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 June 2024

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, Dave McLaughlin, Kyle Meng, Mar Reguant, Michael Roberts, Paul Scott, Joe 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, 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 Sloan Foundation (via the E2e Project), and 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 June 2024
JEL No. Q15,Q25,Q41

ABSTRACT

How do agents respond to policy when investments have high up-front costs and lasting payoffs? We estimate farmers' short- and long-run responses to changes in groundwater pumping costs in California—where perennial crops with these features are prevalent—using both fixed effects and dynamic discrete choice models that leverage quasi-experimental variation. In the short run, farmers' groundwater demand elasticity is -0.76, and they do not change crops. In contrast, the long-run elasticity is -0.38, driven in part by meaningful reductions in water-intensive perennial cropping. Meeting California's sustainability targets would require reallocation of 9% of acres, including a 50% increase in fallowing.

Fiona Burlig
Harris School of Public Policy
University of Chicago
1307 East 60th Street
Chicago, IL 60637
and NBER
burlig@uchicago.edu

Louis Preonas
Department of Agricultural and
Resource Economics
University of Maryland
2200 Symons Hall
7998 Regents Drive
College Park, MD 20742
lpreonas@umd.edu

Matt Woerman
Agricultural and Resource Economics
Colorado State University
Nutrien Building
301 University Ave
Fort Collins, CO 80523
matt.woerman@colostate.edu

1 Introduction

Water is a precious natural resource, which economists have studied for over a century (Coman (1911)). Agriculture contributes 90% of global freshwater consumption (Carleton, Crews, and Nath (2024)); irrigation enables both high-value crop production and farmer adaptation to climate shocks (Schlenker, Hanemann, and Fisher (2005); Hultgren et al. (2022)). Groundwater is a major source of irrigation, supplying over 35% of irrigated acres worldwide (Carleton, Crews, and Nath (2023)). As a textbook common-pool resource (Ostrom (1990); Provencher and Burt (1993)), groundwater is being extracted more quickly than it is naturally recharged in key farming regions—leading to rapid aquifer depletion (Jasechko et al. (2024)), which will necessitate long-run groundwater management.

The efficacy of policies to curb aquifer depletion depends on how farmers respond in both the short and long run. These responses could differ substantially for different time horizons if farmer choices include dynamic considerations. Perennial crops are inherently dynamic, requiring large upfront investment costs and yielding multiple years of production. As a result, a farmer may respond to a short-run groundwater cost shock (e.g., a drought that raises pumping costs) by applying less water to her perennial crop for a single growing season. However, faced with a long-run cost shock (e.g., a permanent groundwater tax), she may instead choose to switch to a less-water-intensive annual crop. Such long-run adjustments will be difficult to detect in a short-run analysis alone. These features are not unique to agriculture: a wide variety of settings are characterized by agents making dynamic investment decisions. Understanding the effectiveness and broader consequences of policy therefore requires long-run estimates.

This paper generates novel empirical estimates of farmers' short- *and* long-run responses to changes in groundwater pumping costs in California, one of the world's most valuable agricultural regions. California farmers produce 18% of total U.S. crop value, and rely heavily on groundwater for irrigation (Bruno (2017); Liu et al. (2022)). Yet 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 groundwater sustainability by 2042.[1] The effectiveness and economic consequences of any such permanent groundwater regulation hinge on both the extent of farmers' response and their means of adapting to higher irrigation costs over the long run.

We ask two 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? Despite the importance of these questions, answers have proven elusive because (i) groundwater extraction 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 overcome these 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 in groundwater pumping. We assemble data on electricity prices and quantities for all agricultural consumers served by Pacific Gas & Electric (PGE). 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 PGE'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 (2021)) in order to estimate model parameters without making parametric assumptions about farmer expectations.

We begin by estimating the short-run effects of groundwater cost shocks on groundwater extraction using a reduced-form approach and exogenous variation in electricity tariffs. We recover an annual groundwater demand elasticity of -0.762, revealing that farmers indeed make behavioral changes in response to groundwater cost shocks.[2] Next, we estimate

1. While SGMA grants local agencies flexibility over the instruments they use to achieve these reductions, more than half are proposing price-based approaches (Bruno, Hagerty, and Wardle (2022)).

2. 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 identification help us overcome these challenges.

how short-run groundwater cost shocks impact land use across four major crop categories in California: annuals, fruit/nut perennials, hay perennials, and non-crop (or fallowing).[3] Restricting our sample to farms growing the same crop category in consecutive years reduces our annual elasticity estimate only slightly to -0.679. Moreover, when we use crop category shares as the outcome variable for analogous reduced-form regressions, we find precise null results on land use for all four categories. These results suggest that farmers do not switch crops in response to short-run variation in pumping costs.

Importantly, this short-run model mischaracterizes the decisions of farmers in California, where 62% of total crop revenues come from perennials such as almonds, grapes, and alfalfa. 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 are able to reduce water use *both* by switching crops *and* by using less water conditional on crop choice (Boser et al. (2024)).[4] As in our reduced-from approach, we identify model parameters using exogenous variation in groundwater costs driven by changes in regulated electricity prices.

Using these parameter estimates, we compute long-run semi-elasticities of land use by solving for farmers' value functions via a fixed-point algorithm and simulating a long-run steady state under different groundwater costs. We find that farmers *do* switch crops in response to permanent groundwater cost changes, with semi-elasticities of 0.071, -0.138, -0.006, and 0.072 for annuals, fruit/nut perennials, hay perennials, and non-crop, respectively. In contrast with our short-run estimates, appropriately accounting for dynamics shows that permanently higher groundwater costs cause farmers to shift towards less water-intensive crops in the long run.

3. We restrict our analysis to these four categories for two reasons. First, our land use data come from the USDA's Cropland Data Layer. These satellite-derived data can clearly distinguish between categories (e.g., orchard vs. row crops), but are less accurate within categories (e.g. almond vs. orange trees). Second, aggregating crops into categories is necessary for the tractability of our structural model. Our short-run elasticity estimates still capture the effect of within-category crop switches.

4. To incorporate the latter channel, we use long-difference models to calibrate the proportion of the short-run elasticity which is likely to persist over the long run. Our central estimates assume that farmers anticipate making intensive-margin adjustments when choosing crops. If we eliminate these intensive-margin adjustments, we recover larger crop-switching semi-elasticities.

Combining these land-use changes in land use with farmers' intensive-margin adjustments, we estimate long-run demand elasticities of -0.363 for groundwater and -0.373 for electricity. Incorporating forward-looking behavior and state dependence is crucial: if we shut down these features, the analogous static model produces elasticity estimates over an order of magnitude smaller. These long-run elasticity estimates are meaningfully smaller than our short-run estimates—suggesting that over a longer time horizon, the mechanism of farmer responses shifts towards crop switching, and away from unsustainable short-run strategies. Our results further demonstrate how short- and long-run responses to environmental policy may diverge substantially, as different mechanisms—yielding different elasticities—are optimal over different time horizons.

Finally, we simulate farmers' long-run responses to counterfactual groundwater taxes, which internalize open-access externalities and may be used to achieve sustainable levels of groundwater extraction.[5] According to the state's Groundwater Sustainability Plans (GSPs), achieving "sustainable yield" under SGMA in currently overdrafted regions will require groundwater pumping reductions of 19% on average. Our simulations suggest that a 60% tax on groundwater pumping could achieve this sustainability target. Such a stringent tax would cause nearly 9% of cropland to switch category, driven by a 24% drop in fruit/nut perennials and a 50% rise in fallowing (compared to a no-tax scenario). These results imply that SGMA's sustainability goals are achievable, but will likely induce dramatic changes to California's 20 million acres of cropland.

This paper makes four contributions. First, and most importantly, we use both panel fixed effects and 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); Pfeif-

5. Groundwater extraction creates two main externalities (Provencher and Burt (1993)). The "stock externality" arises when agents fail to fully account for the future value of resource, leading to depletion that is faster than the social planner's optimal extraction path. The "pumping cost" externality arises when one agent's extraction lowers water levels and increases pumping costs for other (nearby) users in the short run. Other externalities might include land subsidence and air pollution from soil drying.

fer and Lin (2014); Smith et al. (2017)).[6] Bruno, Jessoe, and Hanemann (Forthcoming) use a reduced-form approach to study land use and groundwater dynamics over five years, in response to 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.

Our findings highlight the value of long-run environmental policy analysis. While our results come from one key context—California groundwater—the lesson that agents may respond differently to short- vs. long-run policies is applicable wherever resource management requires a long time horizon, including forests (Araujo, Costa, and Sant'Anna (2022); Hsiao (2022); Balboni et al. (2023)), fish (Costello et al. (2010)), the global climate (Nordhaus (2019)), and other renewable (Gordon (1954)) and nonrenewable (Hotelling (1931)) resources. More broadly, 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)), and pollution control technologies (Blundell, Gowrisankaran, and Langer (2020)).

Second, we estimate long-run agricultural *electricity* demand, accounting for crop investment dynamics. Long-term investments are major determinants of electricity use across sectors, including durable appliances (McRae (2015)), energy efficiency upgrades (Fowlie, Greenstone, and Wolfram (2018)), and solar panels (Borenstein (2017)).[7] Despite the influence of investment dynamics on electricity use, existing work in this area is limited.[8] We extend a small recent literature, which uses quasi-experimental variation to estimate long-run electricity demand in the residential sector (Deryugina, MacKay, and Reif (2020); Feehan (2018); Buchsbaum (2023)), by combining quasi-random price changes with a structural

6. 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)).

7. Suppose an electricity price shock causes a household to invest in insulation. Properly modeling this investment decision requires forward-looking expectations (i.e., the household will be less likely to invest in insulation if it believes the price increase to be short-lived) and state-dependence (i.e., after investing in insulation, the household is unlikely to make similar investments in future periods).

8. Numerous studies have using quasi-experimental variation to estimate short-run electricity demand, mainly in the residential sector (e.g., Fell, Li, and Paul (2014); Ito (2014)). There are few short-run estimates of commercial/industrial electricity demand (exceptions include Jessoe and Rapson (2015); Blonz (2022)).

model of a major commercial electricity end-use: groundwater pumping. We provide one of the first long-run estimates of electricity demand derived from a model of forward-looking agents that make state-dependent investments (following Rapson (2014)). Our results are among the first rigorous estimates of long-run electricity demand in the agricultural sector, which consumes nearly 8% of California's electricity.

Third, 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 documented how groundwater costs impact crop choice, which has important implications for agricultural output markets.[9] 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.[10] Our work complements recent studies of surface water irrigation (Rafey (2023); Hagerty (2022); Hagerty (2023)), 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 (Forthcoming)) by providing estimates of land use change under groundwater policy—for the majority of California's farmland—in both the short and long run.

Fourth, we extend the literature on groundwater management by simulating 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 (2023)), as the few existing policies are mostly local in scope.[11] 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

9. 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.

10. 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.

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

anticipatory responses to (Bruno and Hagerty (2023)) 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—which will alter the landscape of some of the most valuable crop land on earth.

This paper proceeds as follows. Section 2 provides background on groundwater pumping and energy use in California agriculture. Section 3 describes our data. Section 4 outlines our reduced-form approach and presents results. Section 5 outlines our structural model and presents our dynamic estimates and counterfactual simulations. Section 6 concludes.

2 Background

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)).[12] To water these thirsty crops, farmers rely on two water sources with vastly different governance structures (Sawyers (2007)): 61% of irrigation comes from surface water, while 39% comes from groundwater (California Department of Water Resources (2015)).

12. For comparison, the average California household uses 0.52 acre-feet per year (Hanak et al. (2011)).

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.[13] Individual farmers typically receive water allocations proportional to their acreage within the district (Schlenker, Hanemann, and Fisher (2007)); these allocations fluctuate from year to year depending on scarcity (e.g., the amount of snowpack). Importantly, farmers pay a lower marginal cost for district water allocations than for self-pumped groundwater (Hagerty (2022)). We therefore treat district water consumption as inframarginal to any observed groundwater use—since a farmer is unlikely to incur groundwater pumping costs without exhausting her annual allocation of (cheaper) district water.

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 several times higher than marginal groundwater pumping costs (Hagerty (2023)). We therefore assume that purchased water is not 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.[14] 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)).[15] This has enabled farmers to extract vast amounts of groundwater to irrigate their overlying cropland.

13. 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.

14. Pre- and post-SGMA, overlying rightsholders face few restrictions to drilling new groundwater wells, which cost $75,000 on average and typically reach 300–500 feet (Bruno et al. (2023)). While new wells must be reported to the Department of Water Resources, construction usually lasts less than one week (Central Valley Flood Protection Board (2020)). There are also "appropriative" groundwater rights, for users who do not own land above the aquifer. Users may only exercise appropriative rights in the case of a surplus.

15. There are limited exceptions to this rule, where a few irrigation districts impose a per-unit price on groundwater (e.g., the Pajaro Valley described in Bruno and Jessoe (2021a)).

Nearly all groundwater pumps in California run on electricity, the sole variable input to groundwater production. 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.[16]

2.2 Groundwater depletion and management policy

Due to California's longstanding open-access groundwater extraction 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 to 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 to 2021 (Liu et al. (2022)). The left panel 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. The right panel of Figure 1 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 empow-

16. While previous studies have used variation in energy costs to estimate the price elasticity of groundwater demand (e.g., Hendricks and Peterson (2012); Pfeiffer and Lin (2014); Badiani and Jessoe (2019); among others) Mieno and Brozovic (2017) argue that prior estimates tend to be limited by a combination of (non-classical) measurement error, a lack of micro-level identifying variation, and relatively narrow geographies—challenges we overcome with our instrumental variables approach and data spanning PGE.

Figure 1: Groundwater depletion and perennial crops

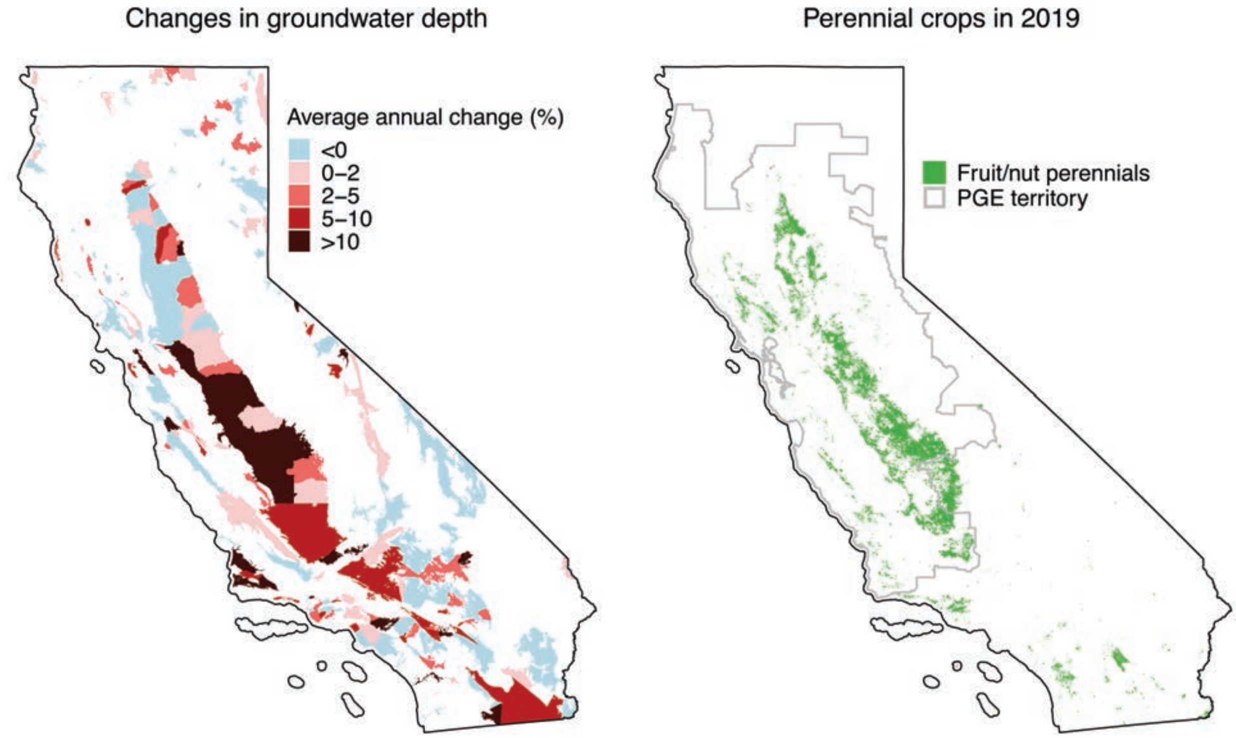

Notes: The left map plots California's groundwater sub-basins, shading based on the average annualized change in depth during our 2008–2019 sample period. A 10% change in depth corresponds to a 10% increase in groundwater pumping costs, holding all else constant. This map averages depth measurements across each sub-basin from April–June of each year, to remove seasonality. The right map 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 PGE's service territory, which encompasses most of this perennial acreage.

ers 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—or having each basin operate within its sustainable yield and avoiding "undesirable results."[17] Critically over-drafted (other medium- and high-priority) basins were required to submit GSPs by 2020 (2022) and are required to achieve sustainability

17. 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)).

by 2040 (2042).[18] Using data from the universe of GSPs, we calculate that the average required reduction in pumping among regions currently experiencing overdraft is 19.2%.[19]

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, 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 Data

3.1 PGE data

We use confidential customer-level microdata from all of PGE's agricultural service points (i.e., electricity meter locations). For each service point, we observe monthly billing data from 2008–2019. These data report the service point's latitude and longitude, monthly electricity consumption (in kilowatt-hours, or kWh), monthly bill amount (in dollars), and electricity tariff. We use PGE's published agricultural tariff schedules to calculate average marginal prices (in \$/kWh) for each service point-month.[20]

In addition, we leverage a unique PGE dataset of agricultural groundwater pump audits. We observe detailed measurements and technical specifications from over 30,000 pump tests from 2011–2019. We match pump tests to service points in our billing data using electricity meter identifiers, isolating a subset of service points where agricultural groundwater pumping is almost certainly the sole electricity end use.[21]

18. All SGMA implementation has occurred after our 2008–2019 analysis period. Bruno and Hagerty (2023) argue that there has not been anticipatory action to reduce groundwater use in response to SGMA's passage.

