

Groundwater, energy, and crop choice

Fiona Burlig, Louis Preonas, and Matt Woerman*

July 13, 2020

Abstract

Groundwater is a key resource for agricultural production globally. Both increasingly rapid drawdowns of aquifers as well the policies intended to increase aquifer sustainability increase costs to agricultural producers, with unknown consequences. In this paper, we provide the first large-scale empirical estimates of how farmers respond to changes in groundwater costs in one of the world's most valuable agricultural areas: California. To do this, we assemble a novel dataset that combines (i) detailed restricted-access microdata on farmers' electricity consumption, (ii) rich data from technical audits of these farmers' pump efficiencies, (iii) measurements of groundwater depths in California aquifers, and (iv) satellite-derived measures of crop cover. For identification, we leverage exogenous variation in the price of electricity, a key marginal input into the groundwater production function. We find that farmers are very price responsive: we estimate large price elasticities of demand for electricity (-1.17) and groundwater (-1.12). We demonstrate that crop switching and fallowing are the main channel through which farmers respond to increases in groundwater costs. Using a discrete choice model, we estimate that a counterfactual \$10 per-acre-foot groundwater tax—approximately the price increase required to meet California's sustainability targets—would lead farmers to reallocate 8 percent of cropland, with increases in fallowing and high-value fruit and nut perennials, and decreases in annual crops and low-value perennials.

Keywords: groundwater, agriculture, electricity

JEL Codes: Q15, Q25, Q41

*Burlig: University of Chicago Harris School of Public Policy and NBER, Keller Center, 1307 E. 60th Street, Chicago, IL 60637. Email: burlig@uchicago.edu. Preonas: Department of Agricultural and Resource Economics, University of Maryland, 2200 Symons Hall, 7998 Regents Drive, College Park, MD 20742. Email: lpreonas@umd.edu. Woerman: Department of Resource Economics, University of Massachusetts Amherst, Stockbridge Hall, 80 Campus Center Way, Amherst, MA 01003. Email: mwoerman@umass.edu. We thank Nick Hagerty for graciously sharing data with us, and thank Amy Ando, Andrew Ayers, Maximilian Auffhammer, Kathy Baylis, Ellen Bruno, Karen Clay, Michael Greenstone, Koichiro Ito, Katrina Jessoe, Erik Lichtenberg, Dave McLaughlin, Kyle Meng, Mar Reguant, Paul Scott, Joe Shapiro, Andrew Stevens, Catherine Wolfram, and seminar participants at UC Energy Camp, Mississippi State, UC Davis, University of Illinois, ITAM, HKUST, University of Chicago, USC, the ASSA Annual Meetings, the AERE Summer Conference, the WEAI annual meetings, and UC Berkeley's POWER Conference for helpful comments and suggestions. Chinmay Lohani, Anna Schmidt, Chen Sui, Yixin Sun, and Xinyi Wang 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, and the Sloan Foundation (via the E2e Project), and the Becker Friedman Institute for Economics. All remaining errors are our own.

1 Introduction

Groundwater is an essential input into agricultural production, responsible for supplying water to 38 percent of irrigated acres worldwide (Siebert et al. (2010)). Recent scientific evidence documents rapid drawdowns of global aquifers, with key agricultural areas seeing their water tables fall by over 4 cm per year.¹ Given that global climate change is projected to increase the frequency and severity of droughts (Famiglietti (2014)), governments will face a greater urgency to enact policies to manage common-pool groundwater resources. For agricultural producers who rely on groundwater for irrigation, both groundwater scarcity itself and groundwater management policies increase the costs of growing crops.

In this paper, we generate novel empirical estimates of farmers' response to changes in groundwater pumping costs in California, one of the world's most valuable agricultural regions. California produces 18 percent of total U.S. crop value, and its farmers rely heavily on groundwater for irrigation. Despite rapidly declining aquifer levels and a series of severe droughts, most California farmers face no meaningful restrictions on groundwater extraction. The state is currently implementing the Sustainable Groundwater Management Act (SGMA), California's first legislation governing groundwater use at scale. The effectiveness and economic consequences of any such groundwater regulation depend on both the extent to which farmers will respond to new pumping regulations and their means of adapting to higher irrigation costs.

We begin by estimating the price elasticity of demand for agricultural groundwater. Estimating this elasticity has historically proven difficult, in large part because groundwater is typically neither priced nor measured. We overcome these challenges by leveraging the fact that electricity is the main variable input in groundwater extraction. Using data on electricity prices and quantities, along with farm-specific mappings from energy inputs to groundwater extraction volumes, we are able to construct measures of groundwater prices and quantities. We assemble a novel dataset that combines (i) confidential electricity consumption data for all agricultural customers served by Pacific Gas & Electric (PGE), California's largest electric utility; (ii) technical pump efficiency audits for nearly 12,000 groundwater extraction points;

1. <https://www.nytimes.com/interactive/2019/08/06/climate/world-water-stress.html>

and (iii) publicly available groundwater measurements across space and time, for all major California aquifers.

Using exogenous variation in PGE’s electricity tariff schedules, we causally identify demand elasticities for both electricity and groundwater. We estimate farmers’ price elasticity of demand for electricity to be -1.17 , which is much more elastic than prior estimates of electricity demand in the residential and commercial/industrial sectors. For groundwater, we estimate a demand elasticity of -1.12 . This is much more elastic than previous estimates from the literature on agricultural groundwater demand. Both elasticity estimates are robust to a variety of sensitivities, and together imply that California’s forthcoming groundwater management efforts have the potential to yield the substantial reductions in extraction necessary to meet the state’s sustainability goals.

