub-basin s governs farmers' costs of irrigation, since deeper aquifer levels increase the energy costs of pulling groundwater up to the surface, leading to demand response (i.e., $\frac{\partial F}{\partial D} < 0$). Unit fixed effects ($\gamma_i$) control for farm-specific factors (e.g., crop choice, irrigation capital), while time fixed effects (δ_t) account for common shocks (e.g., market prices, drought severity). Finally, ε_{it} accounts for idiosyncratic time-varying determinants of groundwater use.

To model the dynamic link between groundwater extraction and groundwater depths, we specify the following aggregate equation of motion (following Carleton, Crews, and Nath (2023)):

$$\overline{D}_{t+1} \;=\; \overline{D}_t \;+\; \rho\Big[(1-\psi)\overline{W}_t \;-\; \overline{R}_t\Big] \tag{9}$$

$\overline{D}_t$ is the average groundwater depth (in feet) in year t across the full aquifer. $\overline{R}_t$ is annual precipitation (in feet) averaged across the full area of the aquifer; rainfall adds water to

the aquifer, decreasing $\overline{D}_{t+1}$. $\overline{W}_t$ is the average groundwater extraction (in feet) across the full basin, equivalent to $\sum_{i=1}^{N} w_{it}$ divided by the area of the basin (in acres). The scalar $\psi \in [0, 1]$ accounts for the return flow of irrigation: for every unit of groundwater extracted, the proportion ψ re-percolates back into the aquifer, while the proportion $1 - \psi$ leaves the aquifer via evapotranspiration. Finally, ρ is the inverse of specific yield, which accounts for soil porosity in translating groundwater displacement to changes in depth (Loheide, Butler, and Gorelick (2005)).[24]

To operationalize Equation (8), we slightly modify our preferred temperature response specification (Equation (7)) to reduce the number of parameters and add flexibility. First, to avoid overfitting we replace our preferred $f(T_{it})$ of 10 temperature bins with a piecewise quadratic in annual degree days. Second, we interact this piecewise quadratic with log depth (as in Figure 4), to allow the cost-benefit tradeoff of adaptive irrigation to shift as temperatures rise. Third, we replace our preferred $g(R_{imt})$ of 12 monthly precipitation coefficients with a quadratic in annual temperature.[25]

To simulate forward, we populate temperature and precipitation using climate projections from Hsiang et al. (2017). These projections combine a suite of 45 climate models, yielding projected counts of days in 1°C temperature bins and in 2mm precipitation bins at a county-year resolution. We aggregate bins to construct projected degree days (linear and squared terms, separately above an below 24.5°C) and projected annual precipitation (in mm) at the county-year level. We present results for two climate scenarios: RCP4.5 and RCP8.5.

Appendix D describes our model calibration and simulation procedure in detail. We present a few salient aspects here. First, we restrict our focus to the San Joaquin Valley groundwater basin, which contains 71% of our estimation sample and is the most economically important groundwater basin in California. The San Joaquin basin comprises 16 sub-basins with spatially heterogeneous climate realizations, irrigation practices, groundwa-

24. Specific yield is the volume of water that can be drained from porous soil by gravity, relative to the total soil volume. Higher porosity is associated with higher specific yield (lower ρ). Extracting 1 AF from (porous) sandy soil increases depth by less than extracting 1 AF from (less porous) sandy clay soil, since the former yields more water per unit of volume.

25. Appendix Figure B3 shows that this alternate functional form on precipitation recovers temperature and depth parameter estimates that are nearly identical to Equation (7).

ter depths, and soil characteristics. We therefore calibrate ρ separately for each sub-basin, following Loheide, Butler, and Gorelick (2005). We set $\psi = 0.2$, following Dewandel et al. (2008). We then aggregate the rightmost term of Equation (9) up from the sub-basin to the basin, which allows for horizontal diffusion across sub-basins of the aquifer.

In each year, we predict $\hat{w}_{it}$ by plugging climate projections and depth into our fitted model of Equation (8). We initialize the model using observed baseline depth, and then endogenize depth in all subsequent years by implementing a calibrated parameterization of Equation (9). We consider three main scenarios, which we simulate under both RCP4.5 and RCP8.5. In our preferred scenario, farmers' groundwater pumping depends both on temperature and on aquifer depth. We also consider alternative scenarios in which farmers respond only to depth (i.e., effectively holding temperatures fixed), or only to temperature (i.e., removing any cost response).

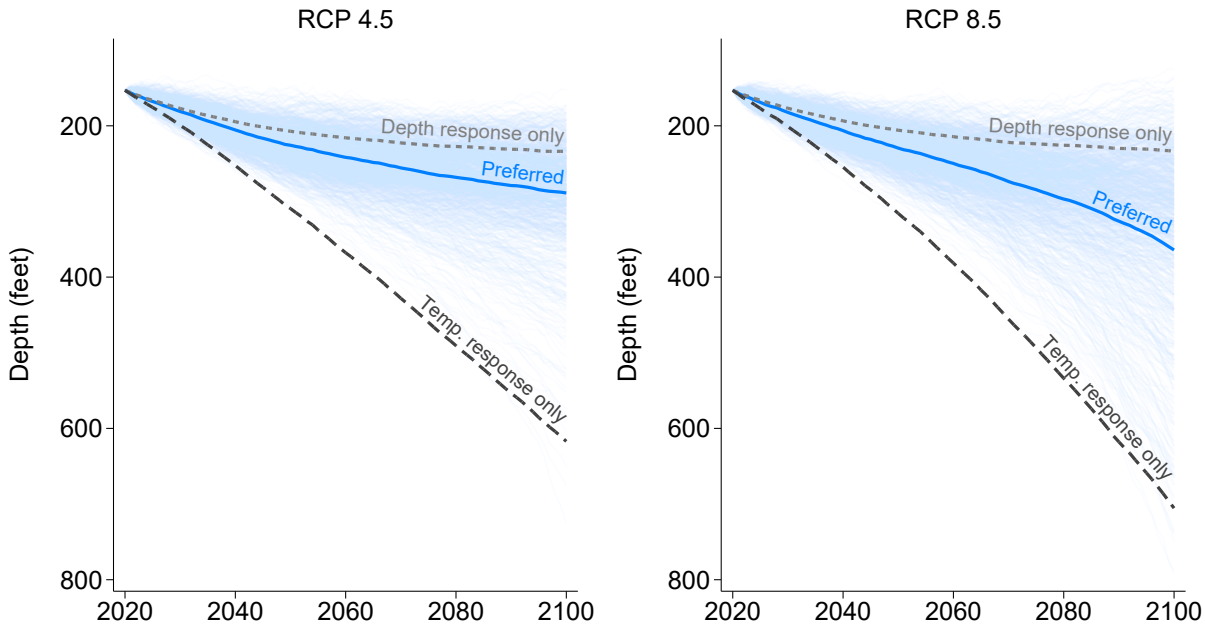
To capture both econometric uncertainty (from our estimate model) and climate uncertainty (from the projections), we simulate 1,000 iterations for each scenario. Each iteration makes joint draws from (i) the sampling distribution of regression model coefficients (including year fixed effects for 2008–2019); (ii) a vector of simulated years drawing from $\{2008, \ldots, 2019\}$ with replacement (to populate year fixed effects out of sample); and (iii) one of the 45 climate models (weighting models following the procedure outlined in Hsiang et al. (2017); Carleton et al. (2022)).

6.2 Simulation results: groundwater depth

Figure 5 presents the results of this simulation exercise. Under our preferred scenario (blue line), we find that groundwater depth increases from 153 feet in 2020 to 289 feet in 2100 under RCP4.5, and to 364 feet in 2100 under RCP8.5. We report the uncertainty in this mean outcome by plotting all 1,000 iterations of our simulations in light blue. This suggests that even accounting for uncertainty, climate change will lead to a significant fall in aquifer levels in the San Joaquin basin, already among the most stressed aquifers on the planet.

How much of this increase in depth is driven by climate change-induced adaptive irrigation? We can compare our preferred model to the "depth response only" model (dash-dotted gray line in Figure 5), where farmers do not change their pumping in response to temperature,

Figure 5: Simulated impact of climate change on groundwater depths



Notes: This figure presents the results of our simulation exercise. In the top panels, we model groundwater depth from 2020 through 2100. The left panel shows results under RCP4.5 (moderate climate change) and the right panel shows results under RCP8.5 (severe climate change). In both panels, the blue line denotes our preferred model, where farmers' groundwater extraction responds to both temperature and depth. We run 1,000 draws of the model. The thick blue line plots the mean over these simulation runs. The individual light blue lines show individual simulations. These draws incorporate both econometric and climate model uncertainty. The two gray lines plot the means over simulation runs where we only have farmers responding to depth (dash-dot) and only have farmers responding to temperature (dashed). The bottom panel (in progress) will plot the welfare consequences of climate change, as measured by revealed preference.

but still reduce pumping in response to greater depths. Under "depth response only," depth increases from 153 feet in 2020 to 234 feet in 2100. This comparison implies that RCP4.5 (RCP8.5) leads to 68% (162%) greater end-of-century depth than what end-of-century would be absent increased temperature response. It also suggests that 41% (62%) of the increase in depth from 2020 to 2100 is explained by temperature response.

Our results also highlight the importance of the "natural scarcity" effect, whereby farmers' depth response moderates the rate of aquifer depletion. Our "temperature response only" model (dashed gray line in Figure 5) removes this moderating effect, such that farmers pump more in response to increasing temperatures, but do not pump less in response to increase depth. This results to a staggering increase in aquifer depth from 153 feet in 2020 to 617 feet (705 feet) in 2100 under RCP4.5 (RCP8.5). This suggests that the "natural scarcity" effect mitigates 71% (62%) of the aquifer depletion likely to occur over this 80-year horizon,

and 85% (72%) of the aquifer depletion likely to occur as a result of climate change. However, even net of this strong natural scarcity mechanism, our simulations imply that climate change will meaningfully alter the trajectory of California's aquifers.

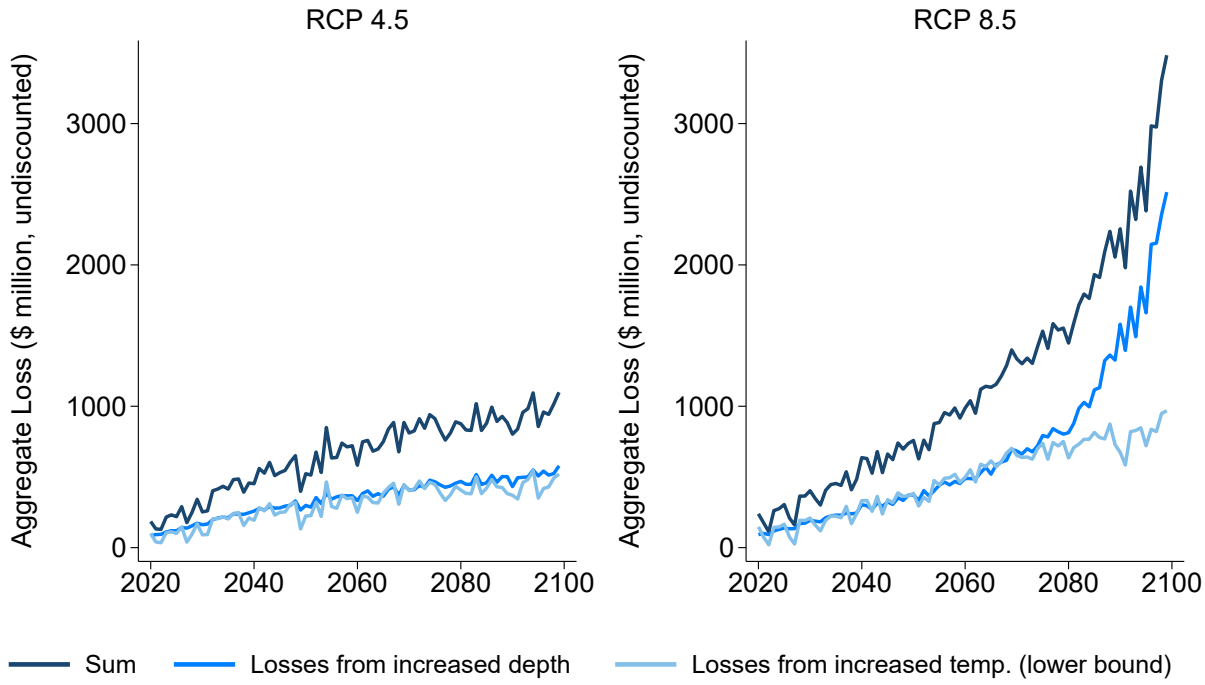
6.3 Simulation results: welfare losses due to climate change

We can bound the welfare losses due to climate change in our preferred scenario using an approach analogous to the willingness to pay analysis in Section 5.3. We describe this approach in detail in Appendix D.3, which departs from our earlier welfare calculations in several ways.

First, since we are assessing the welfare loss due to climate change instead of marginal temperature increases, we consider the difference between projected temperatures in each future year and 2008–19 in-sample average temperatures. In expectation, this accounts for the gradual rightward shift in the temperature distribution depicted in Appendix Figure A2). Second, since the cumulative effect of adaptive irrigation alters the trajectory of groundwater depth, we consider the difference in simulated year-t depth between our preferred vs. "depth response only" scenarios. The latter provides a reference point for what projected depth would be absent shifts in the temperature distribution. Third, we parameterize parcel-specific $\frac{\partial c(D_t)}{\partial D_t}$ using electricity prices and operating pump efficiencies from 2019, our last sample year. This extrapolation holds constant both electricity prices and pumping capital through 2099—a simplifying assumption that is preferable to modeling changes in either factor in the absence of high-quality forecasts. Fourth, we sum losses across all parcels in the San Joaquin basin to arrive at a lower bound on aggregate welfare. Fifth, we account for uncertainty in both estimated model parameters and climate projections by calculating welfare losses within each simulation iteration and taking the mean across all 1,000 iterations.

Figure 6 reports simulated welfare effects under climate change using our preferred scenario, where farmers change their pumping behavior in response to both temperature and depth. Under RCP4.5, our simulations predict that annual losses for farmers in the San Joaquin Valley grow over time, exceeding $1.1 billion per year in 2099. Roughly 47% of this welfare bound comes from the direct effect of hotter temperatures (while holding depth fixed

Figure 6: Simulated welfare losses due to climate change



Notes: This figure plots the average aggregate welfare losses for all agricultural groundwater pumpers in the San Joaquin basin. These estimates all come from our preferred scenario, implementing Equations (D4)–(D5) in Appendix D.3. Losses from increased temperatures represent a lower bound, since they do not account for lost crop revenues. Losses from increase depth are point estimates. The sum of both losses is therefore a lower bound. We report aggregate annual losses in millions of $2019 without discounting.

at the "depth response only" level). The remaining 53% comes from deeper groundwater levels driven by climate change-induced pumping.

Under RCP8.5, our simulations predict that annual losses will increase at an increasing rate, exceeding $3.5 billion per year in 2099. The direct temperature-related losses and the accompanying depth-related losses contribute roughly equally up until 2080, as which point the temperature-related losses level off while the depth-related losses rapidly accelerate—such that by 2099, depth-related losses represent 72% of this aggregate lower bound. The timing of this kink in depth-related losses aligns both with the divergence in the two climate scenarios after 2070 (see the top panels of Appendix Figure A2) and with the inflection point in depth in our preferred scenario under RCP8.5 in Figure 5.

How large are these lower bounds on the welfare loss from climate change? As a point of reference, the total revenue from crop production across the San Joaquin Valley was $24

billion in 2018 (Escriva-Bou et al. (2023)). Our results therefore imply that climate change will reduce aggregate welfare in 2099 by at least 4.5% of current crop revenues under RCP4.5, or by at least 14.5% of crop revenues under RCP8.5.

We can also use these simulated welfare bounds to assess the present discounted welfare loss due to climate change. Applying a 3% annual discount rate and integrating from 2020 to 2099, the total present-discounted welfare loss (as of 2020) has a lower bound of $15.5 billion under RCP4.5 and of $22.6 billion under RCP8.5. If we increase the discount rate to 5%, these lower bounds fall to $8.6 billion under RCP4.5 and to $11.3 billion under RCP8.5.

7 Climate change and the stock externality

The above simulations demonstrate that, under open access, climate change will impose substantial welfare consequences on California farmers. Next, we outline our procedure for empirically quantifying the magnitude of the stock externality and how this externality changes with climate change (work in progress). Our algorithm is as follows:

1. Define the shape of the time path of aggregate groundwater extraction $\{W_t\}$ as a piecewise polynomial in time, defined by parameters from the feasible set $\mathbf{\Theta}$. This feasible set is restricted such that for all periods, $\{W_t\}$ is bounded below by 0 (i.e., zero aggregate groundwater extraction in year t) and bounded above by twice the in-sample maximum of W_t.

2. Define a large number of candidate parameters $\theta^C \in \mathbf{\Theta}$. Each θ^C implies a unique $\{W_t^C\}$.

3. For each θ^C:

   (a) Simulate the preferred model scenario, stepping forward for years $t = 2020, \ldots, 2099$. For each year t:

      i. Solve for the candidate groundwater tax τ_t^C that would induce aggregate pumping equal to W_t^C. (Our functional form is homogeneous across farms, implying a one-to-one mapping between a uniform τ_t and aggregate W_t.)

ii. Use the equation of motion to solve for the corresponding depth D^C_{t+1}, which dynamically feeds into the next year's simulation. (Our linear equation of motion implies a one-to-one mapping between W_t and D_{t+1}).

(b) After having solved for the sequence $\{\tau^C_t\}$, calculate for each year $t \in \{2020, \ldots, 2099\}$:

i. The willingness to pay to have avoided the tax τ^C_t in year t (compared to zero tax), aggregated across all farmers ($COST^C_t$).

ii. The willingness to pay for decreases in groundwater depth induced by the current-period tax τ^C_t, discounted to year-t dollars and summed over all future years, and aggregated across all farmers ($BENEFIT^C_t$).

iii. The total revenue collected in groundwater pumping taxes in year t, given tax τ^C_t and the simulated groundwater pumping in that year ($TAXREV^C_t$).

4. The optimal path $\{W^*_t\}$ is defined by the θ^* that maximizes the discounted sum of welfare across all periods:

$$\theta^* \;=\; \underset{\theta^C}{\text{argmax}} \sum_t \beta^t \Big[BENEFIT^C_t \;-\; COST^C_t \;+\; TAXREV^C_t \Big]$$

The optimal set of $\{\tau^*_t\}$ defined by θ^* are equivalent to the stock externality.

5. Conduct Steps 2–4 twice: once for the scenario with depth response only (i.e., no adaptive irrigation response), and once for the preferred scenario. Comparing the two resulting optimal paths and stock externalities will reveal the impacts of climate change.