19. This weights GSPs by current pumping levels. An unweighted calculation (which gives each GSP equal weight) yields a required reduction of 20.8%. See Appendix C.7 for more details on the GSP data.

20. We drop customers with solar panels from our analysis, since their billed electricity use is net of (unobserved) solar production. Appendix C provides further details on the data described in this section.

21. PGE typically installs a dedicated meter for each groundwater pump. We focus our analysis on this subset of matched service points to ensure that we are measuring energy used for pumping, avoiding other agricultural electricity end uses (e.g., refrigeration, or heating greenhouses).

3.2 Constructing groundwater prices and quantities

Physics governs the relationship between kWh of electricity input and acre-feet (AF) of groundwater output for each pump:

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

PGE's pump audit data report the operating efficiency of each pump. To parameterize lift—the vertical distance from the groundwater source to the surface—we combine PGE's measurements with publicly available data on groundwater depths from California's Department of Water Resources. Since these depth measurements vary across space and time, we condition Equation (1) on contemporaneous groundwater levels at each service point.[22] Using these pump-specific production functions and data on electricity quantities and prices, we compute groundwater use (in AF) and marginal groundwater costs (in $/AF).

3.3 Land use data

We use California county assessor tax parcels as farm boundaries, as in Bruno, Jessoe, and Hanemann (Forthcoming). We spatially merge PGE service points to parcel polygons, linking each groundwater pump to the fields that it most likely irrigates. Then, we 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.[23] 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).

22. We rasterize thousands of separate 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 depends on rate of flow and subsurface characteristics. We observed repeat tests for 64% of service points, enabling us to incorporate within-pump variation in efficiency over time.

23. 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.

3.4 Summary statistics

Panel A of Table 1 compares our matched sample of groundwater pumps to the remaining unmatched agricultural service points. While these two groups face similar marginal electricity prices ($0.13/kWh vs. $0.16/kWh), our matched sample has much greater energy consumption (10,122 vs. 4,962 kWh/month). This is unsurprising since groundwater pumping is far more energy-intensive than other farm end-uses. To the extent that our matching process filters out groundwater pumps that never received PGE pump tests, our sample is skewed towards larger pumps that are most important for groundwater management policy.[24]

Within our matched sample, the average service point consumes 33.5 AF of groundwater per month at a marginal cost of $47.37/AF. Farmers face far more variation in groundwater costs than in electricity prices, due to both dispersion in pumping efficiencies and changing groundwater depths. Panel B aggregates our matched sample up to the parcel-year level. The average parcel has 290 croppable acres, with 22%, 45%, 25%, and 8% of acres in annuals, fruit/nut perennials, hay perennials, and non-crop, respectively. Across all four categories, farmers consume an average of nearly 4 AF of groundwater per acre per year.[25]

4 Reduced form estimation and results

This section presents our reduced-form analysis, using a panel fixed effects approach to estimate short-run responses to changes in groundwater costs. We present three sets of short-run results. First, we estimate monthly cost elasticities of groundwater and electricity demand at the pump level, as a proof-of-concept of our identification strategy. Next, we estimate our main short-run results of interest: annual, parcel-level cost elasticities of groundwater demand. Finally, we present annual crop choice responses to groundwater cost shocks.

24. In Appendix Tables B7 and B11, we use hierarchical clustering to predict which unmatched service points might be groundwater pumps. Including these "predicted pumpers" in our reduced-form analysis does not meaningfully alter our estimates of electricity demand.

25. This average (which comes from applying Equation (1) to our electricity data) aligns with irrigation budgets in agronomic studies. 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).

Table 1: Summary statistics

A. Service point $\times$ month sample	Matched to pumps	Unmatched
Unique service points (SPs)	10,248	90,137
Months observed (2008–2019)	110.7 (41.0)	105.8 (45.8)
Average electricity consumption (kWh/month)	10,122 (13,668)	4,962 (33,390)
Average marginal electricity price ($/kWh)	0.13 (0.03)	0.16 (0.04)
Average electricity bill ($/month)	1,813.16 (1972.45)	836.05 (3718.41)
Average groundwater consumption (AF/month)	33.45 (42.76)	
Average marginal groundwater cost ($/AF)	47.37 (25.49)	
B. Parcel $\times$ year sample	Parcels containing matched SPs	
Unique parcels with SP-pump matches	6,416	
Count of matched SPs per parcel	1.82 (1.50)	
"Croppable" area of parcel (acres)	290.57 (283.59)	
Average share of annual crops	0.217 (0.277)	
Average share of fruit/nut perennial crops	0.453 (0.403)	
Average share of hay perennial crops	0.248 (0.284)	
Average share of non-crop (fallow)	0.082 (0.134)	
Average groundwater use (AF/acre)	3.987 (4.699)	
Parcel within surface water district (1/0)	0.600 (0.490)	

Notes: We report means and standard deviations (in parentheses) across unit-specific averages. Panel A uses monthly billing data for PGE agricultural electricity service points (i.e. meters). We restrict our analysis to the 10,248 SPs that we can match to a vertical agricultural groundwater pump audit. The right column includes all unmatched agricultural SPs. We omit unmatched SPs from our analysis to avoid mistakenly including other agricultural electricity end-uses (e.g. refrigeration). All reported unmatched sample means are statistically different from the matched sample ($p < 0.01$). We cannot calculate kWh/AF, AF/month, or $/AF for unmatched serivce points. Panel B aggregates the matched sample up to the parcel-year level. For consistency with our reduced-form and dynamic discrete choice analyses, we: weight parcels by (time-invariant) croppable acreage; use croppable (rather than total) acreage to denominate crop shares; omit parcels with less than 1 or greater than 5,000 croppable acres; and omit parcels with annual electricity bills exceeding $3,000 per croppable acre.

4.1 Monthly demand elasticities

We estimate groundwater demand using the following two-stage least squares specification:

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

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

The outcome variable is the natural logarithm of groundwater extractedat service point i in month t. The explanatory variable is the natural logarithm of the marginal cost of groundwater. δ_{it} and ψ_{it} are a set of fixed effects, which we describe below. ϵ_{it} and ν_{it} are idiosyncratic errors, which we two-way cluster by service point and month-of-sample.

To econometrically identify the demand elasticity, γ, in Equation (2), we leverage cross-sectional and time-series variation in electricity prices (a component of P_{it}^{water}). PGE's tariff schedules are the outcome of statewide regulatory proceedings, which individual farmers cannot plausibly influence. Moreover, tariff decisions are made 1–3 years in advance, and they do not reflect current conditions (e.g., droughts). We therefore treat the tariff *schedules* as plausibly exogenous.

However, since farmers have some choice over their tariff (conditional on the schedules), the *prices* they face may be endogenous.[26] Specifically, PGE sorts farmers into four mutually exclusive tariff categories based on their pump capacity (smaller vs. larger than 35 horsepower) and their electricity meter type (conventional vs. smart meters). Farmers with smart meters are able to choose tariffs from category-specific menus. We therefore instrument for the marginal cost of water P_{it}^{water} using the "default" marginal electricity price for farmer i's category ($P_{it}^{\text{elecDefault}}$).[27]

26. Unlike PGE's residential tariffs, which have increasing block prices, making a household's marginal price endogenous to its own consumption (Ito (2014)), PGE's agricultural tariffs have linear volumetric prices. This means that unit i's marginal electricity price in a given hour depends solely on its tariff.

27. Small-conventional and large-conventional categories each comprise a single (default) tariff. Small-smart and large-smart categories comprise 8 and 12 tariffs respectively, which we plot in Appendix Figure C2. We assign the simplest (i.e. least time-varying) tariff in each category as the default tariff; instrumenting with each category's modal tariff yields similar results (see Appendix Tables B6 and B10). We omit a fifth category reserved for the 1.7% of farmers transitioning off of internal combustion engines, since they are likely not comparable to the rest of our sample, and we do not observe them before they consumed electricity.

Figure 2: Marginal electricity prices for four default tariffs

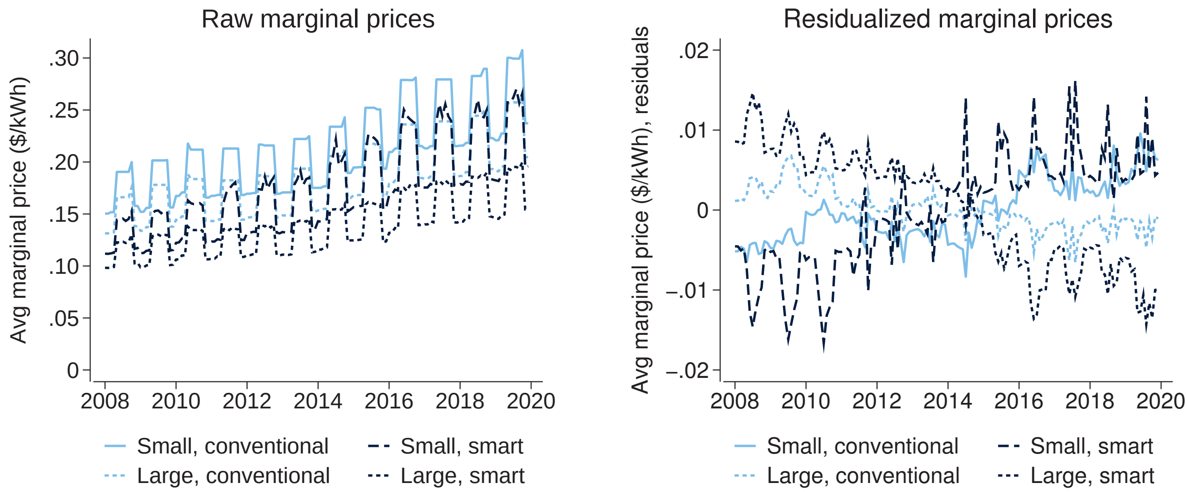

Notes: This figure plots times series of monthly average marginal electricity prices ($/kWh) for PGE's four default agricultural tariffs. The left panel plots raw average marginal prices for each month in our estimation sample, taking unweighted averages across all hours. Marginal prices are systematically higher in summer months, and on non-holiday weekdays. The right panel plots residuals of these four time series after partialing out tariff × month-of-year and month-of-sample 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 (≥ 35 hp) pumps, and conventional vs. smart meters. Our identifying variation comes from differential price changes across default tariffs, as well as PGE's smart meter rollout—which exogenously shifted many customers from conventional to smart categories, lowering their marginal price.

This instrument eliminates selection bias from a high-volume pumper choosing a tariff with advantageously low volumetric prices. It also eliminates: (i) other potentially endogenous variation in P_{it}^{water} (e.g., groundwater demand shifters); (ii) within-pump simultaneity, whereby high Q_{it}^{water} can depress localized groundwater levels and increase P_{it}^{water}; and (iii) measurement error in P_{it}^{water} from imprecision in parameterizing Equation (1).[28]

Having purged within-category variation, we identify γ off of differential changes in electricity prices *across* categories over time. The left panel of Figure 2 plots raw time series of $P_{it}^{\text{elecDefault}}$ during our sample period. Equations (2)–(3) include a series of fixed effects to address potential confounders. We use month-of-sample fixed effects to control for common trends in prices that might be correlated with unobservable determinants of groundwater pumping (e.g., crop prices). We also use unit-by-month-of-year fixed effects to control for seasonality in $P_{it}^{\text{elecDefault}}$, which may also be correlated with unit i's seasonal irrigation patterns. The right panel of Figure 2 shows that, after partialling out these fixed effects, there remains substantial across-category variation in $P_{it}^{\text{elecDefault}}$ over time.

28. In Equation (2), measurement error from converting kWh to AF enters in P_{it}^{water} (i.e., $/kWh ÷ AF/kWh) and in Q_{it}^{water} (i.e., kWh × AF/kWh). Instrumenting with default electricity prices negates the correlation between left-hand-side and right-hand-side measurement error. An un-instrumented OLS regression returns a larger elasticity (in absolute value), which we report in Appendix Table B6.

We also include groundwater-basin-by-year and water-district-by-year fixed effects, to control for differential selection related to changes in depth or surface water allocations. A remaining concern is selection across tariff categories. We observe no bunching at the 35 horsepower cutoff between small vs. large pumps.[29] Moreover, only 4% of service points in our sample switch between small- and large-pump categories; we use unit-by-large-pump fixed effects to control for these shifts (i.e., a small-to-large-pump switch mechanically increases Q_{it}^{water} and decreases $P_{it}^{\text{elecDefault}}$). In contrast, 21% of service points switch from conventional- to smart-meter categories during our sample period due to PGE's smart meter rollout. Since the timing of this rollout reflected institutional factors outside of farmers' control, meter-induced category switches provide additional plausibly exogenous variation in $P_{it}^{\text{elecDefault}}$.[30]

Table 2 presents our demand estimates. We find a monthly groundwater cost elasticity of -0.294 (Column (1); $p < 0.01$), implying that farmers do respond to short-run increases in groundwater costs by reducing groundwater pumping. To rule out the possibility that farmers respond to higher costs by investing in pump upgrades (which would affect our kWh-to-AF conversion), we restrict the sample to the 30% of observations within 12 months of an observed pump test; this yields a similar elasticity estimate (Column (2); $p < 0.10$). Our results are highly robust, including to: using various windows around a pump test; interacting month-of-sample fixed effects with service point characteristics; adding time-varying controls; and alternate kWh-to-AF conversions.[31] We also estimate the analogous model of electricity demand (replacing Q_{it}^{water} and P_{it}^{water} with Q_{it}^{elec} and P_{it}^{elec}), which yields a nearly identical monthly elasticity estimate of -0.285 (Column (3); $p < 0.01$). These results provide strong evidence that farmers' short-run irrigation decisions are sensitive to cost.

