NBER WORKING PAPER SERIES

GROUNDWATER AND CROP CHOICE IN THE SHORT AND LONG RUN

Fiona Burlig
Louis Preonas
Matt Woerman

Working Paper 28706
http://www.nber.org/papers/w28706

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2021, Revised February 2026

We thank Nick Hagerty for graciously sharing data with us. For helpful comments and suggestions, we thank Amy Ando, Andrew Ayers, Maximilian Auffhammer, Kathy Baylis, Ellen Bruno, Karen Clay, Michael Greenstone, Gautam Gowrisankaran, Koichiro Ito, Katrina Jessoe, Ryan Kellogg, Ashley Langer, Erik Lichtenberg, David McLaughlin, Kyle Meng, Mar Reguant, Michael Roberts, Paul Scott, Joseph Shapiro, Andrew Stevens, Anant Sudarshan, Catherine Wolfram, Brian Wright, David Zilberman, and numerous seminar participants. Lauren Beatty, Junshi Han, Chinmay Lohani, Ucindami Mafeni, Shubhalakshmi Nag, Anna Schmidt, Chen Sui, Yixin Sun, Xinyi Wang, Camellia Ye, and Ao Yu provided excellent research assistance. Karen Notsund provided invaluable support in obtaining data. We received generous funding in support of this project from the Giannini Foundation of Agricultural Economics, the Alfred P. Sloan Foundation (via the E2e Project), the Becker Friedman Institute for Economics at the University of Chicago, and the USDA Economic Research Service under Cooperative Agreement 58-6000-0-0052. We circulated a previous version of this paper under the title “Energy, Groundwater, and Crop Choice.” All remaining errors are our own. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Groundwater and Crop Choice in the Short and Long Run
Fiona Burlig, Louis Preonas, and Matt Woerman
NBER Working Paper No. 28706
April 2021, Revised February 2026
JEL No. Q15, Q25, Q41

ABSTRACT

Natural resource management is a dynamic problem, whose analysis is complicated by agents choosing to respond differently in the short- and long-run. Long-run considerations are particularly salient in groundwater management. We estimate farmers' responses to changes in groundwater pumping costs in California, one of the world's most valuable agricultural regions, where perennial crops induce dynamics via upfront costs and long-lived payoffs. Leveraging quasi-experimental variation in groundwater costs driven by regulated electricity tariffs, we estimate a dynamic discrete choice model of land use with state dependence and forward-looking farmers. Farmers' short-run elasticity of groundwater demand is −0.79, with no cropping response. In contrast, their long-run elasticity is −0.46, including meaningful reductions in water-intensive perennial cropping and increased fallowing. Short-run analysis alone therefore yields quantitatively different conclusions. California's flagship groundwater sustainability targets will require a 52% tax in regulated areas, which would lower perennial acreage by 10% and increase fallowing by 21%.

Fiona Burlig
University of Chicago
Harris School of Public Policy
and NBER
burlig@uchicago.edu

Louis Preonas
University of Maryland, College Park
lpreonas@umd.edu

Matt Woerman
Colorado State University
matt.woerman@colostate.edu

1 Introduction

Economists have long recognized that optimal management of both non-renewable (Hotelling (1931)) and renewable (Gordon (1954)) natural resources is a dynamic problem. Long-run policies are required to address the over-exploitation of resources including fish (Costello et al. (2010)), trees (Balboni et al. (2023)), the global climate (Nordhaus (2019)), and water (Carleton, Crews, and Nath (2025)). However, an agent's responses to policies may differ in the short vs. long run, especially if their choice set includes actions with high upfront costs and lasting payoffs. This complicates environmental policy analysis, since understanding the effectiveness and broader consequences of persistent policies requires estimates of agents' long-run responses.

Groundwater is a textbook common-pool resource that necessitates persistent management policies (Ostrom (1990); Provencher and Burt (1993)). Agricultural production consumes 70% of global groundwater withdrawals (UNESCO (2022)), such that absent policy to stem overextraction, groundwater aquifers are being rapidly depleted in farming regions across the globe (Jasechko et al. (2024)), threatening high-value crop production. Many high-value water-intensive crops are perennials that produce multiple years of harvest from a single upfront planting investment (French and Matthews (1971); Sant'Anna (2024); Hsiao (Forthcoming)). When these investment dynamics are present, the efficacy of groundwater management policies hinges on farmers' long-run responses.

We study California, where farmers produce 18% of total U.S. crop value, most of which comes from perennials that rely heavily on groundwater for irrigation (Bruno (2017); Liu et al. (2022)). Despite rapidly declining aquifer levels and a series of severe droughts, most California farmers currently face no meaningful restrictions on groundwater extraction. The state is in the process of implementing the Sustainable Groundwater Management Act (SGMA)—its first comprehensive groundwater regulations, which are designed to achieve aquifer sustainability by 2042. Given the sweeping nature of SGMA, it has the potential to significantly alter California's agricultural sector.

We ask three main research questions. First, what is the elasticity of demand for agricultural groundwater over the short and long run? Second, to what extent do farmers switch crops in response to higher groundwater costs? Third, how stringent would groundwater taxes need to be to achieve California's SGMA policy goals? Answers to these questions are essential for understanding the effectiveness of groundwater management policy, but have proven elusive because (i) groundwater pumping is rarely priced or measured, (ii) there is a dearth of plausibly exogenous variation in groundwater costs, and (iii) modeling how forward-looking farmers respond to cost shocks is a complex dynamic problem.

We are the first to answer these questions, overcoming the empirical challenges using a new measurement strategy, quasi-experimental variation in groundwater costs, and a dynamic discrete choice model. First, we leverage the fact that electricity is the main variable input to groundwater pumping. We assemble data on electricity prices and quantities for all agricultural consumers served by Pacific Gas & Electric (PG&E). Combining these data with newly constructed pump-specific production functions enables us to recover groundwater costs and quantities for farmers across the majority of California's Central Valley. Second, we use exogenous variation in PG&E's regulated electricity tariffs, which change farmers' groundwater costs differentially across space and over time. Third, to understand farmers' long-run responses to changing groundwater costs, we use this exogenous variation in electricity prices to identify a dynamic discrete choice model of farmers' cropping decisions. We use the conditional choice probability (CCP) approach (Scott (2013); Kalouptsidi, Scott, and Souza-Rodrigues (2021b)) in order to estimate model parameters without requiring strong assumptions on farmer beliefs beyond rational expectations and stationarity.

A static model of short-run farm profits would mischaracterize the decisions of farmers in California, where 60% of crop revenues come from perennials such as almonds and grapes (CDFA (2020)). Since these crops have high upfront planting costs and produce multiple years of harvests, accurately characterizing crop choices in this setting requires modeling state dependence and forward-looking farmers. We embed these features into a dynamic discrete choice model of crop choice, in which farmers can reduce water use *both* by switching crops *and* by using less water conditional on crop choice (behavior documented in Boser et al. (2024)). We identify model parameters using exogenous variation in groundwater costs driven by changes in regulated electricity prices. This structural approach uses short-run variation to quantify long-run responses while also facilitating counterfactual policy simulations.

We begin by outlining a simple illustrative model to build intuition for how crop choice and water use respond to a temporary vs. permanent groundwater cost shock. For a temporary cost shock, farmers do not incur the fixed costs of switching to a less water-intensive crop; instead, they may respond by applying (close to) no water to their existing crops, sacrificing the current cropping cycle to avoid negative profits. For a permanent cost shock, farmers incur the fixed crop switching costs, shifting to a new equilibrium of less water consumption but lower revenue. This shows how it can be both feasible and economically rational for farmers to exhibit a larger elasticity of groundwater demand in response to a temporary cost shock than to a permanent cost shock.

We then provide reduced-form evidence that farmers respond to pumping costs. Leveraging exogenous year-over-year variation in electricity tariffs in an instrumental variables approach, we find that farmers reduce groundwater consumption in response to pumping cost shocks.[1] This effect is driven by changing water use conditional on crop choice, rather than by crop switching—as evidenced by (i) null reduced-form estimates of land use change and (ii) a quantitatively similar estimate of groundwater response after conditioning on crop choice. We use this conditional reduced-form estimate to calibrate the short-run intensive-margin response in our dynamic model. To calibrate its long-run analog, we estimate the intensive-margin response for farmers who have made the same crop choice in at least six consecutive periods, thereby isolating responses that are likely to persist over the long run.

Next, we use our dynamic model to measure the effects of short-run pumping cost shocks on land use changes. To do so, we solve for farmers' value functions via a fixed-point algorithm, simulating the model forward until it reaches a steady state, and then inject an (unanticipated) one-year cost increase into the model. We find that such a short-run shock would lead to essentially no change in crop choice, with semi-elasticities of -0.0001 for annuals, -0.0004 for fruit/nut perennials, -0.0007 for hay perennials, and 0.001 for non-crop (fallowing). Combining these negligible land use changes with the short-run intensive-margin elasticity yields an overall short-run groundwater elasticity of -0.79.

1. Mieno and Brozovic (2017) point out that prior studies using energy data to estimate groundwater demand often recover biased estimates, due to significant measurement error or a lack of identifying variation. Our detailed microdata and quasi-experimental variation help us overcome these challenges, in the tradition of utility demand estimates that use regulated tariffs for identification (e.g., Olmstead (2009); Ito (2014)).

We then use our model to compute long-run semi-elasticities of land use change. In contrast to a short-run price shock, we find that farmers *do* switch crops in response to permanent groundwater cost changes, with semi-elasticities of 0.071 for annuals, −0.129 for fruit/nut perennials, −0.007 for hay perennials, and 0.065 for non-crop. These results reveal that forward-looking farmers operating state-dependent fields will react very differently to a short-run cost shock vs. a permanent one. In particular, permanent price changes induce farmers to switch into less-water-intensive crops, while short-run shocks do not.

Combining these land-use changes with farmers' intensive-margin adjustments, we find a long-run groundwater demand elasticity of −0.46. Consistent with our illustrative model, this long-run elasticity is meaningfully smaller than our short-run elasticity. These results align with prior evidence that in dynamic settings, long-run elasticities need not be larger than short-run elasticities (e.g., Hall (1991); Hendel and Nevo (2006); Gowrisankaran and Rysman (2012); Castillo (2021); Lemoine (2024)). Our results further show how short- and long-run responses to environmental policy may diverge substantially, as different mechanisms—yielding different elasticities—are privately optimal over different time horizons.

Finally, we use our structural model to simulate farmers' long-run responses to counterfactual groundwater taxes, which can address the open-access externalities associated with groundwater pumping (Provencher and Burt (1993)) by incentivizing sustainable levels of groundwater extraction. According to the state's Groundwater Sustainability Plans (GSPs), achieving "sustainable yield" under SGMA in overdrafted regions of our sample will require groundwater pumping reductions of 16.7% on average, with substantial variation across locations. Our simulations suggest that, on average, a 52.0% tax on groundwater pumping would be required to achieve this sustainability target. The taxes required to achieve these goals—which vary meaningfully over space—would reduce fruit/nut perennial acres by 10.0%, not change hay perennials, increase annual acres by 4.2%, and increase fallowing by 20.7%, compared to their respective acreage in our no-tax scenario. These results imply that non-trivial groundwater taxes can achieve SGMA's sustainability goals, and that doing so will induce meaningful changes to California's 20 million acres of cropland.

This paper makes three main contributions. First, and most importantly, we use dynamic discrete choice methods to generate short- and long-run demand elasticity estimates for agricultural groundwater. We find that these elasticities differ markedly in both their magnitudes and their underlying mechanisms. Our preferred approach captures key dynamics in agricultural land use, in contrast to previous static estimates of groundwater demand (e.g., Hendricks and Peterson (2012); Bruno and Jessoe (2021a); Pfeiffer and Lin (2014); Smith et al. (2017)).[2] Bruno, Jessoe, and Hanemann (2024) use a reduced-form approach to study land use and groundwater dynamics over five years, in response to voluntary water pricing in a single water district in California's Pajaro Valley. Our results build on this work: we capture both short- and long-run land use and groundwater responses to groundwater cost shocks across the majority of California's farming areas, using a fully dynamic modeling approach to simulate counterfactual responses to short- and long-run groundwater policy.

Our findings highlight the value of long-run environmental policy analysis. We demonstrate that if a policymaker assumed the short-run elasticity were persistent, she would anticipate a long-run response that is 70% too large. Moreover, she would erroneously conclude that farmers would achieve groundwater pumping reductions without changing cropping patterns, whereas in reality, a permanent groundwater tax would induce substantial crop switching. The dynamics of perennial cropping, which are driven by high upfront costs followed by a multi-year payoff, mirror a wide range of investments, including vehicles (Dahl (2014)), household appliances (Dubin and McFadden (1984)), pollution control technologies (Blundell, Gowrisankaran, and Langer (2020)), and cattle herds (Scott et al. (2026)). While our results come from one key context—scarce water resources in California, which economists have studied for over a century (Coman (1911))—the lesson that agents may respond differently to short- vs. long-run policies is applicable wherever resource management requires a long time horizon.

Second, we estimate the short- and long-run impacts of water costs on land use. While agricultural economists have long studied the effect of output prices on cropping patterns (e.g., Nerlove (1956); Roberts and Schlenker (2013); Scott (2013)), fewer studies have doc-

2. Scheierling, Loomis, and Young (2006) conduct a meta-analysis of 24 earlier papers estimating agricultural water demand. Most of these studies rely on agronomic models or field crops experiments with restrictive assumptions on farmers' response to changing water costs. In contrast, more recent estimates of agricultural water demand rely on observed farmer responses (e.g., Bruno and Jessoe (2021b)). We also estimate the long-run elasticity of electricity demand in the agricultural sector, which consumes nearly 8% of the state's electricity. This builds on recent work using quasi-experimental variation to estimate long-run residential electricity demand (Deryugina, MacKay, and Reif (2020); Feehan (2018); Buchsbaum (2023)).

umented how groundwater costs impact crop choice, which has important implications for agricultural output markets.[3] We build on Hornbeck and Keskin (2014) by demonstrating that farmers change crops in response to groundwater costs in the long run, but not in the short run.[4] Our work complements recent studies of surface water irrigation (Rafey (2023); Hagerty (2022); Hagerty (2025)), where regulatory and market institutions are far more mature. We extend recent studies of localized groundwater regulations (Ayres, Meng, and Plantinga (2021); Bruno, Jessoe, and Hanemann (2024)) by providing estimates of land use change under groundwater policy—for the majority of California's farmland—in both the short and long run.

Third, we extend the literature on groundwater management by using our structural model to simulate farmer responses to (counterfactual) groundwater policy, in the context of California's landmark SGMA regulation. Natural scientists have uncovered substantial groundwater depletion in key agricultural regions across the globe (Fan, Li, and Miguez-Macho (2013); Rodell et al. (2018)). However, large-scale groundwater regulation remains rare (Carleton, Crews, and Nath (2025)), as the few existing policies are mostly local in scope.[5] In this context, SGMA stands to be one of the world's most consequential groundwater regulations. Early work on SGMA has focused on the political economy of (Bruno, Hagerty, and Wardle (2022)) and anticipatory responses to (Bruno and Hagerty (2025)) the regulation. We contribute novel estimates of the impact of groundwater pricing, demonstrating that stringent policies will be required to achieve SGMA's sustainability goals. This underscores how SGMA is poised to alter the landscape of some of the most valuable cropland on earth.

This paper proceeds as follows. Section 2 provides background on groundwater pumping and energy use in California agriculture. Section 3 outlines our illustrative model. Section 4 describes our data. Section 5 presents our identifying variation and reduced-form estimates.

3. Blakeslee, Fishman, and Srinivasan (2020) and Ryan and Sudarshan (2022) show that groundwater depletion hurts long-run farm profits in India, but there is far less evidence from high-income countries.

4. Dinar (1994) uses a dynamic theory model to show that rising energy costs are likely to impact crop choice. Caswell and Zilberman (1986) analyze the theoretical relationship between energy demand and irrigation technology choice, a separate determinant of irrigation costs.

5. For example, groundwater regulations exist in parts of Kansas (Drysdale and Hendricks (2018)), parts of Colorado (Smith et al. (2017)), and small regions of California (Bruno, Jessoe, and Hanemann (2024); Ayres, Meng, and Plantinga (2021)).

Section 6 outlines our structural model and presents our dynamic estimates and counterfactual simulations. Section 7 concludes.

2 Background and key institutional details

2.1 Agriculture and irrigation in California

California is a major player in global agricultural production. The state produced $32 billion in crop value in 2019, representing 18% of the U.S. total—including 75% of the total value of U.S. fruits and nuts, and 57% of the total value of U.S. vegetables (USDA (2021)). California's 77,000 farms produce over 400 commodities, and they are the exclusive domestic producers of almonds, artichokes, olives, walnuts, and numerous other high-value crops (California Department of Food and Agriculture (2011)).

Irrigation is essential for farming in California due to scant summer precipitation. 95% of the state's 8.3 million harvested acres are irrigated (Johnson and Cody (2015)), and the agricultural sector is responsible for 80% of the state's total water consumption. Many of California's crops use large amounts of water. For example, hay, almonds, grapes, and rice—four of California's top crops by acreage—all require at least 3 acre-feet per acre per year, with rice using 5 acre-feet per acre per year (Bruno (2019)). To water these thirsty crops, farmers rely on two water sources with vastly different governance structures (Sawyers (2007)): in an average year, 61% of irrigation comes from surface water, while 39% comes from groundwater (California Department of Water Resources (2015)).

Surface water Surface water in California is strictly regulated. Almost all farms with access to surface water obtain it via water districts. Most water districts function as cooperatives that divert water from rivers and canals for distribution to farmers in their geographic territory.[6] Individual farmers typically receive water allocations proportional to their acreage within the district (Schlenker, Hanemann, and Fisher (2007)); these allocations fluctuate yearly depending on scarcity (e.g., the amount of snowpack). Importantly, farmers pay a

6. Districts were established between 1860 and 1950, and their boundaries have remained essentially fixed. Though some farms have individual water entitlements, the vast majority of surface water allocations come from districts. Hagerty (2022) provides a detailed description of surface water rights in California.

lower marginal cost for district water allocations than for self-pumped groundwater (Hagerty (2022)). Most farmers are therefore unlikely to incur groundwater pumping costs without also exhausting their annual allocation of cheaper district water; consistent with this, when surface water availability declines, farmers turn to groundwater (Ferguson (2025)).

Farmers also have a limited ability to purchase surface water on the open market. However, such transactions constitute only a very small share of total water deliveries, at prices much higher than marginal groundwater pumping costs (Hagerty (2025)).[7] Purchased water is therefore unlikely to be a viable substitute for agricultural groundwater.

Groundwater Groundwater supplies 30–40% of all water end uses in California in a normal year, and close to 60% in drought years when surface water is scarce (California Department of Water Resources (2014)). Unlike surface water, agricultural groundwater rights in California tend to be far more vague. The typical groundwater right is "overlying," meaning that a landowner whose property sits above an aquifer has the right to extract the underlying groundwater. Overlying rightsholders face few restrictions to drilling new groundwater wells, which cost $75,000 on average and typically reach 300–500 feet (Hadachek et al. (2026)).[8] Historically, the vast majority of groundwater use has been unmetered, with users facing no variable prices beyond the costs of pump operation (Bruno and Jessoe (2021a)). This has enabled farmers to extract vast amounts of groundwater to irrigate their overlying cropland.

Nearly all groundwater pumps in California run on electricity, the sole variable input to groundwater production.[9] This makes groundwater pumping the dominant electricity end use in the agricultural sector, which accounts for nearly 8% of the state's electricity consumption (California Energy Commission (2005)). Our empirical strategy leverages exogenous variation in electricity prices to instrument for groundwater pumping costs. While previous studies have used variation in energy costs to estimate the price elasticity of groundwater

7. In our sample, the 99th percentile of marginal groundwater pumping costs is $137 per acre-foot. By contrast, Hagerty (2025) reports an *average* transaction price of $221 per acre-foot on the open market.

8. New wells must be reported to the Department of Water Resources, and construction usually lasts less than one week (Central Valley Flood Protection Board (2020)). In some cases, users who do not own land above an aquifer hold "appropriative" groundwater rights based on the prior appropriation doctrine (i.e., "first in time, first in right", by seniority of historic extraction). Appropriative groundwater rights are legally subordinate to overlying groundwater rights, and users may only exercise them in the event of a surplus.

9. Per the 2018 Census of Agriculture's Irrigation and Water Management Survey, California farms operate 94,698 pumps. 84,856 are powered by electricity, and only 8,043 are powered by diesel (USDA (2018)).

Figure 1: Groundwater depletion and perennial crops

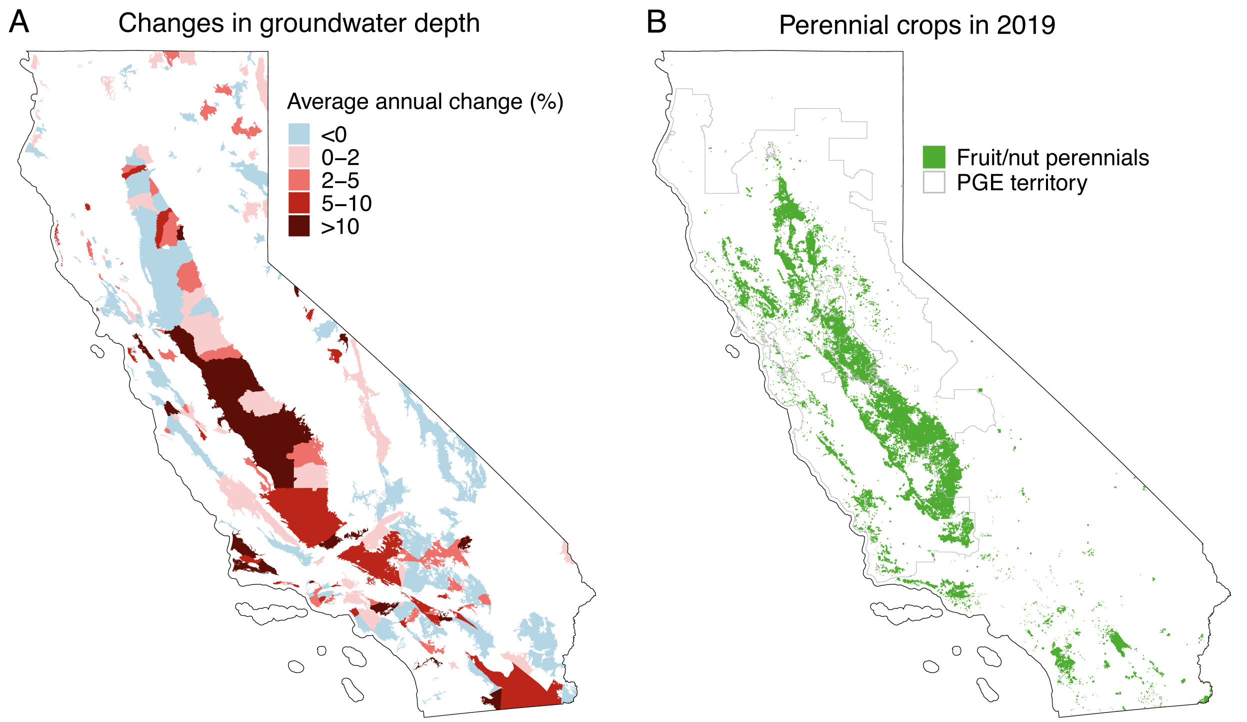

Notes: Panel A plots California's groundwater sub-basins, shading based on the average annualized change in depth during our 2008–2019 sample period. A 10% change in depth corresponds to a 10% increase in groundwater pumping costs, holding all else constant. This map averages depth measurements across each sub-basin from April–June of each year, to remove seasonality. Panel B plots the extent of fruit and nut perennial cropping from 2019, shading parcels for which "fruit/nut perennial" was the modal crop category. We also plot PG&E's service territory in gray, which encompasses most of this perennial acreage.

demand (e.g., Badiani and Jessoe (2019)), our instrumental variables approach overcomes several inherent challenges highlighted by Mieno and Brozovic (2017)—including measurement error and lack of micro-level identifying variation.

2.2 Groundwater depletion and management policy

Due to California's longstanding open-access groundwater regime, many of the state's groundwater basins are "overdrafted"—meaning that withdrawals exceed the pace of replenishment, often by millions of acre-feet each year. The Central Valley has seen substantial groundwater losses, where the "critically overdrafted" Tulare and San Joaquin basins lost a combined 120 million acre-feet of groundwater from 1925–2008 (Konikow (2013)). California's groundwater depletion has been accelerating: while the average depletion rate from 1961–2021 was approximately 1.5 million acre-feet per year, a series of severe droughts increased this rate to 7 million acre-feet per year from 2019–2021 (Liu et al. (2022)). Panel A of Figure 1 shows that much of the Central Valley faced 10% average annual increases in groundwater depths (i.e., reductions in aquifer levels) during our 2008–19 study period, with greater losses in the southern half of the Valley. Panel B illustrates that these same areas are home to concentrated production of (high-value, water-intensive) fruit and nut perennial crops.

A severe drought beginning in 2011 raised serious concerns about the future sustainability of California's groundwater resources. In September 2014, state lawmakers responded by passing the Sustainable Groundwater Management Act (SGMA). This sweeping legislation represented the first statewide effort to regulate groundwater extraction across all agricultural areas in the state, which are responsible for 90% of groundwater pumping (Bruno, Hagerty, and Wardle (2022)). SGMA comprises three separate bills. AB 1739 empowers California's Department of Water Resources (DWR) or local groundwater sustainability agencies (GSAs) to charge fees for groundwater extraction, and it requires GSAs to prepare groundwater sustainability plans (GSPs). SB 1319 authorizes GSAs to implement these GSPs. SB 1168 mandates that groundwater end uses be both reasonable and beneficial, and it enables GSAs and the DWR to require groundwater monitoring.

SGMA represents the future of groundwater management in California, with the goal of long-run sustainability—with each basin operating within its sustainable yield and avoiding "undesirable results."[10] Critically overdrafted (other medium- and high-priority) basins were required to submit GSPs by 2020 (2022) and are required to achieve sustainability by 2040 (2042). Importantly, all SGMA implementation has occurred after our 2008–2019 analysis period. Bruno and Hagerty (2025) argue that there has not been anticipatory action to reduce groundwater use in response to SGMA's passage.