8 Conclusion

Climate change is already impacting economic activity and human health (Burke, Hsiang, and Miguel (2015), Hsiang et al. (2017), and Carleton et al. (2022). Mitigation alone will not be sufficient (Lee et al. (2023)): a growing body of literature demonstrates the necessity of climate change adaptation in moderating current and future costs of an evolving climate (Carleton et al. 2024).

In this paper, we use rich data to empirically characterize adaptation behavior and its costs in the context of agricultural water use. We find that farmers are privately willing to pay at least $775 per parcel-year to avoid high heat, and respond to hot temperatures by substantially increasing irrigation. However, this comes at a social cost: as farmers pump more water, aquifer depth increases, which in turn raises pumping costs for all aquifer users. We therefore document a feedback loop: if farmers respond to hot temperatures this year by pumping additional groundwater, the costs of groundwater pumping rise in the future.

We incorporate these features into a dynamic model of the aquifer. We find that climate change will increase pumping and significantly exacerbate groundwater depletion by end of century. However, we also document that the aquifer's "natural scarcity" effect moderates the overall impact of climate change on groundwater extraction. Failing to take this into account would lead to a substantial *over*estimate of the costs of climate change. We use our model to calculate the welfare consequences of climate change, accounting both for farmers' responses to temperature and to depth, and find that climate change leads to losses of at least $1 to $3 billion annually. In ongoing work, we use our simulation model to estimate the size of the stock externality, and quantify how this changes with climate change.

While we model the dynamics of climate change adaptation in one specific context, agricultural groundwater in California, the lessons from this paper are likely broader. There is also a clear feedback loop in adapting to extreme heat with air conditioning: while operating an air conditioner can keep you cooler today, the electricity this same air conditioner consumers— and thus the greenhouse gases it produces—lead to worsening climate change tomorrow. Even more broadly, whenever there is an upward-sloping supply curve of adaptation actions, taking some adaptation action in the present makes future adaptation more costly. Understanding the impacts of climate adaptation therefore requires modeling its dynamic consequences.

References

Abajian, Alexander, Tamma Carleton, Kyle C. Meng, and Olivier Deschenes. 2025. "Quantifying the global climate feedback from energy-based adaptation." *Nature Communications* 16:3928.

Allen, Richard C, and Micha Gisser. 1984. "Competition versus optimal control in groundwater pumping when demand is nonlinear." *Water Resources Research* 20 (7): 752–756.

Auffhammer, Maximilian. 2022. "Climate Adaptive Response Estimation: Short and long run impacts of climate change on residential electricity and natural gas consumption." *Journal of Environmental Economics and Management* 114:102669.

Boser, Anna, Kelly Caylor, Ashley Larsen, Madeleine Pascolini-Campbell, John T. Reager, and Tamma Carleton. 2024. "Field-scale crop water consumption estimates reveal potential water savings in California agriculture." *Nature Communications* 15:2366.

Bruno, Ellen M., Katrina Jessoe, and W. Michael Hanemann. 2024. "The dynamic impacts of pricing groundwater." *Journal of the Association of Environmental and Resource Economists* 11 (5): 1201–1227.

Burke, Marshall, and Kyle Emerick. 2016. "Adaptation to climate change: Evidence from U.S. agriculture." *American Economic Journal: Economic Policy* 8 (3): 106–140.

Burke, Marshall, Solomon M. Hsiang, and Edward Miguel. 2015. "Global non-linear effect of temperature on economic production." *Nature* 527, no. 7577 (October 21, 2015): 235–239.

Burke, Marshall, Mustafa Zahid, Mariana C. M. Martins, Christopher W. Callahan, Richard Lee, Tumenkhusel Avirmed, Sam Heft-Neal, Mathew Kiang, Solomon M. Hsiang, and David Lobell. 2024. "Are we adapting to climate change?" NBER working paper 32985.

Burlig, Fiona, Louis Preonas, and Matt Woerman. 2025. "Groundwater and crop choice in the short and long run." NBER working paper 28706.

Carleton, Tamma, Levi Crews, and Ishan Nath. 2023. "Agriculture, trade, and the spatial efficiency of global water use." Working paper.

Carleton, Tamma, Esther Duflo, Kelsey Jack, and Guglielmo Zappalà. 2024. "Adaptation to climate change." In *Handbook of the Economics of Climate Change,* edited by Lint Barrage and Solomon Hsiang. Elsevier.

Carleton, Tamma, Amir Jina, Michael Delgado, Michael Greenstone, Trevor Houser, Solomon Hsiang, Andrew Hultgren, et al. 2022. "Valuing the global mortality consequences of climate change accounting for adaptation costs and benefits." *The Quarterly Journal of Economics* 137, no. 4 (November): 2037–2105.

Deschênes, Olivier, and Michael Greenstone. 2007. "The economic impacts of climate change: evidence from agricultural output and random fluctuations in weather." *American Economic Review* 97 (1): 354–385.

Dewandel, B., J.-M. Gandolfi, D. de Condappa, and S. Ahmed. 2008. "An efficient methodology for estimating irrigation return flow coefficients of irrigated crops at watershed and seasonal scale." *Hydrological Processes* 22 (11): 1700–1712.

Escriva-Bou, Alvar, Ellen Hanak, Spencer Cole, and Josué Medellín-Azuara. 2023. *The Future of Agriculture in the San Joaquin Valley.* Public Policy Institute of California policy brief.

Famiglietti, James S, Minhui Lo, Sing L Ho, James Bethune, KJ Anderson, Tajdarul H Syed, Sean C Swenson, Caroline R de Linage, and Matthew Rodell. 2011. "Satellites measure recent rates of groundwater depletion in California's Central Valley." *Geophysical Research Letters* 38 (3).

Foster, Timothy, Nicholas Brozović, and Adrian P Butler. 2014. "Modeling irrigation behavior in groundwater systems." *Water resources research* 50 (8): 6370–6389.

Gisser, Micha, and David A Sanchez. 1980. "Competition versus optimal control in groundwater pumping." *Water resources research* 16 (4): 638–642.

Gordon, H. Scott. 1954. "The economic theory of a common-property resource: The fishery." *Journal of Political Economy* 62 (2): 124–142.

Hagerty, Nick. 2022. "Adaptation to surface water scarcity in irrigated agriculture." Working paper.

———. 2023. "What holds back water markets? Transaction costs and the gains from trade." Working paper.

Hotelling, Harold. 1931. "The economics of exhaustible resources." *Journal of Political Economy* 39 (2): 131–175.

Hrozencik, R Aaron, Dale T Manning, Jordan F Suter, Christopher Goemans, and Ryan T Bailey. 2017. "The heterogeneous impacts of groundwater management policies in the Republican River Basin of Colorado." *Water Resources Research* 53 (12): 10757–10778.

Hsiang, Solomon, Robert Kopp, Amir Jina, James Rising, Michael Delgado, Shashank Mohan, DJ Rasmussen, Robert Muir-Wood, Paul Wilson, Michael Oppenheimer, et al. 2017. "Estimating economic damage from climate change in the United States." *Science* 356 (6345): 1362–1369.

Hultgren, Andrew, Tamma Carleton, Michael Delgado, Diana R Gergel, Michael Greenstone, Trevor Houser, Solomon Hsiang, Amir Jina, Robert E Kopp, Steven B Malevich, et al. 2025. "Impacts of climate change on global agriculture accounting for adaptation." *Nature* 642 (8068): 644–652.

Hurr, R. Theodore, and David W. Litke. 1989. *Estimating Pumping Time and Groundwater Withdrawals Using Energy-Consumption Data.* U.S. Geological Survey, Water-Resources Investigations Report 89-4107.

Jasechko, Scott, Hansjörg Seybold, Debra Perrone, Ying Fan, Mohammed Shamsudduha, Richard G. Taylor, Othman Fallatah, and James W. Kirchner. 2024. "Rapid groundater decline and some cases of recovery in aquifers globally." *Nature* 625:715–721.

Joaquín Morales, Zaira, and Andrew Ayres. 2023. *PPIC Farm Sizes in the San Joaquin Valley: Land Uses.* https://www.ppic.org/data-set/ppic-farm-sizes-in-the-san-joaquin-valley-land-uses/. Dataset.

Johnson, Renée, and Betsy A Cody. 2015. *California agricultural production and irrigated water use.*

Lee, Hoesung, Katherine Calvin, Dipak Dasgupta, Gerhard Krinner, Aditi Mukherji, Peter Thorne, Christopher Trisos, José Romero, Paulina Aldunce, Ko Barret, et al. 2023. "IPCC, 2023: Climate change 2023: Synthesis report, summary for policymakers. Contribution of working groups i, II and III to the sixth assessment report of the intergovernmental panel on climate change [core writing team, h. Lee and j. Romero (eds.)]. IPCC, geneva, Switzerland."

Li, Jiazhen, Wenxu Dong, Oene Oenema, Tuo Chen, Chunsheng Hu, Haijing Yuan, and Liying Zhao. 2019. "Irrigation reduces the negative effect of global warming on winter wheat yield and greenhouse gas intensity." *Science of the Total Environment* 646:290–299.

Li, Yan, Kaiyu Guan, Bin Peng, Trenton E Franz, Brian Wardlow, and Ming Pan. 2020. "Quantifying irrigation cooling benefits to maize yield in the US Midwest." *Global Change Biology* 26 (5): 3065–3078.

Liu, Pang-Wei, James S. Famiglietti, Adam J. Purdy, Kyra H. Adams, Avery L. McEvoy, John T. Reager, Rajat Bindlish, David N. Wiese, Cédric H. David, and Matthew Rodell. 2022. "Groundwater depletion in California's Central Valley accelerates during megadrought." *Nature Communications* 13:7825.

Loheide, Steven P., James J. Butler, and Steven M. Gorelick. 2005. "Estimation of groundwater consumption by phreatophytes using diurnal water table fluctuations: A saturated-unsaturated flow assessment." 41 (7).

Luan, Xiangyu, Riccardo Bommarco, Anna Scaini, and Giulia Vico. 2021. "Combined heat and drought suppress rainfed maize and soybean yields and modify irrigation benefits in the USA." *Environmental research letters* 16 (6): 064023.

Maucieri, Carmelo, Maurizio Borin, Francesco Morbidini, Tjaša Pogačar, Marko Flajšman, Graziano Ghinassi, Leonardo Verdi, Anna Dalla Marta, Roberto Ferrise, et al. 2025. "Projecting the impacts of climate change on soybean production and water requirements using AquaCrop model." *European Journal of Agronomy* 165:127538.

Mendelsohn, Robert O., William Nordhaus, and Daigee Shaw. 1994. "The impact of global warming on agriculture: A Ricardian analysis." *American Economic Review* 84 (4): 753–771.

Nordhaus, William. 2019. "Climate change: The ultimate challenge for economics." *American Economic Review* 109 (6): 1991–2014.

Nordhaus, William D. 1993. "Reflections on the economics of climate change." *Journal of Economic Perspectives* 7 (4): 11–25.

Olson, Mancur. 1965. *The Logic of Collective Action: Public Goods and the Theory of Groups.* Cambridge, MA: Harvard University Press.

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

Pindyck, Robert S. 2013. "Climate change policy: What do the models tell us?" *Journal of Economic Literature* 51 (3): 860–872.

Provencher, Bill, and Oscar Burt. 1993. "The externalities associated with the common property exploitation of groundwater." *Journal of Environmental Economics and Management* 24:139–158.

Rad, Mani Rouhi, Nicholas Brozović, Timothy Foster, and Taro Mieno. 2020. "Effects of instantaneous groundwater availability on irrigated agriculture and implications for aquifer management." *Resource and Energy Economics* 59:101129.

Rasmussen, D.J., M. Meinshausen, and R.E. Kopp. 2016. "Probability-weighted ensembles of U.S. county-level climate projections for climate risk analysis." *Journal of Applied Meteorology and Climatology.*

Schlenker, Wolfram, Michael Hanemann, and Anthony C. Fisher. 2007. "Water availability, degree days, and the potential impact of climate change on irrigated agriculture in California." *Climatic Change* 81 (1): 19–38.

Schlenker, Wolfram, W. Michael Hanemann, and Anthony C. Fisher. 2005. "Will U.S. agriculture really benefit from global warming? Accounting for irrigation in the hedonic approach." *American Economic Review* 95 (1): 395–406.

Schlenker, Wolfram, and Michael J. Roberts. 2009. "Nonlinear temperature effects indicate severe damages to U.S. crop yields under climate change." *Proceedings of the National Academy of Sciences of the United States of America* 106 (37): 15594–15598.

Tack, Jesse, Andrew Barkley, and Nathan Hendricks. 2017. "Irrigation offsets wheat yield reductions from warming temperatures." *Environmental Research Letters* 12 (11): 114027.

Timmins, Christopher. 2002. "Measuring the dynamic efficiency costs of regulators' preferenes: Municipal water utilities in the arid West." *Econometrica* 70 (2): 603–629.

United States Department of Agriculture. 2023. *Census of Agriculture: 2018 Irrigation and Water Management Survey.*

Van Vuuren, Detlef P., Jae Edmonds, Mikiko Kainuma, Keywan Riahi, Allison Thomson, Kathy Hibbard, George C. Hurtt, et al. 2011. "The Representative Concentration Pathways: An Overview." *Climatic Change* 109 (1): 5.

Vanuytrecht, Eline, Dirk Raes, Pasquale Steduto, Theodore C Hsiao, Elias Fereres, Lee K Heng, Marga Garcia Vila, and Patricia Mejias Moreno. 2014. "AquaCrop: FAO's crop water productivity and yield response model." *Environmental Modelling & Software* 62:351–360.

Groundwater in a changing climate

Supplementary Appendix: For online publication

R. Aaron Hrozencik ⓡ Matt Woerman ⓡ Fiona Burlig ⓡ Louis Preonas*

Appendix contents

A Data **A1**

B Reduced-form sensitivities **A3**

C Theoretical model **A14**
C.1 Setup A14
C.2 Social Planner's problem A15
C.3 Farmer's problem A16
C.4 The stock externality A17
C.5 Climate change A17
C.6 Willingness to pay A19

D Simulation model **A22**
D.1 Model setup A22
D.2 Simulation algorithm A24
D.3 Quantifying welfare losses A26

E Additional model results **A28**

F Assessing potential bias in depth coefficient due to dynamic feedback **A31**

G Comparison to engineering estimate **A34**

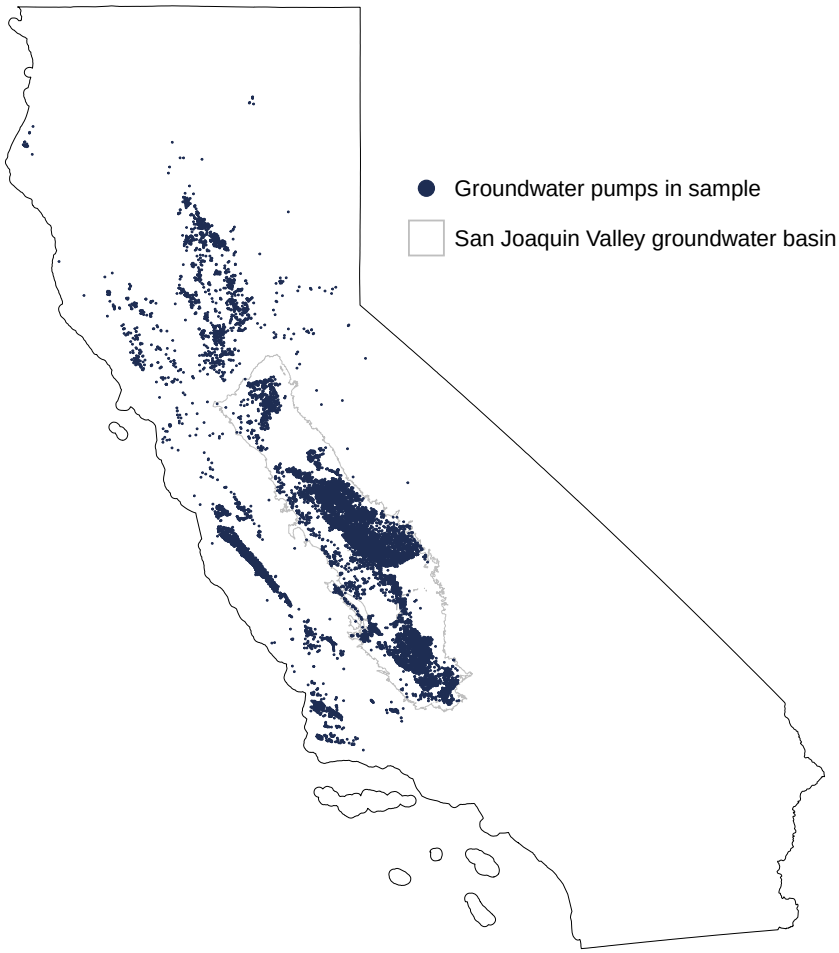
Appendix references **A35**

∗. Hrozencik: University of Georgia. Email: Aaron.Hrozencik@uga.edu. Woerman: Department of Agricultural and Resource Economics, Colorado State University. Email: matt.woerman@colostate.edu. Burlig: Harris School of Public Policy and Energy Policy Institute (EPIC), University of Chicago, and NBER. Email: burlig@uchicago.edu. Preonas: Department of Agricultural and Resource Economics, University of Maryland. Email: lpreonas@umd.edu.

A Data

Details on our estimation sample Appendix Figure A1 shows the geographic distribution of groundwater pumps in our sample. Appendix Table 1 presents summary statistics for key variables.

Figure A1: Sample of groundwater pumpers

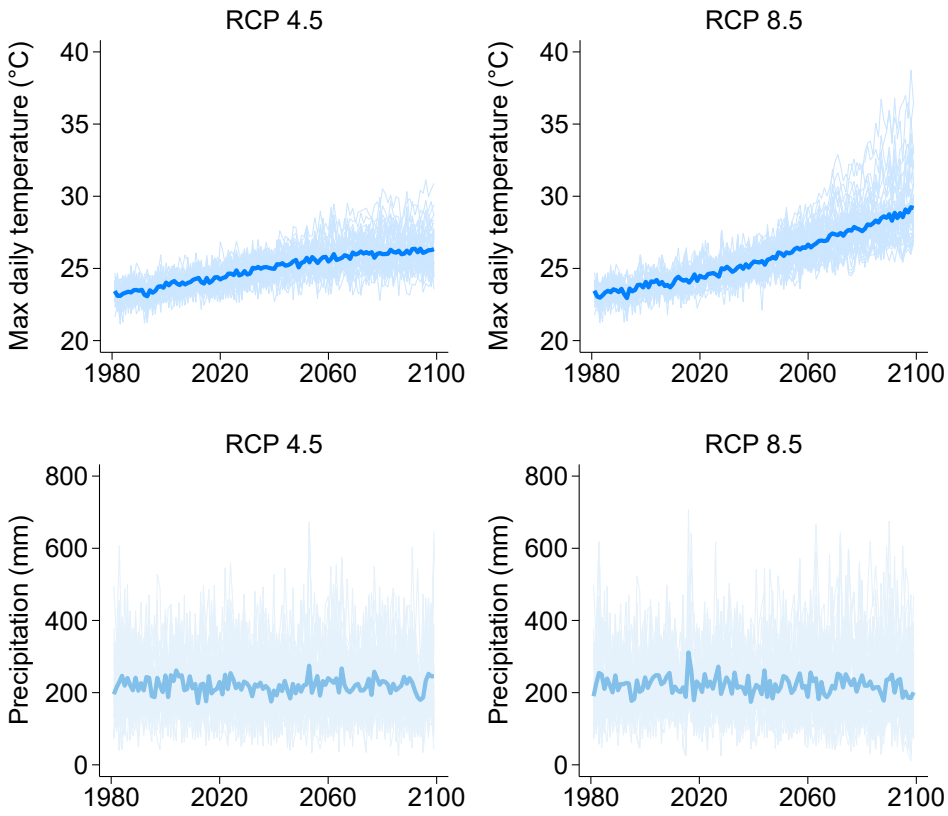


Notes: This figure plots the groundwater pumpers in our sample. As described in Section 4, our sample is made up of groundwater pumpers in the Pacific Gas and Electric service territory. In light gray, we plot the outline of the San Joaquin Valley groundwater basin, home to 71% of our pumps.