29. See Appendix Figure C3. Appendix Table B2 reveals that our results are similar if we interact month-of-sample fixed effects with deciles of horsepower.

30. During our sample period, PGE gradually replaced any remaining conventional (analog) meters with smart (digital) meters, for both agricultural and non-agricultural customers. Previous research has established that PGE 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 they are systematically correlated with unobserved changes in pumping behavior. Our results are robust to interacting month-of-sample fixed effects with predictors of the rollout (Appendix Table B5) and lagging the instrument to mitigate any rollout anticipation effects (Appendix Table B6).

31. We present these sensitivities in Appendix B.1.

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

	(1) $\log(Q^{\text{water}})$	(2) $\log(Q^{\text{water}})$	(3) $\log(Q^{\text{elec}})$
$\log\left(P^{\text{water}}\ (\$/\text{AF})\right)$	-0.294^{***}	-0.260^{*}	
	(0.080)	(0.134)	
$\log\left(P^{\text{elec}}\ (\$/\text{kWh})\right)$			-0.285^{***}
			(0.080)
Pump test within 12 months		Yes	
Service point units	10,091	9,202	10,091
Months	148	124	148
Observations	953,469	276,958	953,469
First-stage estimates			
$\log\left(P^{\text{elecDefault}}\ (\$/\text{kWh})\right)$	1.321^{***}	1.486^{***}	1.306^{***}
	(0.024)	(0.060)	(0.016)
Kleibergen-Paap F-statistic	3079	609	6341

Notes: Column (1) estimates Equations (2)–(3) at the service point (SP) by month level. The outcome variable is the natural logarithm of groundwater consumption, and we instrument for pumping costs with default electricity prices. This recovers an estimate of the cost elasticity of groundwater demand. Column (2) restricts the sample to SP-months within 12 months of an observed pump test, to minimize unobserved changes in pump efficiency. Column (3) uses the quantity and price of electricity (which we observe directly), rather than the quantity and price of groundwater (which we construct). All regressions use two-stage least squares, and include the following fixed effects: SP $\times$ month-of-year (to control for seasonality), SP $\times$ $\mathbf{1}$[large pump] (to control for tariff category switches), groundwater basin $\times$ year (to control for trends in depth), water district $\times$ year (to control for changes in surface water availability), and month-of-sample. Standard errors (in parentheses) are two-way clustered by service point and month-of-sample. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

4.2 Annual elasticities and crop choice

For our main short-run estimates, we next aggregate from service points to parcels (i.e., our definition of farms). This links groundwater pumps to the fields they likely irrigate, while accounting for multiple pumps within the same farm. We also aggregate from months to years to align with the annual cropping cycle. To estimate Equations (2)–(3) at the parcel-year level, we use the following fixed effects: parcel-by-large-pump (to control for unit-specific means and small-to-large category switches), groundwater-basin-by-year (to control for differential trends in depth), and water-district-by-year (to control for changing surface water allocations). We weight these regressions by each parcel's croppable acreage (excluding forests, development, etc.), making our estimates representative per acre of cropland.

We also use this parcel-year panel to estimate an *intensive-margin* cost elasticity of groundwater demand, holding each parcel's crop category (i.e., annuals, fruit/nut perennials, hay perennials, or non-crop) constant. To do this, we restrict our sample to parcels that

Table 3: Short-run groundwater elasticities do not reflect crop switching

	(1)	(2)	(3)	(4)	(5)	(6)
	$\log(Q^{\text{water}})$	$\log(Q^{\text{water}})$	Share annuals	Share fruit/nut	Share hay	Share non-crop
	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS
$\log\left(P^{\text{water}}\ (\$/\text{AF})\right)$	-0.762^{***}	-0.679^{***}	0.022	-0.011	-0.007	-0.003
	(0.219)	(0.230)	(0.017)	(0.020)	(0.018)	(0.017)
Intensive margin		Yes				
Parcel units	6,388	6,287	6,415	6,415	6,415	6,415
County-years	367	334	367	367	367	367
Observations	54,220	41,445	55,363	55,363	55,363	55,363
First-stage estimates						
$\log\left(P^{\text{elecDefault}}\ (\$/\text{kWh})\right)$	1.409^{***}	1.318^{***}	1.414^{***}	1.414^{***}	1.414^{***}	1.414^{***}
	(0.042)	(0.051)	(0.042)	(0.042)	(0.042)	(0.042)
Kleibergen-Paap F-statistic	1121	660	1139	1139	1139	1139

Notes: Column (1) estimates Equations (2)–(3) at the parcel by year level, using two-stage least squares. The outcome variable is the natural logarithm of groundwater consumption, and we instrument for pumping costs with default electricity prices. This recovers an estimate of the short-run cost elasticity of groundwater demand. Column (2) is identical, but restricts the sample to parcel-years with the same modal crop category (i.e., annuals, fruit/nut perennials, hay perennials, non-crop) as the preceding year and interacts parcel fixed effects with the four categories; this shuts down the crop-switching channel that we structurally estimate below. Columns (3)–(6) are identical to Column (1), with shares of acres in a crop category for each parcel-year as outcome variables. All regressions include the following fixed effects: parcel × **1**[large pump] (to capture tariff category switches), 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: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

chose the same crop category as the preceding year, while also adding parcel-by-crop-category fixed effects (to control for parcel-crop-specific irrigation needs). Finally, we explore the crop-switching margin directly by replacing the dependent variable in Equation (2) with the share of the parcel's (croppable) acres allocated to each crop category in each year.

Table 3 reports these annual estimates. In Column (1), we find that the short-run groundwater demand elasticity is -0.762 ($p < 0.01$). This is larger than the monthly elasticity estimate of -0.294, consistent with farmers having more flexibility over longer time intervals. Column (2) reports our intensive-margin elasticity estimate of -0.679 ($p < 0.01$), which is nearly as large as the Column (1) estimate that includes both margins—indicating that farmers' responses to year-on-year changes in groundwater costs are not driven by the crop switching margin. The semi-elasticity estimates in Columns (3)–(6) provide corroborating evidence: for all four crop categories, we recover precise null estimates on crop switching due to year-on-year groundwater cost shocks. Together, these results suggest that we do not detect crop-switching behavior because our panel fixed effects estimator does not incorporate dynamics.

5 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 (Scott (2013); Kalouptsidi, Scott, and Souza-Rodrigues (2021)) to estimate parameters without needing to specify the evolution of individual market-level states. We then use our estimated dynamic model to simulate long-run steady states under counterfactual groundwater taxes. These results generate long-run (semi-)elasticities of crop choice, groundwater use, and electricity use with respect to the marginal cost of groundwater. Our model closely follows Scott (2013), who estimates how changes in crop revenues impact land use. We extend this prior work by (i) estimating land use responses to input costs rather than output prices; (ii) expanding from two to four land-use categories; (iii) allowing for both intensive- and extensive-margin responses to water costs in counterfactual simulations; (iv) incorporating expectations over drought, a key market state variable, in simulations; and (v) going beyond land use to generate elasticities for electricity and groundwater.

5.1 Model of crop choice

We model annual farmer profits on a given field as a function of crop choice, with crop-specific groundwater pumping costs. Each year, a farmer chooses a crop $c \in \mathcal{C} = \{$annuals, fruit/nut perennials, hay perennials, non-crop$\}$ to maximize expected discounted profits over an infinite time horizon.[32] 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

32. 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

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). 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.[33] $\alpha(c_t, k_t)$ represents the time-invariant component of average net returns to 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. $\xi(c_t, k_t, \omega_t)$ 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 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 $\alpha(c_t, k_t)$.[34] 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

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).

33. 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.

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

costs than switching *into* the same perennial crop because the latter comes with an upfront investment cost.

Formally, we assume the field state in year t is Markovian and depends on only the preceding year's crop choice, not on choices in prior years: $k_t = c_{t-1}$. Thus, profit in year t is unaffected by choices made prior to year $t-1$. This assumption captures the salient feature of perennial cropping in our setting—high upfront costs followed by a stream of annual harvests with lower recurring costs.[35] As a result, any crop choice $c \in \mathcal{C}$ is a "renewal action," meaning that choice c_t will yield a particular field state in the following year k_{t+1} regardless of states in prior years (Kalouptsidi, Scott, and Souza-Rodrigues (2021)).[36]

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 also treat each field as independent.[37]

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 (2022)). The resulting conditional choice probabilities (CCPs), or the probability that

35. This assumption is sufficient to generate distinct responses to short-run transient cost shocks vs. long-run permanent cost shocks. Without field state dependence, farmers might unrealistically tear out their almond orchards in response to a single-year shock and then replant almonds the following year. By assuming even a single year of field state dependence, we impose crop switching costs that would likely make this kind of response unprofitable. In doing so, we abstract from second-order dynamic considerations that are less likely to impact the long-run cropping response, such as the time it takes for young perennials to reach maturity and bear fruit, because (i) modeling multiple perennial vintages would be less tractable and (ii) growth profiles vary substantially across crops within our "fruit/nut perennials" category.

36. 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.

37. If one landowner operates multiple fields, this assumption implies that maximizing profits jointly across fields would be equivalent to maximizing profits for each field independently.

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).[38] 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. By Assumption 2, any crop choice in year $t+1$ is a renewal action, which we denote as 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:[39]

$$\begin{aligned}
\ln\left[\frac{p(c_t, k_t, \omega_t)}{p(c'_t, k_t, \omega_t)}\right] + \beta \ln\left[\frac{p(r_{t+1}, c_t, \omega_{t+1})}{p(r_{t+1}, c'_t, \omega_{t+1})}\right] &= \alpha_G\left[G(c_t, \omega_t) - G(c'_t, \omega_t)\right] \\
&+ \alpha(c_t, k_t) - \alpha(c'_t, k_t) + \beta\left[\alpha(r_{t+1}, c_t) - \alpha(r_{t+1}, c'_t)\right] \\
&+ \xi(c_t, k_t, \omega_t) - \xi(c'_t, k_t, \omega_t) + \beta\left[\xi(r_{t+1}, c_t, \omega_{t+1}) - \xi(r_{t+1}, c'_t, \omega_{t+1})\right] \\
&+ \beta\left[e^V(c_t, \omega_t, \omega_{t+1}) - e^V(c'_t, \omega_t, \omega_{t+1})\right]
\end{aligned} \tag{7}$$

Each side of Equation (7) is equivalent to the difference in values between choosing c_t or c'_t in year t, followed by renewal action r_{t+1}, and then choosing optimally in all following years.

5.2 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

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

39. 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.

endogeneity.[40] 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.[41] We construct CCPs by aggregating crop choices within each market, and we use observed groundwater costs in the market to construct crop-specific pumping costs.

5.2.1 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.[42]

Groundwater cost, groundwater use, and electricity use For each parcel, we observe realized groundwater pumping costs, groundwater quantity, and electricity quantity. We

40. For each field, we observe groundwater costs for the *chosen* crop, but not CCPs or groundwater costs for all other possible crop choices.

41. We first split by PGE'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 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.

42. 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.

project these outcomes at the parcel level using the following OLS specification:

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

where O_{ft} is the per-acre outcome (pumping cost, groundwater quantity, or electricity quantity) for parcel f in year t. F_{ft}^c is the fraction of parcel f planted with crop c in year t, omitting non-crop ($c = 0$) to avoid collinearity. η_f are parcel fixed effects, ϕ_{mt} are market-year fixed effects, and ι_{ft} is an idiosyncratic error term. $\zeta_m^c + \kappa_m^c t$ recovers the average per-acre outcome 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 groundwater needs (e.g., due to surface water allocations) and irrigation needs within each crop category c (e.g., grape- vs. almond-growing regions).

Using these fitted regression models, we project per-acre outcomes for each parcel-year under each crop category. Then, we aggregate projections up to the market-crop-year by taking an acreage-weighted median of all parcels within that market-year.[43] This aggregation yields total variable groundwater costs $G_{mct} = G_m(c_t, \omega_{mt})$, groundwater quantities $\hat{Q}_{mct}^{\text{water}}$, and electricity quantities $\hat{Q}_{mct}^{\text{elec}}$, each at the market-crop-year level and measured per-acre.[44]

5.2.2 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., $c_t' = 0$ and $r_{t+1} = 0$), leaving c_t to denote any of the three other crop choices. We estimate this resulting regression equation at the market level, using the data described above and weighting markets by acreage:

43. This procedure yields crop-specific estimates for total variable pumping costs, groundwater quantities, and electricity quantities for a typical acre in each market and year. Taking acreage-weighted means yields similar results (see Appendix Figure A2). We further calculate a time-invariant market-level measure of each crop-specific outcome by taking the weighted median over all parcel-years in a market, which we use as steady-state costs and quantities in our counterfactual simulations. We also aggregate separately by drought vs. non-drought years to incorporate drought expectation in our counterfactual simulations.

44. 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.

$$\ln\left[\frac{p_m(c_t,k_t,\omega_{mt})}{p_m(0,k_t,\omega_{mt})}\right]+\beta\ln\left[\frac{p_m(0,c_t,\omega_{mt+1})}{p_m(0,0,\omega_{mt+1})}\right]=\alpha_G\Delta G_{mct}+\tilde{\Delta}\alpha_{mck}+\tilde{\Delta}\xi_{mckt}+\Delta e^V_{mct} \quad (9)$$

The outcome variable 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, which we construct from our calculated CCPs. Each right-hand-side term represents a component of this difference.[45] 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 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 $\mathbb{E}[\tilde{\Delta}\xi_{mckt}+\Delta e^V_{mct}\mid\Delta G_{mct}]=0$. While the expectational error Δe^V_{mct} is uncorrelated with ΔG_{mct} by construction, unobserved net returns $\tilde{\Delta}\xi_{mckt}$ may be correlated with ΔG_{mct}.[46] 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} as the product of plausibly exogenous quantity and price measures. The quantity is the time-invariant average electricity quantity needed for groundwater pumping for crop c (relative to non-crop) in market m, $\Delta\hat{Q}^{\text{elec}}_{mc}$. The price is the time-varying average default electricity price in market m, $P^{\text{elecDefault}}_{mt}$.[47] This instrument,

45. 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.

46. 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.

47. 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

$\Delta\hat{Q}^{\text{elec}}_{mc} \times P^{\text{elecDefault}}_{mt}$, is strongly correlated with ΔG_{mct} since changes to default electricity prices influence the variable costs of groundwater pumping. It is also plausibly excludable, since exogenous electricity tariffs should only influence pumping costs (as discussed in Section 4.1) and time-invariant $\Delta\hat{Q}^{\text{elec}}_{mc}$ is uncorrelated with annual variation in the market state.

Recovering profit intercept parameters We require estimates of 16 profit intercepts for each market—one $\alpha_m(c_t, k_t)$ for each crop choice-field state pair. However, Equation (9) only includes 12 $\tilde{\Delta}\alpha_{mck}$ intercept terms for each market, since we use the non-crop category as the comparison. Recovering all 16 intercepts therefore requires additional assumptions.[48] 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.[49]

5.3 Counterfactual simulations

We use estimated model parameters to simulate steady-state counterfactuals under different groundwater tax scenarios. In each scenario, we proceed as follows. First, we use Equation (4) to calculate expected annual profit—which is 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. Third, 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.[50] To simulate these counterfactuals, we make two additional assumptions.

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.

48. As described by Scott (2013) and Kalouptsidi, Scott, and Souza-Rodrigues (2021), dynamic discrete choice models are typically not fully identified (Magnac and Thesmar (2002)). See Appendix A.2 for a mathematical statement of these assumptions.

49. $\alpha_m(c_t, k_t)$ incorporates both net returns to crop c_t and any additional transition costs due to field state k_t. For $c_t = k_t$, 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.