Using data from the universe of GSPs, we calculate an average required reduction in pumping of 16.7% among regions in our sample currently experiencing overdraft. GSAs have a variety of tools at their disposal for reducing groundwater pumping, including price instruments (such as taxes or fees), quantity instruments (including both tradable and non-tradable allocations), ad-hoc pumping restrictions, and other conservation incentives (Bruno,

10. Undesirable results include "chronic lowering of groundwater levels," "significant and unreasonable reduction of groundwater storage," "significant and unreasonable seawater intrusion," significant and unreasonable degraded water quality," "significant and unreasonable land subsidence," or "depletions of interconnected surface water" (California Department of Water Resources (2017)).

Hagerty, and Wardle (2022)). Researchers predict these policy instruments will induce a variety of behavioral changes, including reducing irrigation intensity, shifting towards less water-intensive crops, and/or land fallowing (Bruno (2019)).

3 Illustrative model

This section outlines a simple illustrative model to build intuition for how farmers respond to a temporary (i.e., short-run) or permanent (i.e., long-run) shock to the marginal cost of groundwater. We assume that a farmer maximizes profit by choosing which crop to produce and deciding how much water to apply to that chosen crop. Farmers incur three types of costs: a linear marginal cost of water, fixed costs of switching crops, and other input costs that are crop specific (e.g., labor for harvest). Revenues are crop-specific, increasing, and concave with respect to water, consistent with a large literature in agricultural economics (e.g., Caswell and Zilberman (1986)). For simplicity, we assume two crop types: Crop 1 has high value and high water needs (e.g., almonds), while Crop 2 has low value and low water needs (e.g., corn). Panel A of Figure 2 depicts revenues as a function of water input, with Crop 1 in dark blue and Crop 2 in light blue. $\underline{R}_1$ represents the revenue threshold below which it is not profitable to produce Crop 1 because of substantial non-water input costs.[11]

Panel B of Figure 2 plots the corresponding marginal revenue curves and marginal revenue threshold $\underline{MR}_1$. Conditional on crop choice, a farmer will apply the quantity of water that equates marginal cost and marginal revenue, as long as the resulting profit is not negative. If equating marginal cost and marginal revenue would yield negative profit (i.e., producing above MR_1 when growing Crop 1), the farmer will instead cease production, thereby avoiding water costs and other input costs. (In reality, the origin might not imply literally zero water input, as farmers maintain their land for the next growing season.) We assume two types of farmers: Farmer L faces a low marginal cost of water MC_L, while Farmer H faces a high marginal cost of water MC_H. In the status quo, it is profit-maximizing for both farmers to choose the high-value, water-intensive Crop 1, but at different levels of

11. To parameterize this figure, we draw on real input cost and revenue information for almonds and corn in California's Northern San Joaquin Valley (Davis Cost Studies (2014–24)). Almonds have large costs of non-water inputs, implying a high $\underline{R}_1$. We do not plot a similar revenue threshold for Crop 2 for legibility.

Figure 2: Illustrative model of crop choice and water use

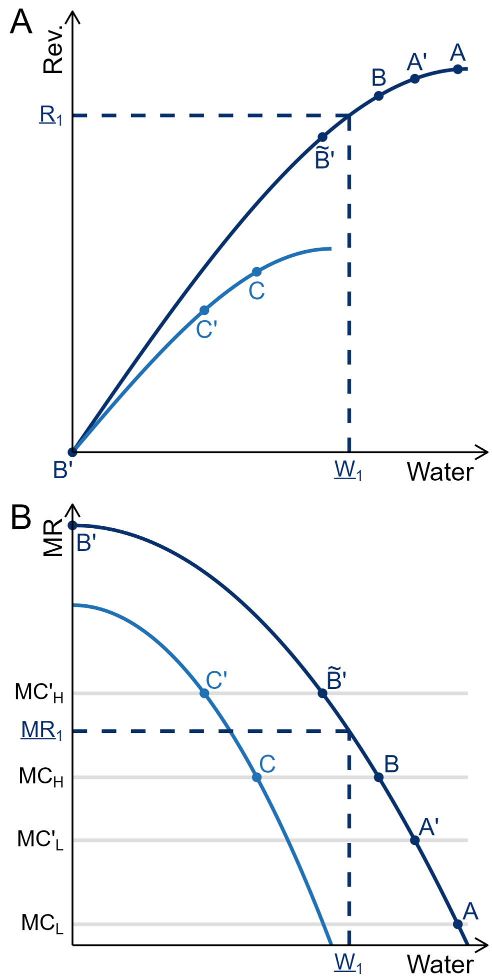

Notes: This figure presents our illustrative model of crop choice and water use. Panel A plots crop revenue against water use. Panel B plots *marginal* revenue against water use. Both plots depict two crops: the high-revenue, high-water-intensity Crop 1 (dark blue; parameterized to approximate the production function for almonds), and the low-revenue, low-water-intensity Crop 2 (light blue; parameterized to approximate the production function for corn). $\underline{R}_1$ denotes the revenue threshold below which it becomes unprofitable to grow Crop 1 due to high non-water input costs; $\underline{MR}_1$ and $\underline{W}_1$ denote the corresponding marginal revenue threshold and water use level, respectively. In the status quo, Farmers L and H produce Crop 1 at points A and B, respectively. A temporary (or permanent) shock to the marginal cost of water induces farmers to shift their production as described in the text.

production and water use: Farmer L chooses to produce at point A, while Farmer H chooses to produce at point B.

We now consider a temporary increase to the marginal cost of water, where MC_L and MC_H rise to MC'_L and MC'_H for the current growing season only (which is known to the farmer *ex ante*). For Farmer L, it is still profit-maximizing to produce Crop 1, but with a lower level of output at point A'. However, Farmer H can no longer profitably produce Crop 1, because $MC'_H > \underline{MR}_1$, and producing at point $\tilde{B}'$ would yield negative profit. Farmer H has two options: remain in Crop 1 and cease production, earning zero profit at point B'; or incur the fixed costs of switching to Crop 2 and produce at point C', necessitating additional fixed costs of switching back to Crop 1 in the future. For sufficiently large switching costs, switching to Crop 2 is not long-run profit-maximizing when marginal water cost reverts back to MC_H in the next growing season. Hence, Farmer H responds to this temporary cost shock by remaining in Crop 1 and earning zero profit at point B' for one year.[12] This pattern produces a sizable reduction in water use from point B to point B', generating a large short-run elasticity of water demand.

Next, we consider a permanent marginal cost increase of the same magnitude. As above, Farmer L remains in Crop 1 and produces at point A'. Farmer H also faces the same two options: cease production of Crop 1 or switch to Crop 2. In this case, however, Farmer H chooses to switch to Crop 2, now the profit-maximizing crop choice in current *and future* years. After incurring switching costs only once, Farmer H will realize the benefit of switching crops—the profit from producing Crop 2 at point C'—in all future years. Hence, Farmer H responds to this permanent cost shock by switching to Crop 2 and producing at point C'. This pattern generates a smaller reduction in water use from point B to point C', amounting to a smaller long-run elasticity of water demand.

This simple model illustrates two key points. First, it is feasible and economically rational for farmers to have a larger elasticity of water demand in response to a temporary cost shock than to a permanent cost shock: the average short-run elasticity includes A to A' and B to B', while the average long-run elasticity includes A to A' and B to C'. Second, *conditional on remaining in a given crop*, farmers' average elasticity of water demand is smaller in the long run than in the short run: whereas the average short-run within-crop response includes both A to A' and B to B', the average long-run within-crop response only includes A to A' (because Farmer H has switched crops).

4 Data

This section provides an overview of our data. For additional details, see Appendix C.

12. Shifts from A to A' and from B to B' would effectively concentrate irrigation and production of Crop 1 onto the most profitable land, aligning with Manning, Goemans, and Maas (2017). Marcheva (2025) documents that ceasing production ("idling", "crop abandonment") is common in this setting.

4.1 Electricity quantities and prices

We use confidential customer-level microdata from all of PG&E's agricultural service points (i.e., electricity meter locations). For each service point, we observe monthly billing data from 2008–2019. These data report each service point's latitude and longitude, electricity tariff, monthly bill amount (in dollars), and monthly electricity consumption Q^{elec} (in kilowatt-hours, or kWh).[13]

PG&E publishes its agricultural tariff schedules, which are the outcome of statewide regulatory proceedings. The California Public Utilities Commission sets rates based on the expected variable cost of electricity and the recovery of reasonable non-marginal costs (e.g., transmission investments). Tariffs are fixed 1–3 years in advance of their implementation, and do not respond to contemporaneous conditions (e.g., drought). Tariffs are geographically uniform, and individual farmers cannot plausibly influence how they are determined. Each tariff comprises a combination of fixed charges ($ per bill and $ per kW of maximum energy draw) and volumetric charges ($ per kWh). These volumetric charges are linear, meaning that a farmer's marginal price is not endogenous to her own electricity consumption—unlike the nonlinear increasing block rates studied by Olmstead (2009) and Ito (2014). By matching historic tariff schedules to our monthly billing data, we construct our main electricity price variable, P^{elec}, as the average marginal (i.e., volumetric) electricity price faced by each customer over all hours in each billing period.

Importantly, PG&E restricts a farmer's tariff eligibility based on two dimensions of physical capital: (i) pump size (smaller vs. larger than 35 horsepower), and (ii) electricity meter type (conventional analog vs. digital smart meters). This creates four mutually exclusive tariff categories: small-conventional, large-conventional, small-smart, and large-smart. Farmers cannot select across categories absent a major capital equipment change or intervention from PG&E. Small-conventional and large-conventional pumps are eligible for one tariff each. Small-smart and large-smart pumps are eligible for eight and twelve tariffs, respectively, each with different time-varying volumetric prices. A farmer can select a tariff only from her category-specific menu. To account for this within-menu selection, we define the

13. We drop the 9% of PG&E agricultural customers that ever have solar panels from our analysis, since their billed electricity use is net of (unobserved) solar production (i.e., we do not observe their actual electricity consumption).

instrumental variable $P^{\text{elecDefault}}$ as the average *default* marginal (i.e., volumetric) electricity price in each customer's tariff category in each billing period.[14] This instrument exploits the tariff category restrictions for identification, as we discuss in Section 5.1 below.

4.2 Groundwater quantities and costs

To construct groundwater quantities and marginal pumping costs, we leverage a unique PG&E dataset of agricultural groundwater pump audits conducted as part of an ongoing energy efficiency program.[15] We observe operating pump efficiencies from over 30,000 pump tests from 2011–2019, along with other detailed measurements and technical specifications. We match pump tests to service points in our billing data using electricity meter identifiers, isolating a subset of 10,146 service points with confirmed agricultural groundwater pumps.[16]

Physics governs the kWh of electricity input required to produce 1 acre-foot (AF) of groundwater (Hurr and Litke (1989)):

$$\frac{\text{kWh}}{\text{AF}} = \frac{[\text{Lift (feet)}] \times 1.0241}{\text{Operating pump efficiency (\%)}} \tag{1}$$

PG&E's audit data report the operating efficiency of each pump. To parameterize lift—the vertical distance from the groundwater source to the surface—we combine PG&E's measurements with publicly available data on groundwater depths from California's Department of Water Resources. These depth measurements vary across space and time, allowing us to construct pump-by-month specific $\frac{\text{kWh}}{\text{AF}}$ that accounts for contemporaneous groundwater depths at each pump's location.[17] Then, we define groundwater consumption $Q^{\text{water}} = Q^{\text{elec}} \div \frac{\text{kWh}}{\text{AF}}$

14. For the small-conventional and large-conventional categories, $P^{\text{elecDefault}} = P^{\text{elec}}$. For the small-smart and large-smart categories, we assign $P^{\text{elecDefault}}$ as P^{elec} of the least-time-varying tariff in each category. Using each category's modal tariff yields similar results (see Appendix Table B8). We omit a fifth category reserved for the 1.7% of farmers transitioning from internal combustion to electric power, since they are likely not comparable to the rest of our sample, and we do not observe them before they consumed electricity.

15. PG&E heavily subsidizes these pump audits, such that farmers bear (close to) zero cost. Contractors are incentivized to test as many pumps as possible, yielding a non-random subset of audited pumps.

16. PG&E typically installs a dedicated meter for each groundwater pump. Nearly all pump tests match to an agricultural service point. Focusing our analysis on this matched subset of confirmed pumps ensures that we are measuring energy used for pumping, avoiding other agricultural electricity end uses (e.g., refrigeration, barn lighting, or heating greenhouses). We drop service points for which a matched pump test reports a non-well water source (e.g., canal), ensuring that our sample comprises confirmed *groundwater* pumps.

17. We rasterize thousands of depth measurements for each sample month. Calculating lift also requires pump-specific measures of drawdown (i.e., how much a pump's extraction impacts its own depth), which

(in AF) and marginal groundwater pumping costs $P^{\text{water}} = P^{\text{elec}} \times \frac{\text{kWh}}{\text{AF}}$ (in \$/AF) for each pump in each month.

Constructing Q^{water} and P^{water} introduces multiple sources of measurement error: infrequent pump tests, spatial and temporal interpolation of depth measurements, and imprecise modeling of how pumps impact their own depth. Our instrumental variable $P^{\text{elecDefault}}$ addresses this measurement error because electricity tariffs are uncorrelated with the parameterization of Equation (1). Section 5.1 describes our instrumental variables approach.

We restrict our analysis to confirmed pumps (i.e., electricity service points with matched groundwater pump tests) for two reasons. First, we require pump audit data to convert from electricity to groundwater. Second, this prevents us from mistakenly incorporating other agricultural electricity uses not directly related to groundwater.[18]

4.3 Land use data

We use county assessor tax parcels as farm boundaries, as in Bruno, Jessoe, and Hanemann (2024). We spatially merge PG&E service points to parcel polygons, linking each groundwater pump to the fields that it most likely irrigates. We aggregate by summing quantities (Q^{water}, Q^{elec}) and averaging prices (P^{water}, P^{elec}, $P^{\text{elecDefault}}$) across all pumps in each parcel and all months of each year, defining our unit of analysis as the parcel (i.e., farm) by year.

We also match parcel polygons to the USDA's Cropland Data Layer (CDL), which reports annual satellite-derived crop coverage for each 30m^2 pixel in California. We classify CDL-reported land types into five mutually-exclusive and exhaustive categories: annuals, fruit/nut perennials, hay perennials, non-crop (i.e., fallow cropland), and not croppable.[19] We also link parcels to groundwater sub-basins (to enable controls for common shocks to groundwater depth) and to water districts (to enable controls for surface water allocations).

depends on rate of flow and subsurface characteristics. We parameterize drawdown using the value reported in each pump test. Appendix C.3 provides further details on how we apply and parameterize Equation (1).

18. This sample restriction is conservative, since it removes any groundwater pumps that did not receive PG&E pump audits. However, we recover similar reduced-form estimates if we expand the sample to include all PG&E agricultural customers (see Table 2, Columns (2)–(3)). We can also rule out cone-of-depression spillover effects from unobserved pumps that we omit from our sample (see Appendix Table B7).

19. Our analysis removes all not-croppable acreage (e.g., development, forests), adjusting the denominator of each parcel to include only cropland. Due to measurement error in the CDL (as discussed in Hagerty (2022)) and to ease computational burden, we use crop categories rather than individual crop classifications.

Table 1: Summary statistics

	Confirmed pumps	Other agricultural users
A. Service point-level statistics		
Unique service points (SPs)	10,146	83,721
Months observed (2008–2019)	102.7	101.9
	(42.0)	(43.2)
Average electricity use (kWh/month)	10,133	4,519
	(13,789)	(27,204)
Average marginal electricity price ($/kWh)	0.13	0.16
	(0.03)	(0.04)
Average electricity bill ($/month)	1,825.78	781.12
	(2,006.48)	(3,181.37)
Average groundwater use (AF/month)	33.26	
	(42.52)	
Average marginal groundwater cost ($/AF)	47.69	
	(25.91)	
B. Parcel-level statistics (acreage-weighted)		
Unique parcels	7,127	41,732
Count of SPs per parcel	1.92	2.00
	(1.71)	(1.56)
"Croppable" area of parcel (acres)	129.96	72.11
	(161.82)	(123.80)
Average share of annual crops	0.218	0.235
	(0.272)	(0.313)
Average share of fruit/nut perennial crops	0.449	0.304
	(0.397)	(0.375)
Average share of hay perennial crops	0.245	0.366
	(0.275)	(0.376)
Average share of non-crop (fallow)	0.089	0.095
	(0.139)	(0.163)
Average groundwater use (AF/acre per year)	4.201	
	(5.127)	
Parcel within surface water district (1/0)	0.600	0.609
	(0.490)	(0.488)

Notes: We report means and standard deviations (in parentheses) across unit-specific averages. Panel A uses monthly billing data for all PG&E agricultural electricity service points (i.e. the physical locations of electricity meters). The vast majority of our analysis focuses on the SPs in the left column, for which we observe PG&E pump audits, thereby (i) confirming that the SP operates a vertical groundwater pump, and (ii) letting us populate Equation (1) to convert from kWh to AF. The right column includes all remaining agricultural SPs, for which we cannot calculate AF/month or $/AF: we cannot distinguish groundwater pumps from other agricultural electricity end uses (e.g. refrigeration), let alone parameterize pump-specific production functions. Panel B aggregates up to the parcel level, our main unit of analysis. The left column includes our core sample for reduced-form and structural estimation: agricultural parcels that contain SPs in the left column of Panel A. The right column includes the additional parcels that contain SPs in the right column of Panel A (but not in the left column). For consistency with our reduced-form and structural analyses, Panel B: weights parcels by (time-invariant) croppable acreage; uses croppable (rather than total) acreage to denominate crop shares; omits parcels with less than 1 or greater than 5,000 croppable acres; and omits parcels with annual electricity bills exceeding $3,000 per croppable acre.

4.4 Summary statistics

Panel A of Table 1 compares our preferred sample, including *only* confirmed pumps, to all other PG&E agricultural service points. While these two groups face similar marginal electricity prices (\$0.13/kWh vs. \$0.16/kWh), confirmed pumps have much greater energy consumption (10,133 vs. 4,519 kWh/month). This is unsurprising, as groundwater pumping is far more energy-intensive than other farm end-uses. To the extent that our matching process filters out (unconfirmed) groundwater pumps that never received PG&E pump tests, our preferred sample is skewed towards larger pumps that are most important for groundwater management policy. The average confirmed pump produces 33.26 AF of groundwater per month at a marginal cost of \$47.69/AF.

Panel B aggregates from the service point-month to the parcel-year, our unit of analysis. This yields a preferred sample (left column) of 7,127 parcels containing confirmed pumps. The average in-sample parcel has 331 croppable acres, with 22%, 45%, 25%, and 9% of acres in annuals, fruit/nut perennials, hay perennials, and non-crop, respectively. Compared to parcels in the right column containing other agricultural users (i.e., parcels where electricity is not necessarily used for irrigation), our sample is selected towards larger parcels with more fruit/nut perennials and less hay perennials. Across all four categories, in-sample parcels use an average of 4.2 AF of groundwater per acre per year. This average aligns with irrigation budgets in agronomic studies.[20]

5 Reduced-form estimation and results

In this section, we use a panel fixed effects approach to measure farmer responses to groundwater pumping costs on an annual time scale. We present two sets of reduced-form results: the effects of year-over-year cost shocks on groundwater and electricity consumption, and the impacts of these same cost shocks on crop choice.

This section serves four purposes. First, we introduce our variation in default electricity price—which also underpins our dynamic discrete choice model—and argue that it

20. For example, almond orchards in California's San Joaquin Valley are estimated to require 3–5 AF per acre per year (https://coststudies.ucdavis.edu/current/commodities/almonds). We cannot calculate a comparable statistic for parcels containing (unconfirmed) pump(s) that PG&E did not test.

is plausibly exogenous. Second, we show a proof-of-concept that farmers indeed respond to changes in groundwater costs by altering their pumping behavior. Third, we demonstrate that year-over-year groundwater cost shocks do not induce crop switching, suggesting that quantifying cropping responses to *persistent* cost shocks requires a dynamic model. Fourth, we generate two key inputs into our structural model: the short- and long-run "intensive-margin" elasticities of groundwater demand with respect to groundwater cost conditional on crop choice.

5.1 Groundwater responses to cost shocks

We estimate the effect of cost shocks on groundwater use via the following two-stage least squares specification:

$$\log\left(Q_{it}^{\text{water}}\right) = \gamma \widehat{\log\left(P_{it}^{\text{water}}\right)} + \psi_i + \delta_t + \epsilon_{it} \tag{2}$$

$$\log\left(P_{it}^{\text{water}}\right) = \theta \log\left(P_{it}^{\text{elecDefault}}\right) + \psi_i + \delta_t + \nu_{it} \tag{3}$$

The outcome variable is the natural logarithm of groundwater extracted at parcel i in year t (Q_{it}^{water}). The explanatory variable is the natural logarithm of the marginal cost of groundwater (P_{it}^{water}). ψ_i and δ_t are a set of cross-sectional and time fixed effects, which we describe below. ϵ_{it} and ν_{it} are idiosyncratic errors, which we two-way cluster by parcel and county-by-year. We weight these regressions by each parcel's croppable acreage (excluding forests, development, etc.), making our estimates representative per acre of cropland.[21]

We instrument for groundwater costs using default marginal electricity prices ($P_{it}^{\text{elecDefault}}$). This leverages plausibly exogenous variation in PG&E's tariff schedules, which are predetermined, are geographically uniform, and assign farms to restrictive tariff categories based on pump size and meter type (as we discuss in Section 4.1). This instrument eliminates multiple sources of potential bias in $\hat{\gamma}$. First, it eliminates selection bias driven by farmers' within-category tariff choices (e.g., choosing the most advantageous tariff within a given category). Second, by predicting P_{it}^{water} using only electricity price variation, it removes simultaneity bias that arises when a pump's own extraction (Q_{it}^{water}) locally displaces groundwater, thereby

21. Appendix Table B9 shows that our estimates attenuate slightly (though retain statistical significance) when we omit these weights, suggesting that larger parcels tend to be more groundwater-cost-responsive.

Figure 3: Marginal electricity prices for four default tariffs

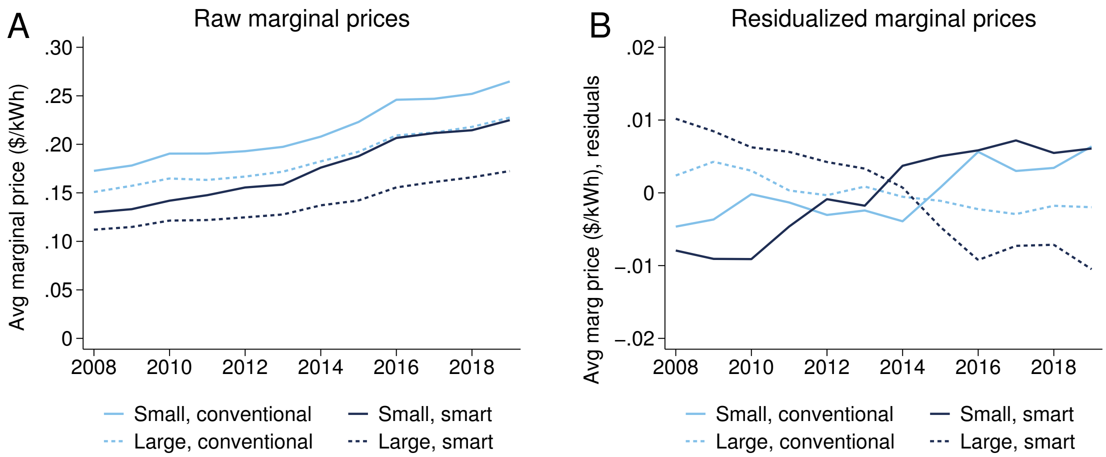

Notes: This figure plots time series of annual average marginal electricity prices ($/kWh) for PG&E's four default agricultural tariffs. Panel A plots raw average marginal prices for each year in our estimation sample, taking unweighted averages across all hours (marginal prices are higher in summer months and on weekdays). Panel B plots residuals of these four time series after partialing out tariff and year fixed effects (aligning with the fixed effects we use when estimating Equations (2)–(3)). The four tariff categories are defined by customers' physical capital: small ($<$ 35 hp) vs. large ($\geq$ 35 hp) pumps, and conventional vs. smart meters. Our identifying variation comes from differential price changes across default tariffs, as well as PG&E's smart meter rollout—which exogenously shifted many customers from conventional to smart categories, lowering their marginal price.

increasing P^{water}_{it} local to parcel i. Third, it purges other potential endogenous variation in P^{water}_{it} (e.g., drought, which increases both groundwater depths and irrigation demand; differences in groundwater depth that may covary with productivity). Fourth, it eliminates attenuation bias in $\hat{\gamma}$ from measurement error in $\frac{\text{kWh}}{\text{AF}}_{it}$, which we use to construct P^{water}_{it}. Finally, it eliminates the mechanical correlation between Q^{water}_{it} and P^{water}_{it}, which are both constructed using $\frac{\text{kWh}}{\text{AF}}_{it}$: since $P^{\text{elecDefault}}_{it}$ is uncorrelated with $\frac{\text{kWh}}{\text{AF}}_{it}$ (conditional on fixed effects), using the fitted values in Equation (2) isolates $\frac{\text{kWh}}{\text{AF}}_{it}$ on the left-hand-side.

We identify γ off of differential changes in default electricity prices *across* tariff categories over time. Panel A of Figure 3 plots raw time series of $P^{\text{elecDefault}}_{it}$ during our sample period. Panel B plots these same time series partialing out tariff and year fixed effects (yielding the same residual variation in $P^{\text{elecDefault}}_{it}$ as the fixed effects in Equations (2)–(3)), illustrating the identifying variation used in both our reduced-form and structural estimation.

We include a series of fixed effects to address remaining potential confounders. Given our instrument, the main selection concern stems from farmers choosing a tariff *category*. Since PG&E controls farmers' electricity meter type, the only way farmers can influence their category is through their choice of pumping capital. We use parcel fixed effects to eliminate bias from such tariff category selection (e.g., parcel A has a larger pump than parcel B, and thus uses more electricity at a lower marginal price), and to control for other time-invariant differences across parcels. Reassuringly, we observe no bunching at PG&E's 35-horsepower cutoff, suggesting that farmers do not strategically size their pumps around PG&E's small vs. large tariff categories.[22]

A related concern is tariff category switches caused by changes in pump size. While only 4% of service points switch between PG&E's small- vs. large-pump tariff categories during our sample, such switches could be endogenous: increasing pump capacity from below to above 35 horsepower both lowers $P_{it}^{\text{elecDefault}}$ (see Panel A of Figure 3) and raises Q_{it}^{water} (since larger pumps produce more groundwater). To address this, we interact parcel fixed effects with a (potentially time-varying) indicator for having a large pump.