Given these large elasticities, a natural question is: *how* are farmers reducing their water use in response to increases in groundwater costs? We consider four potential mechanisms that farmers may be engaging in: (i) applying less water to existing crops; (ii) changing irrigation efficiency; (iii) switching water sources; and (iv) switching crops or fallowing. We present empirical evidence against (i): farmers do not appear to adjust water consumption in response to large short-run fluctuations in price. We can also rule out (ii): we estimate quantitatively similar groundwater elasticities in the short windows around pump tests, where major changes in pumping capital are highly unlikely. We then develop a stylized model of farm crop choice and irrigation costs to provide theoretical guidance on (iii) and (iv). This model reveals that (iii) is only likely at extremely high groundwater pumping costs—which we rarely observe in our data—at which point farmers may substitute groundwater for water purchased on the open market (Hagerty (2018)).

This leaves crop switching as the leading candidate mechanism driving farmers’ groundwater large demand response. We conduct several empirical tests to provide evidence in support of this crop switching mechanism. We estimate annual elasticities that are similar to our monthly elasticities, suggesting that farmers are not arbitraging water sources within a growing season. We find that these annual elasticities can be decomposed into substantial intensive- and extensive-margin responses to increasing groundwater costs, consistent with farmers switching to less-water-intensive crops and fallowing land.

Building on these results, we estimate a discrete choice model of farmers’ cropping decisions to directly quantify the causal effect of groundwater costs on land use. Using a multinomial probit model, identified using exogenous variation in PGE’s electricity tariffs, we find that increases in groundwater costs reallocate cropland across four broad crop types. A 10 percent increase in the cost of groundwater causes farmers to increase the proportion of land in fruit and nut perennials by 1.0 percentage point, increase fallowed land by 0.4 percentage points (almost identical to our extensive-margin semi-elasticity), decrease land in annuals by 0.9 percentage points and reduce land in other perennials by 0.5 percentage points. Because we use a static discrete choice model, these are likely underestimates of the true impacts of increasing groundwater costs on crop reallocation (Scott (2013)).

Finally, we simulate how cropland would be reallocated in response to a counterfactual groundwater tax—a regulation that could internalize open-access externalities and achieve sustainable levels of groundwater extraction.² Under SGMA, California will require high-use basins to reduce extraction by between 20 and 50 percent (Bruno (2019)). We find that a \$10 per acre-foot tax on groundwater, representing approximately a 25 percent increase in pumping costs, would result in a 27 percent reduction in groundwater extraction, near the low end of this targeted curtailment range. We estimate that this same tax would lead farmers to reallocate nearly 8 percent of cropland to a different crop type. While the level of the efficient Pigouvian groundwater tax remains unknown, if the true externality is around \$10 per acre-foot, our results imply that that 8 percent of total cropland in California is currently misallocated from its socially optimal use.

This paper makes three main academic contributions. First, we provide the first large-scale empirical estimates of the impact of groundwater pricing on crop choice in one of the most valuable agricultural regions of the world: California. We demonstrate that, as groundwater extraction costs rise, farmers respond by increasing fallowing, shifting into fruit/nut perennials (among the highest-revenue crops per acre-foot of water), and shifting

2. The water economics literature typically cites two main externalities associated with groundwater extraction (Provencher and Burt (1993)). The first is the “stock externality”, which arises from the finite nature of non-renewable groundwater stocks, driving farmers to collectively extract faster than the social planner’s optimal extraction path. The second is the “pumping cost” externality: when a farmer extracts an acre-foot of groundwater from an aquifer, the water level falls, increasing pumping costs for other (nearby) users. Recently, other potential externalities such as air quality issues associated with soil drying and land subsidence have been raised as well.

out of low-value perennials and annual crops. These results have important implications for agricultural markets, as California is a monopoly producer of many crops that may become less prevalent under groundwater regulation. They also contribute estimates from an understudied sector to the broader literature on the impacts of environmental regulation.

Second, to quantify the impacts of groundwater cost increases on land use, we begin by contributing novel estimates of the price elasticity of electricity demand in agriculture—a major, but heretofore overlooked, electricity end-use sector. While many studies estimate the relationship between electricity prices and consumption in the residential sector (Alberini and Filippini (2011); Fell, Li, and Paul (2014); Ito (2014); Deryugina, MacKay, and Reif (2018)), far fewer have focused on commercial/industrial electricity consumption (Paul, Myers, and Palmer (2009); Jessoe and Rapson (2015); Blonz (2016)). To the best of our knowledge, there exists no comparable study estimating the price elasticity of electricity demand in the agricultural sector. By leveraging microdata for thousands of agricultural consumers across PGE’s service territory, along with plausibly exogenous changes in farmers’ marginal electricity prices, we identify California farmers as unusually elastic electricity consumers (a central elasticity estimate of -1.17), with potential implications for agricultural electricity pricing.

Third, we leverage the central role of electricity use in groundwater extraction to estimate the elasticity of groundwater demand for California farmers—a policy-relevant elasticity that has proven elusive due to both data and identification challenges (Mieno and Brozovic (2017)). Our empirical strategy overcomes many of these challenges by combining comprehensive electricity consumption data with technical audits of groundwater pumps, and by leveraging exogenous variation in electricity prices (a major component of pumping costs) to credibly identify changes in farmers’ effective price of groundwater. Many previous studies have estimated water demand outside the agricultural sector (Hewitt and Hanemann (1995); Renwick and Green (2000); Olmstead, Hanemann, and Stavins (2007)), while others have focused specifically on groundwater demand in agriculture (Hendricks and Peterson (2012); Pfeiffer and Lin (2014); Badiani and Jessoe (2015)). We provide well-identified groundwater demand estimates for thousands of farms in one the most important agricultural regions in the world: California’s Central Valley. We find that farmers are surprisingly elastic, with

a central estimate of -1.12 . We rule out several hypotheses about what might be driving this large elasticity, and demonstrate that crop reallocation is the major mechanism behind farmers' responses to increases in water costs.