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

Assumption CF1: Drought state Drought, which is included in the market state ω_{mt}, is a key determinant of annual groundwater pumping costs.[51] To incorporate drought conditions 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.[52]

Assumption CF2: Intensive-margin response Our reduced-form analysis shows that farmers respond to groundwater cost shocks on the intensive margin (i.e., adjust their water use even conditional on crop choice). We therefore also allow farmers to adjust each crop's groundwater use in response to counterfactual groundwater taxes.[53] To calibrate this long-run within-crop adjustment, we estimate a series of long-differences regressions—analogous to Equations (2)–(3) in our reduced-form analysis—using increasingly longer differences. These estimates imply that roughly 35% of our short-run intensive-margin elasticity estimate is likely to persist in the long run, translating into a long-run intensive-margin elasticity of -0.238.[54] This intensive-margin adjustment reduces groundwater and electricity consumption for a particular crop, thereby lowering the crop's total variable groundwater costs.[55]

51. Drought can increase groundwater needs due to low precipitation levels and curtailed surface water allocations. It can also increase groundwater scarcity, which raises marginal pumping costs. Both effects increase farmers' groundwater expenditures.

52. 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 A2).

53. The profit function (Equation (4)) includes total variable pumping costs, or the product of groundwater quantity and marginal cost: $G = Q^{\text{water}} \times P^{\text{water}}$. A marginal groundwater tax will cause farmers to reoptimize Q^{water} within each crop c (i.e., $\partial Q^{\text{water}}/\partial P^{\text{water}}$). This intensive-margin response is already accounted for when projecting *factual* pumping costs and quantities, which we then use to estimate Equation (9), because observed Q^{water} has been optimized to *factual* P^{water}. However, simulating counterfactual P^{water} necessitates assumptions on this optimal $\partial Q^{\text{water}}/\partial P^{\text{water}}$ response.

54. Our short-run elasticity estimate (-0.679) is identified using short-run cost shocks. However, our counterfactuals simulate a permanent cost shock, and many intensive-margin responses may not persist over the long run. For example, while a farmer may respond to a one-year groundwater cost shock by allowing her crop to fail, she is unlikely let the same crop fail year after year. Alternative simulations vary this assumption of 35% persistence. Appendix A.3 describes these regressions, and Appendix Figure A1 presents their results.

55. Suppose the marginal cost of pumping increases by 20% due to a tax. With a long-run intensive-margin elasticity of -0.238, groundwater quantity falls by 4.25%. As a result, the total variable cost of pumping increases by 14.9%, not the full 20%. For this reason, our (semi-)elasticities are with respect to the *marginal* groundwater cost, even though we use total *variable* groundwater costs to estimate the model.

This additional margin of response alters expected profits for each crop under 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 counterfactual groundwater tax scenarios, we increase the marginal cost of pumping by a specified tax percentage. Since we use crop-specific groundwater costs, a given tax increases total variable costs more for relatively water-intensive crops, which can induce 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. 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 steady-state 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:[56]

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

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 500 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 500 sets of (semi-)elasticities for each tax scenario. Our reported (semi-)elasticities are the means of these

56. 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.

distributions, and our reported 95% confidence intervals span the 2.5th and 97.5th percentiles of the distributions.

5.4 Static model

Following Scott (2013), we also estimate a static model of crop choice for comparison. While this model does not reflect real-world farmer behavior, it provides a useful comparison for both our reduced-form analysis and our dynamic model. This static model follows directly from the dynamic model, with two important changes. First, farmers are not forward-looking, so we set the discount factor to $\beta = 0$; this means that only current-year profits enter the value function. Second, farmers do not account for field state dependence, so we remove all field state conditioning from the model. The static estimating equation is therefore:

$$\ln\left[\frac{p_m(c_t, \omega_{mt})}{p_m(0, \omega_{mt})}\right] = \alpha_G \Delta G_{mct} + \Delta\alpha_{mc} + \Delta\xi_{mct}$$

The outcome variable is the difference in profits between crop choice c and non-crop in year t, while each term on the right-hand side represents a component of this difference.[57] ΔG_{mct} is the difference in pumping costs between the two crop choices. $\Delta\alpha_{mc}$ is a set of intercept terms that reflect the difference in average net returns between the two crop choices. $\Delta\xi_{mct}$ is an unobserved term that reflects the difference in time-varying profits. Importantly, none of these terms depend on year $t+1$ or field state k, unlike in the dynamic estimating equation.

We estimate this model using the same empirical strategy, instrumenting for ΔG_{mct} with the product of plausibly exogenous quantity and price measures. We then use our estimated static model to simulate static counterfactuals following the same structure outlined above, but invoking the modified assumptions on the discount factor and field state dependence.

57. See Appendix A.4 for mathematical definitions of each of these terms.

Figure 3: Long-run elasticities with respect to groundwater pumping cost

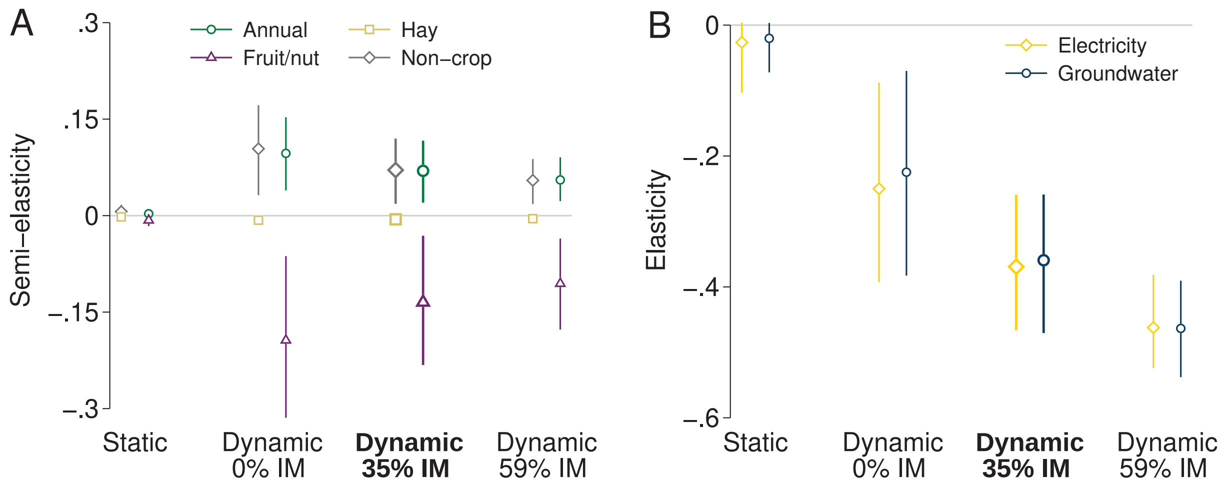

Notes: This figure plots long-run (semi-)elasticities of land use (panel A) and groundwater and electricity use (panel B) with respect to groundwater cost, estimated using our discrete choice model. To recover these (semi-)elasticities, we simulate a 20% increase in groundwater pumping costs and compute a steady state over a 20-year horizon. In the static model, there is no state dependence and farmers are not forward looking. In the myopic model, there is state dependence, but farmers are not forward looking. In the dynamic models, there is state dependence and farmers are forward looking. In "0% IM", farmers can only respond to groundwater cost changes by changing crops. In "35% IM," farmers can also respond on the intensive-margin, by reducing water use conditional on crop choice by 35% of our within-crop reduced-form elasticity, our central estimate based on our preferred long-differences model. We also present a specification in which farmers' intensive-margin response is 59% of our reduced-form estimate, per a long-differences model with fewer fixed effects. Panel A shows semi-elasticities for annual crops (circles), hay perennials (squares), fruit/nut perennials (triangles), and non-crop (diamonds). Panel B shows the electricity of demand for electricity (diamonds) and groundwater (circles). The reported elasticities and semi-elasticities are the means over 500 draws for each model. The plotted 95% confidence intervals (vertical lines) show the 2.5th and 97.5th percentile over draws.

5.5 Model results

Figure 3 presents our main discrete choice results. Panel A plots semi-elasticities of land use, while Panel B presents demand elasticities for groundwater and electricity, all for a permanent 20% tax on the marginal cost of groundwater pumping.[58]

Static model We estimate static semi-elasticities with respect to marginal groundwater costs of 0.003 (s.e. 0.002), −0.007 (s.e. 0.005), −0.002 (s.e. 0.002), and 0.006 (s.e. 0.005) for annual crops, fruit/nut perennials, hay perennials, and non-crop (fallowing), respectively. These results align with our reduced-form results in Table 3, where we find precise null cropping responses to year-on-year changes in marginal groundwater costs. Consistent with these null results, our static model suggests that farmers' groundwater and electricity consumption is close to perfectly inelastic, with elasticity estimates of −0.020 (s.e. 0.019) for groundwater

58. Appendix Table A2 presents these results in tabular form. Appendix Table A1 presents the parameter estimates resulting from estimating Equation (9). Appendix Figure A2 presents robustness to parcel sample selection and aggregation to the market level.

and -0.026 (s.e. 0.027) for electricity. All six estimates are precise and not statistically different from 0 (i.e., our 95% confidence intervals span 0 for all six (semi-)elasticities). These results imply that, in the absence of field state dependence and forward-lookingness, farmers would not switch crops in response to permanent increases in marginal pumping costs.

Dynamic model When we incorporate both forward-lookingness and field state dependence in our dynamic model, we find that farmers exhibit a large cropping response to marginal groundwater costs, driving large groundwater and electricity elasticities.

We first present results assuming no intensive-margin response, meaning that farmers only respond to changes in marginal groundwater costs on the crop-switching margin ("0% IM" in Figure 3). We find cropping semi-elasticities that are roughly 20 times their static counterparts: 0.097 (s.e. 0.030), -0.194 (s.e. 0.067), -0.007 (s.e. 0.004), and 0.104 (s.e. 0.035) for annuals, fruit/nut perennials, hay perennials, and non-crop, respectively. The signs of these effects are intuitive: farmers reduce land in water-intensive fruit/nut perennial crops, while increasing land in annual crops and substantially increasing fallowing. Correspondingly, we find that crop switching alone leads to elasticities of -0.225 (s.e. 0.081) for groundwater and -0.250 (s.e. 0.075) for electricity.

Next, we invoke Assumption CF2 and incorporate the long-run intensive-margin response into our dynamic model. In our central case (35% IM, with an intensive-margin elasticity of -0.238), we estimate cropping semi-elasticities of 0.070 (s.e. 0.024) for annual crops, -0.135 (s.e. 0.052) for fruit/nut perennials, -0.006 (s.e. 0.004) for hay perennials, and 0.071 (s.e. 0.026) for non-crop. These semi-elasticity estimates are slightly attenuated from the 0% IM case, consistent with farmers' ability to reduce within-crop water use as a substitute for crop switching. Adding this intensive-margin channel yields larger elasticities of -0.359 (s.e. 0.058) for groundwater and -0.369 (s.e. 0.055) for electricity—showing that farmers' intensive- vs. extensive-margin substitution is incomplete.[59]

To summarize, our static model corroborates our reduced-form land use regressions, finding no statistical evidence of crop switching in response to marginal groundwater costs.

59. When we increase the intensive-margin response to 59% of our short-run estimate (an intensive-margin elasticity of -0.401)—per our alternative long-differences specification with fewer fixed effects—these two effects are magnified. See Appendix A.3 for details.

As a result, our static model would suggest that farmers' groundwater and electricity use are close to perfectly inelastic, consistent with our finding that crop switching does not explain our reduced-form elasticity estimates in Table 3. In stark contrast, our dynamic model reveals substantial cropping semi-elasticities. In this more realistic model with forward-looking farmers with field state dependence, larger semi-elasticities translate into meaningful long-run elasticities of groundwater and electricity use. Our results demonstrate that both the magnitude and mechanisms underlying farmers' short- and long-run responses to groundwater policy differ substantially, highlighting the importance of estimating both.

5.6 Impacts of counterfactual groundwater policy

Finally, we use our dynamic model (under our central 35% IM assumption) to understand the potential effectiveness of California's landmark groundwater policy, SGMA.[60] In Figure 4, we present simulation results for land use (Panel A) and groundwater and electricity use (Panel B) responses to counterfactual groundwater taxes. Panel A shows that as we increase the groundwater tax from zero, farmers reduce the amount of land planted in fruit/nut perennials (purple triangles) and hay perennials (beige squares), while increasing the amount of land planted in annual crops (green circles) and fallowed (gray diamonds). Panel B shows the resulting reductions in electricity and groundwater use under each counterfactual tax.

Under SGMA, the average Groundwater Sustainability Plan will require groundwater pumping reductions of 19.2% to achieve sustainability.[61] Panel B implies that this will require a groundwater tax of roughly 60%. A tax of this size would have meaningful impacts on land use, leading approximately 8.6% of acres to change crop category: 8.2% of total cropland switching out of fruit/nut perennials, 0.5% switching out of hay, 3.6% switching into annual crops, and 5.0% switching to fallowing. These changes are substantial, representing a 24% decline in fruit/nut perennials, a 2% decline in hay, a 12% increase in annual crops, and a 50% increase in fallowing compared to our no-tax scenario. These results illustrate that,

60. 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 estimates speak directly to these plans, and provide a heuristic for possible responses to non-price instruments.

61. See Appendix C.7 for details on our Groundwater Sustainability Plan data.

Figure 4: Farmer responses to groundwater taxes

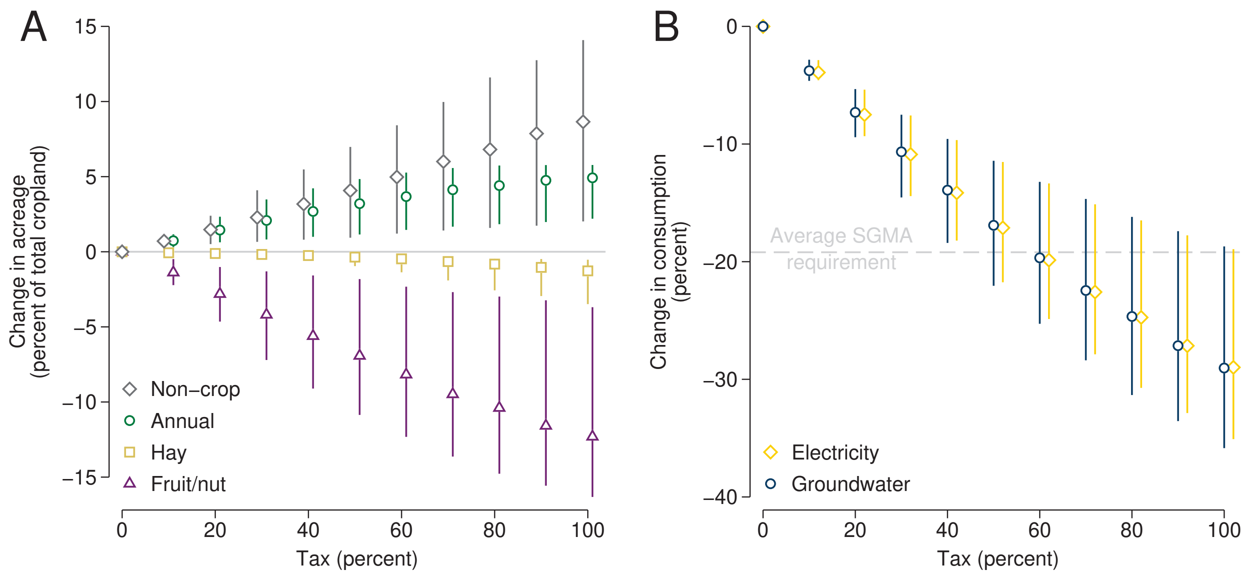

Notes: This figure plots counterfactual changes in crop choice (panel A), and groundwater extraction and electricity use (panel B) resulting from varying groundwater taxes, estimated using our dynamic discrete choice model. Panel A shows acres in our sample in annual crops (circles), hay perennials (squares), fruit/nut perennials (triangles), and non-crop (diamonds). Panel B shows the percent change in electricity (diamonds) and groundwater (circles) for each tax level. 95% confidence intervals, constructed from the 2.5th and 97.5th percentile of the sampling distribution across 500 simulations, are shown with vertical lines. In Panel B, the light gray dashed line shows the average groundwater pumping reduction in overdrafted regions under SGMA, 19.2%.

while SGMA's sustainability targets are likely attainable with a moderate groundwater tax, they are likely to induce a major reallocation of California's 20 million acres of cropland.