In addition, year fixed effects absorb time-varying conditions that are common across all parcels (e.g., crop prices). We also include groundwater-basin-by-year and water-district-by-year fixed effects, which absorb shocks common to relatively small geographic areas, including surface water availability, basin-wide groundwater depths, and productivity.

Table 2 presents our results from estimating Equations (2)–(3). A 1% increase in groundwater costs leads farmers to reduce groundwater use by 0.938% (Column (1); $p < 0.01$), demonstrating that farmers do change pumping behavior in response to groundwater cost shocks. We also estimate the analogous model for electricity (replacing Q_{it}^{water} and P_{it}^{water} with Q_{it}^{elec} and P_{it}^{elec}), which yields a nearly identical 0.899% reduction in electricity usage due to a 1% electricity cost shock (Column (2); $p < 0.01$). When we include *all* parcels with agricultural users (i.e., both columns of Panel B of Table 1, rather than just the left column), we find a similar 0.734% reduction due to a 1% electricity cost shock (Column (3); $p < 0.01$). Even though our sample of confirmed pumps differs from other agricultural users, comparing Columns (2) vs. (3) reveals that this selection does not create a meaningful difference in how farmers respond to cost shocks.

While our instrument and fixed effects eliminate major threats to identification, we also address a series of remaining possible concerns.

22. See Appendix Figure C2. Appendix Table B4 shows that our results are similar if we interact month-of-sample fixed effects with deciles of horsepower.

Table 2: Groundwater use responds to annual variation in pumping costs

	(1) $\log(Q^{\text{water}})$	(2) $\log(Q^{\text{elec}})$	(3) $\log(Q^{\text{elec}})$
$\log(P^{\text{water}}$ (\$/AF))	-0.938^{***}		
	(0.220)		
$\log(P^{\text{elec}}$ (\$/kWh))		-0.899^{***}	-0.734^{***}
		(0.224)	(0.092)
Include all parcels with agricultural SPs			Yes
Parcel units	7,104	7,104	46,563
County-years	367	367	506
Parcel-year observations	60,490	60,490	391,707
First-stage estimates			
$\log(P^{\text{elecDefault}}$ (\$/kWh))	1.418^{***}	1.284^{***}	1.180^{***}
	(0.043)	(0.024)	(0.013)
Kleibergen-Paap F-statistic	1094	1494	8575
OLS estimates			
$\log(P^{\text{water}}$ (\$/AF))	-0.819^{***}		
	(0.094)		
$\log(P^{\text{elec}}$ (\$/kWh))		-0.735^{***}	-0.831^{***}
		(0.133)	(0.079)

Notes: Column (1) estimates Equations (2)–(3) at the parcel-year level using two-stage least squares. The outcome variable is the natural logarithm of groundwater consumption; we instrument for pumping costs with default electricity prices. This recovers an estimate of the cost elasticity of groundwater demand. Column (2) uses the quantity and price of electricity (which we observe directly), rather than the quantity and price of groundwater (which we construct). Columns (1)–(2) use the sample of parcels containing confirmed pumps (i.e., Table 1, Panel B, left column). Column (3) is analogous to Column (2), expanding the sample to include all parcels containing PG&E agricultural customers (i.e., Table 1, Panel B, both columns). Regressions include the following fixed effects: parcel, parcel $\times$ $\mathbf{1}$[large pump] (to capture tariff category switches), year, groundwater basin $\times$ year (to capture trends in depth), and water district $\times$ year (to capture changes in surface water availability). Regressions are weighted by each parcel's "croppable" acreage (excluding forests, development, etc.). Standard errors (in parentheses) are two-way clustered by parcel and county-year. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Unobserved pump upgrades Higher groundwater costs could induce farmers to invest in pump efficiency improvements. Since we do not observe groundwater use directly (instead inferring it from electricity consumption), we could mistake these efficiency upgrades for groundwater pumping reductions (rather than reductions in the amount of electricity needed to pump an acre-foot of water). Appendix Table B1 addresses this concern by restricting the sample to parcel-years proximate to an observed pump test, thereby minimizing the possibility of bias from unobserved efficiency improvements. If anything, this slightly increases the magnitudes of our estimates, suggesting that the pump efficiency channel is not driving our results by biasing our construction of Q_{it}^{water}.

kWh-to-AF conversions As we discuss in Section 4.2, converting from electricity to groundwater introduces multiple sources of measurement error. Our instrument, $P_{it}^{\text{elecDefault}}$, eliminates these sources of measurement error, since electricity tariffs are not correlated with the timing and specifics of pump audits or groundwater depth readings. Still, Appendix Tables B2 and B3 show that our estimates are robust to: alternative parameterizations of lift in Equation (1), omitting pumps with questionable audit data, and omitting pumps far from contemporaneous groundwater readings.

Smart meter rollout During our sample period, PG&E gradually replaced remaining conventional (analog) meters with smart (digital) meters. As part of this broad smart meter rollout, 21% of the service points in our sample were switched from conventional- to smart-meter categories. Since the rollout's timing reflected institutional factors outside of farmers' control, meter-induced category switches provide additional plausibly exogenous variation in $P_{it}^{\text{elecDefault}}$.[23] To assuage any concerns about endogenous meter upgrades, Appendix Tables B4 and B5 show that our estimates are robust to controlling for differential time trends in covariates that could have predicted the smart meter rollout.

Weather realizations PG&E's tariff schedule should be uncorrelated with climatic conditions, especially after controlling for parcel, year, basin-by-year, and water-district-by-year fixed effects. Appendix Table B6 confirms that our estimates are robust to controlling for local weather realizations and drought severity, assuaging any concerns that correlations between weather and electricity prices could be generating omitted variables bias.

Cone of depression spillovers Cones of depression form when extraction from a well temporarily removes groundwater from the surrounding areas of the aquifer (Alley, Reilly, and Franke (1999)). These between-well spillovers could potentially violate our exclusion restriction: if parcel i experiences the same shock to default electricity prices as its neighboring parcel j, and if j responds by altering their pumping behavior, the resulting change in parcel j's cone of depression could impact parcel i's groundwater depth—in turn impacting parcel

23. The rollout affected both agricultural and non-agricultural customers. Previous research has established that PG&E did not design their smart meter rollout to target customers with particular usage patterns (Blonz (2022)). Since farmers could not influence the timing of their meter upgrades, it is highly unlikely that these upgrades are systematically correlated with unobserved changes in pumping behavior.

i's pumping behavior. Appendix Table B7 explicitly models such spillovers, by controlling for average default electricity prices of parcel i's neighbors. This does not meaningfully alter our results, assuaging concerns of bias due to localized spillovers.

5.2 Groundwater responses and crop choice

Having shown that farmers respond to groundwater cost shocks, we next estimate the extent to which they respond on the *intensive margin*—that is, holding constant their existing crop category (i.e., annuals, fruit/nut perennials, hay perennials, or non-crop). This intensive-margin elasticity is a key input to our dynamic structural model of crop choice, since counterfactual groundwater costs will alter *both* farmers' crop choices *and* their irrigation behavior conditional on crop choices.

To estimate the short-run intensive-margin elasticity, we restrict our sample to parcels that chose the same crop category as the preceding year, while also adding a separate parcel fixed effect for each continuous "cropping streak" (i.e., consecutive years when the parcel chose the same crop category) to control for irrigation needs specific to that particular cropping period. We report these results in Columns (1)–(2) of Table 3. Our conditional point estimates of -0.944 for groundwater and -0.906 for electricity (both $p < 0.01$) are nearly identical to our estimates of the unconditional responses shown in Table 2. This indicates that farmers' responses to year-on-year cost shocks are driven by the intensive margin, rather than the crop switching margin.[24]

Columns (3)–(6) of Table 3 provide direct reduced-form tests of this crop switching margin, by replacing the dependent variable in Equation (2) with the share of parcel i's (croppable) acres allocated to a specific crop category in year t. These regressions corroborate our intensive-margin results: for all four crop types, we recover precise null estimates of the effect of year-on-year groundwater costs shocks on crop switching. Together, they suggest that farmers are unlikely to switch crops in response to short-run cost shocks.

These short-run estimates reflect how farmers respond to the year-on-year variation in pumping costs that we observe. We also seek to quantify how farmers would respond to a

24. Appendix Table B10 estimates heterogeneous intensive-margin elasticities for the four crop categories. We cannot statistically reject that these crop-specific elasticities are identical.

Table 3: Short-run responses do not reflect crop switching

	(1) log(Q^{water})	(2) log(Q^{elec})	(3) Share annuals	(4) Share perennials	(5) Share hay	(6) Share non-crop
$\log(P^{\text{water}}$ (\$/AF))	-0.944^{***}		0.002	-0.003	-0.004	0.005
	(0.270)		(0.017)	(0.020)	(0.019)	(0.017)
$\log(P^{\text{elec}}$ (\$/kWh))		-0.906^{***}				
		(0.264)				
Intensive margin	Yes	Yes				
Parcel units	6,997	6,997	7,131	7,131	7,131	7,131
County-years	334	334	367	367	367	367
Parcel-year observations	46,202	46,202	61,910	61,910	61,910	61,910
First-stage estimates						
$\log(P^{\text{elecDefault}}$ (\$/kWh))	1.274^{***}	1.289^{***}	1.425^{***}	1.425^{***}	1.425^{***}	1.425^{***}
	(0.056)	(0.044)	(0.043)	(0.043)	(0.043)	(0.043)
Kleibergen-Paap F-statistic	523	847	1092	1092	1092	1092
OLS estimates						
$\log(P^{\text{water}}$ (\$/AF))	-0.779^{***}		-0.009	-0.008	0.001	0.016^{**}
	(0.121)		(0.010)	(0.012)	(0.010)	(0.008)
$\log(P^{\text{elec}}$ (\$/kWh))		-0.652^{***}				
		(0.155)				

Notes: Columns (1)–(2) are identical to Columns (1)–(2) of Table 2, except that they capture the intensive-margin response within "cropping streaks". We define a parcel's cropping streaks based on the consecutive years that the parcel chose a specific crop category (annuals, fruit/nut perennials, hay perennials, non-crop). For example, if a parcel's modal category was fruit/nut perennials in 2008–13 and annuals in 2014–19, this parcel had two cropping streaks. Columns (1)–(2) restrict the sample to parcel-years in at least the second year of a cropping streak, and also include a separate parcel fixed effect for each cropping streak. This isolates the within-crop intensive margin, shutting down the crop-switching channel that we structurally estimate below. Columns (3)–(6) are identical to Column (1) of Table 2, except that the outcome variables are the share of acres in a crop category for each parcel-year. All regressions include the following fixed effects: parcel, parcel × $\mathbf{1}$[large pump] (to capture tariff category switches), year, groundwater basin × year (to capture trends in depth), and water district × year (to capture varying surface water availability). Regressions are weighted by each parcel's "croppable" acreage (excluding forests, development, etc.). Standard errors (in parentheses) are two-way clustered by parcel and county-year. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

permanent cost shock, such as California's SGMA policy. Since we lack permanent quasi-random variation, we turn to a structural model. Our dynamic discrete choice model uses short-run cost variation to recover long-run responses and lets us simulate counterfactual policies to achieve SGMA targets.

A key input to our dynamic model is a *long-run* intensive-margin elasticity analogous to the short-run elasticity in Column (1) of Table 3. To calibrate this long-run parameter, we estimate a series of intensive-margin regressions that incrementally restrict the sample

Table 4: Intensive-margin response degrades after 5 years in the same crop category

	log(Q^{water})					
	(1)	(2)	(3)	(4)	(5)	(6)
log (P^{water} (\$/AF))	-0.944^{***}	-0.934^{***}	-0.946^{***}	-0.591	-0.380	-0.267
	(0.270)	(0.301)	(0.351)	(0.379)	(0.370)	(0.552)
Year of cropping streak	≥ 2	≥ 3	≥ 4	≥ 5	≥ 6	≥ 7
Parcel units	6,997	6,628	6,035	5,356	4,703	4,084
County-years	334	302	271	240	209	178
Parcel-year observations	46,202	36,492	29,548	23,881	19,024	15,004
First-stage estimates						
log ($P^{\text{elecDefault}}$ (\$/kWh))	1.274^{***}	1.214^{***}	1.185^{***}	1.178^{***}	1.176^{***}	1.172^{***}
	(0.056)	(0.059)	(0.066)	(0.073)	(0.089)	(0.095)
Kleibergen-Paap F-statistic	523	424	321	259	174	152

Notes: These regressions estimate the persistence of intensive-margin groundwater response within "cropping streaks". We define a parcel's cropping streaks based on the consecutive years that the parcel chose a specific crop category (annuals, fruit/nut perennials, hay perennials, non-crop). For example, if a parcel's modal category was fruit/nut perennials in 2008–13 and annuals in 2014–19, this parcel had two cropping streaks. Column (1) is identical to Column (1) of Table 3, restricting the sample to parcel-years in at least the second year of a cropping streak, and also including a separate parcel fixed effect for each cropping streak. This isolates the within-crop intensive margin, shutting down the crop-switching channel that we structurally estimate below. Columns (2)–(6) are identical, except that they incrementally tighten this sample restriction to parcel-years in least the third-to-seventh year of a cropping streak—thereby isolating the intensive-margin response that persists multiple years after the parcel's most recent crop switch. All regressions also include the following fixed effects: parcel, parcel $\times$ $\mathbf{1}$[large pump] (to capture tariff category switches), year, groundwater basin $\times$ year (to capture trends in depth), and water district $\times$ year (to capture varying surface water availability). Regressions are weighted by each parcel's "croppable" acreage (excluding forests, development, etc.). Standard errors (in parentheses) are two-way clustered by parcel and county-year. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

to parcel-years in at least the nth year of a consecutive cropping streak. This identifies the within-crop response that remains available to farmers even after they have grown a particular crop for at least n years. Referring back to the illustrative model in Figure 2, these regressions isolate the persistent within-crop shifts from A to A', excluding the short-run shifts from B to B'.

Table 4 reports these estimates, revealing that the short-run intensive-margin elasticity of -0.944 falls to -0.591 at five years, -0.380 at six years, and -0.267 at seven years in the same crop category. This indicates that the long-run intensive-margin elasticity is meaningfully smaller than its short-run analog—likely because short-run strategies (e.g., ceasing production mid-season, as shown in Figure 2) are not sustainable over multiple years. These estimates lose statistical power due to decreasing sample size, especially in Column (6) where removing two-thirds of observations to isolate "≥ 7" streaks nearly doubles the standard error. Thus, our preferred long-run intensive-margin elasticity is the "≥ 6" estimate of -0.380 from Column (5), which is approximately 40% the size of the short-run intensive-margin elasticity.

6 Structural estimation and results

We specify a dynamic discrete choice model of farmers' cropping decisions that captures two key features of our setting. First, since many California farmers make long-run investments in perennial crops, we incorporate state dependence in annual cropping decisions. Second, because SGMA introduces permanent changes to groundwater policy, we let farmers' annual decisions reflect rational forward-looking expectations.

We use a conditional choice probability approach (following Scott (2013); Kalouptsidi, Scott, and Souza-Rodrigues (2021b)) to estimate model parameters without the need to specify the evolution of individual market-level states. Three notable differences are that we identify the model using a plausibly exogenous groundwater cost shifter (changes in regulated electricity prices) as an instrument, we quantify the impact of groundwater costs rather than expected revenues, and we estimate the model without first-differencing.[25]

We then use our estimated dynamic model to simulate responses to both temporary and permanent counterfactual groundwater taxes. These simulations generate short- and long-run (semi-)elasticities of crop choice, groundwater use, and electricity use with respect to the marginal cost of groundwater. Finally, we simulate a policy counterfactual that achieves the groundwater reduction targets required by California's SGMA legislation.

6.1 Model of crop choice

We model annual profits on a given field as a function of crop choice with crop-specific groundwater pumping costs.[26] Each year, a farmer chooses a crop $c \in \mathcal{C} = \{$annuals, fruit/nut perennials, hay perennials, non-crop$\}$ to maximize expected discounted profits

25. We use groundwater costs, rather than expected returns (as in Scott (2013)), as the driver of profits in our model because (i) we have exogenous variation in groundwater costs and (ii) California's rich landscape of specialty crops makes it difficult to form accurate predictions of returns in our setting. Our instrumental variables strategy obviates the need to first-difference for identification.

26. Our structural model abstracts away from the entry/exit decision of drilling or abandoning wells. This assumption is consistent with the institutional setting. Between 2007–2008 and 2017–2018, total irrigated

over an infinite time horizon.[27] Profits from crop choice c depend on two state variables: the field state and the market state. The field state $k \in \mathcal{K}$ represents field-level characteristics at the start of a growing season, which depend on past cropping decisions. The market state $\omega \in \Omega$ is the set of market-level variables that affect the expected profitability of each crop, such as input prices, output demand, government policies, and widespread weather events (e.g. drought, which affects surface water allocations). The market state is known to all farmers but is not fully observed by the econometrician.

Assumption 1: Profit function Annual profits on a given field in year t depend on crop choice c_t, field state k_t, market state ω_t, and a vector of idiosyncratic shocks ε_t. We define the profit function as:

$$\pi(c_t, k_t, \omega_t, \varepsilon_t) = \alpha_G G(c_t, \omega_t) + \alpha(c_t, k_t) + \xi(c_t, k_t, \omega_t) + \varepsilon_{ct} \tag{4}$$

$G(c_t, \omega_t)$ is the total variable cost of groundwater pumping, which depends on the water requirements of crop c_t and the market state ω_t; we estimate the parameter α_G.[28] Because farmers know the market state, we assume they know this cost of groundwater pumping for each crop type when making their crop choice. $\alpha(c_t, k_t)$ represents the time-invariant component of average net returns to growing crop c_t, excluding groundwater costs and net of the costs of transitioning from field state k_t to crop c_t; we estimate these parameters.

acreage remained largely constant in California, falling by just 2.3%, while the number of farms with irrigation wells rose by just 1.4% (USDA NASS (2019)). Our model also holds pumping capital fixed at observed levels.

27. We aggregate crops into these four broad categories both for model tractability and to avoid concerns about measurement error in the CDL. Fruit/nut crops and hay crops are both perennials, but have different cost structures: whereas hay requires low upfront costs and can be harvested soon after planting, orchards and vineyards require high upfront costs and take longer to reach maturity. We use "fruit/nut" to refer to perennial fruit and nut crops, and "hay" to refer to perennial hay crops (e.g., alfalfa).

28. We use total variable costs, rather than total costs, because fixed fees on electricity bills are crop-choice invariant. Assuming Q^{water} responds only on the extensive (crop choice) margin, a percent change in marginal cost (measured in \$/AF) is equivalent to the same percent change in total variable cost (measured in \$), meaning (semi-)elasticities with respect to both cost measures are identical. If Q^{water} also responds on the intensive margin, a percent change in marginal cost yields a smaller percent change in total variable cost. Our simulations below consider both margins, necessitating this marginal vs. variable cost distinction. Note that surface water availability is a component of the market state, ω_t. Equation (4) does not include a surface water cost term, consistent with surface water being largely inframarginal in this setting (see Appendix Table B11; Ferguson (2025)). Instead, surface water costs are captured as part of $\alpha(c_t, k_t)$ and $\xi(c_t, k_t, \omega_t)$. If surface water is indeed inframarginal and therefore does not differ by crop choice, then these components are differenced out in our Euler equation (Equation (7)). If surface water use does differ with crop choice, such differences are contained in the composite error term; in that case, our groundwater cost-shift instrument, which is uncorrelated with surface water use, yields unbiased estimates of α_G.

$\xi(c_t, k_t, \omega_t)$ represents the time-varying component of average net returns to crop c_t in field state k_t, which depends on the market state ω_t. Finally, ε_{ct} is an idiosyncratic shock to profits for crop c_t in year t, which we assume is independent and identically distributed Type-I extreme value; we denote the joint distribution of the vector ε_t as $F^{\varepsilon}(\varepsilon_t)$.

Assumption 2: State dependence and renewal actions Crop choice dynamics enter through the transition cost component of $\alpha(c_t, k_t)$.[29] Accounting for state dependence is essential given California's abundance of perennial crops, which are harvested across multiple years from a single planting. Growing a perennial crop in consecutive years incurs much lower costs than switching *into* the same perennial crop because the latter comes with an upfront investment cost.

We assume the field state depends on crop choice over the preceding two years. For annuals, hay perennials, or non-crop, one year in that crop type is sufficient to establish the field state. Fruit/nut perennials (FNP), however, typically need multiple years to mature and produce output. Hence, we distinguish between a young fruit/nut perennial crop that was first planted in the prior year and a mature fruit/nut perennial crop that was planted two or more years prior.[30] Formally, we define the field state in year t as:

$$k_t(c_{t-1}, c_{t-2}) = \begin{cases} c_{t-1} & \text{if } c_{t-1} \neq FNP \\ FNP_{young} & \text{if } c_{t-1} = FNP \text{ and } c_{t-2} \neq FNP \\ FNP_{mature} & \text{if } c_{t-1} = c_{t-2} = FNP \end{cases}$$

Any crop choice other than fruit/nut perennials is a "renewal action," meaning that choice $c_t \in \{\text{annuals, hay perennials, non-crop}\}$ will yield a particular field state in the following year k_{t+1} regardless of states in prior years (Kalouptsidi, Scott, and Souza-Rodrigues (2021b)).[31]

29. As Scott (2013) discusses, it is common for dynamic incentives to enter only through an intercept term.

30. While abstracting from the growth profiles of specific perennial crops (e.g., almonds typically reach maturity faster than grapes), this assumption is sufficient to generate distinct responses to temporary vs. permanent cost shocks. Assuming even a single year of field state dependence introduces switching costs that make it unprofitable for farmers to (unrealistically) tear out their almond orchards in response to a one-year cost shock and then replant almonds the following year. Distinguishing between young and mature fruit/nut perennials allows annual returns to vary while the crop is being established, further strengthening the disincentive to switch in or out of almonds in response to temporary shocks.

31. For estimation, we rely on fallowing being a renewal action: choosing c_t = non-crop effectively resets the transition costs in the following year, regardless of the cropping history.

Assumption 3: Small fields We assume that the market state ω_t evolves following a Markov process that is independent of the crop choice on any single field. That is, the distribution of the market state, $F^\omega(\omega_t)$, satisfies $F^\omega(\omega_{t+1} \mid c_t, \omega_t) = F^\omega(\omega_{t+1} \mid \omega_t)$ for all c_t on each field. This assumption implies that fields are small relative to the size of their market, causing farmers to treat ω_t as exogenous. Following from this assumption, we treat each field as independent. This abstracts from any potential externalities due to the common-pool nature of the groundwater resource—which improves model tractability and aligns with our finding that between-well spillovers do not alter contemporaneous groundwater pumping in this setting (see Appendix Table B7). It also implies that if a landowner operates multiple fields, maximizing profits jointly across fields is equivalent to maximizing profits for each field independently (following Scott (2013) in abstracting from any production complementarities).

Value function and conditional choice probabilities Under Equation (4), the expected discounted stream of future profits from a given field is given by the value function:

$$V(k_t, \omega_t, \varepsilon_t) = \max_{c_t \in \mathcal{C}} \{\pi(c_t, k_t, \omega_t, \varepsilon_t) + \beta E\left[V(k_{t+1}, \omega_{t+1}, \varepsilon_{t+1}) \mid c_t, \omega_t\right]\} \tag{5}$$

We assume the common discount factor $\beta = 0.9$ following the literature (e.g., Scott (2013); Hsiao (Forthcoming)). The resulting conditional choice probabilities (CCPs), or the probability that the farmer chooses crop c_t conditional on being in field state k_t, are:

$$p(c_t, k_t, \omega_t) = \frac{\exp\left[v(c_t, k_t, \omega_t)\right]}{\sum_{c'_t \in \mathcal{C}} \exp\left[v(c'_t, k_t, \omega_t)\right]} \tag{6}$$

where $v(c_t, k_t, \omega_t)$ gives the conditional value of selecting crop choice c_t in field state k_t, which follows from the value function in Equation (5).[32] This expression emphasizes that CCPs contain information about the relative value of making different crop choices.

Euler equation To generate an estimating equation, we consider the comparison between two crop choices in year t: c_t and c'_t. Then, suppose crop choice in year $t+1$ is a renewal action (i.e., any choice other than fruit/nut perennials, by Assumption 2), which we denote

32. See Appendix A.1 for mathematical definitions of the *ex ante* and conditional value functions.

r_{t+1}. It follows that the field state in year $t+2$ depends only on r_{t+1}—not on c_t. This means that after choosing r_{t+1} in year $t+1$, continuation values in year $t+2$ will be the same regardless of whether a farmer chooses c_t or c'_t in year t. This comparison produces an Euler equation that can be written as:[33]

$$\ln\left[\frac{p(c_t,k_t,\omega_t)}{p(c'_t,k_t,\omega_t)}\right]+\beta\ln\left[\frac{p(r_{t+1},c_t,\omega_{t+1})}{p(r_{t+1},c'_t,\omega_{t+1})}\right]=\alpha_G\underbrace{\left[G(c_t,\omega_t)-G(c'_t,\omega_t)\right]}_{\Delta G_{ct}}$$
$$+\underbrace{\alpha(c_t,k_t)-\alpha(c'_t,k_t)+\beta\left[\alpha(r_{t+1},c_t)-\alpha(r_{t+1},c'_t)\right]}_{\tilde{\Delta}\alpha_{ck}}$$
$$+\underbrace{\xi(c_t,k_t,\omega_t)-\xi(c'_t,k_t,\omega_t)+\beta\left[\xi(r_{t+1},c_t,\omega_{t+1})-\xi(r_{t+1},c'_t,\omega_{t+1})\right]}_{\tilde{\Delta}\xi_{ckt}}$$
$$+\beta\underbrace{\left[e^V(c_t,\omega_t,\omega_{t+1})-e^V(c'_t,\omega_t,\omega_{t+1})\right]}_{\Delta e^V_{ct}} \tag{7}$$

Each side of Equation (7) is equivalent to the difference in values between choosing c_t or c'_t in year t, taking renewal action r_{t+1}, and then choosing optimally in all following years. As shown in underbraces, the right-hand side of this equation can be summarized by four terms, each representing one component of this difference in values.[34]

6.2 Estimation

To empirically estimate Equation (7), we require CCPs and total variable groundwater costs for each crop in the choice set, as well as an instrument for (potential) groundwater cost endogeneity. We construct these variables using data from all fields in a "market" in which farmers face a similar choice environment—including similar transition costs, crop-specific groundwater costs, and market states. We define markets using three criteria: electricity tariff, surface water availability, and geography.[35] We construct CCPs by aggregating crop

33. The term $e^V(c_t,\omega_t,\omega_{t+1})$ is the expectational error given by the difference between expected and realized *ex ante* value functions in year $t+1$. See Appendix A.1 for a mathematical definition of the expectational error term and a derivation of this Euler equation.