Climate projections Appendix Figure A2 shows how temperature and precipitation in the San Joaquin basin are projected to change by end of century, using data from Hsiang

et al. (2017) and Rasmussen, Meinshausen, and Kopp (2016). We plot individual projections from each climate model in the ensemble, as well as a weighted mean.

Figure A2: Projected temperature and precipitation

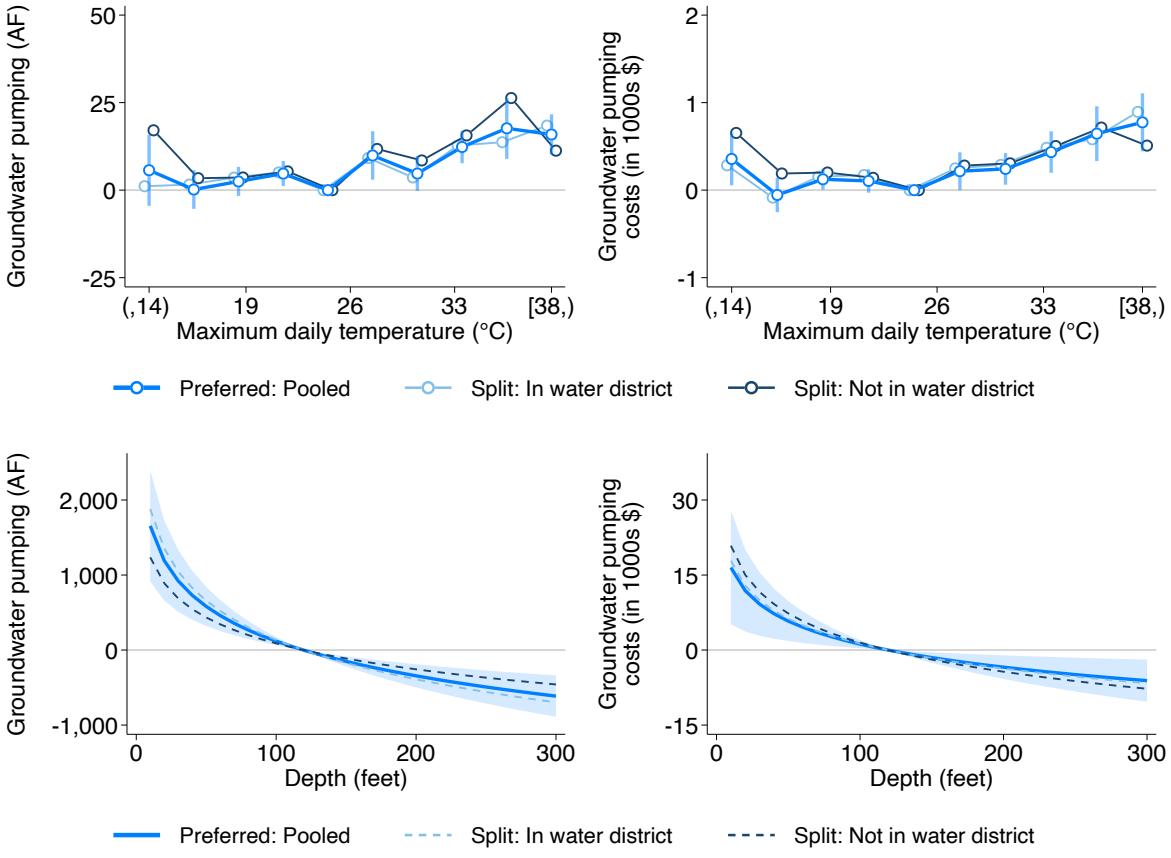


Notes: This figure plots projected temperature (top row) and precipitation (bottom row) in the San Joaquin basin from 1980 to 2100. We plot RCP4.5 (moderate climate change) on the left, and RCP8.5 (severe climate change) on the right. Projection data come from Hsiang et al. (2017) and Rasmussen, Meinshausen, and Kopp (2016), and are derived from an ensemble of 43 climate models for RCP4.5 and 45 models for RCP8.5. The thick line displays the probability-weighted mean, according to weights laid out in Rasmussen, Meinshausen, and Kopp (2016). The light lines present the projection from each individual model.

B Reduced-form sensitivities

Robustness check: surface water It is possible that unobserved within-parcel variation in surface water access covaries with temperature, depth, and groundwater pumping. This has the potential to bias our temperature coefficients, if hotter years tend to have smaller surface water allocations, meaning that farmers need to compensate for lost surface water with additional groundwater. It also has the potential to bias our depth coefficients, if groundwater scarcity (i.e. depth) and surface water scarcity are correlated, generating a similar substitution pattern. While we cannot control for surface water allocations directly, we can test for these confounding effects by splitting the sample into farms that lie within water districts (and therefore receive annual surface water allocations) vs. farms that do not (and therefore are unlikely to have surface water access). Appendix Figure B1 reports these

Figure B1: Robustness – access to surface water

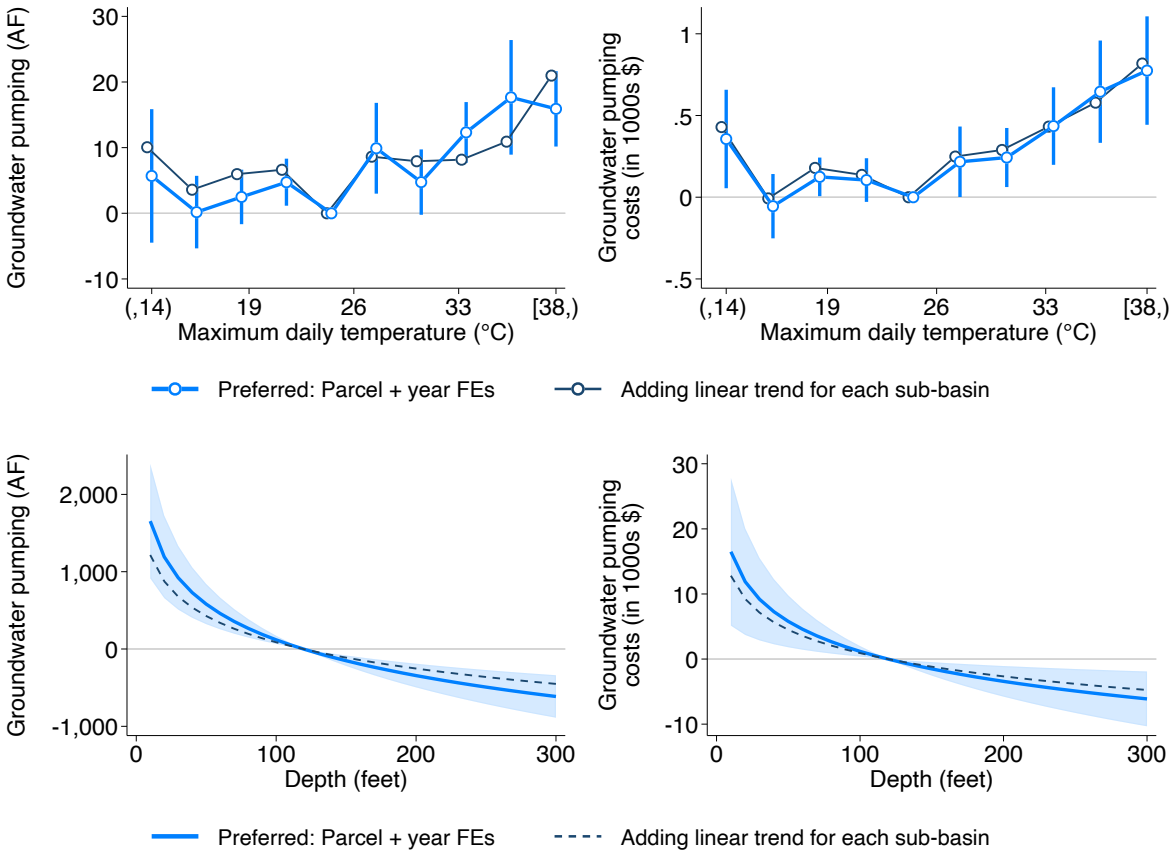


Notes: This figure shows how our main results vary with surface water access, as measured by belonging to a surface water district. We plot our pooled estimates (from Figure 2) in middle blue. In light blue, we plot estimates for farms in water districts. In dark blue, we plot estimates for farms not in water districts. The top panels plot temperature estimates; the bottom panels plot depth estimates (relative to the mean depth of 119.5 feet). The left panels plot effects on annual groundwater extraction; the right panels plot effects on annual groundwater expenditures. We estimate Equation (7) with 10 temperature bins (omitting the 23–26°C bin), log depth, month-specific linear precipitation controls, parcel fixed effects, and year fixed effects. Regressions are weighted by croppable acres. Standard errors are clustered by water district. We plot 95% confidence intervals for our preferred estimates with vertical lines (top panels) and shading (bottom panels).

split sample results, which are remarkably similar for both sets of farms. We cannot reject equality in temperature or depth effects for farms inside vs. outside of water districts. This assuages concerns about omitted variable bias due to surface water.

Robustness check: non-parallel trends It is possible that non-parallel counterfactual trends in groundwater depth across sub-basins is confounding our depth estimates. In other words, sub-basins experiencing a secular decline in aquifer levels may not serve as good "controls" for sub-basins not experiencing such a decline. To test for this, we modify Equation (7) by controlling for a separate linear trend for each sub-basin s. Appendix Figure B2 reports these results, revealing similar temperature estimates and slightly attenuated depth estimates: in the bottom-left panel, $\hat{\beta}$ decreases from -667 ($p < 0.001$) to -490 ($p < 0.001$). This attenuation could reflect the removal of omitted variables bias *or* could simply reflect

Figure B2: Robustness – adding sub-basin-specific linear trends

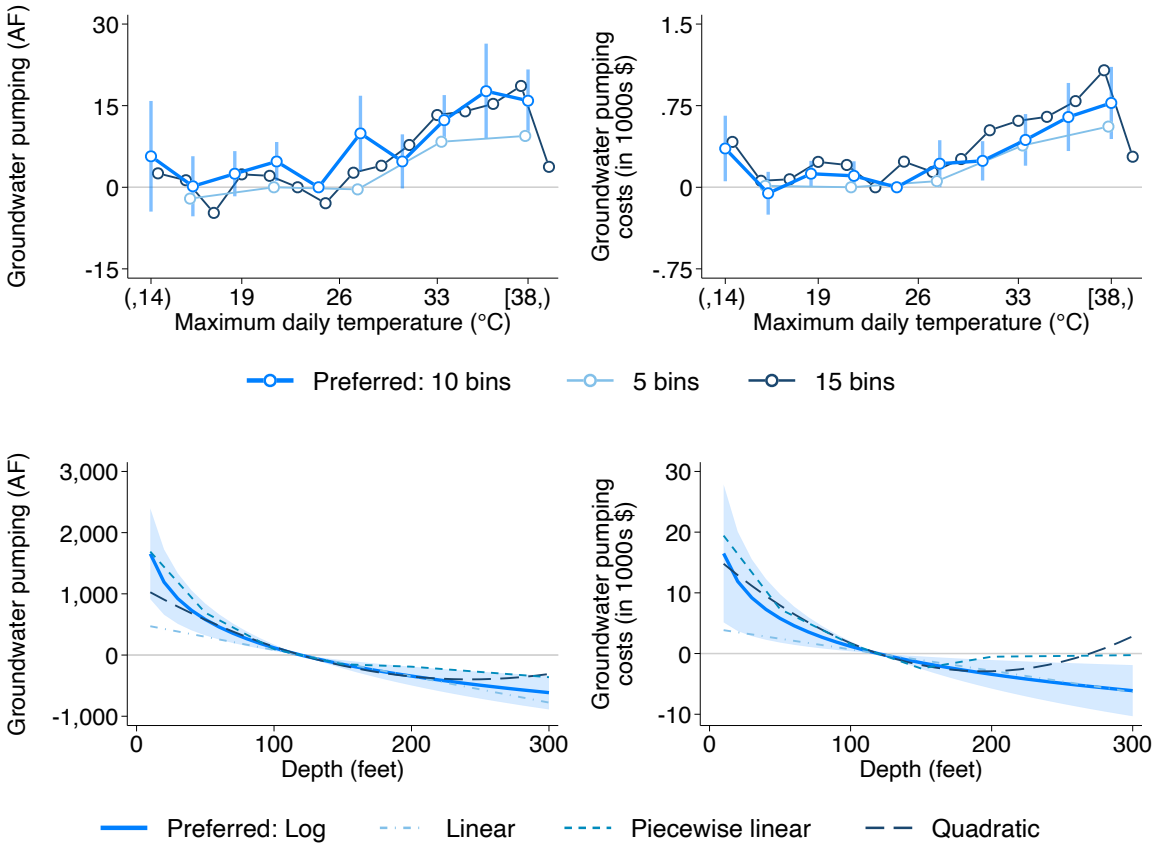


Notes: This figure demonstrates how our main results are not driven by time-varying unobservable trends. We plot our pooled estimates (from Figure 2) in middle blue. In dark blue, we plot estimates for the same specification after adding sub-basin-specific linear trends. The top panels plot temperature estimates; the bottom panels plot depth estimates (relative to the mean depth of 119.5 feet). The left panels plot effects on annual groundwater extraction; the right panels plot effects on annual groundwater expenditures. We estimate Equation (7) with 10 temperature bins (omitting the 23–26°C bin), log depth, month-specific linear precipitation controls, parcel fixed effects, and year fixed effects. Regressions are weighted by croppable acres. Standard errors are clustered by water district. We plot 95% confidence intervals for our preferred estimates with vertical lines (top panels) and shading (bottom panels).

the removal of identifying variation (since after including trends, $\hat{\beta}$ is estimated only off of year-on-year deviations from these trends). Importantly, the attenuated estimate retains strong statistical significance and remains internally consistent with the groundwater cost elasticity estimates in Burlig, Preonas, and Woerman (2025).

Robustness to functional form Appendix Figure B3 shows that our results are robust to alternate numbers of temperature bins, and to alternate parameterizations of depth.

Figure B3: Robustness – functional form of temperature and depth

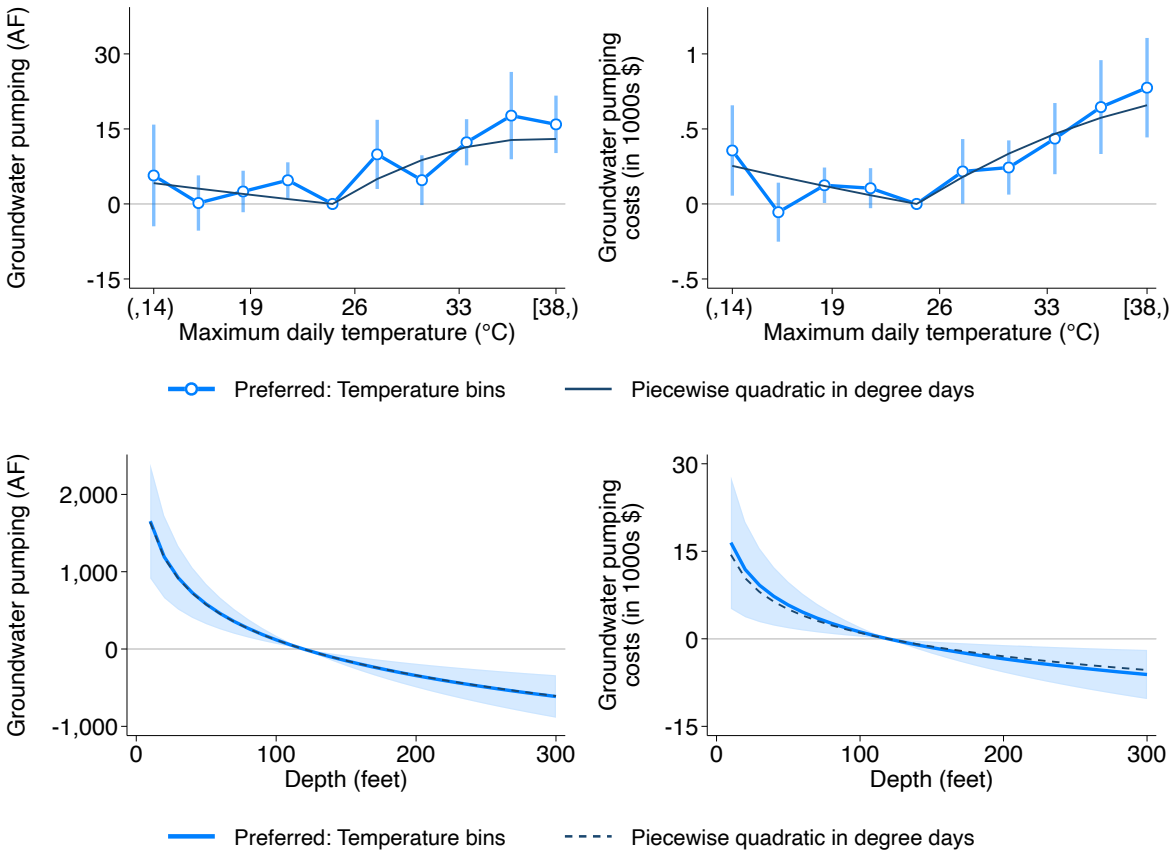


Notes: This figure shows that our results are robust to alternate function forms on temperature and depth. The top panels plot our preferred temperature estimates (from Figure 2) using 10 bins (in middle blue). We also plot the estimated temperature response for two alternate specifications: 5 bins (in light blue) and 15 bins (in dark blue). These six regressions all use our preferred log depth functional form. The bottom panels plot our preferred depth estimates (from Figure 2) using log depth (solid line, plotted relative to the mean depth of 119.5 feet). We also plot the estimated depth response from three alternate specifications: linear depth (light dash), piecewise linear depth (medium dash), and quadratic depth (dark dash). These eight regressions all use our preferred 10-temperature-bin functional form. The left panels plot effects on annual groundwater extraction; the right panels plot effects on annual groundwater expenditures. All regressions include month-specific linear precipitation controls, parcel fixed effects, and year fixed effects. Regressions are weighted by croppable acres. Standard errors are clustered by water district. We plot 95% confidence intervals for our preferred estimates with vertical lines (top panels) and shading (bottom panels).