6 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 in two empirical strategies to estimate the elasticity of demand for groundwater over different time horizons. First, we use a panel fixed effects model, in which farmers respond to year-over-year changes in groundwater costs, to estimate a short-run elasticity. Second, we use a dynamic discrete choice model of land use, which models farmers as forward-looking and accounts for field state dependence, to estimate a long-run elasticity. We find that both the magnitude of and the mechanisms underlying farmer responses differ between the short run and the long run. Our long-run elasticity estimate of -0.359 is smaller than our annual elasticity estimate of -0.762, but is driven by substantial changes in cropping patterns that do not arise in the short run.

Our dynamic model predicts that the equivalent of a 60% tax on groundwater pumping costs will be required to achieve California's sustainability goals. 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- vs. 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

Araujo, Rafael, Francisco Costa, and Marcelo Sant'Anna. 2022. "Efficient forestation in the Brazilian Amazon: Evidence from a dynamic model." Working paper.

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.

Borenstein, Severin. 2017. "Private Net Benefits of Residential Solar PV: The Role of Electricity Tariffs, Tax Incentives, and Rebates." *Journal of the Association of Environmental and Resource Economists* 4 (S1): S85–S122.

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, and Nick Hagerty. 2023. "The anticipatory effects of regulation in open access." Working paper.

Bruno, Ellen, 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, and Katrina Jessoe. 2021a. "Missing markets: Evidence on agricultural groundwater demand from volumetric pricing." *Journal of Public Economics,* no. 196, 104374.

Bruno, Ellen M., Jeffrey Hadachek, Nick Hagerty, and Katrina Jessoe. 2023. *External costs of climate change adaptation: Agricultural wells and access to groundwater.* Working paper.

Bruno, Ellen M., and Katrina Jessoe. 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. Forthcoming. "The dynamic impacts of pricing groundwater." *Journal of the Association of Environmental and Resource Economists.*

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."

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 April 2015.

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

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

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

———. 2024. "Is the world running out of fresh water?" *American Economic Association: Papers and Proceedings* 114:31–35.

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.

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 (C): 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.

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.

Fell, Harrison, Shanjun Li, and Anthony Paul. 2014. "A new look at residential electricity demand using household expenditure data." *International Journal of Industrial Organization* 33:37–47.

Fowlie, Meredith, Michael Greenstone, and Catherine Wolfram. 2018. "Do Energy Efficiency Investments Deliver? Evidence from the Weatherization Assistance Program." *Quarterly Journal of Economics* 133 (3): 1597–1644.

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

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

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

Hanak, Ellen, Jay Lund, Ariel Dinar, Brian Gray, Richard Howitt, Jeffrey Mount, Peter Moyle, and Barton Thompson. 2011. *Managing California's water: From conflict to resolution.* San Francisco, California: Public Policy Institute of California.

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. 2022. "Coordination and commitment in international climate action: Evidence from palm oil." Working paper.

Hultgren, Andrew, Tamma Carleton, Michael Delgado, Diana R. Gergel, Michael Greenstone, Trevor Houser, Solomon Hsiang, et al. 2022. "Estimating global impacts to agriculture from climate change accounting for adaptation." Working paper.

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 groundater decline and some cases of recovery in aquifers globally." *Nature* 625:715–721.

Jessoe, Katrina, and David Rapson. 2015. "Commercial and industrial demand response under mandatory time-of-use electricity pricing." *Journal of Industrial Economics* 63 (3): 397–421.

Johnson, Renee, and Betsy A. Cody. 2015. "California agricultural production and irrigated water use." Congressional Research Service report.

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.

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

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.

McRae, Shaun. 2015. "Infrastructure Quality and the Subsidy Trap." *American Economic Review* 105 (1): 35–66.

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.

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.

Rapson, David. 2014. "Durable goods and long-run electricity demand: Evidence from air conditioner purchase behavior." *Journal of Environmental Economics and Management* 68:141–160.

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 (7707): 651–659.

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

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.

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

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

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

Groundwater and Crop Choice in the Short and Long Run

Fiona Burlig, Louis Preonas, and Matt Woerman

June 2024

Appendix: For online publication

Contents

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 choice 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:

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

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:

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

Because any crop choice is a renewal action in this setting[1]—including the choice of crop r_{t+1} in year $t+1$—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:

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

1. Kalouptsidi, Scott, and Souza-Rodrigues (2021) use the term "renewal action" to denote any choice c_t that yields a particular field state k_{t+1} in the following year regardless of states in years prior to year t. Since we model state-dependence as having only one-year memory, all crop choices $c_t \in \mathcal{C}$ are renewal actions.

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:

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

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 choices of c_t, c'_t, and 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.[2] 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.[3] 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.

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.[4] Third, we assume there is no transition cost to remain in the same crop: $T_m(c_t, c_{t-1}) = 0$.[5]

A.3 Calibrating the intensive-margin elasticity

In our main counterfactual simulations, we assume farmers adjust each crop's groundwater use in response to a counterfactual groundwater tax. Our reduced-form analysis shows that farmers have a short-run intensive-margin elasticity of −0.679 (Column (2) of Table 3). However, since this estimate comes from short-run variation in pumping costs, the long-run within-crop elasticity will likely differ from this short-run elasticity.

We calibrate long-run within-crop adjustment using a series of long-difference regressions of increasing length. We estimate first-difference analog of Equations (2)–(3) at the parcel-year level:

$$\Delta \log (Q_{it}^{\text{water}}) = \gamma \Delta \widehat{\log (P_{it}^{\text{water}})} + \delta_{it} + \epsilon_{it} \tag{A4}$$

$$\Delta \log (P_{it}^{\text{water}}) = \theta \Delta \log (P_{it}^{\text{elecDefault}}) + \psi_{it} + \nu_{it} \tag{A5}$$

2. 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.

3. 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).

4. 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.

5. 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.

where the Δ operator denotes a within-parcel first difference over $\{1, 2, 3, 4\}$ years. Increasing the length of the first difference allows to understand the extent to which the intensive-margin elasticity persists over longer time frames.

To ensure that we capture only within-crop adjustments—rather than the effect of crop switching—we apply sample restrictions comparable to the "intensive-margin" restriction in Column (2) of Table 3. That is, we restrict the sample of observations to parcels that do not change crop type over the period of long-differencing. For example, for a four-year difference, we only include parcel-years for which the modal crop choice satisfies $c_t = c_{t-1} = c_{t-2} = c_{t-3} = c_{t-4}$.[6] As a result, this restriction becomes increasingly strict for longer differences.

Figure A1: Persistence of intensive-margin elasticity over longer time scales

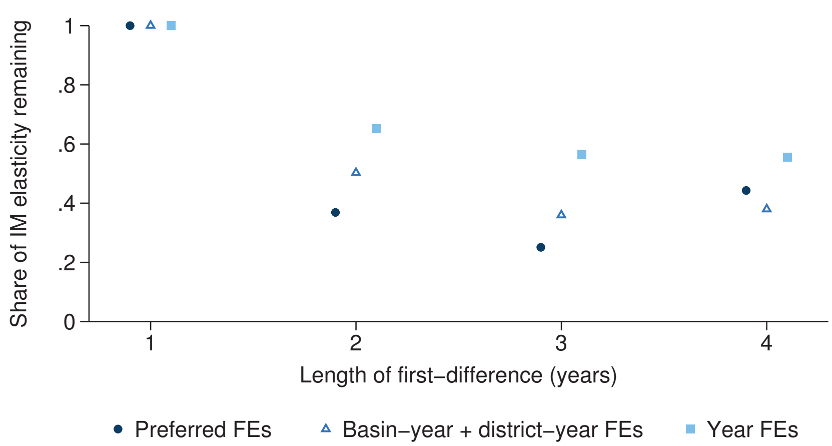

Notes: We estimate the first-difference analog of the annual versions of Equations (2)–(3), as shown in Equations (A4)–(A5) using 1-, 2-, 3-, and 4-year differences. Each regression applies the "intensive-margin" restriction as in Column (2) of Table 3, which becomes increasingly restrictive for longer differences (e.g., for a 4-year difference, $c_t = c_{t-1} = c_{t-2} = c_{t-3} = c_{t-4}$). This figure plots the ratio of each first-difference point estimate and its corresponding 1-year first-difference point estimate, to show how intensive-margin demand response attenuates over multiple years of groundwater cost shocks. Our preferred fixed effects are exactly analogous to the fixed effects used in Table 3; averaging these three ratios motivates our use of 35% IM as a central scenario. More parsimonious fixed effects imply a more persistent intensive-margin response; these ratios motivate our 59% IM scenario. We do not report confidence intervals, as they are not identified for ratios of coefficients.

Our preferred fixed effects are analogous to the regressions in Table 3: parcel-by-$\mathbf{1}$[small/large pump switch] (to capture tariff category switches), groundwater-basin-by-year (to capture trends in depth), and water-district-by-year (to capture varying surface water availability). In alternate specifications, we remove the parcel-switch fixed effect or use only a year fixed effect.

Figure A1 reports results of these regressions. To focus on how the intensive-margin elasticity evolves over longer time horizons, we report our $\hat{\gamma}$ estimates normalized by the 1-year difference effect. Using our preferred fixed effects, the average of the 2-, 3-, and 4-year differences is 35% of the 1-year difference. Therefore, we use "35% IM" as our central simulations. In a more parsimonious specification with only year fixed effects, the average of

6. As discussed in Appendix C.4, we use the modal crop choice because our intensive-margin restrictions necessitate imposing discreteness in crop choices.

the 2-, 3-, and 4-year differences is 59% of the 1-year difference. This informs our alternative "59% IM" simulations.

A.4 Static model

Our static model of crop choice—which we estimate for comparison with the dynamic model—follows from the dynamic model with two important changes. First, we remove farmers' forward-looking behavior by setting the discount rate to zero: $\beta = 0$. Second, we remove field state dependence from all terms in the model. With these changes, our static estimating equation is:

$$\ln\left[\frac{p_m(c_t, \omega_{mt})}{p_m(0, \omega_{mt})}\right] = \alpha_G \Delta G_{mct} + \Delta\alpha_{mc} + \Delta\xi_{mct}$$

where

$$\begin{aligned}
\Delta G_{mct} &= G_m(c_t, \omega_{mt}) - G_m(0, \omega_{mt}) \\
\Delta\alpha_{mc} &= \alpha_m(c_t) - \alpha_m(0) \\
\Delta\xi_{mct} &= \xi_m(c_t, \omega_{mt}) - \xi_m(0, \omega_{mt})
\end{aligned}$$

All terms in this static model are comparable to those in the dynamic and myopic models, but with field state dependence removed. For example, the static term $p_m(c, \omega_{mt})$ is an unconditional choice probability that does not depend on field state, rather than a CCP. Other static terms are defined similarly. As with the dynamic version, we estimate this model by instrumenting for ΔG_{mct} with the product of plausibly exogenous quantity and price measures.

A.5 Model results

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 find that remaining in the same crop type ($c = k$) yields weakly positive annual returns net of groundwater cost, while switching to a different crop type yields a negative annual return. These results are expected in this setting, in which switching crop type requires a large upfront investment—such as planting or removing an entire orchard of trees—that pays 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. These results are the same as those depicted in Figure 3 in the main text.

Table A1: Dynamic discrete choice parameter estimates

A. Groundwater cost parameter: α_G

	-0.012^{***}
	(0.004)

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

	Field state (k)			
	Annual	Fruit/nut	Hay	Non-crop
Crop choice (c)				
Annual	1.64	−1.07	−0.07	−0.08
	[1.10, 2.22]	[−1.61, −0.50]	[−0.61, 0.50]	[−0.62, 0.49]
Fruit/nut	−1.28	2.09	−0.35	−0.70
	[−1.98, −0.55]	[1.40, 2.83]	[−1.04, 0.39]	[−1.39, 0.04]
Hay	−0.38	−0.59	1.92	−0.72
	[−1.02, 0.30]	[−1.23, 0.10]	[1.27, 2.60]	[−1.36, −0.04]
Non-crop†	−0.86	−1.39	−1.32	0.00

Notes: Panel A displays our estimated groundwater cost parameter, α_G, which we obtain from estimating Equation (9). 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 500 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 16 $\alpha(c, k)$ terms, the intercept terms for the non-crop choice have no variation across draws from the sampling distribution.

A.6 Robustness

Figure A2 plots robustness test for our long-run (semi-)elasticities, reporting results for a variety of alternate model specifications. Each of these counterfactual simulations omits the intensive-margin response (0% IM) to focus on the crop-switching margin. Thus, the comparable "main" results in the top row reproduce our 0% IM results depicted in Figure 3. (Semi-)elasticities of land use and groundwater and electricity use are robust to parcel sample selection ("drop high cost" and "all size"), market-level variable construction ("weighted mean"), independent variable choice ("total cost"), estimation weighting ("weight by m-k"), and drought assumptions ("low drought pr." and "high drought pr.").

Table A2: Long-run elasticities with respect to groundwater pumping cost

	(1)	(2)	(3)	(4)
	Static	**Dynamic**		
		0% IM	35% IM	59% IM
Crop semi-elasticities				
Annual	0.003	0.097***	0.070**	0.056***
	[−0.001, 0.006]	[0.039, 0.153]	[0.020, 0.116]	[0.023,0.091]
Fruit/nut	−0.007	−0.194***	−0.135**	−0.106***
	[−0.016, 0.002]	[−0.314, −0.063]	[−0.232, −0.031]	[-0.177,-0.036]
Hay	−0.002	−0.007***	−0.006**	−0.005**
	[−0.007, 0.000]	[−0.013, −0.002]	[−0.010, −0.000]	[-0.009,-0.000]
Non-crop	0.006	0.104***	0.071**	0.055***
	[−0.001, 0.017]	[0.032, 0.172]	[0.018, 0.120]	[0.018,0.088]
Elasticities				
Groundwater	−0.020	−0.225***	−0.359***	−0.463***
	[−0.072, 0.003]	[−0.383, −0.070]	[−0.471, −0.259]	[-0.538,-0.391]
Electricity	−0.026	−0.250***	−0.369***	−0.462***
	[−0.103, 0.004]	[−0.393, −0.088]	[−0.466, −0.259]	[-0.524,-0.382]

Notes: This table reports the long-run (semi-)elasticities from simulating a 20% increase in groundwater pumping costs over a 20-year horizon in a steady state. In the static model (Column (1)), there is no state dependence and farmers are not forward looking. In the dynamic models (Columns (2)–(4)), there is state dependence and farmers are forward looking. In Column (3), farmers can only respond to groundwater cost changes by changing crops. In Columns (3)–(4), farmers also respond on the margin by reducing water use conditional on crop choice, applying 35% (preferred), or 59% of our within-crop reduced-form elasticity, per our long-difference estimates, respectively. The reported semi-elasticities and elasticities are the means over 500 draws for each model. The 95% confidence intervals (in brackets) are constructed from the 2.5th and 97.5th percentile over draws. Significance: *** 99% of simulation draws have the same sign; ** 95% of draws, * 90% of draws.