Beyond our contributions to the academic literature, our findings have direct and immediate policy relevance for California agriculture, as SGMA goes into effect. If we extrapolate beyond our sample to all agriculture in California agriculture, our simulations suggest that the \$10 per acre-foot tax required to meet the lower end of SGMA's groundwater sustainability targets would reallocate nearly two million acres of cropland across the state. Our elasticity estimates also suggest that meeting the approximately 50 percent groundwater use reduction targets in severely overdrawn basins will require a tax greater than \$15 per acre-foot—more than 37 percent of the average cost per acre-foot we observe in our sample.

This paper proceeds as follows. Section 2 provides background on groundwater pumping, California agriculture, and energy use in farming. Sections 3 and 4 describe our data and empirical strategy. Section 5 presents our demand elasticity estimates. In Section 6, we explore mechanisms, and demonstrate the impacts of groundwater costs on crop switching. Section 7 concludes.

2 Background

2.1 Agriculture in California

California is a major player in global agricultural production, and its nearly \$33 billion in crop value in 2018 represented 18 percent of the U.S. total (USDA (2020)). California's 77,000 farms produce over 400 commodities, including more than half of all fruits, nuts, and vegetables grown in the United States. In fact, California is the sole domestic producer of many high-value crops, including almonds, artichokes, olives, and walnuts (California Department of Food and Agriculture (2011)).

Water is an essential input for California's agricultural production. Nearly 80 percent of the state's annual water consumption occurs in the agricultural sector, where crop irrigation is the primary end use. California has nearly 8.3 million harvested acres of cropland, 7.9

million of which are irrigated (Johnson and Cody (2015)). Many of California’s crops require 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)).³ At the same time, droughts of increasing severity have raised serious concerns about the (over)use of water for agriculture in California.

A simple time-series analysis provides suggestive evidence that drought is associated with substantial cropping changes: Figure 1 shows that after the 2011–2016 drought, farmers had substantially reduced land in water-intensive but relatively low-value crops such as alfalfa and winter wheat, substituting similarly water-intensive but high-value crops, such as almonds and grapes. In order to water these thirsty crops, farmers rely on groundwater and/or surface water—two water sources with significantly different governance structures (Sawyers (2007)).

Surface water Approximately 40 percent of California’s surface water is used in agriculture. 61 percent of irrigation water comes from surface sources, with groundwater making up the remaining 39 percent (California Department of Water Resources (2015)). Surface water rights in California follow strict rules. Most farms with access to surface water obtain it via irrigation districts.⁴ In addition to obtaining surface water from individual rights or irrigation districts, farmers have a limited ability to purchase water on the open market. However, these trades constitute only a very small share of total water deliveries, and the prices are extremely high (Hagerty (2018)).

Groundwater In normal weather conditions, groundwater supplies 30 to 40 percent of all water end uses in California. However, this rises to close to 60 percent in drought years, when surface water is unusually scarce (California Department of Water Resources (2014)). In contrast to the strictly defined surface water rights, agricultural groundwater rights in

3. The average California household uses 0.52 acre-feet per year (Hanack et al. (2011)).

4. Irrigation districts were established between 1860 and 1950, and their boundaries have remained essentially fixed over time. Though some individual farms do have their own water entitlements, the vast majority of these allocations belong to districts. These agricultural cooperatives divert water from large rivers and canals, and distribute this water to farmers. Individual farmers receiving water proportional to their acreage within the district (Schlenker, Hanemann, and Fisher (2007)). Hagerty (2019) provides a detailed description of surface water rights in California.

California tend to be far more vague. The typical groundwater right is “overlying,” meaning that landowners whose property sits above an aquifer have the right to extract the underlying groundwater.⁵ The vast majority of groundwater use is unmetered, and users face no variable costs of extraction beyond the energy costs of pumping (Bruno and Jessoe (2018)).⁶ Hence, a farmer may extract as much groundwater as he chooses, conditional on owning the overlying property rights.

Many of California’s groundwater basins are “overdrafted,” meaning that withdrawals exceed the pace of replenishment. As of 2017, some agricultural regions faced overdraft of 2 million acre-feet annually. This has led to a substantial decline in groundwater levels in the Central Valley, most notably in the Tulare and San Joaquin groundwater basins—which have lost a combined 135 million acre-feet of groundwater since 1925.⁷ The state faced a severe drought in 2014, with groundwater levels reaching historic lows in many portions of the state. 21 of the state’s 515 groundwater basins are now considered “critically overdrafted.”

In September 2014, California lawmakers responded to drought conditions by passing the Sustainable Groundwater Management Act (SGMA). This sweeping groundwater legislation is the first statewide regulatory effort to mitigate over-extraction of groundwater. SGMA comprises three separate bills. AB 1739 empowers California’s Department of Water Resources (DWR) or local groundwater sustainability agencies (GSAs) to charge fees for groundwater extraction, and 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 enables GSAs and the DWR to require groundwater monitoring.

This legislation represents the future of groundwater management in California, with the goal of achieving long-run sustainability by 2042. In order to meet these sustainability targets, some basins will have to reduce groundwater use by between 20 and 50 percent. Farmers are expected to meet these targets with a combination of reducing irrigation intensity

5. There are also “appropriative” groundwater rights, for users who do not own land above the aquifer. These rights are lower-priority than the overlying rights, and users may only exercise appropriative rights in the case of a surplus.

6. There are limited exceptions to this rule: a few irrigation districts impose a per-unit price on groundwater, but this remains rare (Bruno and Jessoe (2018)).