Appendix Figure B4 shows that the temperature specification that we use for our forward simulation model produces nearly identical results. This more parsimonious specifi-

cation replaces the 10 temperature bins with a piecewise quadratic in annual degree days, estimating four coefficients on (i) the sum of annual degree days above 24.5°C; (ii) the sum of annual degree days squared above 24.5°C; (iii) the sum of annual degree days below 24.5°C; and (iv) the sum of annual degree days squared below 24.5°C.

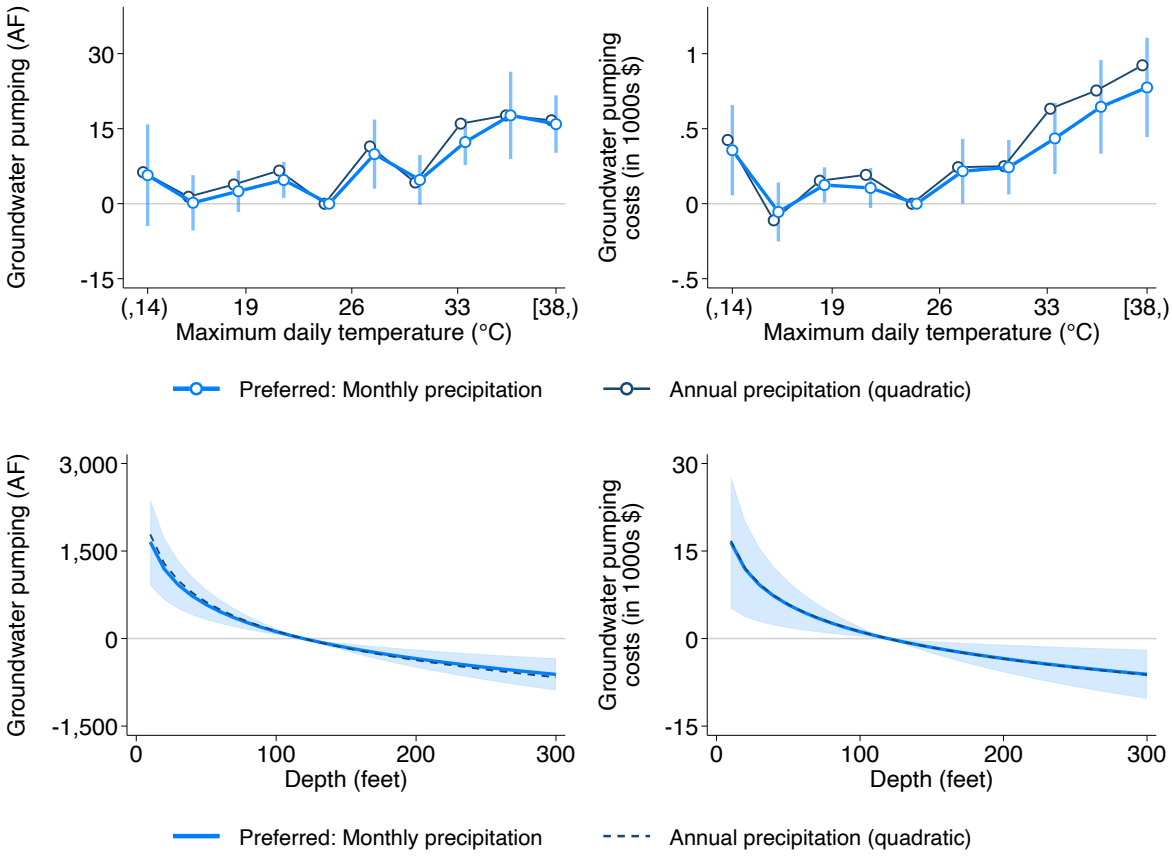
Figure B4: Robustness – parsimonious degree-day specification



Notes: This figure shows that our results are robust to the more parsimonious functional form for temperature that we use in our forward simulation model. This parsimonious model is a piecewise quadratic in annual degree days. We estimate four coefficients on the annual sums of degree days and degree days squared, separately above and below the reference temperature of 24.5°C. The top panels plot these degree day estimates in dark blue, and our preferred temperature estimates (from Figure 2) in middle blue. The bottom panels plot the coefficients on log depth for the preferred model (solid line) and for the degree day model (dashed line), both relative to the mean depth of 119.5 feet. The left panels plot effects on annual groundwater extraction; the right panels plot effects on annual groundwater expenditures. All regressions include log depth, month-specific linear precipitation controls, parcel fixed effects, and year fixed effects. Regressions are weighted by croppable acres. Standard errors are clustered by water district. We plot 95% confidence intervals for our preferred estimates with vertical lines (top panels) and shading (bottom panels).

Appendix Figure B5 shows that our temperature and depth response estimates are nearly identical when we replace our preferred precipitation controls (linear in each month's precipitation) with a quadratic in total annual millimeters of precipitation. This more parsimonious functional form for precipitation is necessary for our forward simulations, since the climate projections (from Hsiang et al. (2017); Rasmussen, Meinshausen, and Kopp (2016)) do not have a sub-annual temporal resolution.

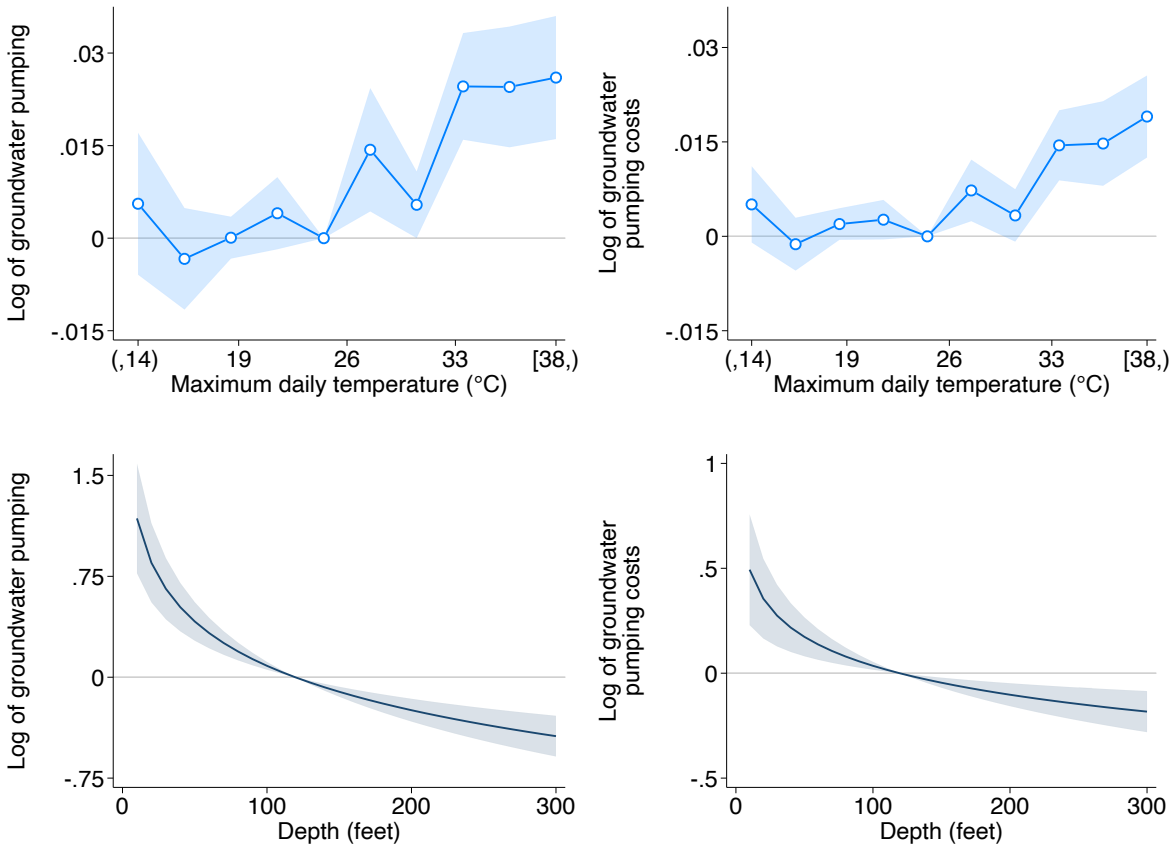
Figure B5: Robustness – functional form of precipitation controls



Notes: This figure shows that our results are robust to the more parsimonious functional form for precipitation that we use in our forward simulation model: quadratic in total annual mm of precipitation. The top panels plot temperature estimates for our preferred model with month-specific linear precipitation controls (from Figure 2; in middle blue) and for the parsimonious model controlling for a quadratic in annual precipitation (in dark blue). The bottom panels plot the coefficients on log depth for the preferred model (solid line) and for the model with a quadratic in annual precipitation (dashed line), both relative to the mean depth of 119.5 feet. The left panels plot effects on annual groundwater extraction; the right panels plot effects on annual groundwater expenditures. All regressions include 10 temperature bins (omitting the 23–26°C bin), log depth, parcel fixed effects, and year fixed effects. Regressions are weighted by croppable acres. Standard errors are clustered by water district. We plot 95% confidence intervals for our preferred estimates with vertical lines (top panels) and shading (bottom panels).

Finally, Appendix Figure B6 shows that our preferred temperature and depth response estimates (estimated with outcome variables in levels) are robust to applying the natural log transformation to annual groundwater extraction and annual groundwater expenditures.

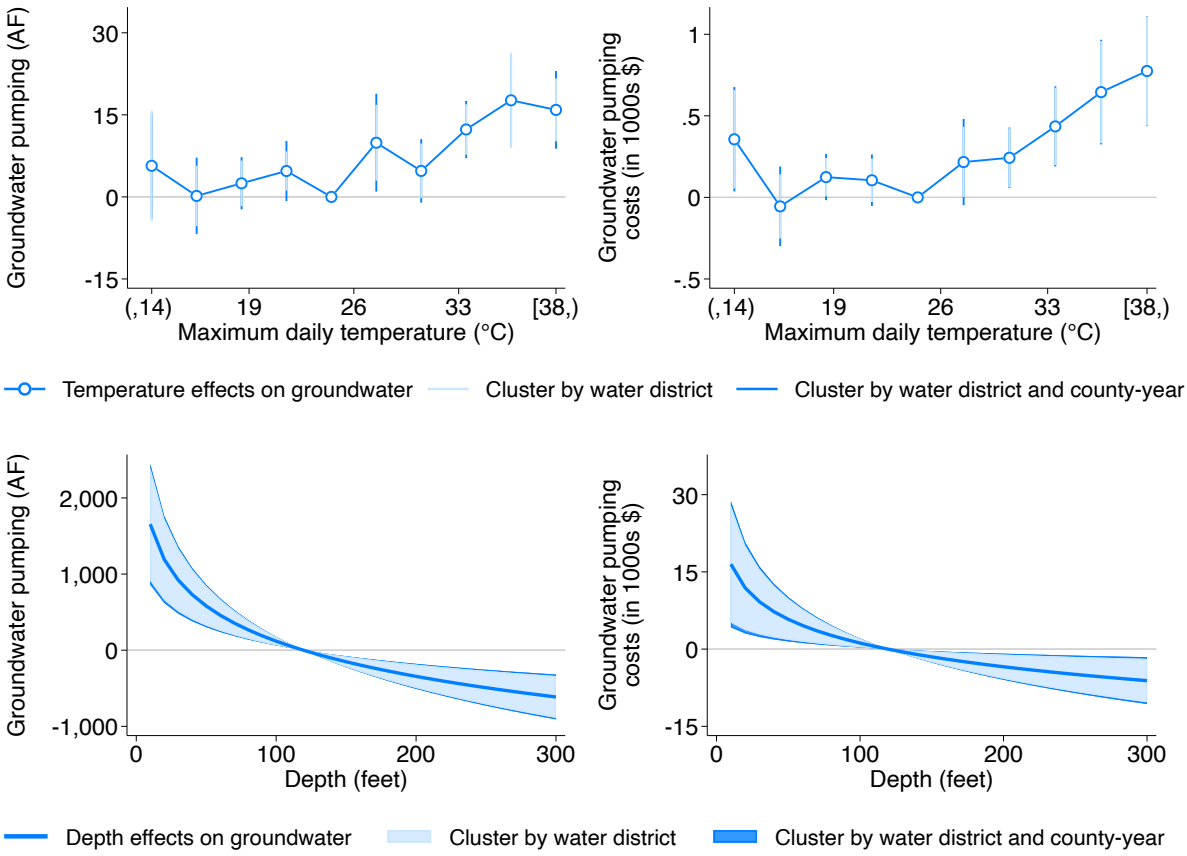
Figure B6: Robustness – log-transformed outcome variables



Notes: This figure presents alternate versions of our main results from Figure 2 where we apply the natural log transformation to our outcome variables (rather than using levels). The top panels plot temperature estimates; the bottom panels plot depth estimates, relative to the mean depth of 119.5 feet. The left panels plot effects on annual groundwater extraction; the right panels plot effects on annual groundwater expenditures. Regressions include 10 temperature bins (omitting the 23–26°C bin), log depth, month-specific linear precipitation controls, parcel fixed effects, and year fixed effects. Regressions are weighted by croppable acres. Standard errors are clustered by water district. Shading denotes 95% confidence intervals.

Robustness to alternate standard errors Appendix Figure B7 demonstrates that our confidence intervals are only slightly changed when we apply more conservative two-way clustering by water district and by county-year. (Our preferred estimates one-way cluster by water district.)

Figure B7: Robustness – alternative standard errors

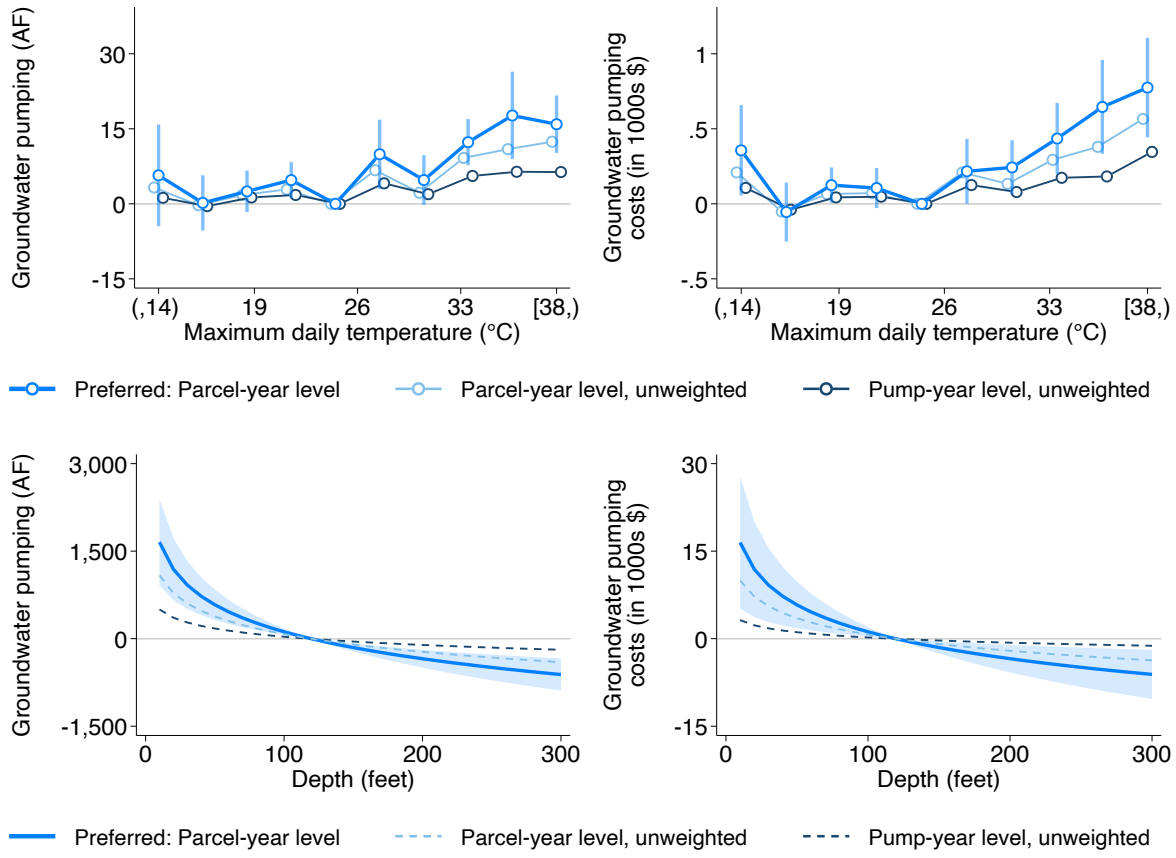


Notes: This figure demonstrates that our results are robust to more conservative standard errors: two-way clustering by water district and by county-year. The top panels plot our preferred temperature estimates (from Figure 2). Light vertical lines denote 95% confidence intervals under our preferred clustering by water district; darker extensions to those vertical lines denote 95% confidence intervals under the more conservative two-way clustering. The bottom panels plot our preferred depth estimates (from Figure 2), relative to the mean depth of 119.5 feet. Light shading denotes 95% confidence intervals under our preferred clustering by water district; darker shading denotes 95% confidence intervals under the more conservative two-way clustering. The left panels plot effects on annual groundwater extraction; the right panels plot effects on annual groundwater expenditures. Regressions include 10 temperature bins (omitting the 23–26°C bin), log depth, month-specific linear precipitation controls, parcel fixed effects, and year fixed effects. Regressions are weighted by croppable acres.

Aggregation and weights In our preferred specification, our cross-sectional unit is the parcel (i.e., the farm). We sum all groundwater extraction across all observed pumps with each parcel's footprint. We also weight our regressions by croppable acres to make our estimates representative per acre of cropland (which avoids overweighting small parcels).

Appendix Figure B8 reproduces our main results (middle blue) and then presents two alternatives: (i) a version estimated at the parcel level, but unweighted (light blue), and (ii0 a version estimated at the PG&E service point (i.e., pump) level, also unweighted (dark blue). Removing the croppable acreage weights slightly attenuates both our temperature and depth estimates—suggesting that larger farms are more responsive than smaller farms (likely because they pump larger volumes of groundwater). When we take the additional step of disaggregating parcels to estimate Equation (7) at the pump level, our temperature and depth estimates attenuate further by approximately half—consistent with the fact that the average parcel in our sample contains 1.9 pumps.