Figure A2: Long-run elasticities: Robustness

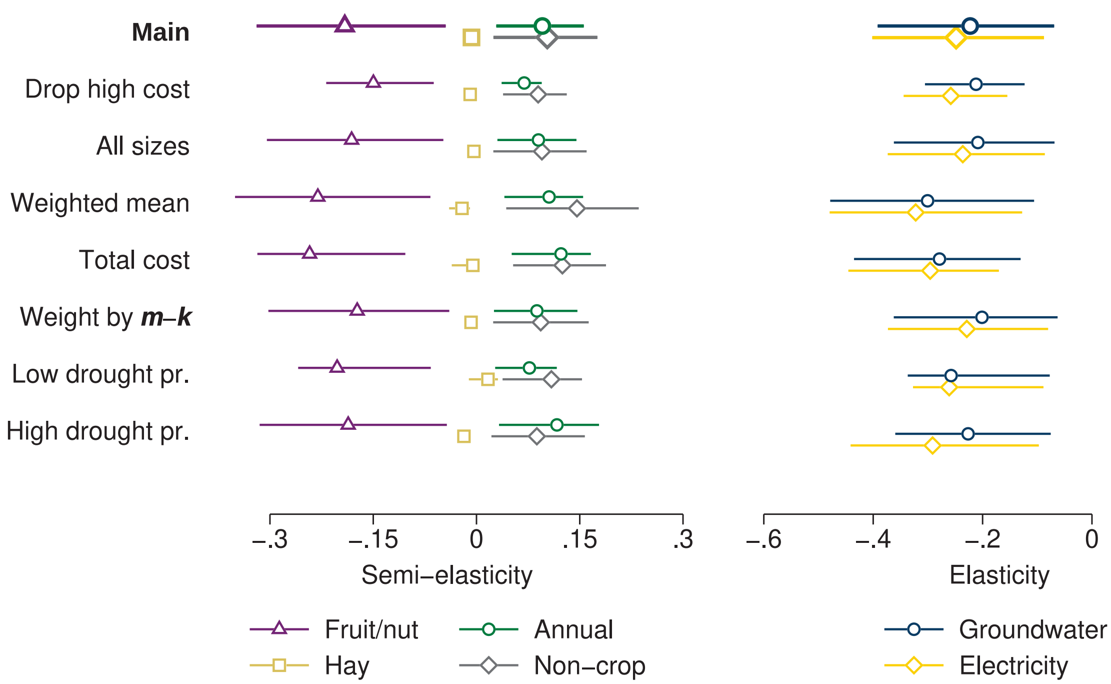

Notes: This figure plots robustness checks on our long-run (semi-)elasticities of land use (left panel) and groundwater and electricity use (right panel) with respect to groundwater cost. The top row reproduces the 0% IM results of Figure 3. In "drop high cost," we drop parcels with groundwater costs > \$5,000 per acre (rather than costs > \$3,000 per acre). In "all sizes," we include parcels of all sizes. In "weighted mean," we aggregate data to the market level using the weighted means (rather than weighted medians) of parcel data. In "total cost," we use the total cost of groundwater pumping (rather than the total *variable* cost of groundwater pumping) as our cost measure. In "weight by m-k," we weight observations by market-field state. In "low drought pr." and "high drought pr.," we assume the probability of a drought is 0.4 and 0.75, respectively. The left panels show semi-elasticities for annual crops (circles), hay perennials (squares), fruit/nut perennials (triangles), and non-crop (diamonds). The right panels show the electricity of demand for electricity (diamonds) and groundwater (circles). The reported semi-elasticities and elasticities are the means over 500 draws for each model. The plotted 95% confidence intervals (horizontal lines) show the 2.5th and 97.5th percentile over draws.

B Reduced-form sensitivity analysis

B.1 Monthly demand regressions

Table B1 provides further evidence that our electricity-to-water conversions are not obscuring unobserved pump efficiency upgrades in a way that meaningfully biases our groundwater demand estimates: our point estimates are quite stable as we tighten the restriction for observations within m months of an observed pump test (i.e., for which we parameterize Equation (1) *measured* operating pump efficiency).

Table B1: Monthly demand sensitivity – months to nearest pump test

	$\log(Q^{\text{water}})$				
	(1)	(2)	(3)	(4)	(5)
$\log\left(P^{\text{water}}\ (\$/\text{AF})\right)$	-0.249^{***}	-0.260^{***}	-0.236^{**}	-0.233^{**}	-0.260^{*}
	(0.094)	(0.098)	(0.105)	(0.117)	(0.134)
Pump test within m months, $m =$	60	48	36	24	12
Service point units	9,982	9,931	9,869	9,741	9,202
Months	148	148	148	136	124
Observations	791,910	712,177	605,892	472,875	276,958
First-stage estimates					
$\log\left(P^{\text{elecDefault}}\ (\$/\text{kWh})\right)$	1.349^{***}	1.366^{***}	1.366^{***}	1.408^{***}	1.486^{***}
	(0.029)	(0.033)	(0.037)	(0.045)	(0.060)
Kleibergen-Paap F-statistic	2095	1755	1374	973	609

Notes: This table conducts sensitivity analysis on Column (1) of Table 2, restricting the sample to SP-months within m months of an observed pump test. Column (5) reproduces Column (2) of Table 2. Progressively restricting the sample in this way does not meaningfully alter our point estimates. This assuages concerns that unobserved changes in pump efficiency (e.g., capital depreciation, efficiency upgrades) are confounding our electricity-to-groundwater conversions. Regressions are otherwise identical to Column (1) of Table 2. All regressions include the following fixed effects: SP $\times$ month-of-year, SP $\times$ $\mathbf{1}$[large pump], groundwater basin $\times$ year, water district $\times$ year, and month-of-sample. Standard errors (in parentheses) are two-way clustered by service point and month-of-sample. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table B2: Monthly demand sensitivity – time-varying confounders

	$\log(Q^{\text{water}})$				
	(1)	(2)	(3)	(4)	(5)
$\log\left(P^{\text{water}}\ (\$/\text{AF})\right)$	-0.314^{***}	-0.295^{***}	-0.291^{***}	-0.297^{***}	-0.293^{***}
	(0.080)	(0.080)	(0.081)	(0.081)	(0.081)
Interact month FEs with	Initial HP	Initial OPE	County	Sub-basin	Water district
Service point units	10,091	10,091	10,090	10,071	10,063
Months	148	148	148	148	148
Observations	953,469	953,469	953,338	951,115	950,346
First-stage estimates					
$\log\left(P^{\text{elecDefault}}\ (\$/\text{kWh})\right)$	1.335^{***}	1.319^{***}	1.319^{***}	1.320^{***}	1.320^{***}
	(0.024)	(0.023)	(0.024)	(0.023)	(0.024)
Kleibergen-Paap F-statistic	3118	3158	3130	3313	3134

Notes: This table conducts sensitivity analysis on Column (1) of Table 2, interacting month-of-sample fixed effects with the following cross-sectional fixed effects: deciles of earliest observed nameplate horsepower (Column (1)); deciles of earliest observed operating pump efficiency (Column (2)); county (Column (3)); groundwater sub-basin (Column (4)); and water district, or county if not in a water district (Column (5)). Regressions are otherwise identical to Column (1) of Table 2. All regressions also include the following fixed effects: SP $\times$ month-of-year, SP $\times$ $\mathbf{1}$[large pump], groundwater basin $\times$ year, and water district $\times$ year. Standard errors (in parentheses) are two-way clustered by service point and month-of-sample. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table B3: Monthly demand sensitivity – adding time-varying controls

	$\log(Q^{\text{water}})$				
	(1)	(2)	(3)	(4)	(5)
$\log\left(P^{\text{water}}\ (\$/\text{AF})\right)$	-0.281^{***}	-0.281^{***}	-0.287^{***}	-0.307^{***}	-0.294^{***}
	(0.080)	(0.080)	(0.080)	(0.081)	(0.080)
Control variables	Precipitaton	Precip & temperature	Precip & temp lags	GW depth in basin	Dist to depth measurement
Service point units	10,062	10,062	10,062	10,075	10,091
Months	144	144	143	148	148
Observations	923,116	923,116	919,660	941,973	953,469
First-stage estimates					
$\log\left(P^{\text{elecDefault}}\ (\$/\text{kWh})\right)$	1.320^{***}	1.320^{***}	1.320^{***}	1.318^{***}	1.321^{***}
	(0.024)	(0.024)	(0.024)	(0.024)	(0.024)
Kleibergen-Paap F-statistic	3090	3089	3078	3072	3079

Notes: This table conducts sensitivity analysis on Column (1) of Table 2, adding the following time-varying controls: precipitation (Column (1)); precipitation and temperature (Column (2)); contemporaneous and monthly lags of precipitation and temperature (Column (3)); average depth in groundwater basin-month (Column (4)); and distance from SP to its nearest groundwater measurement taken in month t (Column (5)). Precipitation is summed to the monthly level at each SP location; we control for the mean, maximum, and minimum daily temperature at each SP location. Regressions are otherwise identical to Column (1) of Table 2. All regressions also include the following fixed effects: SP $\times$ month-of-year, SP $\times$ $\mathbf{1}$[large pump], groundwater basin $\times$ year, water district $\times$ year, and month-of-sample. Standard errors (in parentheses) are two-way clustered by service point and month-of-sample. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table B2 shows that our results are robust to controlling for time-varying confounders, including pump characteristics (horsepower and operating pump efficiency), geography (factors correlated by county and groundwater sub-basin), and surface water institutions (captured by water district). Table B3 shows that our results are robust to controlling for weather realizations (contemporaneous and lagged), groundwater depths, and distance to each SP's nearest contemporaneous groundwater measurement.

Table B4 reports sensitivity analysis on our method for converting from kWh to AF. Our estimates are not sensitive to: using pump-specific drawdown predictions (rather than fixing drawdown at measured levels that don't vary with depth; Column (1)), using average groundwater depths by basin-month (rather than rasterized depth, which may contain localized measurement error; Column (2)), removing pumps with questionable drawdown measurements (Column (3)), removing SPs with multiple pumps tested (where we must average across pumps; Column (4)), and removing SP-months further than 8 miles from the nearest groundwater measurement (Column (5)).

Table B4: Monthly demand sensitivity – conversion from kWh to AF

	$\log(Q^{\text{water}})$				
	(1)	(2)	(3)	(4)	(5)
$\log(P^{\text{water}}$ (\$/AF))	-0.290^{***}	-0.299^{***}	-0.328^{***}	-0.280^{***}	-0.237^{***}
	(0.081)	(0.079)	(0.092)	(0.099)	(0.088)
Predicted drawdown (instead of fixed)	Yes				
Mean depth (insted of rasterized)		Yes			
Drop suspect drawdown measurements			Yes		
Drop SPs with multiple pumps tested				Yes	
GW measurement within 8 miles					Yes
Service point units	10,091	10,091	9,071	7,181	9,957
Months	148	148	148	148	148
Observations	953,469	953,448	822,096	656,286	444,470
First-stage estimates					
$\log(P^{\text{elecDefault}}$ (\$/kWh))	1.314^{***}	1.330^{***}	1.330^{***}	1.274^{***}	1.386^{***}
	(0.024)	(0.024)	(0.025)	(0.024)	(0.028)
Kleibergen-Paap F-statistic	3100	3073	2761	2779	2413

Notes: This table conducts sensitivity analysis on Column (1) of Table 2, focusing on the construction of kWh/AF conversion factors. Column (1) uses time-varying predictions of pump drawdown, rather than drawdown as reported at the time of pump tests. Column (2) uses average groundwater depth by basin-month, rather than rasterized monthly depth at each SP location. Column (3) removes pump tests with questionable drawdown measurements (i.e., extreme values, internal inconsistencies). Column (4) removes SPs with multiple pumps appearing in the PGE's database. Column (5) restricts the sample to SP-months with a contemporaneous groundwater measurement within 8 miles (the median distance). Regressions are otherwise identical to Column (1) of Table 2. All regressions include the following fixed effects: SP $\times$ month-of-year, SP $\times$ $\mathbf{1}$[large pump], groundwater basin $\times$ year, water district $\times$ year, and month-of-sample. Standard errors (in parentheses) are two-way clustered by service point and month-of-sample. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table B5 tests for time-varying confounders related to the timing of PGE's smart meter rollout. It seems plausible that PGE might have prioritized replacing smart meters based on: the age of customers' accounts (proxied by SP start date in Column (1)), the cost of accessing customers' meters (proxied by an SP's distance to the edge of its parcel in Column (2)), customers' expected draw from the grid (proxied by 2008 maximum monthly kWh in Column (3)), or air conditioning demand among nearby households (proxied by climate zone in Column (4)). However, our elasticity estimates are largely unchanged when we interact these proxies with month-of-sample fixed effects.

Table B6 conducts sensitivity analysis on our instrumental variables estimation. Column (1) shows that our results are robust to lagging the default electricity price instrument, which assuages concerns about anticipation of tariff changes or smart meter upgrades. Column (2) shows that our results are nearly identical if we instrument using the modal tariff in each category (rather than the default tariff). Column (3) presents the uninstrumented OLS which reveals a downward bias perhaps due to non-classical measurement error in our kWh-to-AF conversions. Column (4) presents the uninstrumented OLS replacing P_{it}^{water} with P_{it}^{elec}, which reveals that the main benefit of the default price instrument is to remove bias from the kWh-to-AF conversion—not from a within-category tariff selection effect. Column (5) presents the reduced-form OLS, for comparison.

Table B5: Monthly demand sensitivity – smart meter rollout

	$\log(Q^{\text{water}})$			
	(1)	(2)	(3)	(4)
$\log\left(P^{\text{water}}\ (\$/\text{AF})\right)$	-0.299^{***}	-0.293^{***}	-0.182^{**}	-0.314^{***}
	(0.080)	(0.080)	(0.080)	(0.080)
Interact month FEs with	SP start date	Dist to parcel edge	2008 max kWh	Climate zone
Service point units	10,037	10,091	6,828	10,090
Months	148	148	148	148
Observations	952,835	953,469	771,530	953,338
First-stage estimates				
$\log\left(P^{\text{elecDefault}}\ (\$/\text{kWh})\right)$	1.323^{***}	1.323^{***}	1.323^{***}	1.322^{***}
	(0.024)	(0.024)	(0.025)	(0.024)
Kleibergen-Paap F-statistic	3089	3086	2812	3148

Notes: This table conducts sensitivity analysis on Column (1) of Table 2, focusing on factors that might have been correlated with PGE's smart meter rollout. We interact month-of-sample fixed effects with the following cross-sectional fixed effects: the earliest account open date at the SP (Column (1)); the log of distance from the SP to the edge of its assigned parcel, a proxy for distance to road (Column (2)); the log of maximum monthly kWh consumed in 2008, the first year of our sample period (Column (3)); and PGE-designated climate zone (Column (4)). Regressions are otherwise identical to Column (1) of Table 2. All regressions also include the following fixed effects: SP $\times$ month-of-year, SP $\times$ $\mathbf{1}$[large pump], groundwater basin $\times$ year, and water district $\times$ year. Standard errors (in parentheses) are two-way clustered by service point and month-of-sample. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table B6: Monthly demand sensitivity – IV specification

	$\log(Q^{\text{water}})$				
	(1)	(2)	(3)	(4)	(5)
	2SLS	2SLS	OLS	OLS	OLS
$\log(P^{\text{water}}$ (\$/AF))	-0.353^{***}	-0.278^{***}	-0.845^{***}		
	(0.103)	(0.083)	(0.028)		
$\log(P^{\text{elec}}$ (\$/kWh))				-0.380^{***}	
				(0.054)	
$\log(P^{\text{elecDefault}}$ (\$/kWh))					-0.388^{***}
					(0.107)
Service point units	9,971	10,091	10,091	10,091	10,091
Months	136	148	148	148	148
Observations	856,716	953,469	953,469	953,469	953,469
First-stage estimates					
$\log(P^{\text{elecDefault}}$ (\$/kWh), 6-mth lag)	0.917^{***}				
	(0.040)				
$\log(P^{\text{elecDefault}}$ (\$/kWh), 12-mth lag)	0.030				
	(0.021)				
$\log(P^{\text{elecModal}}$ (\$/kWh))		0.604^{***}			
		(0.009)			
Kleibergen-Paap F-statistic	380	4370			