34. We use Δ to denote a contemporaneous difference between c_t and c'_t, and $\tilde{\Delta}$ to denote this contemporaneous difference plus a discounted difference in year $t+1$. Equation (A3) in Appendix A.1 provides definitions for each of these terms.

35. We first split by PG&E's small- and large-pump tariff categories. For surface water availability and geography, we then group fields by water district or (if not in a water district) by county. For water districts

choices within each market, and we use observed groundwater costs in the market to construct crop-specific pumping costs.

6.2.1 Variable construction

Conditional choice probabilities We observe land use at a 30-meter resolution in the CDL. We calculate CCPs from the observed pixel-level sequence of crop choices in parcels in our sample. We aggregate pixel-level conditional choices within each market:

$$p_m(c_t, k_t, \omega_{mt}) = \frac{n_{mckt}}{\sum_{c' \in \mathcal{C}} n_{mc'kt}}$$

where n_{mckt} is the count of pixels in market m with crop c after starting in field state k in year t. The denominator is the count of all pixels in market m in field state k in year t. As in Scott (2013), we smooth CCPs over space to ensure no CCP has a value of zero or one.

Groundwater cost, groundwater use, and electricity use For every parcel-year, we observe the realized total variable cost of groundwater pumping for the chosen crop. To estimate the model, however, we also need projections of what these costs would have been for each alternative crop category.[36] We use the following OLS specification to estimate how parcel-level pumping costs correlate with within-parcel variation in cropping:

$$G_{it} = \sum_{c \neq 0} \left(\zeta_m^c F_{it}^c + \kappa_m^c F_{it}^c \cdot t \right) + \eta_i + \phi_{mt} + \iota_{it} \tag{8}$$

where G_{it} is the per-acre total variable groundwater cost for parcel i in year t (i.e., $Q_{it}^{\text{water}} \times P_{it}^{\text{water}}$). F_{it}^c is the fraction of parcel i planted with crop c in year t, omitting non-crop ($c = 0$) to avoid collinearity. η_i are parcel fixed effects, ϕ_{mt} are market-year fixed effects, and ι_{it} is an idiosyncratic error term. $\zeta_m^c + \kappa_m^c t$ recovers the average per-acre groundwater cost for crop c in market m and year t (relative to choosing non-crop), which is identified from within-parcel crop switches. These market-specific coefficients accommodate geographic variation in both

where we observe fewer than 30 parcels, we instead group by counties to ensure sufficient observations within a market. Appendix A.2 provides more details on market construction.

36. This is a measurement exercise to populate groundwater costs for each market-crop-year. It does not leverage variation in marginal pumping costs and is distinct from our reduced-form regressions in Section 5.

groundwater needs (e.g., due to surface water allocations) and irrigation needs within each crop category c (e.g., grape- vs. almond-growing regions).

This fitted regression model yields projected per-acre pumping costs for each parcel-year under each crop category. We aggregate these projections by taking an acreage-weighted median of all parcels within each market-crop-year.[37] This yields per-acre total variable groundwater costs $G_{mct} = G_m(c_t, \omega_{mt})$ at the market-crop-year level, which is a key input to Equation (7). We follow the same procedure to project comparable measures of groundwater quantities ($\hat{Q}^{\text{water}}_{mct}$) and electricity quantities ($\hat{Q}^{\text{elec}}_{mct}$), each at the market-crop-year level and measured per-acre. We also construct a time-invariant measure of crop-specific average electricity quantity in each market ($\hat{Q}^{\text{elec}}_{mc}$), which is a component of our instrument that we describe below.[38]

6.2.2 Identification

Equation (7) holds for any choice of crops c_t and c'_t in year t followed by any renewal action r_{t+1} in year $t+1$. To generate an estimable regression equation, we set both the comparison crop c'_t and the renewal action r_{t+1} to be the non-crop category (i.e., $c'_t = r_{t+1} = 0$), leaving c_t to denote any of the three other crop choices. We estimate the resulting regression equation at the market level, using the data described above and weighting markets by acreage.

The outcome variable of this regression, which we construct from our calculated CCPs, is the difference in values between a cropping sequence in which crop c is chosen in year t vs. one in which non-crop is chosen in year t. The right-hand side of the regression represents four components that comprise this difference. ΔG_{mct} is the difference in pumping costs between crop choice c and non-crop in year t. $\tilde{\Delta}\alpha_{mck}$ is a set of intercept terms capturing the difference in the present value of average net returns between the two cropping sequences. $\tilde{\Delta}\xi_{mckt}$ is an unobserved term that reflects the difference in time-varying net returns, and Δe^V_{mct} is the unobserved difference in expectational errors; their sum is the regression's

37. Using acreage-weighted means, rather than medians, yields similar results (see Appendix Figure A1).

38. We use $\hat{Q}$ to differentiate these per-acre quantity projections from the *observed* quantities Q^{water} and Q^{elec} used in our reduced-form analysis. Time-invariant $\hat{Q}^{\text{elec}}_{mc}$ takes the weighted median over all parcel-years for market m and crop c. We use analogous time-invariant measures of groundwater quantities and groundwater costs as steady-state values in our counterfactual simulations. We also aggregate separately by drought vs. non-drought years to incorporate drought expectations in our counterfactual simulations.

composite error term. Our main objects of interest are the groundwater cost parameter α_G and the intercepts $\tilde{\Delta}\alpha_{mck}$, which we use to recover the profit intercept parameters $\alpha_m(c_t, k_t)$. We cluster our standard errors at the market-by-year level to allow for correlation across contemporaneous crop choices and field states within each market.

Instrumenting for groundwater cost In order to recover consistent estimates of α_G, we require $\mathbb{E}[\tilde{\Delta}\xi_{mckt} + \Delta e^V_{mct} \mid \Delta G_{mct}] = 0$. Under rational expectations, the expectational error Δe^V_{mct} is uncorrelated with ΔG_{mct} by construction, but unobserved net returns $\tilde{\Delta}\xi_{mckt}$ may be correlated with ΔG_{mct}. For example, weather could affect groundwater pumping costs for crop c (relative to non-crop) by altering water requirements, and weather could also affect (relative) net returns by influencing output prices. As a result, we must instrument for ΔG_{mct}, the groundwater pumping costs for crop c (relative to non-crop) in market m in year t.

We construct an instrument for ΔG_{mct} that uses plausibly exogenous variation in PG&E's electricity tariff schedules, akin to our reduced-form approach in Section 5. Here, we require an instrument for *total* variable pumping cost, which we generate by combining two factors that are exogenous to annual variation in unobserved net returns: (i) time-varying but crop-invariant electricity prices and (ii) time-invariant but crop-varying electricity quantity needs. The price component is the average default electricity price in market m in year t, $P^{\text{elecDefault}}_{mt}$, which is plausibly excludable because exogenous electricity tariffs are only correlated with groundwater demand through pumping costs (as shown in Section 5.1).[39] The quantity component is the average amount of electricity needed for groundwater pumping for crop c (relative to non-crop) in market m, $\Delta\hat{Q}^{\text{elec}}_{mc}$, which is plausibly excludable because this time-invariant cross-sectional measure is uncorrelated with annual variation in the market state. In an approach similar to a shift-share, we instrument with the product of these two components: $P^{\text{elecDefault}}_{mt} \times \Delta\hat{Q}^{\text{elec}}_{mc}$. This instrument is strongly correlated with ΔG_{mct}, since changes to default electricity prices impact the variable costs of groundwater pumping. It also satisfies the exclusion restriction, since each component is plausibly excludable.

39. Our markets partition small vs. large pump categories. $P^{\text{elecDefault}}_{mt}$ collapses from four to two tariff categories, averaging over the composition of conventional and smart meters within each market-year. We assign each parcel's modal pump size before defining markets, such that $P^{\text{elecDefault}}_{mt}$ eliminates variation from any potentially endogenous switches between small- and large-pump tariff categories.

Recovering profit intercept parameters We require estimates of 20 profit intercepts for each market—one $\alpha_m(c_t, k_t)$ for each crop choice-field state pair. However, Equation (7) only includes 15 $\tilde{\Delta}\alpha_{mck}$ intercept terms for each market, since we use the non-crop category as the comparison. Recovering all 20 intercepts therefore requires additional assumptions. First, we normalize $\alpha_m(0, 0) = 0$, where both field state and crop choice are non-crop. Second, we assume that switching from crop c to non-crop costs half as much as switching from non-crop to crop c. Third, we assume there is no transition cost to remain in the same crop.[40]

6.3 Counterfactual simulations

We use the estimated model to simulate counterfactuals under short- and long-run groundwater tax scenarios. For short-run scenarios, we simulate a long-run steady state using baseline groundwater costs and then add a one-year marginal cost shock (i.e., a temporary tax). For long-run scenarios, we simulate the long-run steady state under a persistent marginal cost change (i.e., a permanent tax) that affects every year. In each scenario, we proceed as follows. First, we use Equation (4) to calculate expected annual profit—a function of that scenario's groundwater tax—for each crop choice in each state in every market. Second, we combine these profits with a fixed point algorithm to solve for the continuation values, which follow from Equation (5), for each crop choice at each state in each market. Third, we use these continuation values to calculate CCPs in each market, per Equation (6). Finally, starting from an initial distribution of field states in each market, we iteratively apply these CCPs to solve for crop choices and, therefore, groundwater and electricity use over a 20-year period.[41] To simulate these counterfactuals, we require two additional assumptions.

Assumption CF1: Drought state Drought, which is included in the market state ω_{mt}, can increase annual groundwater pumping costs both by increasing both irrigation needs (higher Q^{water}) and groundwater scarcity (higher P^{water}). To incorporate drought conditions

40. As described by Scott (2013) and Kalouptsidi, Scott, and Souza-Rodrigues (2021b), dynamic discrete choice models are typically not fully identified (Magnac and Thesmar (2002)). See Appendix A.2 for a mathematical statement of these assumptions. These normalizations do not affect the identification of our counterfactuals (Kalouptsidi, Scott, and Souza-Rodrigues (2021a)). Note that $\alpha_m(c_t, k_t)$ incorporates both net returns to crop c_t and any additional transition costs due to field state k_t. When remaining in the same crop, we assume that all costs are recurring (e.g., the yearly cost of replanting an annual crop) and are therefore captured by the net returns component.

41. We initialize the field states using each market's average distribution of field states for 2008–2019.

in our simulations, we project groundwater costs separately for drought vs. non-drought years and calculate profits under each market state. The resulting continuation values and CCPs become functions of both current drought and future expectations of drought. We assume farmers' expectation of drought in any future year is i.i.d. with probability equal to the frequency of drought in our sample.[42]

Assumption CF2: Intensive-margin elasticity Our reduced-form analysis shows that farmers respond to groundwater cost shocks on the intensive margin (i.e., adjust their water use conditional on crop choice). When estimating Equation (7), our construction of $G = Q^{\text{water}} \times P^{\text{water}}$ accounts for *factual* intensive-margin response, since observed Q^{water} has been optimized to *factual* P^{water}. However, simulating *counterfactual* P^{water} necessitates assumptions on out-of-sample within-crop reoptimization. In our temporary tax scenarios, we set this intensive-margin elasticity equal to the intensive-margin response we estimate in our reduced-form analysis: -0.944 (Column (1) of Table 3). In our permanent tax scenarios, we set this intensive-margin elasticity equal to the intensive-margin response in cropping streaks of at least six years: -0.380 (Column (5) of Table 4). This estimate—40% the magnitude of the short-run value—isolates the response that persists when remaining in the same cropping category for many years.[43]

We assume farmers know the magnitude of this intensive-margin adjustment, and hence the resulting groundwater pumping costs for each crop type, when making their crop choices. Because this intensive-margin response reduces groundwater and electricity consumption for a particular crop, the crop's total variable groundwater cost increases by less than the tax rate.[44] This additional margin of response alters expected profits for each crop under

42. California declared severe droughts in 7 of our 12 sample years (2008–2009 and 2012–2016). We also simulate alternate simulations with higher/lower probabilities of drought (see Appendix Figure A1). In simulating a drought (non-drought) year, we model all components of the market state (e.g., weather realizations, surface water allocations) as equal to the average drought (non-drought) year in our sample.

43. We find that crop-specific intensive-margin responses are not statistically different from one another (see Appendix Table B10), so our preferred scenarios apply a common intensive-margin elasticity to all crop categories. In an alternate specification, we instead use crop-specific intensive-margin elasticities and find similar results (see Appendix Figure A1). To generate crop-specific long-run elasticities, we scale the point estimates in Appendix Table B10 by 40%, the long-run to short-run intensive margin ratio from Table 4.

44. Suppose the marginal cost of pumping increases by 20% due to a tax. With an intensive-margin elasticity of -0.380, farmers reduce crop-specific groundwater use by 6.7%. As a result, the total variable cost of pumping increases by only 12%, not the full 20%.

counterfactual groundwater taxes, which subsequently alters continuation values, CCPs, and therefore counterfactual crop choices.

Counterfactual groundwater tax scenarios Our baseline scenario sets total variable groundwater costs equal to the time-invariant projection for that market and crop. For our temporary tax counterfactuals, we increase the marginal cost of pumping only in a single year of the 20-year simulation; we assume farmers do not anticipate this shock but know that it will persist for only one year. For our permanent tax counterfactuals, we increase pumping costs in all years and assume farmers know this tax will persist permanently. Since we use crop-specific groundwater costs, a given tax increases total variable costs relatively more for more water-intensive crops, which can induce crop switching. Further, because our dynamic model captures forward-looking behavior, temporary and permanent taxes can induce different magnitudes and patterns of crop switching.

(Semi-)elasticities Following Scott (2013), we calculate long-run (semi-)elasticities by comparing the final year of each tax scenario to the final year of the baseline scenario. For short-run (semi-)elasticities, we instead compare the tax and baseline scenarios in the year when the short-run tax occurs. The semi-elasticity of crop c with respect to the marginal cost of pumping groundwater is:

$$\frac{\sum_{m\in\mathcal{M}}\left(A'_{mc}-A_{mc}\right)}{\sum_{m\in\mathcal{M}}\sum_{c\in\mathcal{C}}A_{mc}}\Bigg/\tau$$

where A_{mc} is the acreage in market m planted to crop c in the baseline scenario, A'_{mc} is the comparable acreage in the tax scenario, and τ is the percentage tax on marginal groundwater costs. The corresponding pumping cost elasticities of groundwater and electricity are:[45]

$$\frac{\sum_{m\in\mathcal{M}}\sum_{c\in\mathcal{C}}\left(A'_{mc}\hat{Q}'_{mc}-A_{mc}\hat{Q}_{mc}\right)}{\sum_{m\in\mathcal{M}}\sum_{c\in\mathcal{C}}A_{mc}\hat{Q}_{mc}}\Bigg/\tau$$

45. Because a τ% change in electricity price translates to the same τ% change in marginal groundwater pumping costs (see Equation (1)), this expression recovers the elasticity of groundwater use with respect to marginal groundwater cost, and the elasticity of electricity use with respect to marginal electricity price.

where $\hat{Q}_{mc}$ is the time-invariant projection $\hat{Q}^{\text{water}}_{mc}$ or $\hat{Q}^{\text{elec}}_{mc}$, and $\hat{Q}'_{mc}$ is the analogous quantity after any tax-induced intensive-margin adjustments. We conduct inference on these (semi-)elasticities by taking 1,000 draws from the sampling distribution of our estimated groundwater cost parameter α_G. For each draw, we first recover the corresponding $\alpha_m(c_t, k_t)$ parameters, and we then proceed to simulate both the baseline and tax scenarios using the same parameter draw. This sampling yields 1,000 sets of (semi-)elasticities for each tax scenario. Our reported (semi-)elasticities are the means of these distributions, and our reported 95% confidence intervals span the 2.5th and 97.5th percentiles of the distributions.

6.4 Model results

Figure 4 presents our main discrete choice results. Panel A plots semi-elasticities of land use, while Panel B presents demand elasticities for groundwater and electricity, in response to both a temporary and a permanent 20% tax on the marginal cost of groundwater pumping.[46] Panels C and D plot the corresponding time profiles of these land-use, groundwater, and electricity responses.

Short-run (semi-)elasticities From the temporary tax scenario, we recover short-run semi-elasticities with respect to marginal groundwater costs of -0.0001 (s.e. 0.0001) for annuals, -0.0004 (s.e. 0.0002) for fruit/nut perennials, -0.0007 (s.e. 0.0003) for hay perennials, and 0.001 (s.e. 0.001) for non-crop (i.e., fallowing). Farmers instead respond to the one-year tax by reducing water use on existing crops, yielding short-run elasticities of -0.792 (s.e. 0.001) for groundwater and -0.763 (s.e. 0.001) for electricity.[47]

These results imply that a short-run pumping cost shock induces only minimal crop switching, which occurs only in the year of the shock (Panel C). Instead, farmers achieve groundwater reductions by applying less water to the crops they are already growing—

46. Appendix Table A2 presents these results in tabular form. Appendix Table A1 presents the parameter estimates resulting from estimating Equation (7). Appendix Figure A1 presents robustness to parcel sample selection and aggregation to the market level.

47. These short-run groundwater and electricity elasticities are nearly identical to those calculated by applying a 20% non-marginal cost increase to Equation (2) with a marginal elasticity of $\gamma = -0.944$ or -0.906 (our estimates in Table 3). Our short-run crop semi-elasticities align with the reduced-form magnitudes in Table 3, where we find precise null cropping responses to year-on-year changes in marginal pumping costs.

Figure 4: Short- and long-run elasticities with respect to groundwater pumping cost

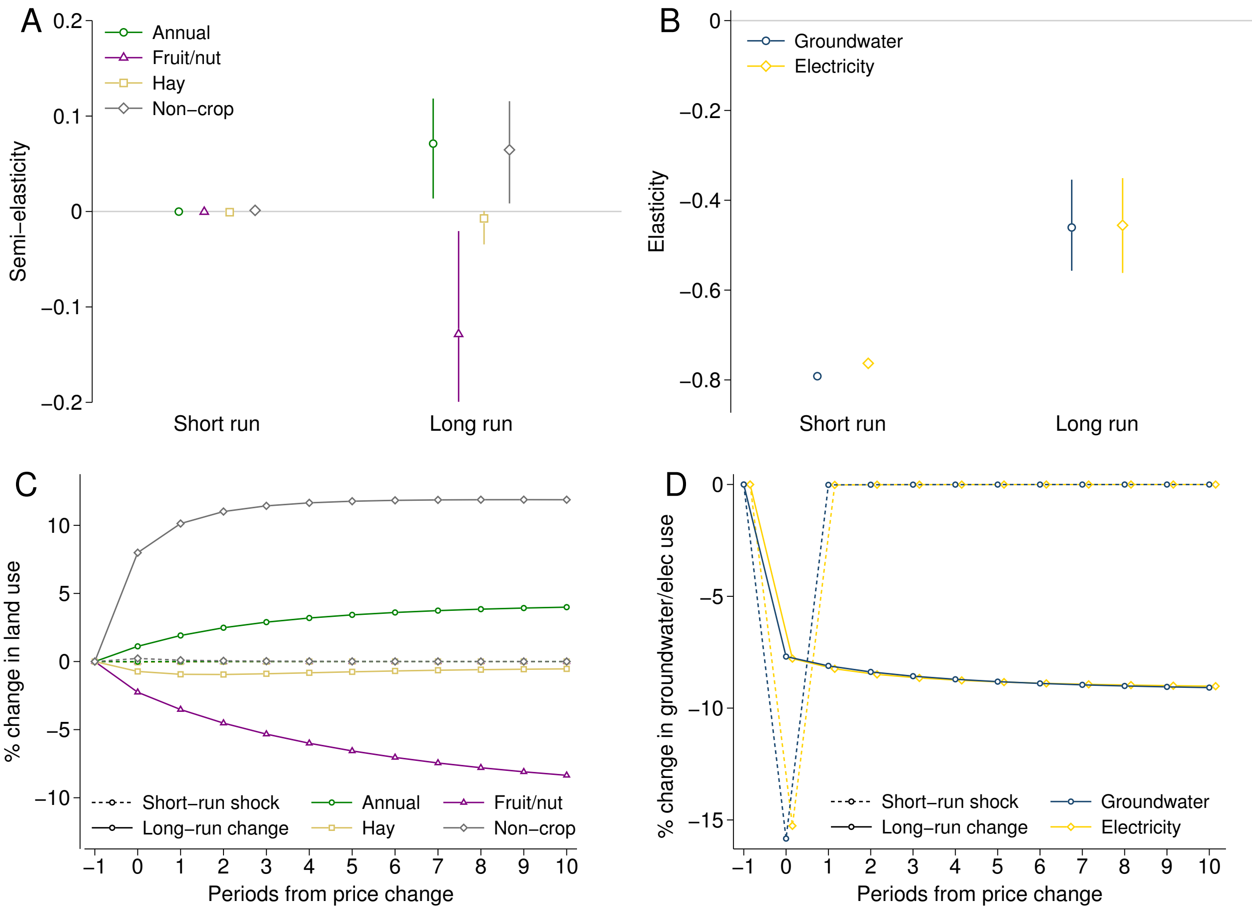

Notes: This figure plots short- and long-run (semi-)elasticities of land use (panel A) and groundwater and electricity use (panel B) with respect to marginal groundwater cost, as well as the time path of these responses (panels C and D). To recover short-run (semi-)elasticities, we simulate the model with baseline costs until it reaches a steady state, and then increase groundwater costs by 20% for one year. Farmers do not anticipate the shock, but they know it will only last for one year. To recover long-run (semi-)elasticities, we increase groundwater pumping costs by 20% and simulate the model forward. Farmers know that this price change is permanent. In both the short- and long-run, farmers can respond to groundwater cost changes on both the crop-switching margin and by reducing water use conditional on crop choice (the "intensive margin"). In the short-run, farmers' intensive-margin elasticity is -0.944 (Table 3, Column (1)). In the long-run, farmers have 40% of this intensive-margin response available to them, per our "≥ 6" cropping streak estimate of -0.380 (Table 4, Column (5)). Panel A shows semi-elasticities for annual crops (circles), hay perennials (squares), fruit/nut perennials (triangles), and non-crop (diamonds). Panel B shows elasticities of demand for electricity (diamonds) and groundwater (circles). In Panels A and B, we report the means over 1,000 draws for each model; 95% confidence intervals (vertical lines) plot the 2.5th and 97.5th percentiles over draws. In Panels C and D, dashed lines plot the time profile of responses for the same short-run scenario, for a one-year 20% cost shock in period 0. Solid lines plot the time profile of response for the same long-run scenario, for a permanent 20% cost shock starting in period 0. In each scenario, we compare outcomes to a baseline with no price change.

and these sharp reductions occur only in the year of the shock, which farmers know to be temporary (Panel D).

Long-run (semi-)elasticities When the groundwater cost change is permanent, and farmers have a smaller intensive-margin response, we find greater crop switching. On net, these effects yield smaller groundwater and electricity elasticities in the long run than in the short run. In our central case (which includes an intensive-margin elasticity of -0.380), we recover cropping semi-elasticities that are orders of magnitude larger than their short-run counterparts: 0.071 (s.e. 0.029) for annuals, -0.129 (s.e. 0.052) for fruit/nut perennials, -0.007 (s.e. 0.009) for hay perennials, and 0.065 (s.e. 0.029) for non-crop. We find overall elasticities of -0.461 (s.e. 0.058) for groundwater and -0.456 (s.e. 0.060) for electricity—showing that substitution between the intensive and extensive margins is incomplete.[48]

Unlike for a temporary cost shock, farmers' responses to a permanent cost shock persist over the long run. Panel C of Figure 4 shows large switches away from fruit/nut perennials and into annuals and fallowing in the first year of the simulation. These switching patterns continue over subsequent years, and eventually the simulation converges to a new long-run steady state with a substantially different crop mix than without the price change. Panel D of Figure 4 reveals that while the initial groundwater/electricity response is more muted than for a temporary cost shock (due to a smaller intensive-margin elasticity), this initial response persists (and even increases over time), giving the new long-run steady state.

Discussion We find that farmers switch crops only minimally in response to temporary groundwater cost shocks. Instead, their primary response to a temporary cost shock is to apply less water to existing crops, which occurs only in the year of the shock. In stark contrast, we find substantial and persistent crop switching in response to permanent cost changes, which translates into meaningful long-run reductions in groundwater and electricity use. However, due to a smaller intensive-margin elasticity over the long run, our overall long-run elasticities of groundwater and electricity demand are smaller than their short-run counterparts. These findings are consistent with our illustrative model in Section 3.

Our results build on prior work measuring agricultural groundwater demand (e.g., Pfeiffer and Lin (2014); Bruno, Jessoe, and Hanemann (2024)) by (i) estimating *both* short- and long-run elasticities in the same policy-relevant setting, (ii) showing that they are meaning-

48. Appendix Table A2 presents additional long-run (semi-)elasticities. Our "No IM" scenario assumes zero long-run intensive-margin elasticity, meaning that farmers only respond to changes in marginal groundwater costs on the crop-switching margin; our "Low IM" and "High IM" scenarios assume long-run intensive-margin elasticities of -0.267 and -0.591, respectively, per alternate cropping streak regressions (Columns (4) and (6) of Table 4). Under No IM, we find larger crop semi-elasticities, as farmers can only respond to the groundwater price increase by switching crops, but smaller groundwater and electricity elasticities, as farmers lose a margin of response. Under Low IM (High IM), crop semi-elasticities are larger (smaller) than in our Preferred IM scenario, while groundwater and electricity elasticities are smaller (larger).

fully different, and (iii) illustrating how this difference reflects mechanisms that vary by time horizon. Our results also connect to Hagerty (2022), who estimates that California farmers have a smaller response to *surface* water scarcity in the long run than in the short run (albeit a narrower long-/short-run gap than we find for groundwater). While we find crop switching only in the long run—driven by the large fixed costs of changing crops—Hagerty (2022) finds that short-run surface water shocks increase fallowing. These results are consistent with surface-water-dependent farmers being less flexible in the short run (i.e., forced to fallow) than farmers with groundwater access, but having greater flexibility to respond to long-run surface water shortages by drilling new groundwater wells (Hadachek et al. (2026)). Thus, the mechanisms underlying farmer responses to both groundwater cost increases and surface water shortages appear to differ across time horizons.