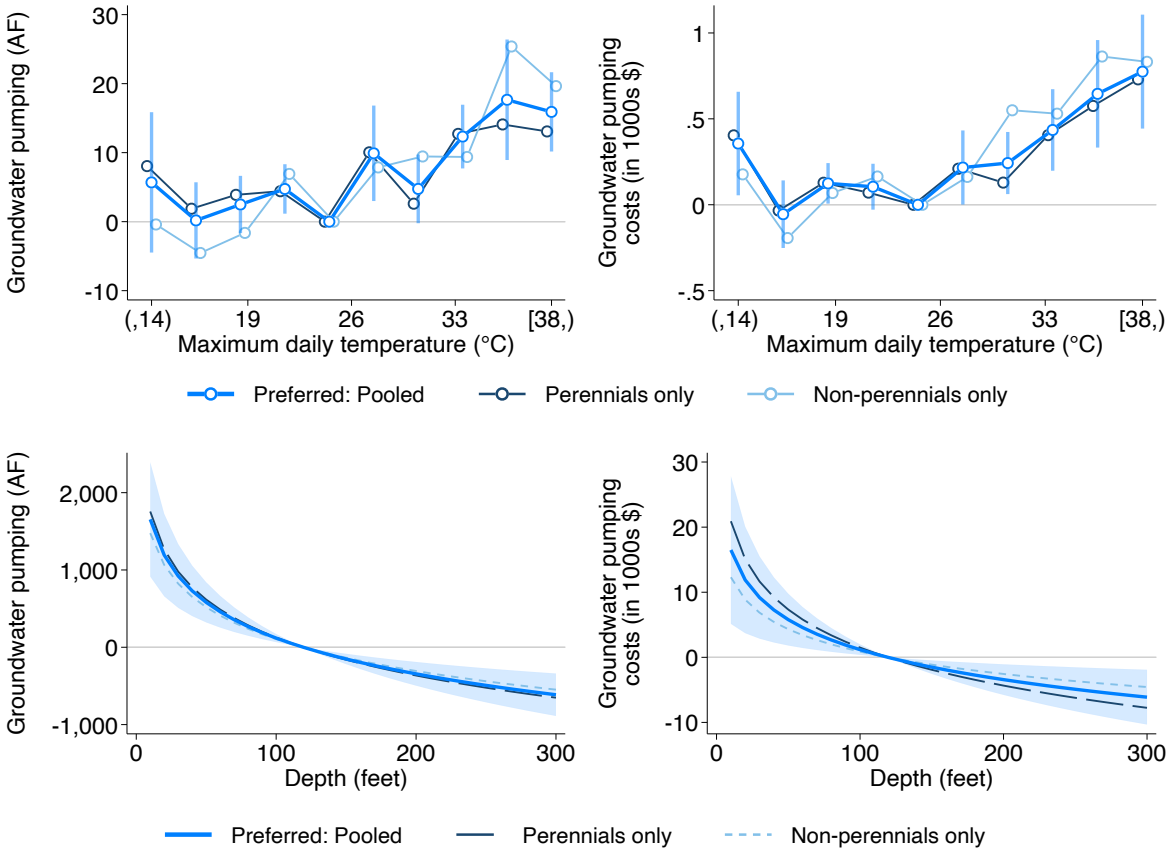
Figure B8: Robustness – regression weights and level of aggregation



Notes: This figure presents our main results from Figure 2, which are estimated at the parcel-year level, weighting parcels by croppable acreage (in middle blue). We also present two alternate versions: (i) remove the acreage weights (in light blue); and (ii) removing the acreage weights and disaggregating to the pump-year level (in dark blue). The top panels plot temperature estimates; the bottom panels plot depth estimates, relative to the mean depth of 119.5 feet. The left panels plot effects on annual groundwater extraction; the right panels plot effects on annual groundwater expenditures. All regressions include 10 temperature bins (omitting the 23–26°C bin), log depth, month-specific linear precipitation controls, unit (i.e., parcel or pump) fixed effects, and year fixed effects. Standard errors are clustered by water district. We plot 95% confidence intervals for our preferred estimates with vertical lines (top panels) and shading (bottom panels).

Accounting for crop choice Burlig, Preonas, and Woerman (2025) underscore the importance of accounting for perennial crops in this setting. We might expect perennial farmers to respond differently than non-perennial farmers to temperature and depth shocks—if, say, adaptive irrigation can mitigate damages from extreme heat that could impact perennial harvests in future years. Appendix Figure B9 splits the sample into parcels where the modal landcover is perennials (including fruit/nut and hay perennials) vs. non-perennials (including annuals and fallowing)—following the crop categories described in Burlig, Preonas, and Woerman (2025). This reveals that while perennials appear to have slightly weaker temperature response and slightly stronger depth response than non-perennials, none of the point estimates for perennials vs. non-perennials are statistically distinguishable from each other.

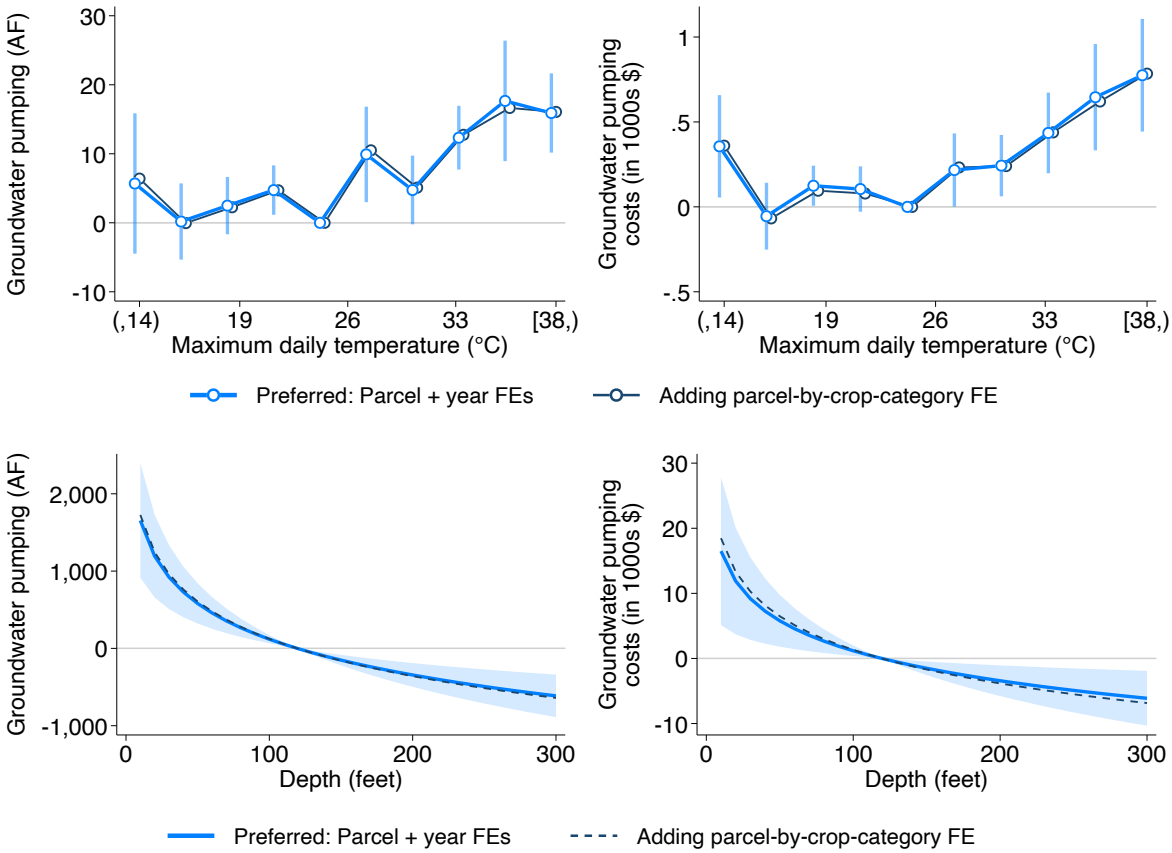
Figure B9: Robustness – splitting by perennials vs. non-perennials



Notes: This figure shows how our main results vary for perennial vs. non-perennial crops. We plot our pooled estimates (from Figure 2) in middle blue. In dark blue, we plot estimates for parcel-years where the modal pixel of cropland is perennial. In light blue, we plot estimates for parcel-years where the modal pixel of cropland is non-perennial (including fallowing). The top panels plot temperature estimates; the bottom panels plot depth estimates, relative to the mean depth of 119.5 feet. The left panels plot effects on annual groundwater extraction; the right panels plot effects on annual groundwater expenditures. We estimate Equation (7) with 10 temperature bins (omitting the 23–26°C bin), log depth, month-specific linear precipitation controls, parcel fixed effects, and year fixed effects. Regressions are weighted by croppable acres. Standard errors are clustered by water district. We plot 95% confidence intervals for our preferred estimates with vertical lines (top panels) and shading (bottom panels).

We can also control for crop choice by interacting parcel fixed effects with the modal crop category (annuals, fruit/nut perennials, hay perennials, and noncrop) for each parcel-year. Appendix Figure B10 compares our preferred estimates from Figure 2 to estimates from a version of Equation (7) that adds parcel-by-crop-category fixed effects. This has virtually no impact on our temperature or depth effect estimates.

Figure B10: Robustness – interacting parcel fixed effects with crop categoreis



Notes: This figure demonstrates how our main results are nearly identical when we control for farmers' crop choices. We plot our pooled estimates (from Figure 2) in middle blue. In dark blue, we plot estimates for the same specification after adding parcel-by-crop-category fixed effects (using four categories: annuals, fruit/nut perennials, hay perennials, noncrop/fallowing). The top panels plot temperature estimates; the bottom panels plot depth estimates, relative to the mean depth of 119.5 feet. The left panels plot effects on annual groundwater extraction; the right panels plot effects on annual groundwater expenditures. We estimate Equation (7) with 10 temperature bins (omitting the 23–26°C bin), log depth, month-specific linear precipitation controls, and year fixed effects. Regressions are weighted by croppable acres. Standard errors are clustered by water district. We plot 95% confidence intervals for our preferred estimates with vertical lines (top panels) and shading (bottom panels).

Main results in table format Finally, we report our main results from Figure 2 numerically in Appendix Table B1.

Table B1: Groundwater use responds to variation in temperature and groundwater depth

	(1) Annual groundwater extraction (AF)	(2) Annual groundwater expenditures ($)
Temperature bin 1: $\geq 14°C$	5.693	356.290**
	(5.146)	(152.439)
Temperature bin 2: $(14°C, 17°C]$	−54.969	−54.969
	(99.518)	(99.518)
Temp bin 3	2.489	124.397**
	(2.100)	(59.738)
Temp bin 4	4.738***	104.947
	(1.814)	(67.577)
Temp bin 6	9.907***	216.768**
	(3.500)	(109.093)
Temp bin 7	4.751*	242.780***
	(2.519)	(91.448)
Temp bin 8	12.328***	435.306***
	(2.336)	(120.034)
Temp bin 9	17.664***	645.622***
	(4.419)	(158.330)
Temp bin 10	15.910***	774.870***
	(2.906)	(167.798)
log(Depth)	−667.104***	−6646.761***
	(151.231)	(2317.242)

Notes: This table reports our main regression results from Figure 2 in numerical form. We estimate Equation (7) at the parcel-year level using two outcomes variables: groundwater consumption (in AF, column (1)), and groundwater expenditures (in $, column (2)). We report coefficients on temperature bins (plotted in the top panels of Figure 2), and on the natural log of average January depth in the groundwater sub-basin (plotted in the bottom panels of Figure 2). Regressions also include month-specific linear precipitation controls, parcel fixed effects, and year fixed effects. Regressions are weighted by each parcel's croppable acres. Standard errors (in parentheses) are clustered by water district. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

C Theoretical model

We develop a stylized model of agricultural production, groundwater use, and groundwater stocks, which we present briefly in Section 3. We provide the full model here.

C.1 Setup

Farm profits There are N identical farmers indexed by i. Farmer i's profits in year t are:

$$\pi(w_t; D_t, T) \;=\; p\,Q(w_t; T) \;-\; c(D_t)\,w_t \tag{C1}$$

The single input w_t is water used for irrigation in year t by farmer i (we suppress the i subscript, since $w_{it} = w_t$ for all N identical farmers). D_t is groundwater depth at the beginning of year t's growing season. Marginal groundwater pumping costs $c(D_t)$ increase linearly in depth D_t, consistent with the physical relationship in Equation (6):

$$\frac{\partial c(D_t)}{\partial D_t} > 0 \quad , \quad \frac{\partial^2 c(D_t)}{\partial D_t^2} = 0$$

T is a stylized temperature parameter, which we use to model a shift in the temperature distribution caused by climate change; this parameter abstracts from uncertainty of temperature realizations, nonstationarity of the temperature distribution, and the timing of climate change. $Q(\cdot)$ is the crop production function, which is increasing and concave in water input w_t and decreasing in temperature T, and has a positive cross-partial (such that the marginal product of water increases with temperature):

$$\frac{\partial Q(w_t; T)}{\partial w_t} > 0 \quad , \quad \frac{\partial^2 Q(w_t; T)}{\partial w_t^2} < 0 \quad , \quad \frac{\partial Q(w_t; T)}{\partial T} < 0 \quad , \quad \frac{\partial^2 Q(w_t; T)}{\partial w_t\, \partial T} > 0$$

Finally, the crop output price p is fixed and exogenous.

We assume that all farmers have an identical profit function in order to simplify the presentation of this stylized model. Importantly, our empirical application relaxes this homogeneity assumption to simulate changes to farmer welfare and quantify the stock externality. This stylized model also abstracts from farmers' crop choices (the focus of Burlig, Preonas, and Woerman (2025)) for simplicity, which appears reasonable in this setting because we find no evidence that farmers' temperature and depth responses vary with crop choice (see Appendix Figure B9).

Aquifer dynamics We model a stylized equation of motion to capture how aggregate groundwater pumping in year t (i.e., Nw_t) impacts depth in year $t+1$ (i.e., D_{t+1}):

$$D_{t+1} \;=\; D_t \;+\; \Phi N w_t$$

The constant $\Phi > 0$ subsumes specific yield of the soil and the effect of evapotranspiration (i.e., ρ and $(1-\psi)$ in Equation (9), respectively), and we assume it is homogeneous across the full aquifer. It also ignores recharge through rainfall (i.e., $-\rho\overline{R}_t$ in Equation (9)), which is not expected to change markedly in our geographic setting due to climate change (see Appendix Figure A2). We make these simplifications only for ease of presentation below.

C.2 Social Planner's problem

The Planner dynamically optimizes groundwater extraction by choosing $\{w_t\}$ to maximize total discounted profits of all farmers over an infinite horizon, subject to the equation of motion:

$$\max_{\{w_t\}} \sum_{i=0}^{N} \sum_{t=0}^{\infty} \beta^t \, \pi_{it} \;=\; N \sum_{t=0}^{\infty} \beta^t \Big[p\, Q(w_t; T) \;-\; c(D_t)\, w_t \Big]$$
$$\text{s.t.} \quad D_{t+1} \;=\; D_t \;+\; \Phi w_t N$$

We specify the following discrete-time Hamiltonian:

$$\mathcal{H}^P \;=\; N\Big[p\, Q(w_t; T) \;-\; c(D_t)\, w_t \Big] \;+\; \beta \lambda^P_{t+1} \Big[D_t \;+\; \Phi N w_t \Big]$$

where λ^P_{t+1} is the shadow value of depth. Since greater groundwater depth increases pumping costs and thereby reduces profits, $\lambda^P_t < 0$ for all t.

Planner's FOC (control variable w_t):

$$N \left[p \frac{\partial Q(w_t; T)}{\partial w_t} \;-\; c(D_t) \;+\; \beta \lambda^P_{t+1} \Phi \right] \;=\; 0$$
$$\Rightarrow \quad p \frac{\partial Q(w_t; T)}{\partial w_t} \;=\; c(D_t) \;-\; \beta \lambda^P_{t+1} \Phi$$

Planner's FOC (state variable D_t):

$$-N w_t \frac{\partial c(D_t)}{\partial D_t} \;+\; \beta \lambda^P_{t+1} \;=\; \lambda^P_t$$

The first FOC shows how the Planner sets each farmer's marginal value of water equal to their marginal pumping cost *plus* the marginal cost imposed on the future value of the aquifer by increasing its depth. The second FOC shows how the shadow value evolves over time, incorporating the fact that period-t pumping will increase future pumping costs.

We can combine these FOCs to yield:

$$p \frac{\partial Q(w_t; T)}{\partial w_t} \;=\; c(D_t) \;+\; N\Phi \sum_{s=t+1}^{\infty} \beta^{s-t} w_s \frac{\partial c(D_s)}{\partial D_s} \tag{C2}$$

This expression more clearly shows that the Planner sets each farmer's marginal value of water equal to their marginal pumping cost *plus* the marginal increase in discounted pumping costs imposed on all N farmers over all future time periods.

C.3 Farmer's problem

We assume that Farmer i is small relative to the aquifer, maximizes her own profits, and is forward-looking with respect to future groundwater pumping costs. It follows that Farmer i solves a dynamic problem similar to that of the Planner's, but only internalizes $1/N$ of the value of future groundwater stocks:

$$\begin{aligned} \max_{\{w_{it}\}} \sum_{t=0}^{\infty} \beta^t \, \pi_{it} \;&=\; \sum_{t=0}^{\infty} \beta^t \Big[p \, Q(w_{it}; T) \;-\; c(D_t) \, w_{it} \Big] \\ \text{s.t.} \quad D_{t+1} \;&=\; D_t \;+\; \Phi[w_{it} \;+\; (N-1) w_{-it}] \end{aligned}$$

Here, the equation of motion decomposes the change in depth due to Farmer i's extraction vs. due to extraction by all other farmers (i.e., $(N-1)w_{-it}$). This yields the following discrete-time Hamiltonian:

$$\mathcal{H}^F \;=\; p \, Q(w_{it}; T) \;-\; c(D_t) \, w_{it} \;+\; \beta \lambda^F_{t+1} [D_t \;+\; \Phi w_{it} \;+\; \Phi(N-1) w_{-it}]$$

Farmer's FOC (control variable w_{it}):

$$\begin{aligned} p \frac{\partial Q(w_{it}; T)}{\partial w_{it}} \;-\; c(D_t) \;+\; \beta \lambda^F_{t+1} \Phi \;&=\; 0 \\ \Rightarrow \quad p \frac{\partial Q(w_{it}; T)}{\partial w_{it}} \;&=\; c(D_t) \;-\; \beta \lambda^F_{t+1} \Phi \end{aligned}$$

Farmer's FOC (state variable D_t):

$$-w_{it} \frac{\partial c(D_t)}{\partial D_t} \;+\; \beta \lambda^F_{t+1} \;=\; \lambda^F_t$$

Combining these FOCs yields:

$$p \frac{\partial Q(w_{it}; T)}{\partial w_{it}} \;=\; c(D_t) \;+\; \underbrace{\Phi \sum_{s=t+1}^{\infty} \beta^{s-t} w_{is} \frac{\partial c(D_s)}{\partial D_s}}_{\approx 0} \tag{C3}$$

This final term is $1/N$ the size of the analogous term in the Planner's FOC, since Farmer i only internalizes her *own* impact on future groundwater depths. This term becomes trivially

small with large N, such that we can rewrite Equation (C3) as:

$$p\frac{\partial Q(w_t;T)}{\partial w_t} \;=\; c(D_t) \tag{C4}$$

Farmer i sets the marginal value of water equal to her own (private) pumping cost. This is equivalent to static profit maximization ignoring future impacts on aquifer depths.