7. See: <https://www.ppic.org/publication/groundwater-in-california/>

and/or technology adoption, water trade with urban areas, and land fallowing or shifts towards less water-intensive crops (Bruno (2019)). However, SGMA’s GSPs may not begin to bind for many years, leaving groundwater pumping effectively unregulated in the interim.

2.2 Electricity for pumping

Electricity is an essential input to groundwater pumping. The California Energy Commission reports that water use accounts for 19 percent of California’s electricity consumption, and close to 8 percent of the state’s energy is used on farms (California Energy Commission (2005)). The state’s investor-owned utilities invest nearly \$50 million annually in agricultural energy efficiency. This makes water use—in particular, agricultural water use—a key component of California’s energy policy goals.

To estimate groundwater demand, we exploit the fact that electricity is a major determinant of pumping costs. Several previous papers have used variation in energy costs to estimate the price elasticity of groundwater demand (Hendricks and Peterson (2012); Pfeiffer and Lin (2014); Badiani and Jessoe (2015); and Mieno and Brozovic (2017)). Mieno and Brozovic (2017) point out that these estimates may exhibit bias due to (non-classical) measurement error and/or poor identification. Furthermore, data limitations have restricted previous studies to relatively narrow geographies.

We build on this existing literature with a sample that covers thousands of farms throughout California’s Central Valley, one of the most productive agricultural regions in the world. We are able to overcome the standard measurement issues in groundwater demand estimation via detailed technical audits that precisely characterize pump-specific electricity-to-groundwater conversion factors. Using exogenous variation from Pacific Gas & Electric’s (PGE) agricultural electricity tariffs, we are also able to overcome the standard identification challenges in this literature.

3 Data

3.1 Electricity data

We begin by estimating how farmers’ electricity consumption responds to changes in electricity price. 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 the years 2008–2017. We aggregate service accounts up to 108,172 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 also observe several key covariates for each service point: its latitude and longitude; an indicator for accounts with solar panels on net-energy metering, which we drop from our estimation sample; and an identifier to link service point locations to physical electricity meters. Figure 2 maps all agricultural service points in our dataset.

PGE offers 23 distinct agricultural tariffs, and our billing data report the particular 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 rules and eligibility criteria. We use these data to construct a 10-year 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 its consumption.

Variation in average volumetric price 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.⁹ Fixed charges (per kW) also play an important role in

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 often turn over within a given service point.

9. Critical peak pricing is a form of electricity price in which farms are offered a slightly lower electricity price throughout the year. In exchange, PGE can raise the price substantially on 15 days throughout the summer, with one or two days’ notice. These event days typically occur on the hottest days of the summer.

offsetting marginal prices (per kWh): rates with higher marginal prices tend to have lower fixed charges, and vice versa. On top of these cross-sectional differences in tariffs pricing schemes, PGE adjusts tariffs’ marginal prices differentially over time.

PGE’s 23 agricultural tariffs are divided into five mutually-exclusive categories. These five categories have strict eligibility criteria, defined both by physical pumping capital (small pumps, large pumps, or auxiliary internal combustion engines) and by type of electric meter (conventional meter, or smart meter).¹⁰ Each of these five categories has a “default” tariff; other tariffs within each category are highly correlated with these default tariffs. Figure 3 plots the five time series of default monthly average marginal prices. Our identification strategy relies on the fact that average marginal prices do not move in parallel across these five default tariffs. It also leverages the strict eligibility rules which place customers into categories. Appendix A presents more details on the full set of PGE agricultural tariffs.

3.2 Pump data

To complement our electricity data, we have rich data on agricultural groundwater pumps collected by PGE’s Advanced Pumping Efficiency Program (APEP). These data include the universe of APEP-subsidized pump tests from 2011–2017, and we observe detailed measurements and technical specifications for 21,851 unique tests at 17,107 unique pump locations. Importantly, we also observe identifiers for the electricity meter associated with each pump test, which we use to match pump tests to electricity service points—thereby isolating a sample of 11,849 service points for which agricultural groundwater pumping is confirmed to be a major end-use. We restrict our empirical analysis to this 11 percent subset of agricultural service points, in order to best isolate groundwater pumpers and avoid incorporating other agricultural electricity end uses.¹¹

See Blonz (2016) for more details on critical peak pricing. Appendix Table A1 summarizes the pricing structure for each tariff.

10. Conventional meters record electricity consumption using an analog dial, whereas smart meters can digitally store the full time profile of consumption. During our sample period, PGE gradually phased out conventional meters, replacing them with smart meters capable of supporting time-varying electricity pricing.

11. Pumping is likely the only end use at matched service points, as PGE typically installs a dedicated meter for each groundwater pump. We are currently working on using satellite images to predict whether service points outside the APEP-matched sample are also groundwater pumps. We hope to incorporate these farms into future analysis, as there are likely many groundwater pumps that never received an APEP-subsidized pump test. Note that not every customer on an agricultural tariff is pumping groundwater. We

Table 1 reports summary statistics for this subset of agricultural service points (in the right column). Compared to the full set of PGE’s agricultural customers, APEP-matched service points tend to consume nearly twice as much electricity and tend to pay lower marginal prices. Figure 2 reveals that APEP-matched service points are heavily concentrated in California’s Central Valley and appear to be a geographically representative subset of PGE’s agricultural customers.