Moreover, our results illustrate how incorporating dynamics can undo the classic economic intuition that longer time horizons lead to larger elasticities (e.g., Castillo (2021); Lemoine (2024)). As with previous models showing the potential for smaller long-run elasticities than short-run elasticities (e.g., Gowrisankaran and Rysman (2012) for durable goods; Hall (1991) for labor supply), this pattern stems from differing mechanisms underlying farmers' groundwater demand response.[49] Whereas intensive-margin reductions in irrigation are not sustainable over longer time horizons, the crop-switching mechanism requires longer time horizons to justify incurring the fixed costs of uprooting or planting crops.

6.5 Impacts of counterfactual groundwater policy

Finally, we use our dynamic model to quantify the potential effectiveness of California's landmark groundwater policy, SGMA.[50] While SGMA's stringency varies substantially across Groundwater Sustainability Agencies, the overdrafted areas of our sample will require a

49. For durable goods in Gowrisankaran and Rysman (2012), the short-run mechanism is waiting for an advantageous price shock, whereas in the long run, households must decide whether or not to make a purchase at all. For labor supply in Hall (1991), short-run mechanisms involve intertemporal substitution to work more in high-wage periods, whereas long-run mechanisms relate to secular trends in labor supply.

50. In order to stem rapid aquifer drawdown, the majority of SGMA groundwater sustainability plans are proposing price-based instruments (Bruno, Hagerty, and Wardle (2022)). Our results speak directly to these plans, while providing a heuristic for stringencies of and possible responses to non-price instruments.

Figure 5: SGMA – Required groundwater taxes and farmer responses

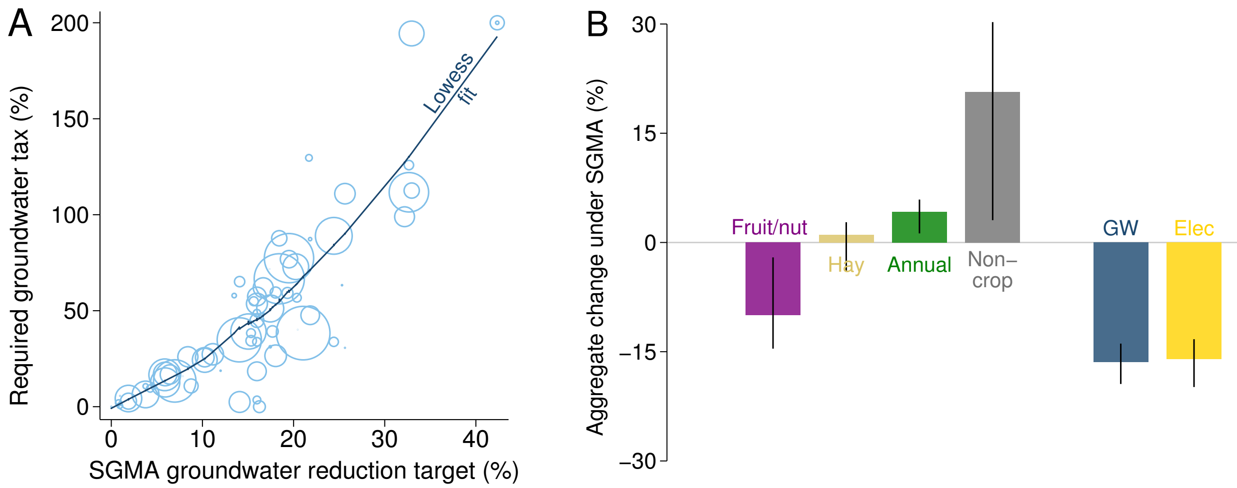

Notes: Panel A plots market-specific SGMA groundwater reduction targets (on the horizontal axis), and our estimates of the long-run groundwater tax stringency that would be required to meet these targets (on the vertical axis). The size of each dot indicates each market's total groundwater extraction in our no-tax baseline, scaled by the share of the market belonging to a SGMA-regulated area; the solid navy line is a Lowess fit, similarly weighted. Panel B plots the aggregate effect of these market-specific taxes on land use, groundwater use, and electricity use. Each bar shows the percent change in land or resource use under a SGMA counterfactual compared to baseline with no taxes. Bars report means over 1,000 draws that are common to all markets; 95% confidence intervals (vertical lines) plot the 2.5th and 97.5th percentiles over these draws.

16.7% reduction in groundwater use on average.[51] We use our model to recover the groundwater tax that would be required to achieve SGMA's spatially heterogeneous targets, and to simulate the resulting impacts on crop choice and groundwater use. To do so, for each of our 82 markets, we calculate the area of the market within overdrafted GSAs and the average reduction in groundwater pumping required by those GSAs. Then, we take our estimated dynamic model (with its market-specific crop switching costs) and, for each market, loop over groundwater tax stringencies until converging on a tax that achieves the reduction target of the market's overdrafted areas.

Panel A of Figure 5 plots these market-specific taxes against the reduction requirements in each market, for markets with groundwater pumping reduction requirements. Achieving long-run sustainability in these markets (i.e., 16.7% for the average market) will require a 52.0% groundwater tax on average. However, there is meaningful variation across markets: the 25th and 75th percentile reductions are 10.4% and 20.6%, and the 25th and

51. Note that PG&E's boundary excludes the southern parts of the Central Valley which are severely overdrafted. SGMA required *all* GSAs in medium- and high-priority basins to submit GSPs. However, 57 of the 120 GSPs in our data report *no* required groundwater pumping reductions. We only report statistics for locations with binding SGMA requirements (i.e., those that require reductions to achieve sustainability).

75th percentile required taxes are 20.4% and 76.9%.[52] Our simulated taxes rise more-than-proportionally with reduction requirements, consistent with increasing marginal costs of groundwater conservation.

Panel B shows how these taxes would impact aggregate crop choice, groundwater use, and electricity use. Our model predicts that achieving SGMA targets would lead to a 10.0% decline in fruit/nut perennials, a 1.0% increase in hay perennials (not statistically different from zero), a 4.2% increase in annuals, and a 20.7% increase in fallowing, each compared to its respective acreage in our no-tax scenario.[53] These correspond to 16.4% and 16.0% reductions in groundwater and electricity, respectively. These results illustrate that SGMA's sustainability targets are likely attainable through groundwater taxes, which in turn will lead to substantial changes in land use.

7 Conclusion

This paper estimates how agents—California farmers—respond to environmental policy—groundwater pricing—in both the short and long run. We leverage quasi-random variation in groundwater costs to estimate the elasticity of demand for groundwater over different time horizons. We provide reduced-form evidence that farmers respond to year-on-year groundwater pumping cost shocks by reducing water consumption. Since long-lived perennial tree crops are common in this setting, we build and estimate a dynamic discrete choice model in which farmers are forward-looking and fields are state-dependent. Using this model, we first recover a short-run elasticity of groundwater demand of -0.79, finding that farmers do not switch crops in response to temporary cost shocks but instead reduce irrigation holding crop fixed. We then recover a long-run elasticity of groundwater demand of -0.46, finding that permanent cost increases cause farmers to switch out of thirsty fruit/nut perennials and into annual crops and fallowing.

52. For all statistics in our SGMA policy counterfactual, we weight markets by their quantity of groundwater pumping in our no-tax baseline scenario. Appendix A.5 shows that 85% of this heterogeneity in tax stringencies is explained by policy variation in SGMA targets across markets.

53. These estimates imply that SGMA will cause 2.5% of total acres to switch into fallowing. This effect is the same order of magnitude as the land use response to water pricing estimated by by Bruno, Jessoe, and Hanemann (2024), and considerably smaller than an engineering estimate from Escriva-Bou et al. (2023) (which only allows farmers to respond by fallowing). Note that our counterfactuals only model fallowing, as opposed to exit from agriculture. We cannot model exit in a way that is identified, as it is rare in sample.

We use our dynamic model to simulate California's flagship groundwater management policy. For the average area where the policy will bind, we find that the equivalent of a 52.0% tax on groundwater pumping costs will be required to achieve the state's sustainability goals—with meaningful heterogeneity in the necessary tax stringency across locations. These results imply that California's Sustainable Groundwater Management Act will alter the landscape of crop production across California by incentivizing large shifts away from fruit and nut perennials and towards exit from agriculture. An important topic for future research will be to quantify the regulation's general equilibrium impacts: to what extent will these land use changes impact crop prices earned by farmers and food prices faced by consumers? Given that California dominates the U.S. market for fruits, nuts, and vegetables, any such price effects could have major welfare consequences.

Our work broadly underscores the importance of using dynamic models to analyze environmental and resource management policies. Our results further highlight that agents' long-run responses need not be larger than their short-run counterparts. These lessons likely apply across a broad range of settings. For example, distinguishing between short- and long-run response margins is crucial in the context of climate policy, where short-run adaptation options may not be feasible (or optimal) in the long run, and vice versa.

References

Alley, William M., Thomas E. Reilly, and O. Lehn Franke. 1999. *Sustainability of Ground-Water Resources.* U.S. Geological Survy Circular 1186.

Ayres, Andrew B., Kyle C. Meng, and Andrew J. Plantinga. 2021. "Do environmental markets improve on open access? Evidence from California groundwater rights." *Journal of Political Economy* 129 (10): 2817–2860.

Badiani, Reena, and Katrina Jessoe. 2019. "Electricity prices, groundwater, and agriculture: The environmental and agricultural impacts of electricity subsidies in India." In *Agricultural productivity and producer behavior,* edited by Wolfram Schlenker, 157–181. Chicago, IL: University of Chicago Press.

Balboni, Clare, Aaron Berman, Robin Burgess, and Benjamin A. Olken. 2023. "The economics of tropical deforestation." *Annual Review of Economics* 15:723–754.

Blakeslee, David, Ram Fishman, and Veena Srinivasan. 2020. "Way down in the hole: Adaptation to long-term water loss in rural India." *American Economic Review* 110 (1): 200–224.

Blonz, Joshua A. 2022. "Making the best of the second-best: Welfare consequences of time-varying electricity prices." *Journal of the Association of Environmental and Resource Economists* 9 (6): 1087–1126.

Blundell, Wesley, Gautam Gowrisankaran, and Ashley Langer. 2020. "Escalation of scrutiny: The gains from dynamic enforcement of environmental regulations." *American Economic Review* 110 (8): 2558–2585.

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. 2017. "California's new groundwater law and the implications for groundwater markets." *ARE Update* 20 (4): 1–4.

———. 2019. *Groundwater Policy and California Agriculture.* Technical report. University of California Cooperative Extension.

Bruno, Ellen M., and Nick Hagerty. 2025. "Anticipatory effects of regulating the commons." *Journal of Environmental Economics and Management* 133:103183.

Bruno, Ellen M., Nick Hagerty, and Arthur R. Wardle. 2022. "The political economy of groundwater management: Descriptive evidence from California." In *American agriculture, water resources, and climate change.* Chicago, IL: University of Chicago Press.

Bruno, Ellen M., and Katrina Jessoe. 2021a. "Missing markets: Evidence on agricultural groundwater demand from volumetric pricing." *Journal of Public Economics* 196:104374.

———. 2021b. "Using price elasticities of water demand to inform policy." *Annual Review of Resource Economics* 13:427–441.

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.

Buchsbaum, Jesse. 2023. "Are consumers more responsive to prices in the long run? Evidence from electricity markets." Working paper.

California Department of Food and Agriculture. 2011. *California agricultural resource directory 2010–2011.* Technical report.

———. 2020. *California County Agricultural Commissioners' Reports, Crop Year 2017–2018.* Technical Report. Prepared in cooperation with the USDA National Agricultural Statistics Service, Pacific Regional Field Office; Gross value of agricultural production by commodity group, County Agricultural Commissioners' Reports. California Department of Food, Agriculture, and USDA National Agricultural Statistics Service, March.

California Department of Water Resources. 2014. *Public update for drought response groundwater basins with potential water shortages and gaps in groundwater monitoring.* Technical report.

———. 2015. *California's groundwater update 2013.* Technical report.

———. 2017. *Best management practices for the sustainable management of groundwater: Sustainable management criteria.* Technical report.

California Energy Commission. 2005. *California's water-energy relationship.* CEC Report 700-2005-011-SF.

Carleton, Tamma, Levi Crews, and Ishan Nath. 2025. "Agriculture, trade, and global water use." Working paper.

Castillo, Juan Camilo. 2021. "Are Elasticities Greater in the Long Run? Beyond Le Chatelier's Principle." Working paper.

Caswell, Margriet F., and David Zilberman. 1986. "The effects of well depth and land quality on the choice of irrigation technology." *American Journal of Agricultural Economics* 68 (4): 798–811.

Central Valley Flood Protection Board. 2020. "Well drilling costs." http://cvfpb.ca.gov/wp-content/uploads/2020/11/8b.-EIS-Attachment-Well-Drilling-Costs.pdf.

Coman, Katharine. 1911. "Some unsettled problems of irrigation." *American Economic Review* 1 (1): 1–19.

Costello, Christopher, John Lynham, Sarah E. Lester, and Steven D. Gaines. 2010. "Economic incentives and global fisheries sustainability." *Annual Review of Resource Economics* 2:299–318.

Dahl, Carol A. 2014. "What do we know about gasoline demand elasticities?" Working paper.

Davis Cost Studies. 2014–24. University of California Davis, https://coststudies.ucdavis.edu/current/commodities.

Deryugina, Tatyana, Alexander MacKay, and Julian Reif. 2020. "The long-run dynamics of electricity demand: Evidence from municipal aggregation." *American Economic Journal: Applied Economics* 12 (1): 86–114.

Dinar, Ariel. 1994. "Impact of energy cost and water resource availability on agriculture and ground water quality in California." *Resource and Energy Economics* 16 (1): 47–66.

Drysdale, Krystal M., and Nathan P. Hendricks. 2018. "Adaptation to an irrigation water restriction imposed through local governance." *Journal of Environmental Economics and Management* 91:150–165.

Dubin, Jeffrey A., and Daniel L. McFadden. 1984. "An econometric analysis of residential electric appliance holdings and consumption." *Econometrica* 52 (2): 345–362.

Escriva-Bou, Alvar, Ellen Hanak, Spencer Cole, and Josué Medellín-Azuara. 2023. *The future of agriculture in the San Joaquin Valley.* Technical report. Public Policy Institute of California.

Fan, Y., H. Li, and G. Miguez-Macho. 2013. "Global patterns of groundwater table depth." *Science* 339 (6122): 940–943.

Feehan, James. 2018. "The long-run price elasticity of residential demand for electricity: Results from a natural experiment." *Utility Policy* 51:12–17.

Ferguson, Billy. 2025. "Trade frictions in surface water markets." Working paper.

French, Ben C., and Jim L. Matthews. 1971. "A supply response model for perennial crops." *American Journal of Agricultural Economics* 53 (3): 478–490.

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

Gowrisankaran, Gautam, and Marc Rysman. 2012. "Dynamics of consumer demand for new durable goods." *Journal of Political Economy* 120 (6): 1173–1219.

Hadachek, Jeffrey, Ellen M. Bruno, Nick Hagerty, and Katrina Jessoe. 2026. "External costs of climate adaptation in common-pool groundwater resources." Working paper.

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

———. 2025. "Transaction costs and the gains from trade in water markets." Working paper.

Hall, Robert E. 1991. "Substitution over time in consumption and work." In *Value and Capital: Fifty Years Later,* 239–267. London: Palgrave Macmillan.

Hendel, Igor, and Aviv Nevo. 2006. "Measuring the implications of sales and consumer inventory behavior." *Econometrica* 74 (6): 1637–1673.

Hendricks, Nathan P., and Jeffrey M. Peterson. 2012. "Fixed effects estimation of the intensive and extensive margins of irrigation water." *Journal of Agricultural and Resource Economics* 37 (1): 1–19.

Hornbeck, Richard, and Pinar Keskin. 2014. "The historically evolving impact of the Ogallala Aquifer: Agricultural adaptation to groundwater and drought." *American Economic Journal: Applied Economics* 6 (1): 190–219.

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

Hsiao, Allan. Forthcoming. "Coordination and commitment in international climate action: Evidence from palm oil." *Econometrica.*

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.

Ito, Koichiro. 2014. "Do consumers respond to marginal or average price? Evidence from nonlinear electricity pricing." *American Economic Review* 104 (2): 537–563.

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

Johnson, Renee, and Betsy A. Cody. 2015. *California agricultural production and irrigated water use.* Congressional Research Service Report R44093.

Kalouptsidi, Myrto, Paul T. Scott, and Eduardo Souza-Rodrigues. 2021a. "Identification of counterfactuals in dynamic discrete choice models." *Quantitative Economics* 12 (2): 351–403.

———. 2021b. "Linear IV regression estimators for structural dynamic discrete choice." *Journal of Econometrics* 222 (1): 778–804.

Konikow, Leonard F. 2013. *Groundwater depletion in the United States (1900–2008).* Scientific Investigations Report 2013-5079. United States Geological Survey.

Lemoine, Derek. 2024. "Estimating the consequences of climate change from variation in weather." Working paper.

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.

Magnac, Thierry, and David Thesmar. 2002. "Identifying dynamic discrete decision processes." *Econometrica* 70 (2): 801–816.

Manning, Dale T., Christopher G. Goemans, and Alexander Maas. 2017. "Producer responses to surface water availability and implications for climate change adaptation." *Land Economics* 93 (4): 631–653.

Marcheva, Antonia. 2025. "The benefits and costs of agricultural adaptation to surface water scarcity." Working paper.

Mieno, Taro, and Nicholas Brozovic. 2017. "Price elasticity of groundwater demand: Attenuation and amplification bias due to incomplete information." *American Journal of Agricultural Economics* 99 (2): 401–426.

Nerlove, Marc. 1956. "Estimates of the elasticities of supply of selected agricultural commodities." *Journal of Farm Economics* 38 (2): 469–509.

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

Olmstead, Sheila M. 2009. "Reduced-form versus structural models of water demand under nonlinear prices." *Journal of Business & Economic Statistics* 27 (1): 84–94.

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

Pfeiffer, Lisa, and C.-Y. Cynthia Lin. 2014. "The effects of energy prices on agricultural groundwater extraction from the High Plains aquifer." *American Journal of Agricultural Economics* 96 (5): 1349–1362.

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.

Rafey, Will. 2023. "Droughts, deluges, and (river) diversions: Valuing market-based water reallocation." *American Economic Review* 113 (2): 430–471.

Roberts, Michael J., and Wolfram Schlenker. 2013. "Identifying supply and demand elasticities of agricultural commodities: Implications for the US ethanol mandate." *American Economic Review* 103 (6): 2265–2295.

Rodell, M., J. S. Famiglietti, D. N. Wiese, J. T. Reager, H. K. Beaudoing, F. W. Landerer, and M. -H. Lo. 2018. "Emerging trends in global freshwater availability." *Nature* 557:651–659.

Ryan, Nicholas, and Anant Sudarshan. 2022. "Rationing the commons." *Journal of Political Economy* 130 (1): 210–257.

Sant'Anna, Marcelo. 2024. "How green is sugarcane ethanol?" *Review of Economics and Statistics* 106 (1): 202–216.

Sawyers, Gary. 2007. "A primer on California water rights." UC Davis Cooperative Extension working paper.

Scheierling, Susanne M., John B. Loomis, and Robert A. Young. 2006. "Irrigation water demand: A meta-analysis of price elasticities." *Water Resources Research* 41 (1): W01411.

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.

Scott, Paul T. 2013. "Dynamic discrete choice estimation of agricultural land use." Working paper.

Scott, Paul T., Eduardo Souza-Rodrigues, Ted Rosenbaum, and Skand Goel. 2026. "Cows and trees." Working paper.

Smith, Steven M., Krister Andersson, Kelsey C. Cody, Michael Cox, and Darren Ficklin. 2017. "Responding to a groundwater crisis: The effects of self-imposed economic incentives." *Journal of the Association of Environmental and Resource Economists* 4 (4): 985–1023.

UNESCO. 2022. *The United Nations World Water Development Report 2022: Groundwater, making the invisible visible.*

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

USDA. 2021. *Crop values: 2020 summary.* United States Department of Agriculture report.

USDA National Agricultural Statistics Service. 2019. *2017 Census of Agriculture: 2018 Irrigation and Water Management Survey.* Farm and Ranch Irrigation Survey Special Studies, Part 1. Tables 1 and 8. United States Department of Agriculture, National Agricultural Statistics Service.

Groundwater and Crop Choice in the Short and Long Run

Supplementary Appendix: For Online Publication

Fiona Burlig Louis Preonas Matt Woerman*

Appendix contents

*. 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. Woerman: Department of Agricultural and Resource Economics, Colorado State University. Email: matt.woerman@colostate.edu.

A Details on dynamic discrete choice estimation

A.1 Model of crop choice

Our model of crop choice and derivation of an estimating equation follow closely from and build upon Scott (2013) and Kalouptsidi, Scott, and Souza-Rodrigues (2021). First, from Assumptions 1–3 in the main text, the value function for a given field is:

$$V(k_t, \omega_t, \varepsilon_t) = \max_{c_t \in \mathcal{C}} \{\pi(c_t, k_t, \omega_t, \varepsilon_t) + \beta E\left[V(k_{t+1}, \omega_{t+1}, \varepsilon_{t+1}) \mid c_t, \omega_t\right]\}$$

as shown in Equation (5) in the main text. This value function gives rise to the *ex ante* value function:

$$\bar{V}(k_t, \omega_t) \equiv \int V(k_t, \omega_t, \varepsilon_t) dF^{\varepsilon}(\varepsilon_t)$$

and the conditional value function:

$$v(c_t, k_t, \omega_t) \equiv \bar{\pi}(c_t, k_t, \omega_t) + \beta E\left[\bar{V}(k_{t+1}, \omega_{t+1}) \mid c_t, \omega_t\right] \tag{A1}$$

where $\bar{\pi}(c_t, k_t, \omega_t) \equiv \pi(c_t, k_t, \omega_t, 0)$ represents an expected profit function with the idiosyncratic shock equal to zero. As shown in Equation (6) in the main text, conditional choice probabilities depend on these conditional value functions:

$$p(c_t, k_t, \omega_t) = \frac{\exp\left[v(c_t, k_t, \omega_t)\right]}{\sum_{c'_t \in \mathcal{C}} \exp\left[v(c'_t, k_t, \omega_t)\right]}$$

We next invoke the Arcidiacono-Miller Lemma (Arcidiacono and Miller (2011)), which follows from the Hotz-Miller inversion (Hotz and Miller (1993)) and yields a new expression for the *ex ante* value function written as a function of the conditional value and conditional choice probability:

$$\bar{V}(k_t, \omega_t) = v(c_t, k_t, \omega_t) - \ln p(c_t, k_t, \omega_t) + \gamma \tag{A2}$$

where γ is the Euler constant. In words, the *ex ante* value equals the conditional value after making any crop choice c_t plus a correction term to account for the relative value of crop c_t compared to the rest of the choice set. This expression further shows that CCPs contain information about the values of making different crop choices.

We continue to follow Scott (2013) and Kalouptsidi, Scott, and Souza-Rodrigues (2021) to derive an Euler equation that will yield an estimating equation for this dynamic discrete choice model. We consider two sequences of crop choices in years t and $t+1$. In the first sequence, the farmer chooses crop c_t in year t followed by a renewal action that we denote r_{t+1} in year $t+1$. In the second sequence, the farmer instead chooses crop c'_t in year t followed by the same r_{t+1} in year $t+1$. In each case, the farmer then chooses optimally in years $t+2$ and beyond. To generate an Euler equation, we compare the value of these two cropping sequences.

We first combine Equations (A1) and (A2) to generate an expression for expected profit of any crop choice in year t:

$$\bar{\pi}(c_t, k_t, \omega_t) = \bar{V}(k_t, \omega_t) - \beta E\left[\bar{V}(k_{t+1}, \omega_{t+1}) \mid c_t, \omega_t\right] + \ln p(c_t, k_t, \omega_t) - \gamma$$

We then decompose the continuation value into its realization and its expectational error, with expectational error defined as the difference between expectation and realization:

$$e^V(c_t, \omega_t, \omega_{t+1}) \equiv E\left[\bar{V}(k_{t+1}, \omega'_{t+1}) \mid c_t, \omega_t\right] - \bar{V}(c_t, \omega_{t+1})$$

This decomposition yields:

$$\bar{\pi}(c_t, k_t, \omega_t) + \beta e^V(c_t, \omega_t, \omega_{t+1}) = \bar{V}(k_t, \omega_t) - \beta \bar{V}(c_t, \omega_{t+1}) + \ln p(c_t, k_t, \omega_t) - \gamma$$

with only realized values (rather than expected values) on the right-hand side.

Next, we eliminate the realized continuation values from this expression, first by differencing the equation across the two different crop choices in year t, c_t and c'_t:

$$\begin{aligned}\ln\left[\frac{p(c_t, k_t, \omega_t)}{p(c'_t, k_t, \omega_t)}\right] &= \bar{\pi}(c_t, k_t, \omega_t) - \bar{\pi}(c'_t, k_t, \omega_t) + \beta\left[e^V(c_t, \omega_t, \omega_{t+1}) - e^V(c'_t, \omega_t, \omega_{t+1})\right] \\ &\quad + \beta\left[\bar{V}(c_t, \omega_{t+1}) - \bar{V}(c'_t, \omega_{t+1})\right]\end{aligned}$$

In words, the CCP term on the left-hand side equals the difference in values from choosing crop c versus crop c' in year t and then choosing crops optimally in all future years.