C.4 The stock externality

Comparing the Planner's FOC (Equation (C2)) vs. Farmer i's FOC (Equation (C4)), the Planner internalizes an additional term:

$$\tau_t \;\equiv\; N\Phi \sum_{s=t+1}^{\infty} \beta^{s-t} w_s \frac{\partial c(D_s)}{\partial D_s} \;>\; 0 \tag{C5}$$

This term τ_t is the stock externality, equivalent to the marginal effect of farmer i's pumping in period t on the present value of aggregate future profits.[2]

We can approximate the total welfare loss from open-access over-extraction as the product of τ_t, the marginal stock externality per unit of water, and the difference in aggregate extraction Nw_t between a no-policy (NP) vs. an optimal policy (OP) scenario. This present discounted sum is:

$$LOSS \;\approx\; N\sum_{t=0}^{\infty} \beta^t \tau_t \Big(w_t^{NP} - w_t^{OP} \Big) \tag{C6}$$

The approximation results from the fact that τ_t is a marginal value along the optimal path $\{w_t^{OP}\}$, whereas $(w_t^{NP} - w_t^{OP})$ is a non-marginal change in extraction levels. Comparing Equation (C2) vs. Equation (C4), it is clear that $w_t^{NP} > w_t^{OP} > 0$. Since $\tau_t > 0$, $LOSS > 0$.

C.5 Climate change

Now, we consider the impact of climate change by taking comparative statics with respect to T. For simplicity, we consider a one-time increase in temperatures (i.e., $dT > 0$) that occurs in year $t = 0$. This shock is permanent and (for the purposes of this stylized model) unanticipated.

No-policy scenario (Farmer's i's comparative static) We can assess the impacts of climate change in a no-policy (NP) scenario by assuming all farmers privately optimize groundwater pumping. Totally differentiating Farmer i's FOC (Equation (C4)) with respect

2. Note that in this expression, w_s is water usage along the Planner's optimal path. Further, if we assume that Farmer i internalizes her own (trivially small) share of future lost profits (as in Equation (C3)), then the stock externally is technically $\tau_t \;\equiv\; (N-1)\Phi \sum_{s=t+1}^{\infty} \beta^{s-t} w_s \frac{\partial c(D_s)}{\partial D_s}$.

to T:

$$\frac{dw_t^{NP}}{dT} \;=\; p\left[\frac{\partial^2 Q(w_t;T)}{\partial w_t^2}\frac{dw_t}{dT} \;+\; \frac{\partial^2 Q(w_t;T)}{\partial w_t dT}\right] \;=\; \frac{\partial c(D_t)}{\partial D_t} N\Phi \sum_{s=0}^{t-1}\frac{dw_s}{dT}$$

The sign of this expression is ambiguous: while hotter temperatures raise the marginal product of water, increasing farmers' preferred amount of pumping, higher past pumping raises groundwater depths, raising the cost of pumping, and reducing farmers' preferred amount of pumping. Which dominates is an empirical question.

The stock externality How will climate change alter the stock externality? We can totally differentiate Equation (C5) with respect to T:

$$\begin{aligned}\frac{d\tau_t}{dT} \;&=\; N\Phi \sum_{s=t+1}^{\infty}\beta^{s-t}\left[\frac{dw_s}{dT}\frac{\partial c(D_s)}{\partial D_s} \;+\; w_s \underbrace{\frac{\partial^2 c(D_s)}{\partial D_s^2}}_{=0} N\Phi \sum_{r=t}^{s-1}\frac{dw_r}{dT}\right] \\ &=\; N\Phi \sum_{s=t+1}^{\infty}\beta^{s-t}\frac{dw_s}{dT}\frac{\partial c(D_s)}{\partial D_s} \qquad \text{(C7)}\end{aligned}$$

This expression is similarly ambiguous, because it depends on the net effect of climate change on farmers' pumping behavior.

Optimal policy scenario (Planner's comparative static) We can assess the impacts of climate change in an optimal policy (OP) scenario using the Planner's FOC (Equation (C2)). Totally differentiating with respect to T:

$$\begin{aligned} p\left[\frac{\partial^2 Q(w_t;T)}{\partial w_t^2}\frac{dw_t}{dT} \;+\; \frac{\partial^2 Q(w_t;T)}{\partial w_t dT}\right] \;&=\; \frac{\partial c(D_t)}{\partial D_t} N\Phi \sum_{s=0}^{t-1}\frac{dw_s}{dT} \\ &\quad + \underbrace{N\Phi \sum_{s=t+1}^{\infty}\beta^{s-t}\left[\frac{dw_s}{dT}\frac{\partial c(D_s)}{\partial D_s} \;+\; w_s \underbrace{\frac{\partial^2 c(D_s)}{\partial D_s^2}}_{=0} N\Phi \sum_{r=t}^{s-1}\frac{dw_r}{dT}\right]}_{=\, d\tau_t/dT} \\ \Rightarrow \quad \frac{dw_t^{OP}}{dT} \;&=\; \frac{p\dfrac{\partial^2 Q(w_t;T)}{\partial w_t dT} \;-\; \dfrac{\partial c(D_t)}{\partial D_t} N\Phi \displaystyle\sum_{s=0}^{t-1}\frac{dw_s}{dT} \;-\; N\Phi \sum_{s=t+1}^{\infty}\beta^{s-t}\frac{dw_s}{dT}\frac{\partial c(D_s)}{\partial D_s}}{-\dfrac{\partial^2 Q}{\partial w_t^2}} \qquad \text{(C8)}\end{aligned}$$

This expression mirrors Farmer i's comparative static except for the addition of the third term in the numerator—which is the negative of the change in the stock externality, $-\frac{d\tau_t}{dT}$. Thus, the effect of climate change on the optimal policy scenario depends on whether

the stock externality shrinks or grows with climate change, since the Planner will account for the greater or smaller future external costs of current-year pumping by reducing or increasing the optimal w_t^* accordingly.

How does climate change impact the welfare losses due to over-extraction? Since optimal groundwater management remains elusive, we can use this stylized model to show how climate change affects welfare losses due to over-extraction. To do this, we can totally differentiate Equation (C6) with respect to T:

$$\frac{dLOSS}{dT} \approx N \sum_{t=0}^{T} \beta^t \left[\frac{d\tau_t}{dT} \left(w_t^{NP} - w_t^{OP} \right) + \tau_t \left(\frac{dw_t^{NP}}{dT} - \frac{dw_t^{OP}}{dT} \right) \right] \tag{C9}$$

The sign of this expression (which is an approximation due to the nonmarginal change between NP and OP), depends both on the effect of climate change on the stock externality, $\frac{d\tau_t}{dT} > 0$, and on the corresponding change in policy, $\frac{dw_t^{NP}}{dT}$ vs. $\frac{dw_t^{OP}}{dT}$. Hence, climate change may either increase or decrease the premium to society from optimal groundwater management policy.

C.6 Willingness to pay

We can differentiate the farmer's profit function (Equation (C1)) to arrive at expressions of the private willingness to pay to avoid increased heat and increased aquifer depth.

WTP for temperature Differentiating Equation (C1) with respect to T at the privately optimal w_t^*:

$$\frac{d\pi_t^*}{dT} = p \left[\frac{\partial Q}{\partial T} + \frac{\partial Q}{\partial w_t} \frac{dw_t^*}{dT} \right] - c(D_t) \frac{dw_t^*}{dT} \tag{C10}$$

To sign this expression, note that $\frac{\partial Q}{\partial T} < 0$ (i.e., yields fall as temperatures rise, per a large literature including our own Figure 1); and that $\frac{\partial Q}{\partial w_t} \frac{dw_t^*}{dT} > 0$, as increases in water improve crop yields, and farmers increase water use as temperatures rise (see Figure 2). Though the overall sign of this term is technically ambiguous, Figure 1 suggests that the overall effect of temperature on crop yields is negative, even accounting for the benefits of adaptive irrigation. Then, note that $c(D_t) \frac{dw_t^*}{dT} > 0$, as marginal pumping costs are positive and, again, farmers apply more water as temperatures rise. Thus, both terms in Equation (C10) are negative.

This implies that our reduced-form estimates from the expenditure regressions can serve as a lower bound on farmers' willingness to pay to avoid an increase in temperature. The temperature coefficients from the top-right panel of Figure 2 are empirical analogs to the second term in Equation (C10). Since the first term has the same sign, and since we do not

see farmers substituting across other inputs (see null effects in Figure 3, which our single-input profit function implicitly assumes), farmers' willingness to pay to avoid a marginal hot day must be weakly larger than the corresponding temperature coefficient would suggest. This logic extends to the model that adds a temperature-depth interaction: the willingness-to-pay to avoid a marginal day in bin 10 is bounded below by the sum of the uninteracted bin-10 coefficient and the interacted bin-10 coefficient evaluated at a specific depth D_t.

To bound farmers' willingness to pay to avoid a non-marginal change in the *climate* (e.g., comparing temperatures in 2050 to temperatures in 2019), we can extrapolate the above approach by multiplying the linear coefficients from the expenditure regression by the change in temperature. For example, $\widehat{f()} \cdot (\mathbf{T}_{i,2050} - \mathbf{T}_{i,2019})$.

WTP for depth Differentiating Equation (C1) with respect to D_t at the privately optimal w_t^*:

$$\begin{aligned}
\frac{d\pi^*}{dD_t} &= p\frac{\partial Q}{\partial w_t}\frac{\partial w_t^*}{\partial D_t} - \left[w_t^*\frac{\partial c(D_t)}{\partial D_t} + c(D_t)\frac{\partial w_t^*}{\partial D_t}\right] \\
&= \left[p\frac{\partial Q}{\partial w_t} - c(D_t)\right]\frac{\partial w_t^*}{\partial D_t} - w_t^*\frac{\partial c(D_t)}{\partial D_t} \\
&= -w_t^*\frac{\partial c(D_t)}{\partial D_t} \qquad \text{(C11)}
\end{aligned}$$

where the first term in the second line equals zero by the farmer's first-order condition (Equation (C4)). Thus, farmers' marginal willingness to pay to avoid an increase in depth is equal to the optimal quantity of groundwater pumped scaled by the (constant) rate at which greater depths increase pumping costs.

To evaluate farmers' willingness to pay to avoid a non-marginal change in depth (e.g., comparing factual depth D_t to counterfactual depth D_t'), we can integrate this expression:

$$\Delta\pi_t^* = \int_{D_t}^{D_t'} -w_t^*\frac{\partial c(D_t)}{D_t}\,dD$$

Note that the optimal w_t^* change is a function of the variable of integration D. We can use our fitted estimates from Equation (7) with water use as the outcome variable as an empirical analog for w_t^*, which accounts depth reoptimization. Substituting, suppressing the i subscript, and pulling out the constant $\frac{\partial c(D_t)}{\partial D_t}$:

$$\begin{aligned}
\Delta\pi_t^* &\approx -\frac{\partial c(D_t)}{D_t}\int_{D_t}^{D_t'}\hat{w}_t(D)\,dD \\
&= -\frac{\partial c(D_t)}{D_t}\int_{D_t}^{D_t'}\left[\hat{\beta}\log(D) + \hat{\Gamma}\right]dD
\end{aligned}$$

where $\hat{\Gamma}$ represents all components of the fitted model that are independent of depth.

$$\begin{aligned} &= -\frac{\partial c(D_t)}{D_t}\left[D\Big(\hat{\beta}\log(D)+\hat{\Gamma}\Big)-\hat{\beta}D\right]_{D_t}^{D_t'} \\ &= -\frac{\partial c(D_t)}{D_t}\left[D_t' w_t^*(D_t') - D_t w_t^*(D_t) - \hat{\beta}(D_t' - D_t)\right] \end{aligned} \tag{C12}$$

This expression simplifies to $-\frac{\partial c(D_t)}{D_t}\hat{\beta}[D_t'(\log(D_t')-1) - D_t(\log(D_t)-1)]$, since Equation (7) treats depth as linearly separable from other determinants of $\hat{w}_t$. For the model with a (nonlinear) temperature-depth interaction, the last term inside the brackets in Equation (C12) becomes $(\hat{\beta} + \sum_k \hat{\beta}^k T^k)(D_t' - D_t)$, where $\hat{\beta}^k$ is the kth interaction coefficient and T^k is the observed temperature that corresponds with that coefficient.

D Simulation model

D.1 Model setup

Fitted regression model To operationalize Equation (8), we modify our preferred temperature response specification (Equation (7)) to reduce the number of parameters increase flexibility. First, we replace our preferred $f(T_{it})$ of 10 temperature bins with a piecewise quadratic in annual degree days. We anchor our degree day sums at 24.5°C, the midpoint of the omitted temperature bin in Figure 2. We estimate four temperature parameters, as the OLS coefficients on: the annual sum of degree days conditional on being below 24.5°C, the annual sum of degree days conditional on being above 24.5°C, the annual sum of degree days squared conditional on being below 24.5°C, and the annual sum of degree days squared conditional on being above 24.5°C. Appendix Figure B4 shows that this modified functional form closely matches our preferred binned specification.

Second, we replace our preferred $g(R_{imt})$ of 12 monthly precipitation coefficients with a quadratic in annual precipitation. Besides the benefits of parsimony, this modification is necessary for parameterizing future precipitation, which the Hsiang et al. (2017) climate projections (described next) do not provide subannually. Appendix Figure B5 shows that this simplified functional form for temperature does not meaningfully alter our estimated temperature or depth coefficients.

Third, we estimate an interaction between log depth and the four degree day variables, as in Figure 4. This accommodates large out-of-sample shifts in either the temperature distribution (due to climate change) or depth (due to the dynamic propagation of the effects of adaptive irrigation over multiple years).

Together, these steps yield the fitted regression model $\widehat{w}\left(\mathbf{T}_{it}\,,\,R_{it}\,,\,D_{st}\,;\,\widehat{\boldsymbol{\beta}}\,,\,\widehat{\gamma}_i\,,\,\widehat{\delta}_t\right)$, with estimated coefficients $\widehat{\boldsymbol{\beta}}$ on temperature (piecewise quadratic degree day), precipitation (quadratic), and depth (log); estimated parcel fixed effects $\widehat{\gamma}_i$; and estimated year fixed effects $\widehat{\delta}_t$. We estimate these coefficients using the subsample of 4,788 parcels located within the San Joaquin Valley groundwater basin (the focus of our simulations) that were present in the sample in 2019. We extrapolate this sample forward for simplicity, abstracting from entry and exit of groundwater pumping (as in Burlig, Preonas, and Woerman (2025)), changes to pumping capital (which would be difficult to model), and crop-switching (supported by our finding in Appendix Figure B9 of statistically indistinguishable responses for perennials vs. non-perennials).

Climate projections To simulate forward, we populate temperature and precipitation using climate projections from Hsiang et al. (2017). These projections combine a suite of 45 climate models, yielding projected counts of days in 1°C temperature bins and in 2mm precipitation bins at a county-year resolution. We aggregate bins to construct projected degree days (linear and squared terms, separately above an below 24.5°C) and projected annual precipitation (in mm) at the county-year level. We model two climate scenarios:

RCP4.5, which is consistent with a stabilization of global fossil fuel emissions; and RCP8.5, which is consistent with intensive growth of global fossil fuel emissions.

Equation of motion To operationalize Equation (9), we restrict our focus to the San Joaquin Valley groundwater basin, which represents 71% of our estimation sample and is the most economically important groundwater basin in California. The San Joaquin basin contains 16 sub-basins with spatially heterogeneous climate realizations, irrigation practices, groundwater depths, and soil characteristics. To capture this inter-sub-basin heterogeneity close to the point of groundwater extraction and recharge, we populate the rightmost term of Equation (9) separately for each sub-basin s and average up to the basin level (weighting by sub-basin acreage A_s):

$$\overline{D}_{t+1} \;=\; \overline{D}_t \;+\; \frac{1}{\sum_s A_s} \sum_s A_s \rho_s \Big[(1-\psi)\alpha \widehat{\overline{W}}_{st} \;-\; \overline{R}_{st}^{\text{proj}} \Big] \tag{D1}$$

$$\text{where} \quad \widehat{\overline{W}}_{st} \;=\; \frac{\alpha_s}{A_s} \sum_{i \in s} \widehat{w} \left(T_{it}^{\text{proj}} \,,\, R_{it}^{\text{proj}} \,,\, \tilde{D}_{st} \,;\, \widehat{\boldsymbol{\beta}} \,,\, \widehat{\gamma}_i \,,\, \widehat{\delta}_{\tilde{t}} \right) \tag{D2}$$

$$\text{and} \quad \tilde{D}_{st} \;=\; \overline{D}_t \;+\; \gamma_s \tag{D3}$$

We assume $\psi = 0.2$, which is consistent with relatively efficient drip irrigation (Dewandel et al. (2008)). We parameterize ρ_s for each sub-basin by combining data on soil characteristics from World Soil Information Service[3] with specific yield estimates from Loheide, Butler, and Gorelick (2005).

We calibrate the α_s parameters by dividing the total number of agricultural/irrigation pumps in the sub-basin (as reported by the DWR's Well Completion Reports) by the number of pumps in our sample. This assumes that our estimation sample is representative of the full population of agricultural groundwater pumpers within each sub-basin. Scaling up from sample to population is essential for simulating the equation of motion, since observed groundwater depths are impacted by *total* groundwater extraction, not just groundwater extracted by pumps in our sample.

We also scale up total groundwater extraction by $\alpha = 1.38$, a calibration factor that ensures Equation (D1) holds as an accounting identity within our sample. This additional adjustment accounts for factors outside of our model that impact groundwater depths, including groundwater extracted for non-agricultural uses, errors in specifying α_s, and recharge due to precipitation outside of the basin (e.g., snowmelt).

We separately calibrate γ_s using time-invariant average sub-basin depths during our 2008–2019 sample period. These sub-basin-specific depth adders account for persistent differences in depth levels across sub-basins within the San Joaquin Valley basin. However, we simplify the hydrology (and the state space) by modeling a common equation of motion

3. https://isric.org/explore/wosis

across the full groundwater basin. This assumes full horizontal diffusion of groundwater across the basin before stepping forward to the next year.