After identifying a subset of agricultural consumers who pump groundwater, we use APEP data to characterize pump-specific groundwater production functions. The amount of groundwater extracted is a linear function of the electricity consumed by the pump; pump characteristics and groundwater depth determine how much water (measured in acre-feet (AF)) is produced by each kWh of electricity consumed. This kWh per AF relationship is governed by physics:

$$\frac{\text{kWh}}{\text{AF}} = \text{kW} \div \frac{\text{AF}}{\text{hour}} = \frac{[\text{Lift (feet)}] \times [\text{Flow (gallon/minute)}]}{[\text{Operating pump efficiency (\%)}] \times [\text{Constant}]} \div \frac{\text{AF}}{\text{hour}} \quad (1)$$

The power (kW) required to pump 1 acre-foot is directly proportional to both 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). We can simplify Equation (1) by converting from gallons to acre-feet:

$$\Rightarrow \frac{\text{kWh}}{\text{AF}} = \frac{[\text{Lift (feet)}] \times [\text{Constant}]}{\text{Operating pump efficiency (\%)}} \quad (2)$$

For each APEP pump test, we observe measurements of kWh/AF, operating pump efficiency, flow, and lift. We also observe the standing water level, or the baseline groundwater depth in the absence of pumping. Because pumping lowers the water level at a given location, standing water levels help us more accurately calibrate how changes in aquifer depth impact lift for each pump.¹²

restrict our sample to the APEP-matched customers in order to avoid including service points with other agricultural end uses, such as operating machinery or heating greenhouses.

12. Lift is (approximately) the sum of the standing water level, drawdown (i.e., how much pump i impacts its own depth), and other pump-specific factors (e.g., discharge pressure, gauge corrections, height of the

3.3 Water data

While a given farm’s pumping technology is relatively constant in the short run, its kWh/AF conversion factor is sensitive to short-run changes in groundwater levels. In order to capture these short-run shocks in pumping costs, we use publicly available groundwater data from California’s Department of Water Resources collected under the California Statewide Groundwater Elevation Monitoring (CASGEM) Program.¹³ These data report depth below the surface at 16,490 unique monitoring stations during our 2008–2017 sample period, with an average of 27 measurements at each location at different points in time. We rasterize all measurements within each month (and quarter), using inverse distance weighting to interpolate a gridded two-dimensional surface of average depth at each point in space. This allows us to construct a monthly (and quarterly) panel of estimated groundwater depths at each electricity service point.

We assign each service point to a groundwater basin, using publicly available shapefiles maintained by the California Department of Water Resources.¹⁴ Groundwater basins are broadly defined by stratigraphic barriers through which water does not travel horizontally. We control for annual changes in water levels that impact all farms within the same water basin. We also obtained shapefiles of irrigation districts in California from the California Department of Water Resources, the California Atlas, and the California Environmental Health Tracking Program, following Hagerty (2019). We spatially match PGE service points to these shapefiles to determine to which irrigation district (if any) each service point belongs.¹⁵ Irrigation districts (a.k.a. water districts) are administrative entities that govern farmers’ annual allocations of surface water. Because groundwater and surface water are obvious substitutes, we non-parametrically control for annual shocks to farms’ surface water allocations at the water district level. This helps to isolate variation in pumping behavior

pump above the surface). 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.

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

14. Water basin shapefiles are available from <https://water.ca.gov/Programs/Groundwater-Management/Bulletin-118>.

15. In ongoing work, we are working to incorporate data on each irrigation district’s water allocations for a heterogeneity analysis.

driven by variation in pumping costs, rather than variation in the availability of groundwater substitutes.

3.4 Groundwater prices and quantities

We merge the above data sources to create a panel of groundwater prices and quantities at the service point by month level. To convert from electricity (kWh or \$/kWh) to groundwater (AF or \$/AF), we simply need to populate a kWh/AF conversion factor for every panel observation. We construct estimates of kWh/AF by parameterizing Equation (2) using (i) monthly (or quarterly) rasters of groundwater depths at each service point; (ii) pump-specific conversions between standing water level and lift, as calculated from APEP pump tests; and (iii) APEP-measured operating pump efficiencies. We take unweighted averages of APEP variables across multiple pumps within a single service point; we also 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.

Table 2 reports summary statistics for this merged panel dataset. We observe 1.67 unique pump tests for the average APEP-matched service point, and APEP data reveal substantial cross-sectional variation in operating pump efficiencies and kWh/AF conversion factors. Our constructed kWh/AF estimates tend to moderate extreme values, which compresses the right tail of measured kWh/AF (while also slightly shifting this distribution left). Interestingly, implied marginal groundwater prices exhibit far less seasonal variation than marginal electricity prices. This is because groundwater levels tend to be higher in summer months (compared to winter months), which tends to reduce (constructed) kWh/AF in months when electricity prices are highest.

3.5 Agricultural data

Common Land Units In order to match electricity meters to cropland, we use the USDA Farm Service Agency’s Common Land Unit (CLU) data. We obtained the universe of USDA

CLUs in California.¹⁶ A CLU is the smallest contiguous unit of agricultural land under common land cover, land management, and ownership. We link PGE service points to CLUs via a spatial match. We also use shapefiles of tax parcels from California county assessors to aggregate CLUs (\sim fields) into groups with a common owner (\sim farms).¹⁷

Cropland Data Layer We obtain data on cropped acreage from the U.S. Department of Agriculture’s (USDA) Cropland Data Layer (CDL). 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 starting in 2007. The CDL is generated using satellite imagery in conjunction with a machine learning algorithm and is ground-truthed against the USDA’s Farm Service Agency’s farm surveys. In California during our sample period, the CDL reports 83 distinct landcover classifications. We classify these landcover categories into four broad groups: annual crops, fruit and nut perennial crops, other perennial crops, and fallowed land. The major annual crops in our sample are winter wheat, cotton, tomatoes, corn, rice, and strawberries. The major fruit and nut perennial crops are almonds, grapes, walnuts, pistachios, and oranges. The other perennials category is primarily cropped in alfalfa. The fallow category consists of land that is fallowed according to the CDL, as well as grass and pastureland.