We then use Equation (A2) to substitute for the continuation values in year $t+1$. That equality holds for all crop choices, including choice r_{t+1} from the cropping sequences described above:

$$\begin{aligned}\ln\left[\frac{p(c_t, k_t, \omega_t)}{p(c'_t, k_t, \omega_t)}\right] &= \bar{\pi}(c_t, k_t, \omega_t) - \bar{\pi}(c'_t, k_t, \omega_t) + \beta\left[e^V(c_t, \omega_t, \omega_{t+1}) - e^V(c'_t, \omega_t, \omega_{t+1})\right] \\ &\quad + \beta\left[v(r_{t+1}, c_t, \omega_{t+1}) - v(r_{t+1}, c'_t, \omega_{t+1})\right] - \beta \ln\left[\frac{p(r_{t+1}, c_t, \omega_{t+1})}{p(r_{t+1}, c'_t, \omega_{t+1})}\right]\end{aligned}$$

Because crop choice r_{t+1} is a renewal action, the field state in year $t+2$ will depend only on the choice of r_{t+1} in year $t+1$ and not on the crop choice in year t. In that case, the continuation values in year $t+2$ will be the same regardless of whether crop c_t or crop c'_t is chosen in year t, so the difference in conditional values reduces to:

$$v(r_{t+1}, c_t, \omega_{t+1}) - v(r_{t+1}, c'_t, \omega_{t+1}) = \bar{\pi}(r_{t+1}, c_t, \omega_{t+1}) - \bar{\pi}(r_{t+1}, c'_t, \omega_{t+1})$$

Then the above expression simplifies to:

$$\ln\left[\frac{p(c_t,k_t,\omega_t)}{p(c_t',k_t,\omega_t)}\right] = \bar{\pi}(c_t,k_t,\omega_t) - \bar{\pi}(c_t',k_t,\omega_t) + \beta\left[e^V(c_t,\omega_t,\omega_{t+1}) - e^V(c_t',\omega_t,\omega_{t+1})\right]$$
$$+\,\beta\left[\bar{\pi}(r_{t+1},c_t,\omega_{t+1}) - \bar{\pi}(r_{t+1},c_t',\omega_{t+1})\right] - \beta\ln\left[\frac{p(r_{t+1},c_t,\omega_{t+1})}{p(r_{t+1},c_t',\omega_{t+1})}\right]$$

We next expand the profit terms, as in Equation (4) in the main text, which yields an Euler equation. We rearrange the expression to get Equation (7) in the main text:

$$\begin{aligned}\ln\left[\frac{p(c_t,k_t,\omega_t)}{p(c_t',k_t,\omega_t)}\right] + \beta\ln\left[\frac{p(r_{t+1},c_t,\omega_{t+1})}{p(r_{t+1},c_t',\omega_{t+1})}\right] &= \alpha_G\left[G(c_t,\omega_t) - G(c_t',\omega_t)\right]\\ &+ \alpha(c_t,k_t) - \alpha(c_t',k_t) + \beta\left[\alpha(r_{t+1},c_t) - \alpha(r_{t+1},c_t')\right]\\ &+ \xi(c_t,k_t,\omega_t) - \xi(c_t',k_t,\omega_t) + \beta\left[\xi(r_{t+1},c_t,\omega_{t+1}) - \xi(r_{t+1},c_t',\omega_{t+1})\right]\\ &+ \beta\left[e^V(c_t,\omega_t,\omega_{t+1}) - e^V(c_t',\omega_t,\omega_{t+1})\right]\end{aligned}$$

Each side of this expression is equivalent to the difference in values between the two sequences of crop choices that we describe above: choosing c_t or c_t' in year t, followed by renewal action r_{t+1}, and then choosing optimally in all following years.

This equality holds for any crop choices c_t and c_t' and for any renewal action r_{t+1}. To generate our estimating equation, we set both c_t' and r_{t+1} equal to the non-crop category—which we denote with 0—while c_t represents one of the other three crop choices. Then our estimating equation is a simple linear regression:

$$Y_{ckt} = \alpha_G\Delta G_{ct} + \tilde{\Delta}\alpha_{ck} + \tilde{\Delta}\xi_{ckt} + \Delta e^V_{ct} \tag{A3}$$

where

$$\begin{aligned}Y_{ckt} &= \ln\left[\frac{p(c_t,k_t,\omega_t)}{p(0,k_t,\omega_t)}\right] + \beta\ln\left[\frac{p(0,c_t,\omega_{t+1})}{p(0,0,\omega_{t+1})}\right]\\ \Delta G_{ct} &= G(c_t,\omega_t) - G(0,\omega_t)\\ \tilde{\Delta}\alpha_{ck} &= \alpha(c_t,k_t) - \alpha(0,k_t) + \beta\left[\alpha(0,c_t) - \alpha(0,0)\right]\\ \tilde{\Delta}\xi_{ckt} &= \xi(c_t,k_t,\omega_t) - \xi(0,k_t,\omega_t) + \beta\left[\xi(0,c_t,\omega_{t+1}) - \xi(0,0,\omega_{t+1})\right]\\ \Delta e^V_{ct} &= \beta\left[e^V(c_t,\omega_t,\omega_{t+1}) - e^V(0,\omega_t,\omega_{t+1})\right]\end{aligned}$$

A.2 Estimation

Market construction We estimate the above regression at the market level, grouping farmers who face a similar choice environment. We define a market according to three criteria: electricity price, surface water availability, and geographic proximity. Because small pumps and large pumps face different marginal electricity prices, we first partition parcels based on whether their groundwater pump is on a small- or large-pump tariff.[1] To account for the final two criteria, we further split parcels by water districts—thereby grouping farms with comparable surface water allocations within contained geographic areas.[2] For parcels located outside of any water district (a.k.a., in "white areas"), we define county-level pseudo-water districts; these units retain the small vs. large tariff split, while also grouping farms with comparable surface water access (i.e., allocations of zero) in contained geographic areas.

Calculating conditional choice probabilities In order to ensure no conditional choice probability (CCP) has a value of zero or one, we smooth a market's CCPs using other markets with similar surface water availability (i.e., in vs. out of a water district) and electricity tariff (i.e., small- vs. large-pump tariffs). Smoothing weights are inversely proportional to the square of the distance between the market centroids. Formally, the smoothed CCPs are:

$$\hat{p}_m(c_t, k_t, \omega_{mt}) = \frac{\sum_{m' \in \mathcal{M}} w_{mm'} p_{m'}(c_t, k_t, \omega_{m't})}{\sum_{m' \in \mathcal{M}} w_{mm'}}$$

where weight $w_{mm'} = (1 + d_{mm'})^{-2}$ if markets m and m' have similar electricity tariffs and surface water availability, and 0 otherwise. $d_{mm'}$ is the distance between centroids of m and m' in kilometers.

Recovering profit intercept parameters Estimating Equation (A3) returns 12 $\tilde{\Delta}\alpha_{mck}$ regression intercept terms for each market, which we use to recover the 16 profit intercepts for each market, $\alpha_m(c_t, k_t)$. To do this, we must make additional assumptions. First, we normalize $\alpha_m(0, 0) = 0$, where both field state and crop choice are non-crop. Second, we decompose $\alpha_m(c_t, k_t) = R_m(c_t) - T_m(c_t, k_t)$, where $R_m(c_t)$ is time-invariant net returns to crop c_t (excluding groundwater costs), and $T_m(c_t, k_t)$ is the time-invariant cost of transitioning from field state k_t to crop c_t. We assume $T_m(0, c_{t-1}) = 0.5 \times T_m(c_t, 0)$, such that switching from crop c to fallow costs half as much as switching from fallow to crop c.[3] Third, we assume there is no transition cost to remain in the same crop: $T_m(c_t, c_{t-1}) = 0$.[4]

1. For parcels containing both small and large pumps (multiple pumps and/or pumps that change capacity), we assign the modal category based on observed groundwater use.

2. Appendix C.5 provides more information on our use of water districts. For some smaller water districts, we observe too few fields to be confident in the construction of our market-level variables (especially after having already split by small vs. large tariff categories). We consider a water district-tariff group to be too small if it contains fewer than 30 (in-sample) parcels. In this case, we create a composite water district-tariff group comprising all water districts in the county (still retaining the small vs. large tariff split).

3. Different coefficients relating $T_m(0, c_{t-1})$ to $T_m(c_t, 0)$ yield nearly identical $\alpha_m(c_t, k_t)$ parameters.

4. Any recurring costs, such as the cost of replanting an annual crop every year, are captured by $R_m(c_t)$. Then, $T_m(c_t, c'_{t-1})$ for $c_t \neq c'_{t-1}$ reflects the additional costs incurred when switching crops.

A.3 Model results

Table A1: Dynamic discrete choice parameter estimates

A. Groundwater cost parameter: α_G

-0.015^{**}
(0.006)

B. Profit intercept parameters: $\alpha(c,k)$

	Field state (k)				
	Annual	Young fruit/nut	Mature fruit/nut	Hay	Non-crop
Crop choice (c)					
Annual	1.63	−1.45	−1.87	−0.07	−0.20
	[1.06, 2.17]	[−2.03, −0.91]	[−2.45, −1.33]	[−0.64, 0.48]	[−0.77, 0.34]
Fruit/nut	−1.59	−0.11	2.45	−0.64	−0.87
	[−2.49, −0.74]	[−1.00, 0.74]	[1.55, 3.30]	[−1.54, 0.21]	[−1.77, −0.03]
Hay	0.04	−0.77	−0.83	2.08	−0.37
	[−0.93, 0.95]	[−1.74, 0.15]	[−1.80, 0.09]	[1.11, 3.00]	[−1.34, 0.55]
Non-crop†	−0.91	−1.66	−1.66	−1.23	0.00

Notes: Panel A displays our estimated groundwater cost parameter, α_G, which we obtain from estimating Equation (7). The standard error (in parentheses) is clustered at the market-year level. Panel B displays our average profit intercept parameters. The parameters are recovered at the market level, and we average over all markets to generate this table. 95% confidence intervals (in brackets) are constructed from the 2.5th and 97.5th percentile of the α_G sampling distribution across 1,000 draws. Because of our normalizations, comparing the profit intercept parameters to a null hypothesis of zero is not appropriate, so we do not report significance on these estimates. Significance of α_G: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.
† Because of our normalizations to recover all 20 $\alpha(c,k)$ terms, the intercept terms for the non-crop choice have no variation across draws from the sampling distribution.

Parameter estimates Table A1 reports the results of our dynamic model estimation. The groundwater cost parameter α_G is common to all markets. As expected, the estimated value is negative, indicating that greater groundwater costs reduce profits. The profit intercept parameters $\alpha(c,k)$ are market-specific, and we report the average values over all markets. We normalize this parameter to equal zero when both field state and crop choice are non-crop: $\alpha(0,0) = 0$. Thus, the other $\alpha(c,k)$ parameters are *relative* to remaining in non-crop for consecutive years. For annual and hay crops, remaining in the same crop type for consecutive years ($c = k$) yields positive returns as expected (e.g., see the upper left entry of Panel B for annual returns when remaining in an annual crop). For fruit/nut perennials, remaining in a young crop yields approximately zero returns, as the young crop has not begun producing yet. Once a fruit/nut crop reaches the mature field state, remaining in that crop yields the largest positive returns. When switching crops ($c \neq k$), our parameter estimates imply that current-year returns are either the same as or worse than remaining in non-crop, indicating that switching costs typically offset or dominate current-year returns from the new crop. These results are expected in our setting, where switching crop type requires large upfront investments—such as planting or removing an entire orchard of trees—that pay out over a longer time horizon.

(Semi-)elasticities Table A2 reports the land-use semi-elasticities and the groundwater and electricity elasticities that result from our counterfactual simulations. Figure 4 in the main text depicts the short-run results (Column (1)) and the long-run, Preferred IM results (Column (4)). This table reports three additional long-run scenarios: No IM (in which farmers cannot respond on the intensive margin and can only change crops), Low IM, and High IM. These alternate long-run intensive-margin elasticities come from selecting cropping streaks of seven or more years (Low IM) or five or more years (High IM)—rather than our preferred choice of cropping streaks with six or more years in a particular crop—when calibrating the response. As expected, allowing for a greater intensive-margin effect mutes the crop-switching response while intensifying the groundwater/electricity response.

Table A2: Short- and long-run elasticities with respect to groundwater pumping cost

	Short run	**Long run**			
		No IM	Low IM	**Preferred IM**	High IM
	(1)	(2)	(3)	(4)	(5)
Modeled IM elasticity	−0.944	0.000	−0.267	−0.380	−0.591
Crop semi-elasticities					
Annual	−0.000	0.112**	0.082**	0.071**	0.042**
	[−0.000, 0.000]	[0.023, 0.181]	[0.019, 0.135]	[0.014, 0.118]	[0.007, 0.075]
Fruit/nut	−0.000**	−0.210**	−0.150**	−0.129**	−0.075**
	[−0.001, −0.000]	[−0.314, −0.034]	[−0.232, −0.028]	[−0.199, −0.021]	[−0.124, −0.011]
Hay	−0.001**	−0.013	−0.008	−0.007	−0.004**
	[−0.001, −0.000]	[−0.059, 0.002]	[−0.039, 0.001]	[−0.034, 0.000]	[−0.020, −0.000]
Non-crop	0.001**	0.111**	0.076**	0.065**	0.036**
	[0.000, 0.002]	[0.015, 0.192]	[0.012, 0.137]	[0.008, 0.116]	[0.005, 0.068]
Elasticities					
Groundwater	−0.792***	−0.225**	−0.387***	−0.461***	−0.592***
	[−0.793, −0.791]	[−0.379, −0.040]	[−0.502, −0.266]	[−0.557, −0.354]	[−0.655, −0.533]
Electricity	−0.763***	−0.240**	−0.388***	−0.456***	−0.578***
	[−0.764, −0.762]	[−0.408, −0.055]	[−0.513, −0.269]	[−0.562, −0.351]	[−0.645, −0.519]

Notes: This table reports the short- and long-run (semi-)elasticities of land use and groundwater and electricity use with respect to the marginal cost of groundwater, estimated using our dynamic discrete choice model. To recover short-run (semi-)elasticities (Column (1)), we simulate the model with baseline costs until it reaches a steady state and then increase the marginal cost of groundwater by 20% for one year. While the shock is unanticipated by farmers, once it arrives, they are aware that it only lasts for one year. For this short-run scenario, we assume farmers' within-crop intensive-margin elasticity is −0.944 (from Column (1) of Table 3). To recover long-run (semi-)elasticities (Columns (2)–(5)), we permanently increase the marginal cost of groundwater by 20% and simulate the model forward. Farmers are aware that this price change is permanent. In "No IM", farmers can only respond to groundwater cost changes by changing crops. In "Low IM", "Preferred IM", and "High IM", farmers also respond on the intensive-margin by reducing water use conditional on crop choice; we assume these responses have cost elasticities of estimates of −0.267, −0.380, and −0.591, respectively (from Columns (4)–(6) of Table 4). The reported semi-elasticities and elasticities are the means over 1,000 draws for each model. The 95% confidence intervals (in brackets) are constructed from the 2.5th and 97.5th percentile over draws. We report Columns (1) and (4) graphically in Figure 4. Significance: *** 99% of simulation draws have the same sign; ** 95% of draws, * 90% of draws.

A.4 Robustness

Figure A1 plots robustness test for our long-run (semi-)elasticities, reporting results for seven alternate model specifications: heterogeneous intensive-margin elasticities ("Crop-specific IM"), parcel sample selection ("Drop high cost" and "All size"), market-level variable construction ("Weighted mean"), estimation weighting ("Weight by m-k"), and drought assumptions ("Low drought pr." and "High drought pr."). Our structural estimates are broadly robust to each of these alternate specifications, which the notes under Figure A1 describe in detail. The only notable departure from our main results is the sensitivity "High drought pr.", where we increase farmers' subjective probability of experiencing a drought from 58% to 75%. This increased risk of drought intensifies the crop switching effect from fruit/nut perennials into annuals, which increases the magnitudes of the long-run elasticities for groundwater and electricity.

Figure A1: Robustness of long-run (semi-)elasticities

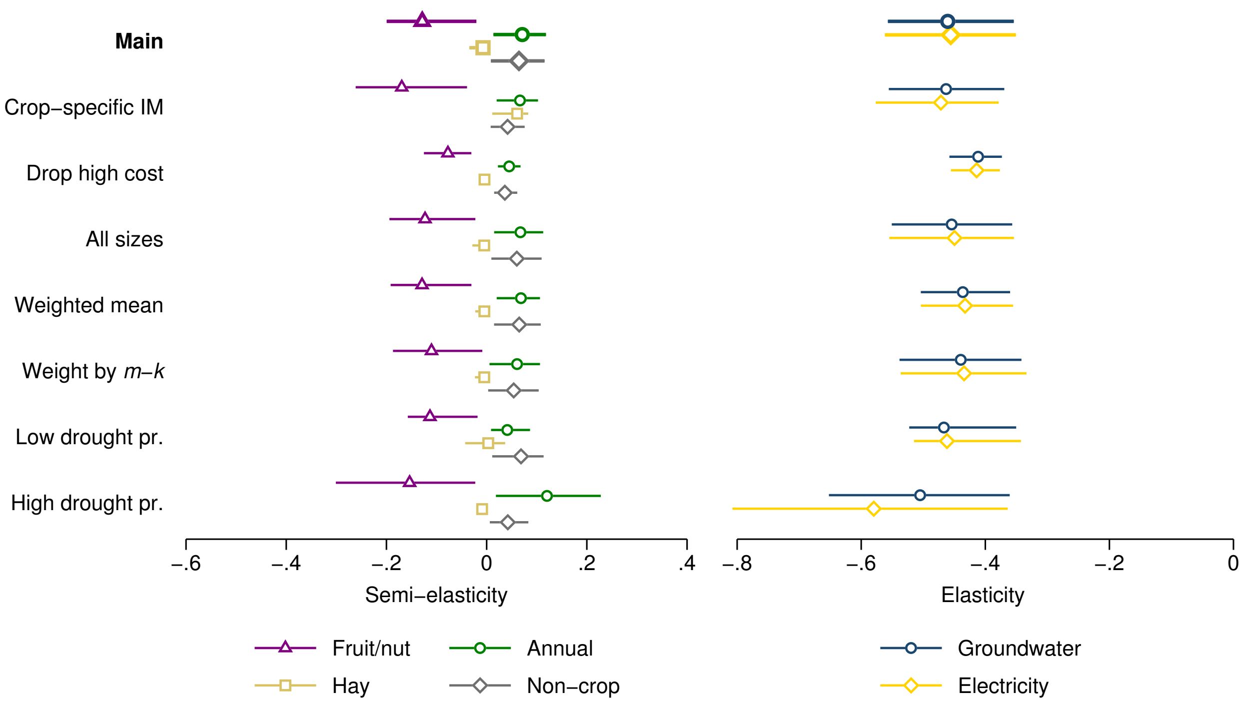

Notes: This figure plots robustness checks on our long-run (semi-)elasticities of land use (left plot) and groundwater and electricity use (right plot) with respect to groundwater cost. The top row reproduces the results of our preferred model: Preferred IM from Column (4) of Table A2. In "Crop-specific IM," we apply crop-specific heterogeneous intensive-margin responses (scaling the point estimates in Table B10 by 40%, rather than our preferred homogeneous intensive-margin response). In "Drop high cost," we drop parcels with groundwater costs $> \$5,000$ per acre (rather than our preferred threshold of $> \$3,000$ per acre). In "All sizes," we include parcels of all sizes (rather than our preferred exclusion of parcels smaller than 1 acre and greater than 5,000 acres). In "Weighted mean," we aggregate data to the market level using the weighted means (rather than our preferred weighted medians) of parcel data. In "Weight by m-k," we weight observations by croppable acres of the market-field state (rather than by our preferred weights by croppable acres of the market). In "Low drought pr." and "High drought pr.," we assume the probability of a drought is 0.4 and 0.75, respectively (rather than our preferred probability of 0.58). The left plot shows semi-elasticities for annual crops (circles), hay perennials (squares), fruit/nut perennials (triangles), and non-crop (diamonds). The right plot shows the electricity of demand for groundwater (circles) and electricity (diamonds), for the same six sensitivities. The reported semi-elasticities and elasticities are the means over 1,000 draws for each model. The plotted 95% confidence intervals (horizontal lines) show the 2.5th and 97.5th percentile over draws.

A.5 Heterogeneity in necessary tax stringencies

Our policy simulations imply substantial heterogeneity across markets in the tax stringency that would be required to achieve SGMA's groundwater reduction targets (see Panel A of Figure 5). This heterogeneity has two possible sources, which are not mutually exclusive: (i) policy-driven heterogeneity in the stringency of SGMA targets across GSAs, with stricter targets necessitating larger taxes; and (ii) spatial heterogeneity in factors that influence farmers' responsiveness to groundwater costs, with less cost-responsive markets necessitating larger taxes to achieve a given target. Here, we explore both drivers of heterogeneity by estimating descriptive cross-sectional regressions at the market (m) level.

Since the SGMA targets are endogenously related to local groundwater irrigation practices, we first regress SGMA targets on a set of market-level covariates.[5] These covariates include groundwater depth, average pump efficiency, crop shares by category (i.e., annual, fruit/nut perennial, and hay), and groundwater use—each averaged across all parcels in each market (weighted by parcel i's croppable acreage multiplied by the number of years it appears in the sample). We report these regression results in Column (1) of Table A3. We find that 10-foot deeper groundwater levels are associated with a 1.2 pp greater SGMA target, consistent with the policy's goals. Greater shares of all three crop categories (compared to non-crop, the omitted category) are also strongly associated with greater SGMA targets, which is consistent with agriculture being a major source of the aquifer drawdown that motivated SGMA. Interestingly, we find that surface water districts tend to have smaller SGMA targets in percentage terms, holding other factors constant. Together, these ten covariates predict 49% of the variation in SGMA targets.

Column (2) of Table A3 estimates the relationship between our estimates of the required groundwater tax stringency and SGMA's pumping reduction target (Panel A of Figure 5). We regress the market-specific required tax stringencies, which we estimated through our policy counterfactual simulations, on a quadratic in SGMA targets. Heterogeneity in the targets alone explains 85% of the variation in tax stringencies across markets. Column (3) replaces the quadratic in SGMA targets with the ten covariates from Column (1). While seven of these covariates are strong predictors of the market-specific tax stringencies, together they explain only 44% of the variation in our policy simulations.

Columns (2)–(3) are not straightforward to interpret, since the SGMA targets in Column (2) are endogenous to the covariates in Column (3). Because of this, the 0.44 R^2 in Column (3) could indicate either that the covariates proxy imperfectly for the SGMA targets *or* that the covariates directly explain 44% of the variation in tax stringencies. This motivates Column (4) of Table A3, which includes both the SGMA targets and the other covariates. This reveals that the statistical correlations in Column (3) appear to be proxy effects—we cannot reject zero correlation for any of the ten covariates after directly controlling for the target. The R^2 of 0.87 is only slightly higher than in Column (2), indicating that

5. We extract these SGMA targets from GSPs, as we discuss in Appendix C.6.

between-market heterogeneity in tax stringencies in our policy simulations is driven almost exclusively by policy variation in the SGMA targets.

Table A3: SGMA targets explain most of the variation in market-specific taxes

	(1) SGMA target (%)	(2) Market-specific tax (%)	(3)	(4)
SGMA target (%)		1.36**		1.62**
		(0.58)		(0.67)
SGMA target (%) squared		0.08***		0.07***
		(0.01)		(0.02)
Average groundwater depth (feet)	0.12**		0.63**	−0.04
	(0.06)		(0.31)	(0.14)
Average pump efficiency (%)	−0.08		−0.98	0.04
	(0.30)		(1.40)	(0.56)
Average farm size (croppable acres)	0.00		−0.08	−0.05
	(0.02)		(0.08)	(0.05)
Average share of annual crops	47.38**		260.55**	23.51
	(21.67)		(101.97)	(41.31)
Average share of fruit/nut perennial crops	40.90**		194.46**	−3.69
	(16.05)		(80.26)	(35.73)
Average share of hay perennial crops	40.41***		189.50***	−1.38
	(14.60)		(70.24)	(31.41)
Average groundwater use (AF/acre-year)	0.12		−1.35	−1.14
	(0.78)		(3.27)	(1.52)
1[in surface water district]	−8.42***		−45.36***	−4.11
	(3.18)		(16.40)	(7.92)
1[in San Joaquin Valley basin]	−8.60*		−51.54**	−7.53
	(4.76)		(22.46)	(10.35)
1[in Sacramento Valley basin]	−13.52***		−51.62***	−16.17
	(4.52)		(18.84)	(15.07)
Market-level observations	82	82	82	82
R^2	0.49	0.85	0.44	0.87

Notes: These cross-sectional OLS regressions describe the variation in SGMA targets and market-level taxes depicted in the Panel A of Figure 5. Column (1) regresses the market-level SGMA targets (i.e., the horizontal axis of Figure 5, Panel A) on market-level covariates. Columns (2)–(4) regress the required market-level groundwater taxes (i.e., the vertical axis of Figure 5, Panel A) on these targets and market-level covariates. We construct these market-level covariates by averaging across parcels in each market, weighting by the product of each parcel's croppable acreage and its count of years in that market. Regressions are weighted by the sum of these weights across their constituent parcels. Removing these regression weights and/or dropping markets with non-binding SGMA targets yields similar results. Heteroskedasticity-robust standard errors are in parentheses. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

B Reduced-form sensitivity analysis

Tables B1–B9 present a series of robustness checks for our reduced-form estimate of groundwater cost response in Column (1) of Table 2 (odd columns), and for our intensive-margin groundwater demand estimate in Column (1) of Table 3 (even columns).

B.1 Addressing measurement error

Unobserved pump upgrades If higher electricity costs incentivized farmers to invest in (unobserved) pump efficiency improvements, we could mistakenly interpret these energy efficiency improvements as reductions in groundwater use (rather than reductions in the electricity needed to produce the same quantity of groundwater). Such unobserved "drift" in our kWh-to-AF conversion factors would bias our groundwater estimates away from zero. In Table B1, we minimize the potential for unobserved drift in AF/kWh ratios by restricting our sample to parcel-years for which every constituent SP has an observed pump test within m months of each month in that calendar year (i.e., observations with more contemporaneous parameterizations of operating pump efficiency and lift in Equation (1)). If anything, this increases the magnitudes of our point estimates, suggesting that unobserved pump efficiency improvements are not biasing our estimates away from zero.