Notes: This table conducts sensitivity analysis on Column (1) of Table 2, focusing on factors related to our instrumental variables specification. Column (1) lags the default tariff instrument by 6 and 12 months, in case default tariff changes happened to coincide with factors affecting electricity or groundwater consumption. Column (2) uses an alternate instrument: the modal tariff in each category, rather than the default tariff. Column (3) presents the uninstrumented OLS estimate corresponding to our preferred 2SLS regression. Column (4) is also an OLS regression, but uses the log of mean electricity price as an intermediate step between Column (3) (i.e. mean electricity price transformed from \$/kWh into \$/AF) and the reduced form presented in Column (5) (i.e. default mean electricity price). Regressions are otherwise identical to Column (1) of Table 2. All regressions include the following fixed effects: SP $\times$ month-of-year, SP $\times$ $\mathbf{1}$[large pump], groundwater basin $\times$ year, water district $\times$ year, and month-of-sample. Standard errors (in parentheses) are two-way clustered by service point and month-of-sample. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Finally, Table B7 presents additional electricity demand estimates. Column (1) reports the uninstrumented OLS. Columns (2)–(3) expand the sample using hierarchical clustering to predict which unmatched agricultural service points are likely to be groundwater pumps. For Column (2), our clustering algorithm is based on the monthly time series of electricity consumption at each service point. For Column (3), we use an alternative clustering algorithm based on cross-sectional covariates including: summary statistics on monthly kWh consumed, county, water district, distance to the edge of each SP's assigned parcel polygon, and proximity to a well in the DWR's Well Completion Reports. In each case, we construct clusters separately for each county, pooling matched and unmatched SPs. Then, we classify all unmatched SPs within a cluster as "predicted pumpers" if the share of the county's

Table B7: Monthly demand estimates: electricity consumption

	(1)	(2)	(3)
	$\log(Q^{\text{elec}})$	$\log(Q^{\text{elec}})$	$\log(Q^{\text{elec}})$
	OLS	2SLS	2SLS
$\log\left(P^{\text{elec}}\ (\$/\text{kWh})\right)$	-0.375^{***}	-0.271^{***}	-0.233^{***}
	(0.053)	(0.051)	(0.054)
Including predicted pumpers		Yes (kWh time series)	Yes (covariates)
Service point units	10,091	40,045	50,910
Months	148	148	148
Observations	953,469	3,876,641	4,638,030
First-stage estimates			
$\log\left(P^{\text{elecDefault}}\ (\$/\text{kWh})\right)$		1.048^{***}	1.042^{***}
		(0.006)	(0.006)
Kleibergen-Paap F-statistic		30200	31237

Notes: This table reports additional monthly demand estimates for electricity used to pump groundwater. Column (1) estimates an uninstrumented OLS of log Q_{it}^{elec} on endogenous P_{it}^{elec}. Columns (2)–(3) expand the sample to include agricultural SPs without observed APEP pump tests, but with a high likelihood of operating vertical groundwater pumps. We employ a hierarchical clustering algorithm to predict latent pumpers using both SP-specific time series of monthly electricity consumption (Column (2)) and cross-sectional covariates (Column (3)). Regressions are otherwise identical to Column (3) of Table 2. All regressions include the following fixed effects: SP $\times$ month-of-year, SP $\times$ $\mathbf{1}$[large pump], groundwater basin $\times$ year, water district $\times$ year, and month-of-sample. Standard errors (in parentheses) are two-way clustered by service point and month-of-sample. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

matched SPs in that cluster is greater than share of the county's total SPs in that cluster. Both expanded samples yield similar monthly electricity demand elasticity estimates.

B.2 Annual demand regressions

Table B8 is analogous to Table B2, revealing that time-varying geographic confounders are unlikely to be biasing our annual elasticity estimates. Table B9 is analogous to Table B3, showing that our annual estimates are not sensitive to the inclusion of weather controls.

Table B10 explores our instrumental variables specification. Columns (1)–(2) show that the model price instrument produces similar annual elasticity estimates (as in Column (2) of Table B6). Columns (3)–(4) show that the uninstrumented OLS estimate is biased away from zero (as in Column (3) of Table B6). Columns (5)–(6) reveal much larger reduced-form annual estimates (in contrast with Column (5) of Table B6).

Table B11 presents annual demand elasticities for electricity. Columns (1)–(2) find elasticity estimates that are similar to our preferred groundwater estimates (consistent with the monthly electricity result in Column (3) of Table 2). Columns (3)–(6) show that our demand estimates are similar if we expand our sample to include parcels containing unmatched service points that we predict to be likely groundwater pumpers (as in Columns (2)–(3) of Table B7).

Table B8: Annual demand sensitivity – time-varying confounders

	log(Q^{water})					
	(1)	(2)	(3)	(4)	(5)	(6)
log (P^{water} (\$/AF))	-0.749^{***}	-0.682^{***}	-0.755^{***}	-0.659^{***}	-0.746^{***}	-0.672^{***}
	(0.202)	(0.224)	(0.218)	(0.231)	(0.217)	(0.232)
Interact year FEs with	Initial HP		Initial OPE		County	
Intensive margin		Yes		Yes		Yes
Parcel units	6,388	6,287	6,388	6,287	6,388	6,287
County-years	367	334	367	334	367	334
Observations	54,220	41,445	54,220	41,445	54,220	41,445
First-stage estimates						
log ($P^{\text{elecDefault}}$ (\$/kWh))	1.420^{***}	1.392^{***}	1.392^{***}	1.374^{***}	1.400^{***}	1.365^{***}
	(0.042)	(0.044)	(0.041)	(0.043)	(0.042)	(0.044)
Kleibergen-Paap F-statistic	1159	688	1151	708	1125	652

Notes: This table conducts sensitivity analysis on Columns (1)–(2) of Table 3, interacting month-of-sample 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) (Column (2)) of Table 3. All regressions also include the following fixed effects: parcel × **1**[large pump], groundwater basin × year, and water district × year. 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: Annual demand sensitivity – adding time-varying controls

	log(Q^{water})					
	(1)	(2)	(3)	(4)	(5)	(6)
log (P^{water} (\$/AF))	-0.765^{***}	-0.687^{***}	-0.758^{***}	-0.684^{***}	-0.899^{***}	-0.691^{***}
	(0.218)	(0.229)	(0.219)	(0.228)	(0.201)	(0.228)
Control variables	Precipitaton		Precip & temp		Precip & temp lags	
Intensive margin		Yes		Yes		Yes
Parcel units	6,388	6,287	6,388	6,287	6,373	6,287
County-years	367	334	367	334	336	334
Observations	54,220	41,445	54,220	41,445	50,121	41,445
First-stage estimates						
log ($P^{\text{elecDefault}}$ (\$/kWh))	1.409^{***}	1.381^{***}	1.381^{***}	1.378^{***}	1.388^{***}	1.358^{***}
	(0.042)	(0.044)	(0.042)	(0.044)	(0.044)	(0.047)
Kleibergen-Paap F-statistic	1120	660	1123	655	988	654

Notes: This table conducts sensitivity analysis on Columns (1)–(2) of Table 3, adding the following time-varying controls: precipitation (Columns (1)–(2)); precipitation and temperature (Columns (3)–(4)); and contemporaneous and yearly lags of precipitation and temperature (Column (5)–(6)). Odd (even) columns are otherwise identical to Column (1) (Column (2)) of Table 3. All regressions include the following fixed effects: parcel × **1**[large pump], groundwater basin × year, and water district × year. 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 B10: Annual demand sensitivity – IV specification

	log(Q^{water})					
	(1)	(2)	(3)	(4)	(5)	(6)
	2SLS	2SLS	OLS	OLS	OLS	OLS
log (P^{water} ($/AF))	-0.674^{***}	-0.584^{**}	-0.819^{***}	-0.808^{***}		
	(0.229)	(0.232)	(0.084)	(0.099)		
log ($P^{\text{elecDefault}}$ ($/kWh))					-1.074^{***}	-0.895^{***}
					(0.314)	(0.306)
Intensive margin		Yes		Yes		Yes
Parcel units	6,388	6,287	6,388	6,287	6,388	6,287
County-years	367	334	367	334	367	334
Observations	54,220	41,445	54,220	41,445	54,220	41,445
First-stage estimates						
log (Modal P^{elec} ($/kWh))	0.625^{***}	0.616^{***}				
	(0.017)	(0.018)				
Kleibergen-Paap F-statistic	1326	828				

Notes: This table conducts sensitivity analysis on Columns (1)–(2) 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 estimate. Columns (5)–(6) present the reduced form of our preferred specification. Odd (even) columns are otherwise identical to Column (1) (Column (2)) of Table 3. All regressions include the following fixed effects: parcel × **1**[large pump], groundwater basin × year, and water district × year. 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 B11: Annual demand estimates: electricity consumption

	log(Q^{elec})					
	(1)	(2)	(3)	(4)	(5)	(6)
log (P^{elec} ($/AF))	-0.719^{***}	-0.641^{***}	-0.752^{***}	-0.674^{***}	-0.688^{***}	-0.618^{***}
	(0.222)	(0.230)	(0.106)	(0.126)	(0.115)	(0.128)
Including predicted pumpers			Yes (kWh time series)		Yes (covariates)	
Intensive margin		Yes		Yes		Yes
Parcel units	6,388	6,287	32,133	31,776	32,850	32,284
County-years	367	334	384	352	378	346
Observations	54,220	41,445	283,744	219,688	280,449	212,465
First-stage estimates						
log ($P^{\text{elecDefault}}$ ($/kWh))	1.281^{***}	1.235^{***}	1.235^{***}	1.077^{***}	1.092^{***}	1.075^{***}
	(0.024)	(0.030)	(0.010)	(0.011)	(0.010)	(0.012)
Kleibergen-Paap F-statistic	1509	1021	4881	4199	5241	4564

Notes: This table reports annual demand estimates for electricity used to pump groundwater. Columns (1)–(2) are analogous to Columns (1)–(2) of Table 3, replacing quantity and price of groundwater with quantity and price of electricity. Columns (3)–(6) expand the sample to include parcels containing agricultural SPs without observed APEP pump tests, but with a high likelihood of operating vertical groundwater pumps. We employ a hierarchical clustering algorithm to predict latent pumpers using both SP-specific time series of monthly electricity consumption (Columns (3)–(4)) and cross-sectional covariates (Columns (5)–(6)). Odd (even) columns are otherwise identical to Column (1) (Column (2)) of Table 3. All regressions include the following fixed effects: parcel × **1**[large pump], groundwater basin × year, and water district × year). 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: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Finally, Table B12 conducts sensitivity analysis on our parcel-year sample. In Columns (1)–(2), we include the following outlier parcels that we omit from our main analysis (both reduced-form and structural): parcels with less than 1 croppable acre (for which reported acreages is prone to measurement error), parcels with greater than 5000 croppable acres (all of which are unlikely to be irrigated by our observed groundwater pumps), and parcels with monthly electricity bills exceeding \$3000 per croppable acre (indicating either highly abnormal groundwater use or measurement error in the denominator). Including these outlier parcels produces similar results. Columns (3)–(4) remove our (preferred) croppable-acreage weights, yielding attenuated elasticity estimates; this suggests that larger parcels tend to be relatively more groundwater-cost-responsive than smaller parcels. Finally, Column (5) uses a service-point-by-year panel, which produces an estimate similar to Column (3).[7] This suggests that our decision to aggregate up from SPs to parcels is not meaningfully altering our elasticity estimates.

Table B12: Annual demand sensitivity – parcels and acreage weights

	$\log(Q^{\text{water}})$				
	(1)	(2)	(3)	(4)	(5)
$\log\left(P^{\text{water}}\ (\$/\text{AF})\right)$	-0.760^{***}	-0.677^{***}	-0.574^{***}	-0.439^{**}	-0.526^{***}
	(0.218)	(0.229)	(0.158)	(0.217)	(0.134)
Include outlier parcels	Yes	Yes			
Remove acreage weights			Yes	Yes	Yes
SP-year panel					Yes
Intensive margin		Yes		Yes	
Parcel units	6,913	6,797	6,079	5,867	
SP units					9,558
County-years	367	334	367	332	367
Observations	58,816	44,463	53,518	40,647	83,531
First-stage estimates					
$\log\left(P^{\text{elecDefault}}\ (\$/\text{kWh})\right)$	1.410^{***}	1.319^{***}	1.319^{***}	1.229^{***}	1.327^{***}
	(0.042)	(0.051)	(0.030)	(0.039)	(0.029)
Kleibergen-Paap F-statistic	1131	666	1841	1002	2113

Notes: Columns (1)–(2) include parcels with bills over \$3000 per croppable acre and with croppable areas $\notin [1, 5000]$ acres, all of which we drop from our preferred parcel-year specifications; these regressions are otherwise identical to Columns (1)–(2) of Table 3. Columns (3)–(4) remove the regression weights by each parcel's "croppable" acreage, but are otherwise identical to Columns (1)–(2) of Table 3. Column (5) is analogous to Column (3) but estimates the annual groundwater elasticity at the SP-level, rather than the (more aggregated) parcel level. All regressions include the following fixed effects: unit $\times$ $\mathbf{1}$[large pump], groundwater basin $\times$ year, and water district $\times$ year. Standard errors (in parentheses) are two-way clustered by unit and county-year. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

7. When the panel unit is the service point, it is not straightforward to isolate the intensive-margin (since SPs don't have modal crops *per se*) or weight by croppable acreage (since SPs don't have areas *per se*). Hence, the appropriate comparison here is between Column (5) and Column (3).

C Data

C.1 PGE data

PGE monthly billing data We use confidential customer-level electricity datasets, which PGE's data management team prepared for us under a non-disclosure agreement. These data comprise the universe of agricultural electricity consumers in PGE'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 an indicator 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.

PGE's Advanced Pumping Efficiency Program Under the same non-disclosure agreement, we also obtained rich audit data on agricultural groundwater pumps. PGE collected these data as part of its Advanced Pumping Efficiency Program (APEP), which subsidized pump tests for agricultural consumers across PGE 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 use (i.e. irrigation), and the electricity meter identifier. The latter lets us match pump tests to electricity service points, thereby isolating a sample of 15,732 service points for which agricultural groundwater pumping is the confirmed end-use.[10] We restrict our empirical analysis to this 14% subset of agricultural service points (plotted in dark blue in Figure C1), in order to avoid incorporating other agricultural electricity end uses.[11] We drop all matched service

8. PGE'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 matched service points, as PGE typically installs a dedicated meter for each groundwater pump.

11. This limits our sample by dropping the (many) agricultural groundwater pumps that did not receive APEP pump tests. While we cannot precisely identify which unmatched agricultural service points are in fact groundwater pumps, Tables B7 and B11 report sensitivity results where we expand the sample to include the unmatched service points that we can classify as "predicted pumpers."

points where the pump test data report a non-agricultural end use (e.g., "municipal") or indicate a non-vertical pump (i.e., small horizontal booster pumps).

Figure C1: PGE agricultural customers

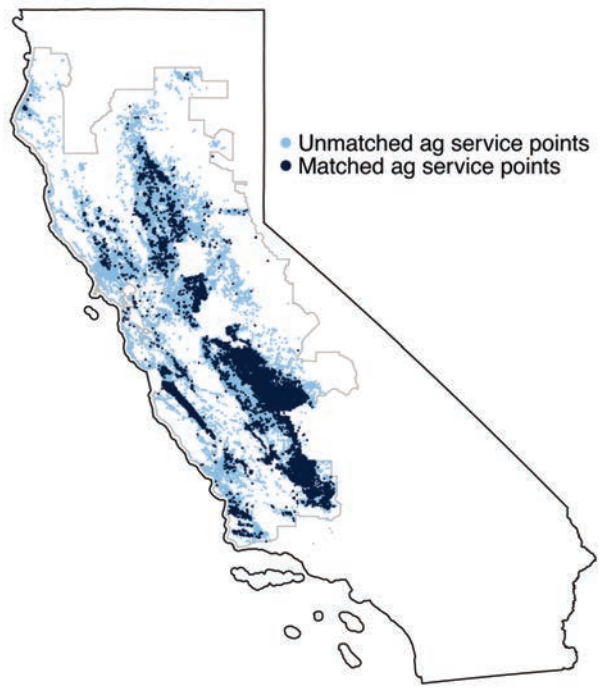

Notes: This figure maps the locations of all agricultural service points served by PGE, from 2008–2019. Dark blue dots indicate the 15,732 service points that we match to an APEP pump test. Light blue dots indicate unmatched agricultural service points. The light grey outline indicates PGE's service territory.