We construct an annual panel of landcover at the CLU level based the modal CDL pixel within the CLU. If this modal crop type does not cover over 50 percent of pixels in the CLU, we label the CLU as having missing data. Using our CLU to service point concordance, we are able to assign individual service points a landcover type for each growing season.

3.6 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

16. In the 2008 Farm Bill, the CLU data were made restricted-access. We therefore use the 2008 CLUs for our full sample. The USDA provides more detail on CLUs here: <https://www.fsa.usda.gov/programs-and-services/aerial-photography/imagery-products/common-land-unit-clu/index>.

17. We collected tax parcel shapefiles from each county assessor’s office in California. Aggregating CLUs within the same tax parcel lets us test for sensitivity to within-farm spillovers.

and Roberts (2009) and Burke and Emerick (2016)).¹⁸ Using gridded data with a 4km-by-4km resolution, we extract daily maximum temperature, minimum temperature, and precipitation at each SP location.

4 Empirical strategy

This section outlines our empirical strategy for estimating farmers’ demand for groundwater pumping. First, we estimate price elasticities of demand for electricity, for the full sample of agricultural consumers where we can match an electricity meter to a groundwater pump. Next, we estimate price elasticities of demand for groundwater by translating prices and quantities of electricity into prices and quantities of water using data on (i) technical pumping production functions and (ii) groundwater depths across space and time.

4.1 Electricity demand

We estimate monthly electricity demand using the following specification:

$$\sinh^{-1}(Q_{it}^{\text{elec}}) = \beta \log(P_{it}^{\text{elec}}) + \gamma_i + \delta_t + \varepsilon_{it} \quad (3)$$

The dependent variable is kWh of electricity consumed at service point i in month t , transformed using the inverse hyperbolic sine function, which closely approximates the natural log transformation but includes zero in its support (Bellemare and Wichman (2020)).¹⁹ P_{it}^{elec} is unit i ’s marginal electricity price (in \$/kWh), averaged across all hours in month t . We include unit-by-month-of-year fixed effects (γ_i) to non-parametrically control for seasonality – including the average agricultural cycle – at every groundwater pump. We also include month-of-sample fixed effects (δ_t) to control for average market-wide time effects in both electricity prices (which rise over time) and pumping behavior, as well as changes in of the market environment (e.g., crop prices). Alternative specifications include groundwater-

18. These data are available at <https://prism.oregonstate.edu/>.

19. Since 14 percent of observations in this panel are zeros, we apply the inverse hyperbolic sine transformation to avoid dropping months when a farm consumes zero kWh for groundwater pumping. Appendix Table D10 presents alternative functional forms, including log-log and (log +1)-log.

basin-by-year fixed effects (to control for time-varying trends in groundwater depth across basins), water-district-by-year fixed effects (to control for annual shocks to surface water allocations), and service-point-specific linear time trends. We two-way cluster standard errors by service point and month-of-sample, which accommodates both arbitrary within-unit serial correlation and arbitrary spatial correlation across units within a month.

To econometrically identify the demand elasticity in Equation (3), we leverage both cross-sectional and time-series variation in electricity prices. Our primary source of exogenous variation comes from the fact that PGE’s agricultural tariff schedules are the outcome of statewide regulatory proceedings. This means that individual farmers cannot plausibly influence how PGE sets prices. Moreover, tariff decisions are made 1–3 years in advance of their implementation, reducing concerns that prices are set in response to real-time events such as droughts.

While the tariff schedules are themselves exogenous, many farmers are able to select a menu of tariffs—effectively choosing which marginal electricity price they face.²⁰ It is therefore important that we purge the resulting endogenous variation in unit i ’s marginal electricity price. To do this, we take advantage of eligibility restrictions that prevent farmers from choosing across the full menu of 23 tariffs. As we discuss above, PGE classifies all agricultural consumers into 5 disjoint categories:

- small pumps (< 35 hp) with conventional meters
- large pumps (≥ 35 hp) with conventional meters
- small pumps (< 35 hp) with smart meters
- large pumps (≥ 35 hp) with smart meters
- pumps with auxiliary internal combustion engines

While farmers may choose among tariffs *within* a category, they may not choose tariffs from other categories. To ensure that this within-category selection is not biasing our estimates,

20. As described in Section 3 above, marginal prices are constant in the amount of electricity consumed. Constant marginal prices simplify our estimation of agricultural electricity demand, because farm i ’s marginal price is determined solely by its tariff schedule. This is in contrast to PGE’s residential electricity tariffs, which have increasing block pricing, wherein a household’s marginal price is endogenous to its own consumption (Ito (2014)).

we instrument for unit i 's marginal price with the marginal price of the “default” tariff within unit i 's category.²¹ This eliminates selection bias from a high-volume pumper choosing a tariff with advantageously low volumetric prices.

For farmers to move *across* tariff categories, they must either adjust their physical pumping capital or have their electricity meter replaced by PGE. Pumping capital-induced category changes have the potential to introduce simultaneity bias: for example, upgrading from a < 35 hp pump to a ≥ 35 hp pump would lead to both a decrease in default marginal price and a mechanical increase in electricity consumption. Figure 4 demonstrates that there is no “bunching” in installed capital around this 35 hp threshold, suggesting that farmers are not endogenously choosing their pumping capital to manipulate their tariff category.²² Nevertheless, we control for potential endogenous changes in price by interacting unit fixed effects with dummies for each type of physical capital: small pumps, large pumps, and auxiliary internal combustion engines.

Moreover, Table 1 shows that 90 percent of changes in a unit's tariff category come from PGE replacing unit i 's conventional meter with a smart meter. It is highly unlikely that such meter upgrades coincided with any other changes in a farmer's pumping behavior.²³ Hence, meter-induced category changes are unlikely to lead to endogenous changes in unit i 's marginal electricity price. As a robustness check, we instrument with lagged default prices to purge potential endogeneity in the timing of unit i 's smart meter installation.