Table B1: Reduced-form sensitivity – months to nearest pump test

	$\log(Q^{\text{water}})$					
	(1)	(2)	(3)	(4)	(5)	(6)
$\log\left(P^{\text{water}}\ (\$/\text{AF})\right)$	-0.793^{***}	-0.820^{***}	-1.021^{***}	-1.212^{***}	-1.240^{***}	-1.505^{***}
	(0.242)	(0.301)	(0.208)	(0.306)	(0.240)	(0.355)
Pump tests within	48 months		24 months		12 months	
Intensive margin		Yes		Yes		Yes
Parcel units	7,008	6,893	6,921	6,725	6,823	6,478
County-years	356	329	316	313	283	279
Parcel-year observations	47,495	37,861	33,895	27,816	24,562	20,118
First-stage estimates						
$\log\left(P^{\text{elecDefault}}\ (\$/\text{kWh})\right)$	1.446^{***}	1.314^{***}	1.314^{***}	1.380^{***}	1.596^{***}	1.568^{***}
	(0.055)	(0.067)	(0.075)	(0.096)	(0.101)	(0.139)
Kleibergen-Paap F-stat	692	322	398	192	248	120

Notes: This table conducts sensitivity analysis on Column (1) of Table 2 and on Column (1) of Table 3, restricting the sample to parcel-years for which all constituent SP-months occur within m months of an observed pump test. Progressively restricting the sample in this way does not systematically attenuate our point estimates. This assuages concerns that unobserved pump efficiency upgrades (incentivized by higher costs) are biasing our electricity-to-groundwater conversions away from zero. Odd (even) columns are otherwise identical to Column (1) of Table 2 (Column (1) of Table 3). Regressions include the following fixed effects: parcel, parcel $\times$ $\mathbf{1}$[large pump], year, groundwater basin $\times$ year, water district $\times$ year, and (for even columns) parcel $\times$ cropping streak. Regressions are weighted by each parcel's "croppable" acreage (excluding forests, development, etc.). Standard errors (in parentheses) are two-way clustered by parcel and county-year. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

kWh-to-AF conversions Table B2 reports sensitivity analyses on our parameterization of lift in Equation (1). Columns (1)–(2) use pump-specific drawdown predictions for months without an observed pump test, rather than our preferred approach of fixing drawdown at measured levels that don't vary with depth. Columns (3)–(4) remove pumps with missing or questionable drawdown measurements, rather than our preferred approach of populating drawdown for these pumps using between-pump predictions. Columns (5)–(6) parameterize groundwater depth using the average of measurements across the full groundwater basin for each sample month, rather than our preferred approach of using month-specific rasters to estimate localized depths (which could be prone to between-pump spillovers from "cones of depression"). All three sensitivities yield similar point estimates, which is unsurprising given that our electricity price instrument is unlikely to be correlated with drawdown or depth.

Table B3 tests three sample restrictions related to our kWh-to-AF conversions. We find similar point estimates for parcels containing SPs with either exactly one vs. multiple observed pump tests (Columns (1)–(2) vs. (3)–(4)). We also find similar estimates for parcel-years with below-median distance to their nearest groundwater measurement in both summer and winter months (Columns (5)–(6)). As with Table B2, these results are unsurprising given that our instrument should be uncorrelated with the number of pump tests, between-test extrapolation/interpolation, or the accuracy of our groundwater depth rasters.

Table B2: Reduced-form sensitivity – parameterization of kWh-to-AF conversion

	$\log(Q^{\text{water}})$					
	(1)	(2)	(3)	(4)	(5)	(6)
$\log\left(P^{\text{water}}\ (\$/\text{AF})\right)$	−0.936***	−0.945***	−0.868***	−0.746**	−0.941***	−0.946***
	(0.221)	(0.268)	(0.256)	(0.295)	(0.221)	(0.274)
Sensitivity	Predicted drawdown (instead of fixed)		Drop suspect drawdown measurements		Basin-month avg depth (instead of rasterized)	
Intensive margin		Yes		Yes		Yes
Parcel units	7,104	6,997	6,302	6,187	7,104	6,997
County-years	367	334	367	334	367	334
Parcel-year observations	60,490	46,202	51,391	39,828	60,490	46,202
First-stage estimates						
$\log\left(P^{\text{elecDefault}}\ (\$/\text{kWh})\right)$	1.417***	1.322***	1.322***	1.291***	1.410***	1.290***
	(0.043)	(0.050)	(0.042)	(0.046)	(0.042)	(0.049)
Kleibergen-Paap F-stat	1068	562	1115	632	1148	565

Notes: This table conducts sensitivity analysis on Column (1) of Table 2 and on Column (1) of Table 3, focusing on components of the kWh-to-AF conversion. Columns (1)–(2) use time-varying predictions of pump-specific drawdown (i.e., the translation from depth to lift), rather than our preferred parameterization that fixes drawdown at the level reported in each pump test. Columns (3)–(4) remove pump tests where the reported drawdown measurement is questionable (e.g., extreme values, internal inconsistencies). Columns (5)–(6) use monthly average groundwater depths across each basin, rather than our preferred parameterization that rasterizes depth at each SP location. Odd (even) columns are otherwise identical to Column (1) of Table 2 (Column (1) of Table 3). Regressions include the following fixed effects: parcel, parcel $\times$ $\mathbf{1}$[large pump], year, groundwater basin $\times$ year, water district $\times$ year, and (for even columns) parcel $\times$ cropping streak. Regressions are weighted by each parcel's "croppable" acreage (excluding forests, development, etc.). Standard errors (in parentheses) are two-way clustered by parcel and county-year. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table B3: Reduced-form sensitivity – kWh-to-AF-related sample restrictions

	$\log(Q^{\text{water}})$					
	(1)	(2)	(3)	(4)	(5)	(6)
$\log\left(P^{\text{water}}\ (\$/\text{AF})\right)$	-0.861^{***}	-0.600^{*}	-1.038^{***}	-1.393^{***}	-0.837^{**}	-1.000^{***}
	(0.304)	(0.311)	(0.241)	(0.465)	(0.371)	(0.346)
Sample restriction	SPs with exactly one pump test		SPs with multiple pump tests		GW measurements w/in 10 miles (in season)	
Intensive margin		Yes		Yes		Yes
Parcel units	4,403	4,306	2,862	2,829	4,160	3,946
County-years	355	325	342	310	265	241
Parcel-year observations	35,837	27,949	24,653	18,253	23,882	19,045
First-stage estimates						
$\log\left(P^{\text{elecDefault}}\ (\$/\text{kWh})\right)$	1.359^{***}	1.268^{***}	1.268^{***}	1.333^{***}	1.532^{***}	1.440^{***}
	(0.046)	(0.052)	(0.076)	(0.091)	(0.066)	(0.078)
Kleibergen-Paap F-stat	874	456	376	171	537	299

Notes: This table conducts sensitivity analysis on Column (1) of Table 2 and on Column (1) of Table 3, for sample restrictions related to our kWh-to-AF conversions. Columns (1)–(2) include only parcels containing SPs with exactly one observed pump test. Columns (3)–(4) include only parcels containing SPs with multiple observed pump tests. Columns (5)–(6) include parcel-years for which the nearest groundwater measurement averaged less than 10 miles (the median) in both summer and winter of year t. Odd (even) columns are otherwise identical to Column (1) of Table 2 (Column (1) of Table 3). Regressions include the following fixed effects: parcel, parcel $\times$ $\mathbf{1}$[large pump], year, groundwater basin $\times$ year, water district $\times$ year, and (for even columns) parcel $\times$ cropping streak. Regressions are weighted by each parcel's "croppable" acreage (excluding forests, development, etc.). Standard errors (in parentheses) are two-way clustered by parcel and county-year. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

B.2 Identification checks

Time-varying confounders Table B4 interacts year fixed effects with baseline pump characteristics (horsepower and operating pump efficiency) and county, showing that our estimates are robust to time-varying confounders correlated with these cross-sectional factors.

Table B5 tests for time-varying confounders related to the timing of PG&E's smart meter rollout. It seems plausible that PG&E might have prioritized replacing smart meters based on: the age of customers' accounts (proxied by the earliest SP start date in Columns (1)–(2)), customers' expected load on the grid (proxied by 2008 maximum monthly kWh of a constituent SP in Columns (3)–(4)), or air conditioning demand among nearby households (proxied by climate zone in Columns (5)–(6)). Our reduced-form estimates are largely unchanged when we interact these proxies with year fixed effects, which assuages concerns about selection in the timing of smart meter switches.

Table B6 shows that our results are robust to controlling for weather (separately for each month of the current and preceding years), drought severity (by county-year, for five separate severities: "abnormal", "moderate", "severe", "extreme", and "exceptional"), localized groundwater depth, and the distance from parcel i to the nearest contemporaneous groundwater measurement (averaged across constituent SPs and over all months of the year).

Table B4: Demand sensitivity – time-varying confounders

	log(Q^{water})					
	(1)	(2)	(3)	(4)	(5)	(6)
log (P^{water} (\$/AF))	−0.880***	−0.878***	−0.935***	−0.931***	−0.914***	−0.949***
	(0.200)	(0.263)	(0.221)	(0.269)	(0.220)	(0.274)
Interact year FEs with	Initial HP		Initial OPE		County	
Intensive margin		Yes		Yes		Yes
Parcel units	7,104	6,997	7,104	6,997	7,104	6,997
County-years	367	334	367	334	367	334
Parcel-year observations	60,490	46,202	60,490	46,202	60,490	46,202
First-stage estimates						
log ($P^{\text{elecDefault}}$ (\$/kWh))	1.434***	1.325***	1.325***	1.299***	1.408***	1.296***
	(0.043)	(0.050)	(0.042)	(0.049)	(0.043)	(0.051)
Kleibergen-Paap F-stat	1127	547	1131	557	1081	505

Notes: This table conducts sensitivity analysis on Column (1) of Table 2 and on Column (1) of Table 3, interacting year fixed effects with the following cross-sectional fixed effects: earliest observed nameplate horsepower of pump (Columns (1)–(2)); earliest observed operating pump efficiency (Columns (3)–(4)); and county (Columns (5)–(6)). Odd (even) columns are otherwise identical to Column (1) of Table 2 (Column (1) of Table 3). Regressions also include the following fixed effects: parcel, parcel × **1**[large pump], year, groundwater basin × year, water district × year, and (for even columns) parcel × cropping streak. Regressions are weighted by each parcel's "croppable" acreage (excluding forests, development, etc.). Standard errors (in parentheses) are two-way clustered by parcel and county-year. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table B5: Reduced-form sensitivity – smart meter rollout

	log(Q^{water})					
	(1)	(2)	(3)	(4)	(5)	(6)
log (P^{water} (\$/AF))	−0.928***	−0.937***	−0.757***	−0.888***	−0.985***	−0.943***
	(0.220)	(0.270)	(0.205)	(0.247)	(0.198)	(0.269)
Interact year FEs with	Earliest SP start date in PG&E data		Max monthly kWh for SP in 2008		Climate zone	
Intensive margin		Yes		Yes		Yes
Parcel units	7,078	6,974	5,511	5,457	7,104	6,997
County-years	367	334	360	328	367	334
Parcel-year observations	60,458	46,175	52,164	39,393	60,490	46,202
First-stage estimates						
log ($P^{\text{elecDefault}}$ (\$/kWh))	1.426***	1.313***	1.313***	1.330***	1.421***	1.316***
	(0.043)	(0.051)	(0.045)	(0.052)	(0.043)	(0.050)
Kleibergen-Paap F-stat	1085	522	998	514	1102	548

Notes: This table conducts sensitivity analysis on Column (1) of Table 2 and on Column (1) of Table 3, focusing on factors that might have been correlated with PG&E's smart meter rollout. We interact year fixed effects with following cross-sectional continuous variables: the earliest account open date at an SP within the parcel (Columns (1)–(2)); the log of maximum monthly kWh consumed in 2008, the first year of our sample period (Columns (3)–(4)); and climate zone (Columns (5)–(6)). Odd (even) columns are otherwise identical to Column (1) of Table 2 (Column (1) of Table 3). Regressions include the following fixed effects: parcel, parcel × **1**[large pump], year, groundwater basin × year, water district × year, and (for even columns) parcel × cropping streak. Regressions are weighted by each parcel's "croppable" acreage (excluding forests, development, etc.). Standard errors (in parentheses) are two-way clustered by parcel and county-year. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table B6: Reduced-form sensitivity – adding time-varying controls

	$\log(Q^{\text{water}})$					
	(1)	(2)	(3)	(4)	(5)	(6)
$\log(P^{\text{water}}$ (\$/AF))	-1.053^{***}	-0.911^{***}	-0.921^{***}	-0.924^{***}	-0.937^{***}	-0.949^{***}
	(0.213)	(0.270)	(0.219)	(0.270)	(0.220)	(0.271)
Control variables	Monthly precipitation and temperature		Drought severity by county-year		Depth, and distance to depth measurement	
Intensive margin		Yes		Yes		Yes
Parcel units	6,996	6,896	7,104	6,997	7,096	6,988
County-years	336	334	367	334	367	334
Parcel-year observations	53,868	44,557	60,490	46,202	60,231	46,034
First-stage estimates						
$\log(P^{\text{elecDefault}}$ (\$/kWh))	1.398^{***}	1.284^{***}	1.284^{***}	1.308^{***}	1.417^{***}	1.306^{***}
	(0.046)	(0.053)	(0.043)	(0.050)	(0.043)	(0.051)
Kleibergen-Paap F-stat	917	520	1102	526	1091	522

Notes: This table conducts sensitivity analysis on Column (1) of Table 2 and on Column (1) of Table 3, adding time-varying controls. Columns (1)–(2) control for month-specific precipitation and temperature (e.g. 12 variables for precipitation in each month of year t, and for each month of year $t-1$). Columns (3)–(4) control for the share of each county-year classified as a drought, separately for five severities ("abnormal", "moderate", "severe", "extreme", and "exceptional"). Columns (5)–(6) control for depth at the location of parcel i and parcel i's average distance to the nearest groundwater depth measurement in year t. Odd (even) columns are otherwise identical to Column (1) of Table 2 (Column (1) of Table 3). Regressions include the following fixed effects: parcel, parcel $\times$ $\mathbf{1}$[large pump], year, groundwater basin $\times$ year, water district $\times$ year, and (for even columns) parcel $\times$ cropping streak. Regressions are weighted by each parcel's "croppable" acreage (excluding forests, development, etc.). Standard errors (in parentheses) are two-way clustered by parcel and county-year. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Neighbor spillovers Table B7 addresses a potential violation of our IV exclusion restriction: if a decrease in $P_{it}^{\text{elecDefault}}$ is simultaneously experienced by parcel i's neighbors, these neighbors may increase their groundwater consumption, which can generate contemporaneous between-pump interference (i.e., cones of depression) because multiple farmers share the same aquifer, which in turn lowers the groundwater level at parcel i, increasing i's marginal pumping costs. To address this possible spillover channel, we control for the average default electricity price of parcel i's neighbors within a certain geographic radius—including separate controls for the average price of neighboring in-sample parcels (i.e., confirmed pumpers most likely to impact parcel i's costs) and for the average price of neighboring agricultural service points (i.e., all other agricultural users—including any unconfirmed pumpers—that are not included in our estimation sample). We also interact these two average-price-of-neighbors controls with the count of neighbors (of each type) to model the intensity of potential spillovers, as more neighbors should lead to more between-well interference. Appendix Table B7 demonstrates that including these controls does not meaningfully alter our reduced-form estimates, defining neighbors using radii of 1, 2, or 10 miles.[6] This provides strong evidence that between-well interference is not driving an exclusion violation in this setting, and is therefore unlikely to bias our IV estimates.

6. These radii are informed by the cone of depression analysis in Alley, Reilly, and Franke (1999).

Table B7: Reduced-form sensitivity – controlling for spillovers via neighbors' electricity prices

	$\log(Q^{\text{water}})$					
	(1)	(2)	(3)	(4)	(5)	(6)
$\log\left(P^{\text{water}}\ (\$/\text{AF})\right)$	-1.013^{***}	-1.000^{***}	-1.060^{***}	-1.027^{***}	-0.934^{***}	-0.932^{***}
	(0.199)	(0.263)	(0.202)	(0.259)	(0.217)	(0.266)
Control for neighbors' prices	within 1-mile radius		within 2-mile radius		within 10-mile radius	
Intensive margin		Yes		Yes		Yes
Parcel units	6,331	6,233	6,921	6,814	7,097	6,990
County-years	360	325	360	326	360	327
Parcel-year observations	52,070	39,784	58,366	44,560	60,407	46,133
First-stage estimates						
$\log\left(P^{\text{elecDefault}}\ (\$/\text{kWh})\right)$	1.391^{***}	1.308^{***}	1.308^{***}	1.299^{***}	1.421^{***}	1.316^{***}
	(0.044)	(0.054)	(0.044)	(0.052)	(0.043)	(0.051)
Kleibergen-Paap F-stat	982	457	1036	503	1097	534

Notes: This table conducts sensitivity analysis on Column (1) of Table 2 and on Column (1) of Table 3, controlling for potential spillovers through neighbors' electricity prices (which could influence neighbors' pumping behavior, thereby indirectly impacting farmer i's groundwater depth). Each regression includes four time-varying controls: the average $P_{it}^{\text{elecDefault}}$ of neighboring in-sample parcels (i.e., confirmed pumpers), the average $P_{it}^{\text{elecDefault}}$ of neighboring agricultural service points (i.e., including all latent pumpers), and the interaction of each of these averages with the number of neighbors of the respective type (to control for the intensity of this potential spillover channel). We define neighbors using three distance radii: 1 mile, 2 miles, and 10 miles. Odd (even) columns are otherwise identical to Column (1) of Table 2 (Column (1) of Table 3). Regressions include the following fixed effects: parcel, parcel $\times$ $\mathbf{1}$[large pump], year, groundwater basin $\times$ year, water district $\times$ year, and (for even columns) parcel $\times$ cropping streak. Regressions are weighted by each parcel's "croppable" acreage (excluding forests, development, etc.). Standard errors (in parentheses) are two-way clustered by parcel and county-year. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

B.3 Additional robustness checks

Table B8 conducts sensitivity analysis on our IV specification. Columns (1)–(2) show that our results are similar if we instrument using the modal tariff in each category (rather than the default tariff). Columns (3)–(4) present the uninstrumented OLS, which is biased slightly towards zero likely due to measurement error in our kWh-to-AF conversions. Columns (5)–(6) present the reduced-form OLS, for comparison.

Table B9 conducts sensitivity analysis on our parcel-year sample. Columns (1)–(2) include the following outliers omitted from both our reduced-form and structural analyses: parcels with (i) < 1 croppable acre (reported acreage is prone to measurement error); (ii) $> 5,000$ croppable acres (unlikely to be irrigated by our observed pumps); and (iii) monthly electricity bills $> \$3,000$ per croppable acre (either highly abnormal groundwater use or measurement error in the denominator). Including these outliers yields similar results. Columns (3)–(4) remove croppable-acreage weights, yielding attenuated elasticity estimates; this suggests that larger parcels are more groundwater-cost-responsive than smaller parcels. Column (5) uses a service-point-by-year panel, which produces an estimate similar to Column (3).[7] This suggests that aggregating from SPs to parcels does not meaningfully alter our results.

7. The appropriate comparison here is between Column (5) and Column (3)—since SPs do not have croppable acreage *per se*, and do not easily map to a definition of the intensive margin.

Table B8: Reduced-form sensitivity – IV specification

	$\log(Q^{\text{water}})$					
	(1)	(2)	(3)	(4)	(5)	(6)
	2SLS	2SLS	OLS	OLS	OLS	OLS
$\log\left(P^{\text{water}}\ (\$/\text{AF})\right)$	-0.865^{***}	-0.853^{***}	-0.819^{***}	-0.779^{***}		
	(0.233)	(0.270)	(0.094)	(0.121)		
$\log\left(P^{\text{elecDefault}}\ (\$/\text{kWh})\right)$					-1.330^{***}	-1.202^{***}
					(0.320)	(0.349)
Intensive margin		Yes		Yes		Yes
Parcel units	7,104	6,997	7,104	6,997	7,104	6,997
County-years	367	334	367	334	367	334
Parcel-year observations	60,490	46,202	60,490	46,202	60,490	46,202
First-stage estimates						
$\log\left(\text{Modal } P^{\text{elec}}\ (\$/\text{kWh})\right)$	0.626^{***}	0.594^{***}				
	(0.017)	(0.020)				
Kleibergen-Paap F-stat	1386	691				

Notes: This table conducts sensitivity analysis on Column (1) of Table 2 and on Column (1) of Table 3, focusing on our instrumental variables specification. Columns (1)–(2) use an alternate instrument: the modal tariff in each category, rather than the default tariff. Columns (3)–(4) present the uninstrumented OLS estimates. Columns (5)–(6) present the reduced form of our preferred specification. Odd (even) columns are otherwise identical to Column (1) of Table 2 (Column (1) of Table 3). Regressions include the following fixed effects: parcel, parcel × **1**[large pump], year, groundwater basin × year, water district × year, and (for even columns) parcel × cropping streak. Regressions are weighted by each parcel's "croppable" acreage (excluding forests, development, etc.). Standard errors (in parentheses) are two-way clustered by parcel and county-year. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table B9: Reduced-form sensitivity – parcels and acreage weights

	$\log(Q^{\text{water}})$				
	(1)	(2)	(3)	(4)	(5)
$\log\left(P^{\text{water}}\ (\$/\text{AF})\right)$	-0.942^{***}	-0.945^{***}	-0.667^{***}	-0.593^{***}	-0.531^{***}
	(0.220)	(0.269)	(0.159)	(0.224)	(0.133)
Sensitivity	Include outlier parcels		Remove acreage weights		SP-year panel
Intensive margin		Yes		Yes	
Parcel units (or SP in Col (5))	7,742	7,619	6,786	6,366	9,575
County-years	367	334	367	330	367
Observations	66,031	49,825	59,736	42,426	83,675
First-stage estimates					
$\log\left(P^{\text{elecDefault}}\ (\$/\text{kWh})\right)$	1.418^{***}	1.273^{***}	1.273^{***}	1.196^{***}	1.332^{***}
	(0.043)	(0.056)	(0.030)	(0.039)	(0.029)
Kleibergen-Paap F-stat	1104	525	1856	940	2044

Notes: Columns (1)–(2) include parcels with bills over \$3,000 per croppable acre and with croppable areas $\notin [1, 5000]$ acres, all of which we drop from our preferred parcel-year specifications. Columns (3)–(4) remove the regression weights by each parcel's "croppable" acreage. Columns (1) and (3) are otherwise identical to Column (1) of Table 2. Columns (2) and (4) are otherwise identical to Column (1) of Table 3. Column (5) is analogous to Column (3) but estimates the groundwater elasticity at the SP-level, rather than the (more aggregated) parcel level. All regressions include the following fixed effects: unit, unit × **1**[large pump], year, groundwater basin × year, water district × year, and (for even columns) parcel × cropping streak. Standard errors (in parentheses) are two-way clustered by unit and county-year. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

B.4 Heterogeneous intensive-margin responses

Heterogeneity by crop type Table B10 reproduces Columns (1)–(2) of Table 3 interacting price (and the instrument) with indicators for the four crop categories. This reveals economically similar estimates for the three crop categories. We fail to reject the null hypothesis that all of these elasticities are identical. Hence, we assume that the intensive-margin elasticity is homogeneous across crops in our counterfactual simulations.

Table B10: Intensive-margin response is not statistically different across crop categories

	(1) $\log(Q^{\text{water}})$	(2) $\log(Q^{\text{elec}})$
$\log\left(P^{\text{water}}\ (\$/\text{AF})\right) \times \mathbf{1}[\text{Annual}]$	-0.962^{***}	
	(0.357)	
$\log\left(P^{\text{water}}\ (\$/\text{AF})\right) \times \mathbf{1}[\text{Fruit/nut}]$	-0.869^{***}	
	(0.292)	
$\log\left(P^{\text{water}}\ (\$/\text{AF})\right) \times \mathbf{1}[\text{Hay}]$	-1.146^{***}	
	(0.373)	
$\log\left(P^{\text{water}}\ (\$/\text{AF})\right) \times \mathbf{1}[\text{Non-crop}]$	-0.715	
	(0.588)	
$\log\left(P^{\text{elec}}\ (\$/\text{kWh})\right) \times \mathbf{1}[\text{Annual}]$		-0.928^{*}
		(0.482)
$\log\left(P^{\text{elec}}\ (\$/\text{kWh})\right) \times \mathbf{1}[\text{Fruit/nut}]$		-0.767^{**}
		(0.310)
$\log\left(P^{\text{elec}}\ (\$/\text{kWh})\right) \times \mathbf{1}[\text{Hay}]$		-1.282^{***}
		(0.455)
$\log\left(P^{\text{elec}}\ (\$/\text{kWh})\right) \times \mathbf{1}[\text{Non-crop}]$		-0.676
		(0.766)
p-value on joint F-test: $\gamma^{\text{A}} = \gamma^{\text{F}} = \gamma^{\text{H}}$	0.618	0.492
p-value on joint F-test: $\gamma^{\text{A}} = \gamma^{\text{F}} = \gamma^{\text{H}} = \gamma^{\text{N}}$	0.788	0.694
Parcel units	6,997	6,997
County-years	334	334
Parcel-year observations	46,202	46,202
Kleibergen-Paap F-statistic	130	204

Notes: These regressions are identical to Columns (1)–(2) of Table 3, except that we interact both the endogenous price variable and the instrument with indicators for the four crop categories. Since these indicators are defined by the modal category of each parcel-year, they are exhaustive and there is no omitted category. We report the p-values for two joint F-tests: that the intensive-margin elasticity estimates are identical for annuals, fruit/nut perennials, and hay perennials; and for all four categories. All regressions isolate the intensive margin by restricting the sample to parcel-years in at least the second year of a cropping streak. Regressions include the following fixed effects: parcel, parcel $\times$ $\mathbf{1}$[large pump], year, groundwater basin $\times$ year, water district $\times$ year, and parcel $\times$ cropping streak. Regressions are weighted by each parcel's "croppable" acreage (excluding forests, development, etc.). Standard errors (in parentheses) are two-way clustered by parcel and county-year. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Heterogeneity by water district Table B11 reproduces Columns (1)–(2) of Table 3 interacting price (and the instrument) with indicators for being in vs. not in a water district. If farmers systematically responded to groundwater cost shocks by substituting towards their surface water allocations, we would expect to find a smaller elasticity for parcels in water districts. However, we find similar intensive-margin elasticity estimates for both groups, and we fail to reject the null hypothesis that they are identical. This further indicates that surface water is unlikely to confound our groundwater elasticity estimates.