PGE agricultural tariffs PGE offers 23 distinct agricultural tariffs, and 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.[12] We use these data to construct a 2008–2019 panel of hourly volumetric (marginal) electricity prices, which we collapse to the monthly level by taking an unweighted average across hours. Importantly, unlike PGE's residential electricity prices, its agricultural tariffs are not tiered: a farm's marginal price does not depend on its consumption.

Variation in average volumetric prices arises from several features of PGE's tariff structure. All 23 tariffs have higher marginal prices during summer months (May–October). Time-varying tariffs have higher marginal prices on weekdays, during peak hours (12–6pm), and on critical peak event days.[13] Fixed charges (per kW) also play an important role in

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

13. Critical peak pricing offers farmers a slightly lower volumetric price throughout the year. In exchange, PGE can raise volumetric prices substantially on 15 days throughout the summer (typically the hottest days of the year). See Blonz (2022) for more details on critical peak pricing.

Figure C2: Average marginal electricity prices, by category

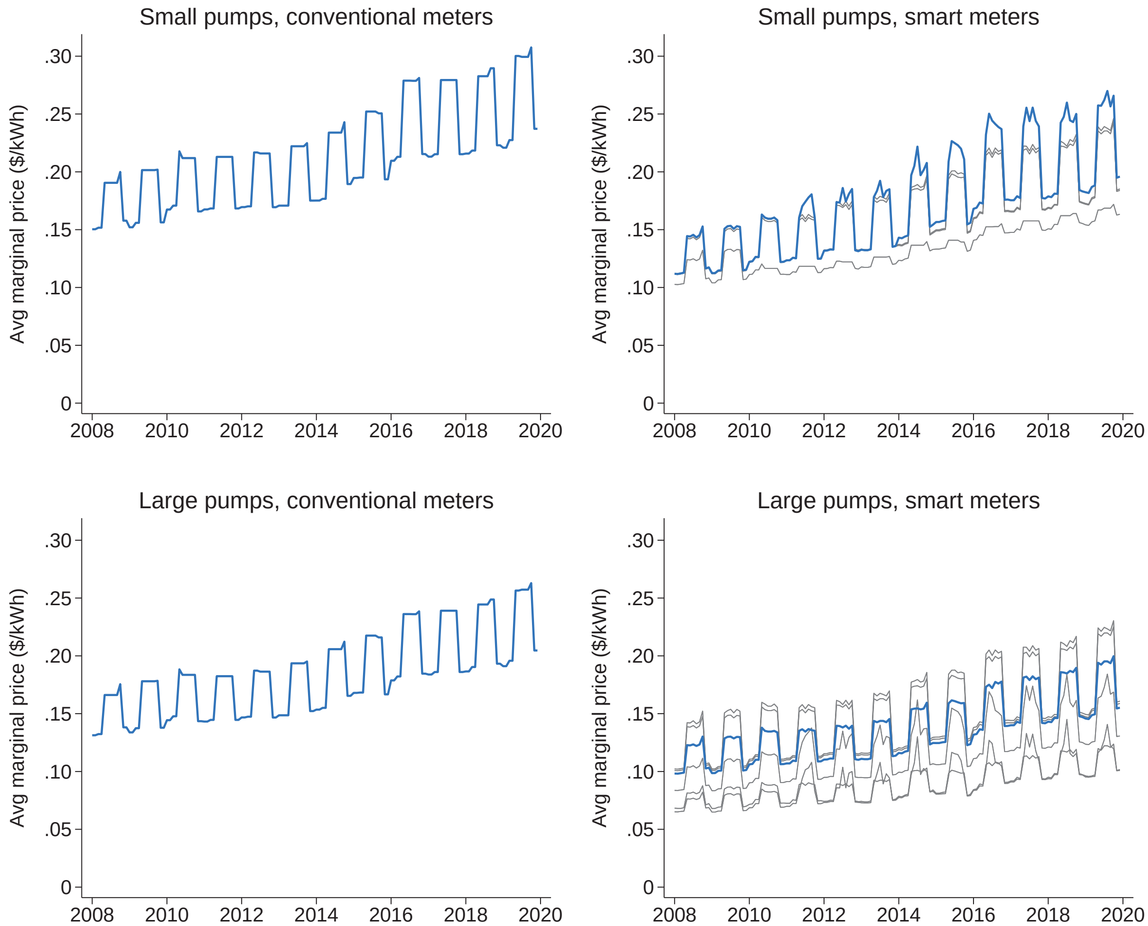

Notes: The thick blue line in each panel plots the time series of average marginal monthly electricity price for the default tariff in each of PGE's four tariff categories. (These are identical to the four series in the left panel of Figure 2.) For the two conventional-meter categories, the default tariff is the sole option. For the two smart-meter categories, customers may choose from a menu of additional tariffs, whose marginal price time series are plotted in thin grey. See surrounding text for further details.

offsetting marginal prices (per kWh): rates with higher marginal prices tend to have lower fixed charges, and vice versa. PGE adjusts tariffs' volumetric prices differentially over time.

22 of PGE'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).[14] 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 simplest (i.e., least time-varying) tariff as the default in each of these categories, which serves as our instrument. Figure C2 plots the time series of average monthly marginal prices for these 22

14. To be precise, the 35 horsepower cutoff applies to pumps with a single motor. For the few pumps with multiple motors, PGE defines a pump as "large" if the multiple motors sum to at least 15 horsepower of load. Conventional meters record electricity consumption using an analog dial, whereas smart meters can digitally store the full time profile of consumption.

Figure C3: Histogram of pump horsepower

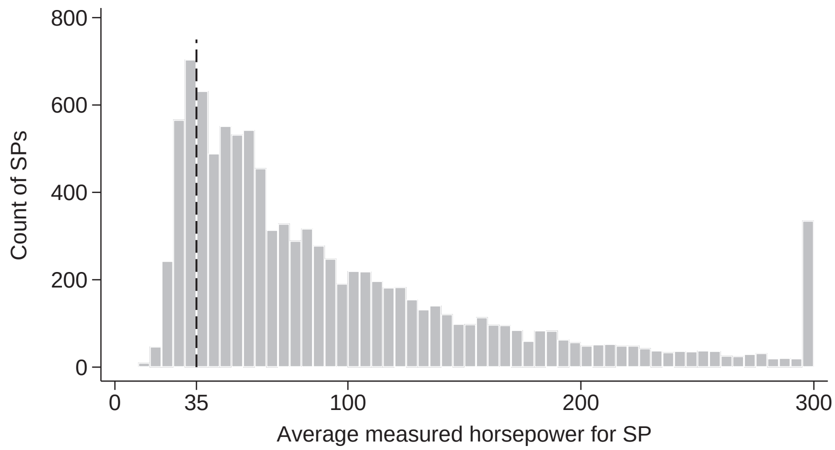

Notes: This is a histogram of measured horsepower from APEP pump tests, averaged for each PGE service point in our estimation sample. This reveals no evidence of bunching at the 35 hp cutoff that defines PGE's small- vs. large-pump tariff categories. Bunching would be consistent with farmers' optimizing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

tariffs by category. Figure C3 reveals no evidence of bunching at this 35 horsepower cutoff, which assuages concerns about tariff-induced selection in pump size.

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 matched 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 (or years) to construct the average marginal electricity price ($/kWh) for each tariff. For our structural analysis we construct the average total variable costs for each tariff by subtracting off fixed charges from each tariff (i.e., non-trivial charges assessed per day or per billing period, 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 (CASGEM) Program.[15] 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

15. These data are available from: https://water.ca.gov/Programs/Groundwater-Management/Groundwater-Elevation-Monitoring--CASGEM

of average depth at each point in space.[16] 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 measurment site in each month; this facilitates a robustness check where we remove observations with a high degree of spatial interpolation in groundwater depths.

We assign each service point to a groundwater basin and sub-basin, using publicly available shapefiles from the DWR.[17] Groundwater (sub-)basins are broadly defined by stratigraphic barriers through which water does not travel horizontally. California has 425 basins and 517 sub-basins (only 6% of basins contain more than one sub-basin); our main estimation sample includes farms from 54 basins and 104 sub-basins. Our reduced-form analysis controls for annual changes in depth that impact all farms within the same basin.

Finally, we link parcels to well-level observations in the DWR's Well Completion Reports.[18] Since these data contain a high degree of measurement error, we only use them as an input for classifying "predicted pumpers" in Tables B7 and B11.

C.3 Constructing groundwater quantities and prices

Energy is the sole variable input to groundwater production, and virtually all agricultural groundwater pumps in California are powered by electricity. Holding pump characteristics and groundwater depths fixed, the relationship between the quantity of groundwater extracted (measured in acre-feet, or AF) the the quantity of electricity (in kWh) consumed is governed by physics:

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

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.

To parameterize Equation (1) for each matched service point-month, we use OPE as reported in the PGE 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 can likewise parameterize lift using observed pump test measurements. However, populating lift for service point-months without an observed pump test is more complicated, since lift is (approximately) the sum of the standing water level (i.e., the baseline groundwater depth in the absence of pumping), drawdown (i.e., how much pump i impacts its own depth),

16. Before rasterizing, we drop depth measurements that are flagged as having questionable accuracy.
17. See here: https://water.ca.gov/Programs/Groundwater-Management/Bulletin-118
18. These data are available at: https://data.cnra.ca.gov/dataset/well-completion-reports

and other pump-specific factors (e.g., discharge pressure, gauge corrections, height of the pump above the surface).[19] Our preferred approach models drawdown as a function of the standing water level and location fixed effects (to account for properties of the stubstrata); then we combine these pump-specific drawdown predictions with (i) our service point-month panel of groundwater depths (i.e. observed variation in the standing water level; described above), and (ii) fixed pump characteristics, in order to populate lift for each service point-month. As with OPE, we extrapolate beyond the first/last pump tests and interpolate between pump tests (for all inputs of lift except standing water level). Table B4 presents sensitivity analysis on how we construct lift.

C.4 Constructing crop choice at the parcel level

Our data on cropped acreage are from the U.S. Department of Agriculture's (USDA) Cropland Data Layer (CDL).[20] This product provides annual information on what crop is being grown 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 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., non-crop land uses including 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 matched 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 (Column (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.[21] In all cases, we remove "not croppable" acreage from the denominator of each parcel-year.

19. Drawdown depends on rate of extraction (i.e. flow) and the physical properties of the substrata. Greater flow increases 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.

20. https://www.nass.usda.gov/Research_and_Science/Cropland/SARS1a.php

21. We use all pixels contained within parcels in market m that merge to matched 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.

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 Environmental Health Tracking Program.[22] Water districts are administrative entities that govern farmers' annual allocations of surface water.[23] 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 C4: Water districts and counties used to construct markets

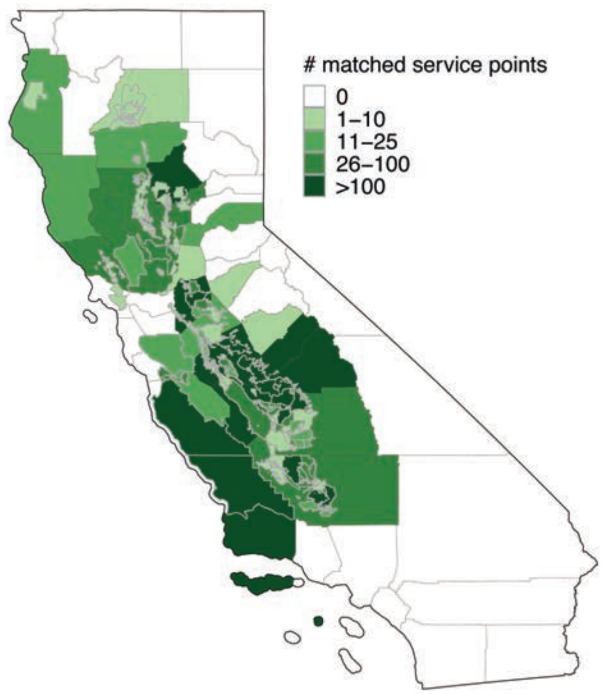

Notes: This figure maps the number of parcels with matched service points by water district and (for parcels not in a water district) by county. We use water districts (plotted with thick grey borders) and counties-less-water-districts to construct markets in our dynamic discrete choice analysis. Note 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).

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.[24] In our reduced-form analysis, we non-parametrically control for annual shocks to surface water allocations at the water district

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

23. 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.

24. A third source of water for irrigation is the open market. However, Hagerty (2023) suggests that purchased water is almost always more expensive than the groundwater costs for farmers in our dataset.

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 matched parcels that are not in a water district (i.e., not receiving surface water allocations), we use counties to define "markets." Figure C4 maps water districts (with thick grey borders) and counties, where shading indicates the number of parcels with matched 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. Appendix A.2 provides further details on how we construct markets.

C.6 Weather data

Weather is a key input into agricultural production, which directly impacts groundwater consumption. 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 (2009) and Burke and Emerick (2016)).[25] Using gridded data with a 4km-by-4km resolution, we extract daily maximum temperature, minimum temperature, and precipitation at each SP location.

C.7 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.[26] 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

25. These data are available at https://prism.oregonstate.edu/.
26. GSPs are available from the Department of Water Resources: https://sgma.water.ca.gov/portal/gsp/all

Figure C5: GSP sustainability targets under SGMA

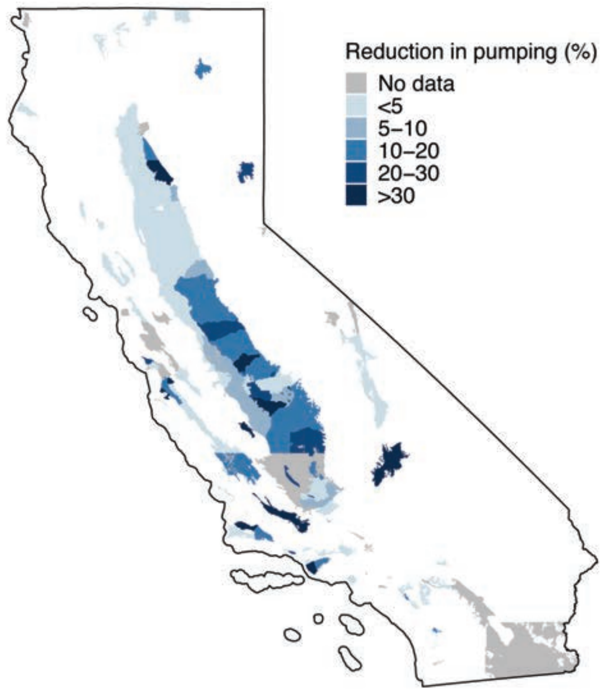

Notes: This figure maps the Groundwater Sustainability Agencies (GSAs) within California's medium- and high-priority basins, which were required to submit Groundwater Sustainability Plans (GSPs) to the Department of Water Resources. The shading reflects the percentage reduction in groundwater pumping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

of Water Resources (2017)).[27] We were able to populate these two numbers for 111 out of the 120 available GSPs.[28]

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 C5 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 levels at or below sustainable yield, thereby already achieving sustainability.[29] Bringing the

27. GSPs are detail documents, frequently over 1,000 pages long. Where possible, we draw these numbers from the executive summary. If these numbers are not in the executive summary, we extract these numbers from the GSP's water budget section.

28. 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 climate.

29. It is possible that the GSPs *understate* the true magnitude of overdraft. Bruno, Jessoe, and Hanemann (Forthcoming) compares reported overdraft to the results from running the C2VSim hydrology model, and

average overdrafted GSP into sustainability under this definition will require reductions in pumping of 20.7% (weighting all GSPs equally) or 19.2% (weighting by current pumping).[30]

Appendix references

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

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.

Bruno, Ellen M., Katrina Jessoe, and W. Michael Hanemann. Forthcoming. "The dynamic impacts of pricing groundwater." *Journal of the Association of Environmental and Resource Economists.*

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

California Department of Water Resources. 2017. *Best Management Practices for the Sustainable Mangement of Groundwater.* Technical report. Technical Report November 2017.

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

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

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

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.

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

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.

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.

30. These averages exclude GSPs for which (current pumping) $\leq$ (sustainable yield).