Finally, our dynamic simulations impose a discrete time step equal to the calendar year. While groundwater extraction, recharge, and diffusion occur continuously year round in reality, discretizing by year aligns with the annual cropping cycle, our annual temperature and depth response estimates, and the format of the Hsiang et al. (2017) climate projections. As we discuss in Section 5.1, January serves as a natural break between discrete years, as it is both the start of the calendar year and the calendar month with the least groundwater pumping in our sample.

D.2 Simulation algorithm

Armed with the fitted regression model, balanced sample of 2019 pumpers in the San Joaquin Valley basin, and calibrated equation of motion (all presented above), we simulate according to the following steps:

1. Start with basin-wide average depth $\overline{D}_t$.

2. Adjust $\overline{D}_t$ by γ_s to arrive at $\tilde{D}_{st}$ (following Equation (D3)), which accounts for persistent differences in depth across San Joaquin's 16 sub-basins.

3. Estimate $\widehat{w}\left(\mathbf{T}_{it}^{\text{proj}}, R_{it}^{\text{proj}}, \tilde{D}_{st}\right)$ for each unit i in year t, by plugging climate projections ($\mathbf{T}_{it}^{\text{proj}}$, R_{it}^{proj}) and sub-basin-specific depths ($\tilde{D}_{st}$) into our fitted regression model, using estimated coefficients $\widehat{\boldsymbol{\beta}}$, estimated unit fixed effects $\widehat{\gamma}_i$, and estimated year fixed effect $\widehat{\delta}_{\tilde{t}}$. Since we are forward-simulating beyond our 2008–2019 sample period, we draw $\tilde{t} \in \{2008, \ldots, 2019\}$ for out-of-sample year t.[4]

4. Aggregate following Equation (D2). First, sum $\widehat{w}_{ist}$ (in AF) up to the sub-basin-year level. Second, scale up by α_s to extrapolate our estimation sample to the full population of agricultural groundwater pumpers in sub-basin s. Third, divide by sub-basin acreage A_s to arrive at $\widehat{\overline{W}}_{st}$, the average projected feet (or acre-feet per acre) of groundwater extracted in year t across sub-basin s.

5. Construct the final term of Equation (D1). First, calculate $\left[(1-\psi)\alpha\widehat{\overline{W}}_{st} - \overline{R}_{st}^{\text{proj}}\right]$, averaging projected precipitation across sub-basin s. This step yields the projected feet of net groundwater extraction across the full area of sub-basin s. Second, apply the multiplier ρ_s to translate groundwater extraction into changes in depth. Third, average changes in depth up to the basin level, weighting by sub-basin acreage A_s.

4. By drawing year fixed effects from within our sample period, we are effectively assuming stationarity. While this is a strong assumption, we prefer it to the alternative of extrapolating parametric time trends. Because we are using the model mainly to make comparisons across scenarios, this assumption is relatively innocuous.

6. Calculate $\overline{D}_{t+1}$, step forward one year, and repeat.

Uncertainty enters the simulations in three ways:

(i) Simulated parameters $\{\widehat{\boldsymbol{\beta}}^S, \widehat{\gamma}_i^S, \widehat{\delta}_t^S\}$ jointly drawn from the sampling distribution of the fitted regression model.

(ii) Simulated assigned years, drawing $\tilde{t} \in \{2008, \ldots, 2019\}$ with replacement 80 times to construct the vector $\vec{\tilde{t}}$. This vector assigns the (jointly drawn) year fixed effects $\widehat{\delta}_{\tilde{t}}^S$ for each year from 2020 to 2099.

(iii) Simulated climate projections, draw 1 of the 45 climate models from Hsiang et al. (2017) and weighting model draws probabilistically according to the sampling weights provided by Hsiang et al. (2017) (following Carleton et al. (2022) to represent uncertainty in climate projections within each RCP scenario). After drawing a climate model, we use that model's projections to construct $\{\mathbf{T}_{it}^{\text{proj},S}, R_{it}^{\text{proj},S}, \overline{R}_{st}^{\text{proj},S}\}$ for each year $t \in \{2020, \ldots, 2100\}$.

To capture this uncertainty, we perform 1,000 iterations of Steps 1–6. Each iteration proceeds as follows:

- Make joint draws (i)–(iii) to arrive at $\{\widehat{\boldsymbol{\beta}}^S, \widehat{\gamma}_i^S, \widehat{\delta}_t^S, \vec{\tilde{t}}, \mathbf{T}_{it}^{\text{proj},S}, R_{it}^{\text{proj},S}, \overline{R}_{st}^{\text{proj},S}\}$.
- Initiate the model using observed $\overline{D}_{2020}$.
- Looping forward over years $t = \{2020, \ldots, 2099\}$, execute Steps 1–6, plugging in $\{\widehat{\boldsymbol{\beta}}^S, \widehat{\gamma}_i^S, \widehat{\delta}_{\tilde{t}}^S, \mathbf{T}_{it}^{\text{proj},S}, R_{it}^{\text{proj},S}, \overline{R}_{st}^{\text{proj},S}\}$, and dynamically updating $\overline{D_t}$.

Our results report means over all 1,000 simulations. We also estimate three alternate scenarios for comparison with our preferred scenario:

- "Depth response only": identical to our preferred scenario, except that we replace $\{T_{it}^{\text{proj}}, R_{it}^{\text{proj}}\}$ in Step 3 with the within-sample averages $\{T_{i,\text{2008-19}}^{\text{proj}}, R_{i,\text{2008-19}}^{\text{proj}}\}$. This effectively holds the climate fixed and "turns off" the effect of climate change in increasing adaptive irrigation.[5]
- "Temperature response only": identical to our preferred scenario, except that we replace $\tilde{D}_{st}$ in Steps 2–3 with the initial values $\tilde{D}_{s,2020}$. This "turns off" the dynamic feedback effect of greater depths increasing pumping costs and reducing groundwater use, while still "adding up" the incremental groundwater depletion in Steps 4–6.

5. For internal consistency, we construct $\{T_{i,\text{2008-19}}^{\text{proj}}, R_{i,\text{2008-19}}^{\text{proj}}\}$ as the 2008–19 averages of the mean climate *projections*, rather than using climate realizations. This introduces trivial differences in $\{T_{it}^{\text{proj}}, R_{it}^{\text{proj}}\}$ in RCP4.5 vs. RCP8.5. Also, note that while we "turn off" changes in temperature and precipitation, we continue to simulate the equation of motion in Step 5 using $\overline{R}_{st}^{\text{proj}}$.

- "Neither response": identical to our preferred scenario, but combining the above two scenario to "turn off" any shocks to groundwater pumping demand. This cases is only useful as an internal calibration, and Figure E3 demonstrates that this case comes close to replicating the extrapolated linear trend in average depth for the San Joaquin Valley basin.

D.3 Quantifying welfare losses

We can use the derivations in Appendix C.6 to quantify the welfare consequences of climate change and greater groundwater depths in our preferred scenario.

Willingness to pay to avoid hot temperatures Following our discussion of Equation (C10), we can use our estimated coefficients from the groundwater expenditure regressions to infer a lower bound on farmers' willingness to pay to have avoided extreme heat. We estimate the identical regression to our fitted model $\widehat{w}(\,\cdot\,)$ for the simulations (described above), but using annual groundwater pumping costs as the outcome variable. This yields four coefficient estimates $\{\widehat{\beta}^{\$}_1, \widehat{\beta}^{\$}_2, \widehat{\beta}^{\$}_3, \widehat{\beta}^{\$}_4, \widehat{\beta}^{\$}_5, \widehat{\beta}^{\$}_6, \widehat{\beta}^{\$}_7, \widehat{\beta}^{\$}_8\}$ on the four uninteracted degree-day regressors and the four degree-day regressors interacted with depth, respectively. To calculate a lower bound on the welfare loss due to hotter temperatures, we calculate:

$$LOSS^T_t \;=\; \sum_s \alpha_s \sum_{i \in s} \left[\sum_{k=1}^{4} \left(T^{k,\text{proj}}_{it} - T^{k,\text{proj}}_{i,\text{2008-19}} \right) \widehat{\beta}^{\$}_k + \sum_{k=5}^{8} \left(T^{k,\text{proj}}_{it} - T^{k,\text{proj}}_{i,\text{2008-19}} \right) \tilde{D}_{st} \widehat{\beta}^{\$}_k \right] \quad \text{(D4)}$$

For each parcel i and each degree-day regressor (indexed by k), we calculate the difference between the projected value of that variable in year t and its 2008–19 baseline (i.e., "pre-climate-change") value. Scaling by $\widehat{\beta}^{\$}_k$, multiplying by the simulated depth (for the interacted regressors $k = \{5, 6, 7, 8\}$), and summing across all eight regressors yields a lower bound on the temperature-related welfare loss for parcel i in simulated year t. Summing across parcels in each sub-basin, scaling up using sub-basin-specific parameters α_s, and then summing across sub-basins yields a basin-wide aggregate lower bound on the welfare lost to due to hotter temperatures (relative to 2008–19 temperatures).

Willingness to pay to avoid depth increases Following our discussion of Equation (C12), we can use our estimated coefficients from our fitted model $\widehat{w}(\,\cdot\,)$ to quantify farmers' willingness to pay to have avoided increases in depth. The fitted model $\widehat{w}(\,\cdot\,)$ has five estimated depth coefficients: $\widehat{\beta}$ is the uninteracted coefficient on log depth, and $\{\widehat{\beta}_5, \widehat{\beta}_6, \widehat{\beta}_7, \widehat{\beta}_8\}$ are the coefficients on the interaction of log depth with the four degree-day temperature variables (as indexed above). Using these coefficients, along with the corresponding temperature projections and the simulated depth in the preferred scenario ($\tilde{D}_{st}$) and in the depth-response-only ($\tilde{D}^0_{st}$), we calculate:

$$LOSS_t^D = \sum_s \alpha_s \sum_{i \in s} \left[-\frac{\partial c(D_t)}{D_t} \right]_i \times \left[\tilde{D}_{st}\, \widehat{w}(\tilde{D}_{st})_{it} - \tilde{D}^0_{st}\, \widehat{w}(\tilde{D}^0_{st})_{it} - \left(\widehat{\beta} + \sum_{k=5}^{8} \widehat{\beta}_k T_{it}^{k,\text{proj}} \right) \left(\tilde{D}_{st} - \tilde{D}^0_{st} \right) \right] \quad \text{(D5)}$$

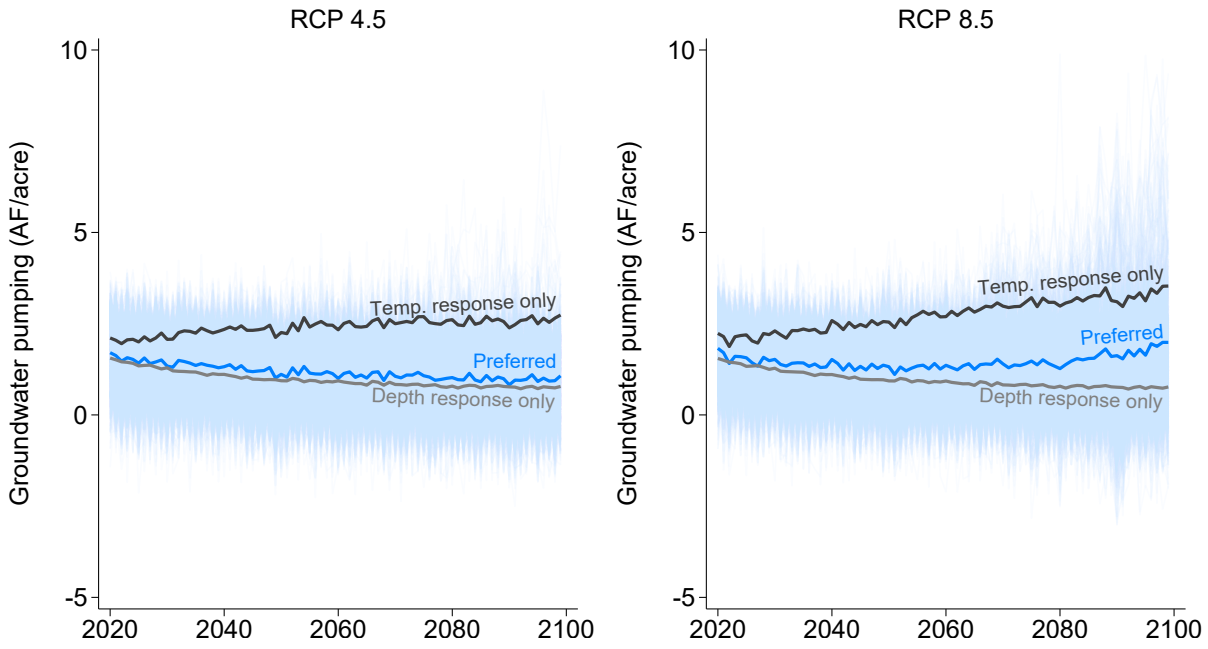
We can parameterize $\frac{\partial c(D_t)}{D_t}$ for each parcel by applying Equation (6) using their average 2019 electricity prices and their average 2019 operating pumping efficiency.[6] By comparing simulated depths across the preferred vs. depth-response-only scenarios, this expression evaluates the extent to with the cumulative effects of climate-induced adaptive irrigation reduce parcel i's welfare by increasing depth. We aggregate across parcels and sub-basins (as described above), to arrive at the basin-wide aggregate welfare lost due to greater depths resulting from hotter temperatures (relative to 2008–19 temperatures).

6. This holds end-of-sample electricity prices and pumping capital constant through 2100.

E Additional model results

Impacts on groundwater pumping Appendix Figure E1 presents simulation results showing groundwater extraction from 2020 through 2100 under both RCP4.5 (left panel) and RCP8.5 (right panel). As with our main simulations, we present results for three scenarios: our preferred model, in which farmers respond both to temperature and to groundwater depth; an alternative model where farmers respond only to depth; and an alternative where farmers respond only to temperature. We find that when farmers respond only to depth, groundwater pumping falls over time, as the slowly deepening aquifer leads costs to rise. When farmers respond only to temperature, pumping increases substantially, increasing by approximately 50% under RCP4.5 and nearly doubling by end of century under RCP8.5. In our preferred model, groundwater extraction stays relatively stable under RCP4.5 and increases somewhat under RCP8.5.

Figure E1: Impact of climate change on groundwater pumping

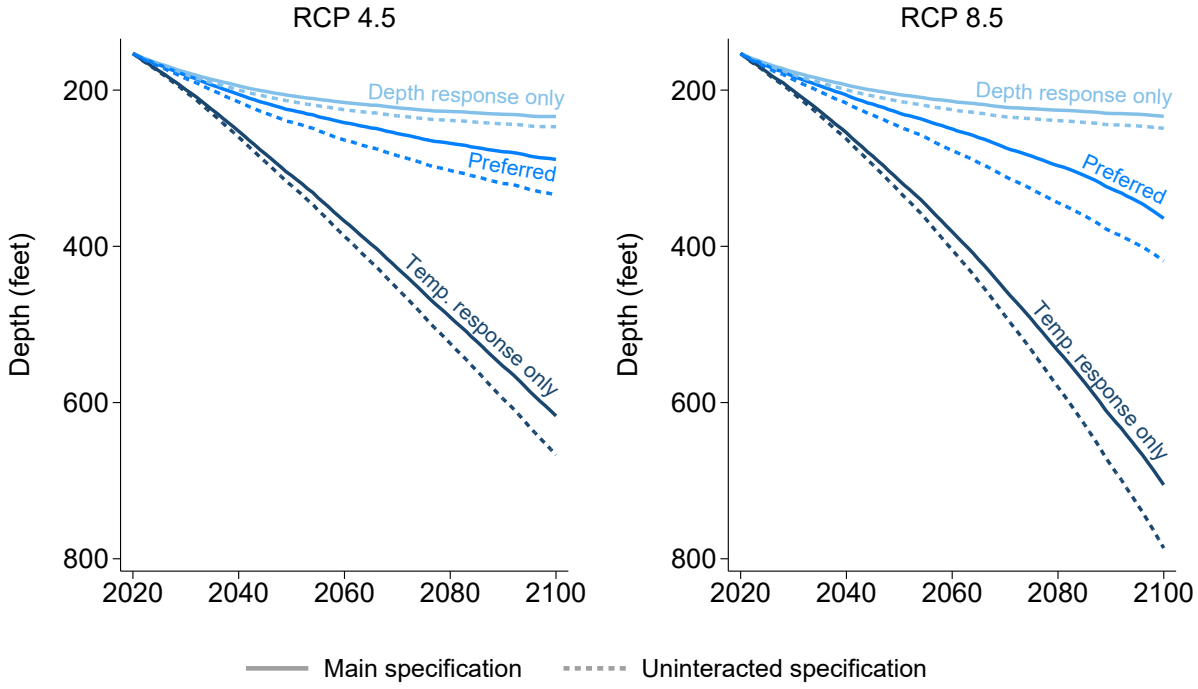


Notes: This figure presents additional simulation results not reported in Figure 5: predicted groundwater use, reported in average acre-feet per acre across the full basin. The left panel shows results under RCP4.5 (moderate climate change) and the right panel shows results under RCP8.5 (severe climate change). In both panels, the blue line denotes our preferred model, where farmers' groundwater extraction responds to both temperature and depth. The thick blue line plots the mean over the same 1,000 simulation draws reported in Figure 5. The individual light blue lines show individual simulations for the preferred scenario only. These draws incorporate both econometric and climate model uncertainty. The two gray lines plot the means over simulation runs where we only have farmers responding to depth (dash-dot) and only have farmers responding to temperature (dashed).