Table B11: Intensive-margin response is not statistically different in vs. not in water district

	(1) $\log(Q^{\text{water}})$	(2) $\log(Q^{\text{elec}})$
$\log\left(P^{\text{water}}\ (\$/\text{AF})\right) \times \mathbf{1}[\text{In water district}]$	-0.974^{***} (0.332)	
$\log\left(P^{\text{water}}\ (\$/\text{AF})\right) \times \mathbf{1}[\text{Not in water district}]$	-0.887^{**} (0.432)	
$\log\left(P^{\text{elec}}\ (\$/\text{kWh})\right) \times \mathbf{1}[\text{In water district}]$		-0.942^{***} (0.316)
$\log\left(P^{\text{elec}}\ (\$/\text{kWh})\right) \times \mathbf{1}[\text{Not in water district}]$		-0.831^{*} (0.451)
p-value on test of equality: $\gamma^{\text{WD}} = \gamma^{\text{NWD}}$	0.870	0.837
Parcel units	6,997	6,997
County-years	334	334
Parcel-year observations	46,202	46,202
Kleibergen-Paap F-statistic	74	93

Notes: These regressions are identical to Columns (1)–(2) of Table 3, except that we interact both the endogenous price variable and the instrument with indicators for whether the parcel is vs. is not in a water district. We report the p-values for the t-test of equality across both groups. All regressions isolate the intensive margin by restricting the sample to parcel-years in at least the second year of a cropping streak. Regressions include the following fixed effects: parcel, parcel $\times$ $\mathbf{1}$[large pump], year, groundwater basin $\times$ year, water district $\times$ year, and parcel $\times$ cropping streak. Regressions are weighted by each parcel's "croppable" acreage (excluding forests, development, etc.). Standard errors (in parentheses) are two-way clustered by parcel and county-year. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

C Data

C.1 PG&E data

PG&E monthly billing data We use confidential customer-level electricity data, which PG&E's data management team prepared for us under a non-disclosure agreement. These data comprise the universe of agricultural electricity consumers in PG&E's service territory, and we observe each customer's monthly bills at the service account level for 2008–2019. We aggregate service accounts up to 112,032 unique service points (i.e. the physical location of an electricity meter) and construct a "monthified" panel of electricity consumption (in kWh) at the service point (SP) level.[8] We observe several key covariates for each service point: its latitude and longitude, its climate zone, its electricity tariff, and indicators for accounts with solar panels on net-energy metering (which we drop from our estimation sample). Our data also include meter identifiers to link service point locations to physical electricity meters. Figure C1 maps all agricultural service points in our dataset.

PG&E's Advanced Pumping Efficiency Program PG&E also provided rich audit data on agricultural groundwater pumps, collected as part of the utility's Advanced Pumping Efficiency Program (APEP), which subsidized pump tests for agricultural consumers across PG&E service territory. We observe the universe of APEP-subsidized pump tests from 2011–2019: 33,747 unique tests at 24,642 unique pump locations. For each test, the data report detailed measurements including: operating pump efficiency, horsepower, standing water level, drawdown, lift (a.k.a. total dynamic head), flow (in gallons per minute), and kWh/AF.[9] We also observe pump make/model, water source (e.g., well, canal, reservoir, etc.), water use (i.e., irrigation vs. commercial or residential), and the electricity meter identifier. The latter lets us match pump tests to electricity service points, thereby isolating a sample of 12,419 service points for which agricultural groundwater pumping is the confirmed end-use.[10] We restrict our empirical analysis to this 11% subset of agricultural service points (plotted in dark blue in Figure C1), in order to avoid incorporating other agricultural electricity end uses.[11] We further drop any APEP-matched service points where the pump test data

8. PG&E's monthly bill cycles are customer-specific, and most billing periods do not line up with calendar months. We "monthify" billed kWh for each SP by splitting/weight-averaging multiple bills in a single calendar month, in order to create a SP by month panel. This is standard practice in the economics literature on electricity demand (e.g. Ito (2014)). Most service points have a single service account at each point in time, but service accounts frequently turn over within a given service point.

9. Measured kWh/AF serves as an important cross-check for computing groundwater quantities. Whereas the pump test data report kWh/AF at the time of each test, our electricity-to-water conversions account for variation in groundwater depth over time.

10. Pumping is almost certainly the only end use at APEP-matched service points, as PG&E typically installs a dedicated meter for each groundwater pump.

11. This limits our sample by dropping agricultural groundwater pumps that did not receive APEP pump tests (Table 1, Panel A, right column). Absent a pump test, we cannot precisely identify which of these service points are in fact groundwater pumps. However, Column (3) of Table 2 show that expanding our sample to include both columns of Table 1 yields a similar short-run elasticity estimate.

indicate a non-well pump (e.g., horizontal booster pumps reporting water source = "canal") or list a non-agricultural end use (e.g., "municipal") from our estimation sample.[12]

Figure C1: PG&E agricultural customers

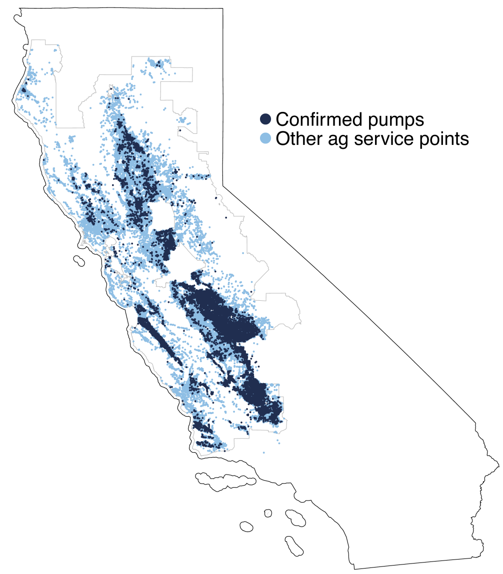

Notes: This figure maps the locations of all agricultural service points served by PG&E, from 2008–2019. Dark blue dots indicate the 12,419 service points that are confirmed pumps (i.e. matched to an APEP test for a groundwater pump). Light blue dots indicate all other agricultural service points. The light grey outline indicates PG&E's service territory.

PG&E agricultural tariffs PG&E offered 23 distinct agricultural tariffs during our sample period. Our billing data report the specific tariff associated with each monthly bill. Prices on each tariff are updated multiple times per year, and historic prices are publicly available, along with information on tariff-specific eligibility criteria.[13] We use these data to construct a 2008–2019 panel of hourly volumetric (i.e., marginal) electricity prices, which we collapse to the monthly level by taking an unweighted average across hours.[14]

22 of PG&E's 23 agricultural tariffs are divided into four mutually exclusive categories, based on pump size ("small" pumps < 35 horsepower, and "large" pumps ≥ 35 horsepower) and electricity meter type (conventional analog meters, and digital smart meters).[15] The small-conventional and large-conventional categories comprise a single tariff. The small-smart and large-smart categories comprise 8 and 12 tariffs respectively; we define the least time-varying tariff as the default in each of these categories, which serves as our instrument.

The remaining (23rd) tariff comprises a fifth category: farmers who have recently transitioned from internal combustion engines to electricity. We omit this 1.7% subset of confirmed

12. We relax the latter restriction when aggregating from SPs to parcels, since our parcel-level analyses weight parcels by croppable acreage. These weights effectively remove non-agricultural water uses with far fewer false positives than the (often miscoded) "water end-use" variable provided by the APEP database.

13. See here: https://www.pge.com/tariffs/en/rate-information.html

14. Importantly, unlike PG&E's residential electricity prices, its agricultural tariffs are not tiered: a farm's marginal price does not depend on its consumption.

15. The 35 horsepower cutoff applies to pumps with a single motor. The few pumps with multiple motors are defined as "large" if all motors sum to at least 15 horsepower of load. Conventional meters record electricity consumption using an analog dial, while smart meters digitally store the full time profile of consumption.

Figure C2: Histogram of pump horsepower

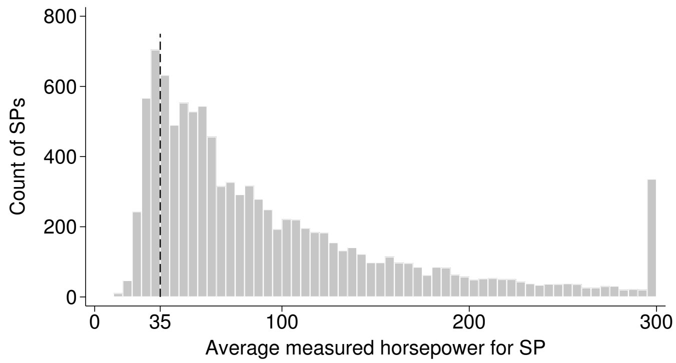

Notes: This is a histogram of measured horsepower from APEP pump tests, averaged for each PG&E service point in our estimation sample. This reveals no evidence of bunching at the 35 hp cutoff that defines PG&E's small- vs. large-pump tariff categories. Bunching would be consistent with farmers' optimizing against tariff categories when making pump investments.

pumps from our analysis entirely, for two reasons: (i) they likely represent an idiosyncratic group of pumps that is less likely to be comparable to pumps in the other four categories, and (ii) we do not observe groundwater consumption prior to switching to electricity, and we worry about selection in the timing of these switches. Our results are not sensitive to this decision to exclude this fifth category.

For our reduced-form analysis, we take unweighted averages over all sample months to construct the annual average marginal (i.e., volumetric) electricity price ($/kWh) for each tariff. For our structural analysis, we construct average total variable costs for each tariff by subtracting off fixed charges from each tariff (i.e., non-trivial charges assessed per day, per billing period, or per kW, all of which are independent of farmers' level of consumption).

C.2 Groundwater data

We use publicly available groundwater data from California's Department of Water Resources (DWR) collected under the California Statewide Groundwater Elevation Monitoring Program.[16] These data report depth below the surface at 16,852 unique monitoring stations during our 2008–2019 sample period, with an average of 33 measurements at each location at different points in time. We rasterize all measurements from each sample month, using inverse distance weighting to interpolate a gridded two-dimensional surface of average depth at each point in space.[17] Using these monthly rasters and service point geocoordinates, we construct a service point-month panel of groundwater depths. We also store the distance from each service point to its nearest measurement site in each month; this facilitates a robustness check where we remove observations with a high degree of spatial interpolation in groundwater depths (see Columns (5)–(6) of Appendix Table B3).

We assign each service point to a groundwater basin and sub-basin, using publicly available shapefiles from the DWR.[18] (Sub-)basins are defined by stratigraphic barriers that

16. https://water.ca.gov/Programs/Groundwater-Management/Groundwater-Elevation-Monitoring--CASGEM
17. Before rasterizing, we drop depth measurements that are flagged as having questionable accuracy.
18. See here: https://water.ca.gov/Programs/Groundwater-Management/Bulletin-118

limit the horizontal movement of groundwater. California has 425 basins and 517 sub-basins (only 6% of basins contain more than one sub-basin); our main sample includes farms located within 54 basins and 104 sub-basins. Our reduced-form analysis controls for changes in depth that impact all farms within the same basin (via basin-by-year fixed effects).

C.3 Constructing groundwater quantities and prices

Energy is the sole variable input to groundwater production, and the vast majority of agricultural groundwater pumps in California are powered by electricity (United States Department of Agriculture (2018)). Holding pump characteristics and groundwater depths fixed, the relationship between the quantity of groundwater extracted (in acre-feet, or AF) and the quantity of electricity (in kWh) consumed is governed by physics (Hurr and Litke (1989)):

$$\frac{\text{kWh}}{\text{AF}} = \text{kW} \div \frac{\text{AF}}{\text{hour}} = \frac{[\text{Lift (feet)}] \times [\text{Flow (gallon/minute)}]}{[\text{Operating pump efficiency (\%)}] \times 5310} \div \frac{\text{AF}}{\text{hour}} \tag{C1}$$

The power (kW) needed to pump 1 acre-foot is directly proportional to the vertical distance the water must travel to the surface (i.e., lift) and the speed at which the water travels (i.e., flow). It is inversely proportional to the rate at which the pump converts electric energy into the movement of water (i.e., operating pump efficiency, or OPE). We can simplify Equation (C1) by converting from gallons to acre-feet, arriving at Equation (1) in the main text.

We parameterize Equation (1) for confirmed pumps at the service point-month level, such that we capture the within-year time profile of kWh-per-AF conversion rates *before* aggregating up to the parcel-year.[19] We use OPE as reported in the PG&E pump test data. We extrapolate each service point's first pump test backwards, extrapolate its last pump test forwards, and interpolate between multiple pump tests using a triangular kernel in time.

We parameterize lift by combining DWR groundwater depths and PG&E pump test measurements. Lift is the sum of three components: (i) standing water level (i.e., baseline groundwater depth in the absence of pumping), (ii) drawdown (i.e., how much pump i impacts its own depth), and (iii) minor pump-specific correction factors (i.e., discharge pressure, gauge corrections, pump height above the surface).[20] We populate the standing water level (the largest component of lift) using the monthly groundwater rasters described above. We parameterize drawdown using the values reported in the pump test data.[21] Finally, we apply discharge pressure, gauge corrections, and other pump-specific adjustments as reported in the APEP database.[22]

19. Parameterizing Equation (1) annually would lead to systematically inaccurate Q^{water} conversions, since months with shallower groundwater depths (i.e., more AF per kWh) tend to have greater groundwater use.

20. Drawdown depends on rate of extraction (i.e. flow) and the physical properties of the substrata. Greater flow means greater drawdown, as water levels fall with faster extraction. More transmissive (or porous) rock formations have lower drawdown, because water levels are able to horizontally reequilibrate more quickly.

21. Where reported drawdown is missing or internally inconsistent, we populate drawdown by modeling it as a function of the standing water level and location fixed effects (to account for properties of the substrata).

22. We extrapolate beyond the first/last pump tests and interpolate between tests for these characteristics.

Appendix Table B2 presents sensitivity analysis on how we construct lift. Columns (1)–(2) replace reported drawdown with our drawdown predictions. Columns (3)–(4) drop pumps where reported drawdown is missing/suspect and thus populated using our predictions (as described above). Columns (5)–(6) parameterize lift using basin-wide average monthly groundwater depths, rather than depths extracted from monthly rasters—thereby removing any potential for localized feedback effects of pumping on depth (e.g., "cones of depression"). Our results are similar in all cases, which is unsurprising given that our electricity price instrument is not correlated with the components of Equation (1).

C.4 Constructing crop choice at the parcel level

Our data on cropped acreage come from the U.S. Department of Agriculture's (USDA) Cropland Data Layer (CDL).[23] This product provides annual crop coverage at every 30-by-30 meter pixel in the United States from 1997 to 2019. California was added to the CDL in 2007. The CDL is generated using satellite imagery in conjunction with a machine learning algorithm, and its land classifications are ground-truthed against the USDA's Farm Service Agency's farm surveys. The CDL reports 97 distinct crops that were grown in California during our sample period. We classify these 97 crops into three broad categories: annual crops, fruit and nut perennial crops, and hay perennial crops. The major annual crops in our sample are winter wheat, cotton, tomatoes, corn, and rice. The major fruit and nut perennial crops are almonds, grapes, walnuts, pistachios, and oranges. The hay perennials category is dominated by alfalfa. Two additional categories are non-crop (which the CDL reports as "fallow/idle cropland"), and not croppable (i.e., forest, shrubland, and development).

Using parcel shapefiles obtained from California county tax assessors' offices, we cookie-cutter each annual CDL image to construct parcel polygons. This yields a parcel-year panel of the shares of land cover by category (e.g., the fraction of acres in parcel f that were crop category c in year t). For parcels that are spatially merged to our sample of confirmed-pump service points, these fractions serve as outcome variables in our reduced-form analysis (i.e., Columns (3)–(6) of Table 3). They also enter our structural analysis as F^c_{ft} in Equation (8). However, for our intensive-margin regressions (Columns (1)–(2) of Table 3), we restrict the sample such that the parcel's *modal* crop choice is the same in adjacent years. Finally, we use year-on-year transitions at the pixel level to calculate conditional choice probabilities.[24] In all cases, we remove "not croppable" acreage from the denominator of each parcel-year.

C.5 Defining markets using surface water districts

Following Hagerty (2022), we spatially merge parcels to water districts. Shapefiles for California's water districts come from the California DWR, the California Atlas, and the California

23. See here: https://www.nass.usda.gov/Research_and_Science/Cropland/SARS1a.php

24. We use all pixels contained within parcels in market m that merge to confirmed-pump service points, dropping not croppable pixels. Calculating CCPs at the pixel level (as opposed to using parcel-specific modes) helps to increase coverage across all possible switches within a market.

Environmental Health Tracking Program.[25] Water districts are administrative entities that govern farmers' annual allocations of surface water.[26] Individual water districts typically offer their constituent farmers a common per-acre allocation at substantially lower marginal price than farmers' marginal cost of groundwater pumping.

Figure C3: Water districts and counties used to construct markets

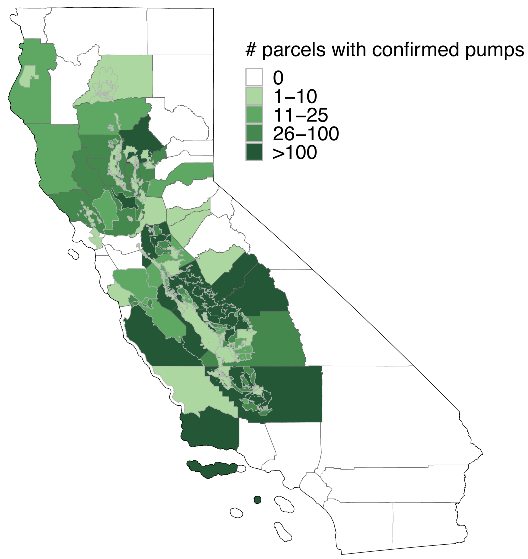

Notes: This figure maps the number of parcels with confirmed-pump service points by water district and (for parcels not in a water district) by county. We use water districts (plotted with thick grey borders) and counties-less-water-districts to construct markets in our dynamic discrete choice analysis. Note that our market definitions further subdivide these polygons by small- vs. large-pump tariffs, and then aggregate water districts with fewer than 30 parcels up to the county level (preserving the small- vs. large-pump split). We drop markets with fewer than 15 parcels after this within-county, within-pump-size aggregation.

Since groundwater and surface water are obvious substitutes, this cost advantage for surface water is key: we can assume that farmers exhaust their (inframarginal) surface water allocations before pumping groundwater, rendering any positive observed groundwater pumping the marginal source of irrigation.[27] In our reduced-form analysis, we non-parametrically control for annual shocks to surface water allocations at the water district level. This helps to isolate changes in pumping behavior driven by variation in pumping cost shocks, rather than by variation in the availability of groundwater substitutes.

For our structural analysis, we use water districts to define "markets." This grouping combines farmers who are geographically proximate and likely to have similar conditional value functions for a given field state and crop choice. It also absorbs heterogeneous surface water allocations and annual shocks to these allocations, which occur at the water district level. For the 40% of confirmed-pump parcels that are not in a water district (i.e., not

25. We thank Nick Hagerty for providing these shapefiles, and for his help in understanding and processing these surface water data.

26. As Hagerty (2022) describes, the term "water district" refers to multiple types of organizations that provide/sell water to irrigators within a defined area, including: irrigation districts, county water agencies, water conservation and flood control districts, reclamation districts, and mutual water companies.

27. A third source of water for irrigation is the open market. However, Hagerty (2025) suggests that purchased water is almost always more expensive than the groundwater costs for farmers in our dataset.

receiving surface water allocations), we use counties to define "markets." Figure C3 maps water districts (with thick grey borders) and counties, where shading indicates the number of parcels with confirmed-pump service points in each polygon. These polygons do not directly correspond to the markets used in our structural analysis, since (i) we further subdivide parcels by small vs. large pump categories, and (ii) we then aggregate water-district-by-pump-size units with fewer than 30 in-sample parcels up to the county level. We drop markets with fewer than 15 in-sample parcels after this within-county, within-pump-size aggregation—since small samples in these markets generate extremely unstable projections of groundwater use, electricity use, and cost.[28]

C.6 SGMA data

To quantify the reductions in groundwater pumping that will be required under SGMA, we collect data from the Groundwater Sustainability Plans (GSPs) that Groundwater Sustainability Agencies (GSAs) submitted to the California Department of Water Resources. All GSAs in the 90 high- and medium-priority basins were required to submit GSPs by January 31, 2022 (California Department of Water Resources (2024)). As of the time of writing, there were 120 available GSPs.[29] We downloaded all available GSPs and extracted two pieces of information from each: (i) annual average groundwater pumping; and (ii) "sustainable yield," or "the maximum amount of water calculated over a base period representative of long-term conditions in the basin and including any temporary surplus that can be withdrawn annually from a groundwater supply without causing an undesirable result" (California Department of Water Resources (2017)).[30] We were able to populate these two numbers for 111 out of the 120 available GSPs.[31]

Our measure of interest is the percent reduction in groundwater pumping that will be needed to meet each GSP's SGMA target, which we define as:

$$\frac{\text{current pumping} - \text{sustainable yield}}{\text{current pumping}} \times 100$$

Figure C4 plots this statistic for all (available) GSAs. 63 GSPs report overdraft conditions, or sustainable yield that is below current pumping levels; 57 GSPs report current pumping

28. Appendix A.2 provides further details on how we construct markets.

29. GSPs are available from the Department of Water Resources: https://sgma.water.ca.gov/portal/gsp/all

30. GSPs are detailed documents, frequently over 1,000 pages long. Where possible, we draw these numbers from the executive summary. Otherwise, we extract these numbers from the GSP's water budget section.

31. Of the 90 basins where GSPs were required, 71 basins' GSPs were fully approved as of January 2024. 13 basins' GSPs were deemed incomplete, and 6 basins' GSPs were deemed inadequate. We include all available GSPs—whether approved or not—in our GSP data, as these are the best available representation of groundwater pumping reductions required under SGMA. We expect that, if anything, the final approved GSPs will be more stringent than the proposals, making our summary statistics underestimates of the ultimate regulatory stringency.

Figure C4: GSP sustainability targets under SGMA

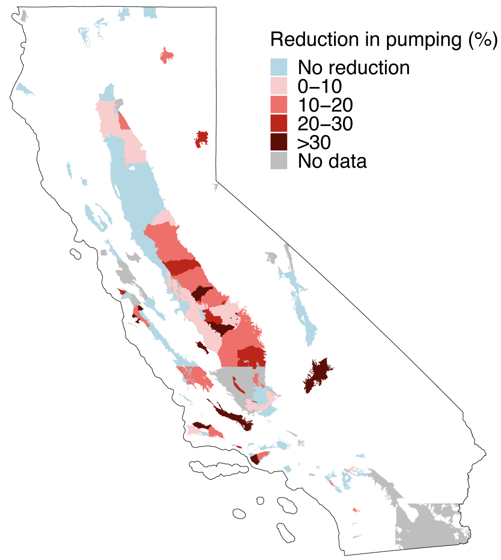

Notes: This figure maps the Groundwater Sustainability Agencies (GSAs) within California's medium- and high-priority basins, which were required to submit Groundwater Sustainability Plans (GSPs) to the Department of Water Resources. The shading reflects the percentage reduction in groundwater pumping that will be required to reach sustainability according to each GSP. See surrounding text for details.

levels at or below sustainable yield, thereby already achieving sustainability.[32] Bringing the average overdrafted GSP into sustainability under this definition will require reductions in pumping of 17.4% (weighting GSPs by historical groundwater pumping levels). Looking at only the overdrafted GSPs within PG&E's territory, this weighted average falls to 16.9%. We convert to the market level using the average required reductions in overdrafted GSPs that overlap with our market boundaries, weighting by the fraction of the market's croppable acreage in each GSP.[33] This yields a weighted average reduction target of 16.7% in our sample.

C.7 Weather data

We obtained daily temperature and precipitation rasters from the PRISM climate group, a standard source in the agriculture economics literature (see, e.g., Schlenker and Roberts

32. It is possible that the GSPs *understate* the true magnitude of overdraft. Bruno, Jessoe, and Hanemann (2024) compares reported overdraft to the results from running the C2VSim hydrology model, and finds they are broadly similar (average reported overdraft: 0.085 AF/acre; average modeled overdraft: 0.094 AF/acre). To the extent that the GSPs are underestimates, our policy estimates will be conservative.

33. This conversion effectively drops cropland that falls either outside of any GSA or inside a non-overdrafted GSA, since neither faces any SGMA compliance obligation. 11 of our 82 markets contain no overdrafted SGMA-regulated areas. Thus, our policy counterfactual applies only to SGMA-regulated areas of the remaining 71 markets. These adjustments are necessary since our market footprints (which importantly align with water district borders to eliminate any confounding effects of changing surface water allocations) do not cleanly map to GSA boundaries.

(2009)).[34] Using gridded data with a 4km-by-4km resolution, we extract daily maximum temperature, minimum temperature, and precipitation at each SP location.

C.8 Drought data

We use historic drought data from the National Oceanic and Atmospheric Administration (NOAA), which publishes weekly maps of drought intensity across the U.S. NOAA uses six drought categories of increasing severity: "No Drought", "Abnormally Dry", "Moderate", "Severe", "Extreme", and "Exceptional".[35] We aggregate these data to the county-year, quantifying the average share of each county's area that falls into each category in each year.

Appendix references

Alley, William M., Thomas E. Reilly, and O. Lehn Franke. 1999. *Sustainability of Ground-Water Resources.* U.S. Geological Survy Circular 1186.

Arcidiacono, Peter, and Robert A. Miller. 2011. "Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity." *Econometrica* 79 (6): 1823–1867.

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.

California Department of Water Resources. 2017. *Best management practices for the sustainable management of groundwater: Sustainable management criteria.* Technical report.

———. 2024. *Groundwater sustainability plans.*

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

———. 2025. "Transaction costs and the gains from trade in water markets." Working paper.

Hotz, V. Joseph, and Robert A. Miller. 1993. "Conditional choice probabilities and the estimation of dynamic models." *Review of Economic Studies* 60 (3): 497–529.

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.

Ito, Koichiro. 2014. "Do consumers respond to marginal or average price? Evidence from nonlinear electricity pricing." *American Economic Review* 104 (2): 537–563.

34. These data are available at https://prism.oregonstate.edu/.
35. These maps are available here: https://www.drought.gov/data-maps-tools/us-drought-monitor

Kalouptsidi, Myrto, Paul T. Scott, and Eduardo Souza-Rodrigues. 2021. “Linear IV regression estimators for structural dynamic discrete choice.” *Journal of Econometrics* 222 (1): 778–804.

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.

Scott, Paul T. 2013. “Dynamic discrete choice estimation of agricultural land use.” Working paper.

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