Figure 3 plots both raw and residualized time series of monthly average marginal prices for the five default tariff categories during our sample period.²⁴ The right panel partials

21. Three categories (conventional meters and internal combustion engines) comprise a single tariff; for these categories, assigning a “default” tariff is trivial. The two smart meter categories comprise 8 and 12 separate time-varying tariffs, respectively; for these categories, we choose as “defaults” the tariffs with the least time-varying marginal prices that most closely resemble their non-time-varying counterparts (AG-4A, AG-4B). Appendix Table A1 summarizes all 23 tariffs by category, with default tariffs in bold, and Appendix Figure A1 shows the time series of each tariff. We find similar results if we instrument using the modal tariff in each category (see Appendix Table D3).

22. As a robustness check, we estimate Equation (3) interacting month-of-sample fixed effects with deciles of pump horsepower. Our results are nearly identical to those in our main specification, assuaging concerns of differential trends between small vs. large pumpers (see Appendix Table D4).

23. During our 2008–2017 sample period, PGE gradually installed smart meters for the vast majority of its customers. The timing of PGE's smart meter rollout was determined by institutional and geographic factors, which were outside of customers' control. Previous research has established that PGE did not design their smart meter rollouts to target customers with particular usage patterns (Blonz (2016)).

24. Appendix Figure A1 presents an extended version of this same figure, with all 23 PGE tariffs.

out both tariff \times month-of-year fixed effects and common month-of-sample fixed effects, thereby illustrating the the main source of exogenous variation we use identify the demand elasticity in Equation 3: that the residualized time series do not move in parallel. We also leverage variation from meter-induced shifts across tariff categories: farmers who received smart meters during our sample saw a systematic decrease in their average monthly marginal prices (i.e., from AG-1A to AG-4A, or from AG-1B to AG-4B).

4.2 Groundwater demand

To estimate the causal effect of groundwater price on groundwater consumption, we construct an electricity-to-water conversion ratio, $\widehat{\frac{\text{kWh}}{\text{AF}}}_{it}$, for each observation in our dataset. While we do not directly observe groundwater prices or quantities, we can use these conversion factors to transform the electricity variables we do observe into their water equivalents:

$$Q_{it}^{\text{water}} = Q_{it}^{\text{elec}} \div \widehat{\frac{\text{kWh}}{\text{AF}}}_{it} \quad P_{it}^{\text{water}} = P_{it}^{\text{elec}} \times \widehat{\frac{\text{kWh}}{\text{AF}}}_{it} \quad (4)$$

Appendix B explores this conversion in detail, and describes how we can decompose the implied values of Q_{it}^{water} and P_{it}^{water} to separately identify farmer responses to changes in electricity price versus changes in pumping costs.

Here, we present a more parsimonious specification for estimating groundwater demand:

$$\sinh^{-1}(Q_{it}^{\text{water}}) = \beta \log(P_{it}^{\text{water}}) + \gamma_i + \delta_t + \varepsilon_{it} \quad (5)$$

This model uses the same fixed effects as in Equation (3) above, and we instrument for $\log(P_{it}^{\text{water}})$ using the same instrument: logged average marginal electricity price of unit i 's within-category default tariff. This isolates changes in the effective price of groundwater driven only by plausibly exogenous changes in the marginal electricity price. The instrument also eliminates the within-pump feedback effect of Q_{it}^{water} on P_{it}^{water} , whereby extraction lowers a pump's own water level and mechanically increases its effective marginal groundwater price.

Finally, instrumenting with default electricity price removes right-hand-side measurement error in $\widehat{\frac{\text{kWh}}{\text{AF}}}_{it}$, which has the potential to bias our elasticity estimates.²⁵

5 Results

5.1 Electricity demand

Table 3 reports results from estimating Equation (3). Column (1) presents the OLS results without instrumenting for marginal electricity price, resulting in a potentially biased $\hat{\beta}$ estimate. Column (2) instruments using unit i 's within-category default marginal price, which eliminates bias from farmers choosing their own electricity tariffs. The direction of this bias is not obvious *ex ante*, because farmers are choosing between tariffs with both volumetric (\$/kWh) and fixed (\$/kW) price components.²⁶ Comparing Columns (1) vs. (2), we see that on average, farmers with higher electricity consumption tend to choose tariffs with relatively low fixed charges per kW and relatively high prices per kWh.

Column (3) eliminates the other potential source of price endogeneity—farmers switching across rate categories—by interacting unit fixed effects with indicators for (i) small pumps (< 35 hp), (ii) large pumps (≥ 35 hp), and (iii) auxiliary internal combustion engines. While only 5 percent of units shift across tariff categories due to changes in their physical capital, the resulting simultaneous changes in Q_{it}^{elec} and P_{it}^{elec} induce substantial bias in Column (2) point estimate. Column (3) reports our preferred estimate of -1.17 , after having purged both sources of price endogeneity.²⁷

25. In Equation (5), measurement error from $\widehat{\text{kWh}/\text{AF}}_{it}$ enters directly on the right-hand side and inversely on the left-hand side. Instrumenting with default electricity prices negates the correlation between left-hand-side vs. right-hand-side measurement error. See Appendix B for further details.

26. PGE tariffs with relatively high volumetric (i.e. marginal) prices tend to have relatively low fixed prices, and vice versa. Two farmers with the same average electricity consumption may optimally choose different tariffs. Suppose farmer A operates a 300 hp pump for 50 hours per month, while farmer B operates 50 hp pump for 300 hours per month. Farmer A should prefer a low fixed price and a high volumetric price, while farmer B should prefer a high fixed price and a low marginal price.