Uninteracted model with additively separable temperature depth We conduct a parallel set of simulations under the assumption that temperature and depth are additively separable, as in our main reduced-form specification (Equation (7). Appendix Figure E2

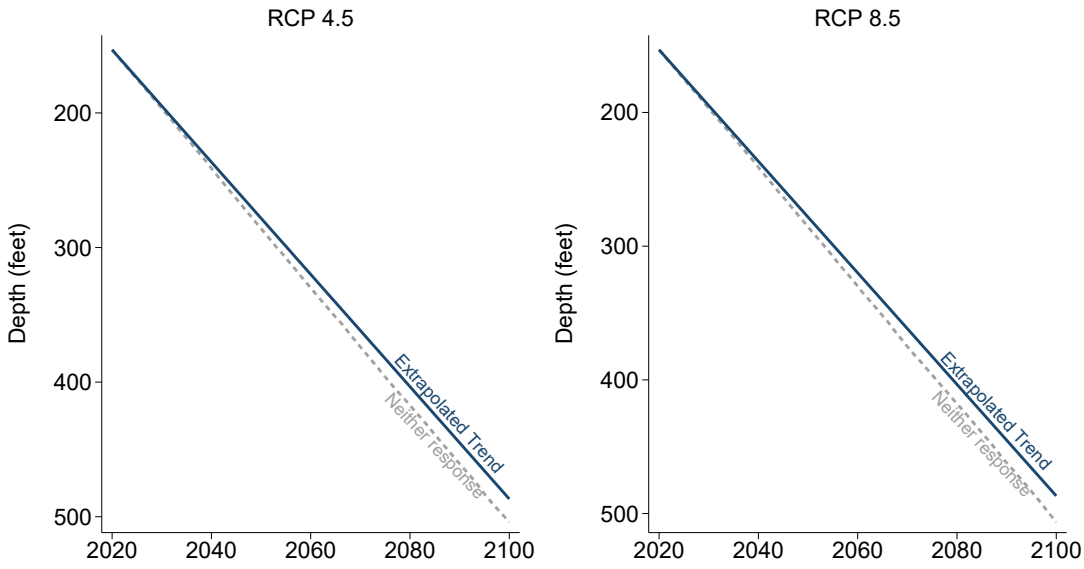
reports the means over 1,000 simulations of this "uninteracted" model (dashed lines), compared with analogous results from our preferred model that interacts temperature and depth (solid lines; reproduced from Figure 5). Compared to our preferred simulation model, the uninteracted model predices slightly greater groundwater extraction across all scenarios.

Figure E2: Impact of climate change on aquifer levels, uninteracted model



Notes: This figure compares our preferred simulation results from Figure 4 (solid lines) with the results from a parallel simulation model that assuming additively separable temperature and depth (dashed lines). Besides this different in the functional form of the fitted groundwater model $\hat{w}(\cdot)$, the two sets of simulations are identical. Lines report means across 1,000 simulation draws.

Figure E3: Impact of climate change on aquifer levels, without demand shocks



Notes: This figure presents the results of an extension of our simulation exercise. We model groundwater depth from 2020 through 2100. The left panel shows results under RCP4.5 (moderate climate change) and the right panel shows results under RCP8.5 (severe climate change). In both panels, the solid lines shows the extrapolated linear trend in average depth for the San Joaquin Valley basin. The dashed line shows our "neither response" model, where we take our preferred model and turn off the shocks to groundwater pumping demand due to temperature and depth responses.

F Assessing potential bias in depth coefficient due to dynamic feedback

Our simulations illustrate the dynamic feedback effect between aggregate groundwater pumping: due to the equation of motion, more pumping in year t serves to increase groundwater depths in year $t+1$. Figure 5 shows that out-of-sample temperature increases induce large aggregate feedback effects that influence out-of-sample depths.

Since Equation (7) does not account for any dynamic feedback from current groundwater use to future depths, our estimated reduced-form coefficient $\hat{\beta}$ on $\log(D_{st})$ depth could be biased. We can treat this as standard omitted variable bias, where the variable omitted from the right-hand side of Equation (7) is lagged groundwater pumping, $Y_{i(t-1)}$. If greater lagged groundwater pumping is correlated with greater contemporaneous depths (i.e., $\text{Cor}(Y_{i(t-1)}, D_{st}) > 0$), and if there positive within-unit serial correlation in the idiosyncratic error ε_{it} (i.e., $\text{Cor}(Y_{i(t-1)}, Y_{it}) > 0$), then $\hat{\beta}$ will have an upward bias (i.e., towards zero). The magnitude of this bias is uncertain a priori, as it depends on both the degree of serial correlation in ε_{it} and the empirical strength of the dynamic feedback within our sample period.

To assess the magnitude of such bias, we conduct a Monte Carlo analysis using simulated data calibrated to our estimation sample. We simulate a balanced panel goverened by the data-generating process:

$$Y_{it} \;=\; \zeta_1 \text{DegreeDayHigh}_{st} \;+\; \zeta_2 \log(D_{st}) \;+\; \gamma_i \;+\; \delta_t \;+\; \varepsilon_{it} \tag{F1}$$

This simplifies our main specification (Equation (7)) such that we can easily simulate data, while capturing any potential bias in $\hat{\zeta}_2$.

Monte Carlo procedure:

1. Create a balanced panel dataset for Monte Carlo iteration M of $N_i = 7,820$ parcels for $N_t = 12$ years, grouped into $N_s = 85$ groundwater sub-basins (all counts matching our estimation sample).

2. Generate parcel-specific means γ_i^M and year-specific means δ_t^M by making N_i and N_t independent draws from Normal distributions $N(\mu_{\hat{\gamma}_i}, \sigma_{\hat{\gamma}_i})$ and $N(\mu_{\hat{\delta}_t}, \sigma_{\hat{\delta}t})$ (with means and standard deviations calibrated to the estimate fixed effects in Equation (F1)).

3. Generate mean-zero idiosyncratic shocks ε_{it}^M following an AR(1) process with serial correlation $\rho = 0.31$, calibrated to the AR(1) correlation of $\hat{\varepsilon}_{it}$ in Equation (F1) (estimated with the restriction $\zeta_2 = 0$, to shut down any serial correlation induced by

persistence in D_{st}).[7] We simulate a random noise parameter calibrated to yield ε_{it}^M with the same variance as $\hat{\varepsilon}_{it}$ in Equation (F1)).

4. Simulate i.i.d. draws of $\text{DegreeDayHigh}_{st}^M$ from the Normal distribution with mean and standard deviation calibrated to the in-sample distribution of $\text{DegreeDayHigh}_{st}$.[8] Whereas our reduced-form model includes parcel-specific temperature realizations (with high spatial correlation), our Monte Carlos simulate a common yearly realization for all parcels within each sub-basin. Conservatively assuming that each parcel $i \in s$ receives the same temperature shock has the effect of concentrating the dynamic feedback effect, thereby erring in favor of biasing $\hat{\zeta}_2$.

5. Initiate first-period depth by drawing D_{s1}^M from the Normal distribution with mean and standard deviation calibrated to in-sample 2008 depths.

6. Construct first-period groundwater use, using the above inputs and the estimated coefficients $\hat{\zeta}_1$ and $\hat{\zeta}_2$:

$$Y_{i1}^M \;=\; \hat{\zeta}_1 \text{DegreeDayHigh}_{s1}^M \;+\; \hat{\zeta}_2 \log(D_{s1}^M) \;+\; \gamma_i^M \;+\; \delta_1^M \;+\; \varepsilon_{i1}^M$$

7. Looping over years $t = \{2, \ldots, 12\}$, construct D_{st}^M and Y_{it}^M:

$$\begin{aligned} D_{st}^M \;&=\; \hat{\chi}_1 D_{s(t-1)}^M \;+\; \hat{\chi}_2 \frac{\alpha}{A} \sum_{i \in s} \left[Y_{i(t-1)}^M \right] \;+\; \nu_{st}^M \\ Y_{it}^M \;&=\; \hat{\zeta}_1 \text{DegreeDayHigh}_{st}^M \;+\; \hat{\zeta}_2 \log(D_{st}^M) \;+\; \gamma_i^M \;+\; \delta_t^M \;+\; \varepsilon_{it}^M \end{aligned}$$

   The scalars α and A are the average α_s and A_s from Equation (D2). The coefficients $\hat{\chi}_1$ and $\hat{\chi}_2$ come from regressing D_{st} on lagged D_{st} and lagged $\frac{\alpha_s}{A_s} \sum_{i \in s} Y_{it}$, at the sub-basin by year level. This regression is a reduced-form analog to our simulated equation of motion (Equation (D1)), at the sub-basin (rather than basin) level. This goal of this regression is to simulate both persistence (through $\hat{\chi}_1$) and dynamics (through $\hat{\chi}_2$) in D_{st}^M with realistic magnitudes, rather than to interpret this relationship or the $\hat{\chi}_1$ and $\hat{\chi}_2$ estimates. Finally, we include a Normally distributed i.i.d. shock to depth ν_{st}^M, with mean and standard deviation calibrated to the error term of this depth regression.

8. Estimate the regression:

$$Y_{it}^M \;=\; \zeta_1^M \text{DegreeDayHigh}_{s1}^M \;+\; \zeta_2^M \log(D_{s1}^M) \;+\; \gamma_i^M \;+\; \delta_1^M \;+\; \varepsilon_{i1}^M$$

   and store the coefficients $\hat{\zeta}_1^M$ and $\hat{\zeta}_2^M$.

7. This AR(1) coefficient is slightly larger than $\rho = 0.25$, the AR(1) coefficient of $\hat{\varepsilon}_{it}$ from our preferred Equation (7). Restricting $\zeta_2 = 0$ has virtually no impact on $\hat{\rho}$.

8. This is defined as linear degree days above 24.5°C, which is one of the four temperature terms in our climate simulation model.

We conduct 1,000 Monte Carlos, repeating Steps 1–8 for each iteration M. One complication is that due to the large (empirical) variance in Y_{it} and ν_{st}, simulated D_{st}^M often becomes negative and breaks the $\log(D_{st})$ functional form. To address this, we inflate initial depths (in Step 5) by a large constant, which prevents successive negative shocks from yielding negative D_{st}^M. This preserves the log functional form and dynamic feedback relationship, but shifts D_{st}^M away from the empirical distribution of D_{st}. We also conduct a parallel set of 1,000 Monte Carlos using a functional form that is linear in depth, aligning D_{st}^M with the empirical distribution of D_{st}. Finally, to assess how stronger serial might induce bias in $\hat{\zeta}_2$, we conduct two additional sets of 1,000 Monte Carlos (for log and linear depth, respectively) where we assume $\rho = 0.9$ in Step 3.

Monte Carlo results: Table F1 reports results from these four sets of 1,000 Monte Carlos. The first row reports $\hat{\zeta}_2$ from Equation (F1), using the actual data from our main estimation sample. The second row reports the mean of $\hat{\zeta}_2^M$ over all 1,000 Monte Carlo iterations, setting $\rho = 0.31$ (to match the degree of serial correlation in ε_{it} in our sample). This mean is slightly less than $\hat{\zeta}_2$ in the log model, and slightly greater than $\hat{\zeta}_2$ in the linear model. These differences are indistinguishable from statistical noise, given the variation in these $\hat{\zeta}_2^M$ estimates (brackets report 5th and 95th percentiles). It follows that the combination of dynamic feedback into D_{st} and serial correlation in ε_{it} is unlikely to meaningfully bias our $\hat{\beta}$ estimates in Equation (7). The bottom row of Table F1 shows that if we simulate unrealistically strong serial correlation in ε_{it} by setting $\rho = 0.9$, the upward bias in $\hat{\zeta}_2$ becomes more pronounced.

Table F1: Monte Carlo simulations indicate negligible upward bias on depth coefficient

	Functional form for depth	
	Log	Linear
Equation (F1) estimate of $\hat{\zeta}_2$ (using actual data)	-426.99	-2.96
Mean $\hat{\zeta}_2^M$ using factual AR(1) $\rho = 0.31$	-432.70 $[-455.90, -410.92]$	-2.94 $[-2.96, -2.92]$
Mean $\hat{\zeta}_2^M$ using extreme AR(1) $\rho = 0.9$	-395.30 $[-430.00, -356.83]$	-2.90 $[-2.92, -2.87]$

Notes: This table reports results from the Monte Carlo simulations described above. We conduct four sets of 10,000 simulations: log vs. linear depth (i.e., $\zeta_2 \log(D_{st})$ vs. $\zeta_2 D_{st}$ in Equation (F1)); and actual vs. extreme AR(1) serial correlation (i.e., $\rho = 0.31$ vs. $\rho = 0.9$). For each set of Monte Carlos, we report the mean of $\hat{\zeta}_2^M$ over all 10,000 iterations, along with the 5th and 95th percentiles (in brackets). This indicates that while dynamic feedback effects induce an upward bias on $\hat{\zeta}_2$ due to positive serial correlation in ε_{it}, the magnitude of this bias in our empirical setting is likely negligible.

G Comparison to engineering estimate

Our reduced-form modeling generates estimates of increased water demand as a function of temperature. This empirical approach differs from approaches often leveraged in the engineering/crop science literature where process-based models are used to simulate water use and crop growth given weather conditions. Here we compare our reduced-form estimates to those generated by a commonly used crop simulation model, Aquacrop, which was developed by the Food and Agricultural Organization of the United Nations (Vanuytrecht et al. 2014). We opt to use Aquacrop over other commonly employed crop growth simulation models (e.g., DSSAT), as the Aquacrop model has been used extensively in the engineering literature measuring irrigation demands (Foster, Brozović, and Butler 2014, 2017).

To estimate increased irrigation water demand resulting from an additional day above 38° C within the framework of Aquacrop we develop two scenarios: 1) a baseline scenario describing average annual water within our study area over the 2008 to 2019 time period and 2) a 'heat shock' scenario where one observed daily temperature realization is altered to be above 38° C. We conduct 1000 iterations of scenario 2), randomly choosing, for each year of the simulation, one day within the growing season with a maximum temperature between 23° C and 26° C (the reference temperature bin in our reduced form estimation) and altering that day's maximum temperature to be 38° C. We randomly choose which day within the growing season to apply the 'heat shock' because the impact of heat exposure on changes in irrigation depends crucially on weather prior to the shock. For example, applying the heat shock to a day that was preceded by a temperate day is likely to result in different water use changes than applying the heat shock to a day that was preceded by a day with high temperatures given that irrigation needs in t are a function of soil moisture in t-1 which is influenced by temperature. This randomization allows us to uncover the average effect of a heat shock. Each iteration of the 'heat shock' scenario yields an estimate of annual total water use given the counterfactual of one day in the growing season having a higher maximum temperature which we use to calculate an average. Comparing this average to 1), the baseline scenario, yields the process-based model equivalent of our reduced-form model estimate i.e., how much water use increases if one additional day's maximum temperature exceeds a given threshold.

We operationalize Aquacrop using Aquacrop-OSPY which facilitates conducting crop growth simulations within Python (Kelly and Foster 2021). Aquacrop requires several data inputs to simulate crop production and irrigation water demands such as daily local weather outcomes (precipitation, maximum/minimum temperature, evapotranspiration), crop parameters, planting date, soil type, and irrigation management parameters. For scenario 1) we use observed daily precipitation and maximum/minimum temperature from PRISM averaging across the whole study area to generate daily averages. We use Hargreave's formula to approximate daily evapotranspiration given minimum and maximum daily temperatures and the extraterrestrial radiation of the Central Valley (Hargreaves and Allen 2003). We conduct our water use simulations using maize (corn) which is commonly used as a reference

crop in crop simulation modeling. We assume that the maize crop is planted on May 1^{st} of each growing season and cultivated on sandy loam soil. Finally, we set irrigation management parameters to reflect a soil moisture target where irrigation events are triggered when soil moisture reaches a given threshold. This threshold, described as a percentage between 0 and 100, reflects soil water availability with 0 representing the crops permanent wilting point (i.e., where it becomes impossible for plants to extract water at a rate fast enough to keep up with their water demand) and 100 representing field capacity (i.e., where the water content of the soil is considered to be ideal for crop growth). When the soil moisture target passes below the predefined threshold an irrigation event is triggered which applies water until the soil water availability reaches 100. For simplicity, we assume that this target is 50%.

We conduct the baseline water use simulation using daily data from our study. The baseline scenario results in annual water use ranging from 2.36 to 3.4 acre feet per acre irrigated. This estimate is slightly lower than water use in our sample, which is to be expected given that corn is slightly less water intensive than many of the annuals and perennials grown in the study area. On average, annual water use for the baseline scenario is 3.12 acre feet per acre. For the heat shock scenario, simulation results indicate that average annual water use is 3.14 acre feet per acre. These simulation results indicate that an additional day with maximum temperature above 38° C compared to a day with maximum temperature between 23° C and 26° C increases total annual water use by approximately 0.54%. This estimate is substantially lower (about a third) than that generated by the reduced-form modeling.

This difference suggests that engineering and process-based modeling approaches to estimating how farmer's respond to higher temperatures may overlook important information and behavioral aspects driving irrigation decision making. For example, Aquacrop, and other models like it, assume that farmers have perfect knowledge of soil moisture when scheduling irrigation events. However, in reality farmers may not have perfect knowledge of their field's soil moisture and decide to irrigated based on less precise heuristics e.g., the 'feel' of the soil. Also, process-based models determine daily irrigation amounts according to a soil moisture target and do not allow for irrigation above the field's water holding capacity. In reality, farmers may apply quantities of water above their field's capacity as insurance against the damages of too little water availability or because they lack information on the water holding capacity of their field's soil.

Appendix references

Burlig, Fiona, Louis Preonas, and Matt Woerman. 2025. "Groundwater and crop choice in the short and long run." NBER working paper 28706.

Carleton, Tamma, Amir Jina, Michael Delgado, Michael Greenstone, Trevor Houser, Solomon Hsiang, Andrew Hultgren, et al. 2022. "Valuing the global mortality consequences of

climate change accounting for adaptation costs and benefits." *The Quarterly Journal of Economics* 137, no. 4 (November): 2037–2105.

Dewandel, B., J.-M. Gandolfi, D. de Condappa, and S. Ahmed. 2008. "An efficient methodology for estimating irrigation return flow coefficients of irrigated crops at watershed and seasonal scale." *Hydrological Processes* 22 (11): 1700–1712.

Foster, Timothy, Nicholas Brozović, and Adrian P Butler. 2014. "Modeling irrigation behavior in groundwater systems." *Water resources research* 50 (8): 6370–6389.

———. 2017. "Effects of initial aquifer conditions on economic benefits from groundwater conservation." *Water Resources Research* 53 (1): 744–762.

Hargreaves, George H, and Richard G Allen. 2003. "History and evaluation of Hargreaves evapotranspiration equation." *Journal of irrigation and drainage engineering* 129 (1): 53–63.

Hsiang, Solomon, Robert Kopp, Amir Jina, James Rising, Michael Delgado, Shashank Mohan, DJ Rasmussen, Robert Muir-Wood, Paul Wilson, Michael Oppenheimer, et al. 2017. "Estimating economic damage from climate change in the United States." *Science* 356 (6345): 1362–1369.

Kelly, TD, and Timothy Foster. 2021. "AquaCrop-OSPy: Bridging the gap between research and practice in crop-water modeling." *Agricultural Water Management* 254:106976.

Loheide, Steven P., James J. Butler, and Steven M. Gorelick. 2005. "Estimation of groundwater consumption by phreatophytes using diurnal water table fluctuations: A saturated-unsaturated flow assessment." 41 (7).

Rasmussen, D.J., M. Meinshausen, and R.E. Kopp. 2016. "Probability-weighted ensembles of U.S. county-level climate projections for climate risk analysis." *Journal of Applied Meteorology and Climatology.*

Vanuytrecht, Eline, Dirk Raes, Pasquale Steduto, Theodore C Hsiao, Elias Fereres, Lee K Heng, Marga Garcia Vila, and Patricia Mejias Moreno. 2014. "AquaCrop: FAO's crop water productivity and yield response model." *Environmental Modelling & Software* 62:351–360.