27. Appendix Table D4 presents results from specifications where we interact month-of-sample fixed effects with bins of horsepower, kW, or operating pump efficiency. We recover extremely similar point estimates, (-1.15 with a standard error of 0.16 for horsepower), indicating that differential trends in pump size are not biasing our results.

Columns (4)–(6) report three additional elasticity estimates, each intended to assuage any remaining concerns over price endogeneity. Column (4) includes separate year fixed effects for each water basin and each water district, to control for potential time-varying confounders related to groundwater depth or surface water availability. The resulting point estimate of -0.95 is similar, albeit slightly attenuated. Column (5) instruments with the 6- and 12-month lags of the default price (rather than the contemporaneous default price), to account for potential endogeneity in the timing of PGE’s smart meter rollout.²⁸ This yields a nearly identical point estimate, implying that farmers’ electricity consumption did not meaningfully change in anticipation of a smart meter installation. Finally, Column (6) adds 11,173 unit-specific linear time trends, to confirm that we are not identifying $\hat{\beta}$ solely off of monotonic trends in price and quantity. The resulting point estimate of -0.76 is attenuated, as linear trends remove much of the (exogenous) variation in electricity prices over time, but still indicates that farmers respond strongly to changes in electricity prices.

In Appendix D, we present a series of sensitivity checks on this main result: instrumenting with modal within-category tariffs, rather than default tariffs (Appendix Table D3); interacting month-of-sample fixed effects with bins of pump horsepower, measured load, and efficiency (Appendix Table D4); interacting month-of-sample fixed effects with geographic fixed effects (Appendix Table D6); and adding weather controls (Appendix Table D8). Across all specifications, we find large, statistically significant elasticities that are quantitatively similar to our preferred specification in Column (3) of Table 3. Appendix Table D10 reveals that while our estimates are not sensitive to our choice of functional form, they are sensitive to the inclusion of observations with zero electricity consumption. If we exclude unit-months with zero electricity consumption, our elasticity estimate attenuates to -0.31 . This is consistent with farmers responding to high electricity prices by changing their cropping patterns—a mechanism that we explore in greater detail in Section 6 below.

Across specifications, we find large, precisely-estimated elasticities. The magnitude of these estimates is surprisingly large, considering that electricity demand tends to be quite inelastic in other contexts. The literature on electricity demand has focused heavily on

28. Recall that farmers may shift across tariff categories (inducing changes to their within-category default price) due to either changes in their physical capital or the installation of a smart meter.

the residential sector, and recent estimates have found elasticities of -0.08 to -0.48 in the short run (Reiss and White (2005); Alberini and Filippini (2011); Fell, Li, and Paul (2014)) and -0.09 to -0.73 in the medium-to-long run (Alberini and Filippini (2011); Ito (2014); Deryugina, MacKay, and Reif (2018)).²⁹ Fewer estimates exist for commercial or industrial electricity demand. Paul, Myers, and Palmer (2009) estimate commercial/industrial elasticities of -0.11 to -0.16 in the short run, and -0.29 to -0.40 in the long run. Jessoe and Rapson (2015) find no demand response to dynamic pricing in these sectors, while Blonz (2016) estimates elasticities of -0.08 to -0.22 using hourly price variation for PGE’s small commercial/industrial customers. To our knowledge, we provide the first large-scale estimates of electricity demand elasticities in the agricultural sector.

5.2 Groundwater demand

Table 4 presents our estimates of the price elasticity of farmer’s groundwater demand. In each column, we estimate Equation (5), including unit-by-month-of-year fixed effects, month-of-sample fixed effects, and interactions between unit fixed effects and physical capital. In Column (1), we present OLS results. As with the electricity results in Table 3, we find a smaller elasticity with the OLS (-0.88) than with our instrumental variables approach in the following columns. In Column (2), our preferred specification, we instrument for the price of groundwater using unit i ’s within-category default marginal price, eliminating bias from farmers choosing their electricity tariff. We estimate a price elasticity of groundwater demand of -1.12 .

Columns (3)–(6) present a series of sensitivity checks around this central estimate. Column (3) restricts the sample to the three largest groundwater basins, each of which has over 1,000 units in our estimation sample.³⁰ The resulting $\hat{\beta}$ estimate is quite similar, which should assuage concerns that the instrument is invalid due to a few large farms located in very

29. These estimates use monthly or annual variation in electricity prices, which aligns with our empirical strategy. Studies that leverage hourly variation in electricity prices have estimated electricity demand elasticities ranging from -0.03 to -0.25 (Wolak (2011); Jessoe and Rapson (2014); Fowlie et al. (2018); Ito, Ida, and Tanaka (2018)).

30. These basins are the San Joaquin Valley, the Sacramento Valley, and the Salinas Valley. The number of agricultural groundwater pumpers in each basin is likely much larger, as our estimation sample comprises only the subset of PGE customers that we can match to an APEP-subsidized pump test.

small groundwater basins. In Column (4), we convert from electricity to groundwater by recalculating $\widehat{\frac{\text{kWh}}{\text{AF}}}_{it}$ using groundwater data rasterized at the quarterly (rather than monthly) level; this addresses the inherent tradeoff between spatial density vs. temporal frequency of groundwater measurements, and has little effect on our results. Column (5) includes water-basin-by-year and water-district-by-year fixed effects, yielding a slightly attenuated point estimate (-0.90). In Column (6), we instrument with 6- and 12-month lags of the default within-category electricity price (rather than contemporaneous prices); we again find a quantitatively similar estimate of -1.14 .

The point estimates in Table 4 are quite similar to our electricity demand estimates.³¹ Appendix B presents the results from an alternative approach where we separately estimate the price elasticity of demand for groundwater with respect to electricity prices vs. kWh-to-AF conversion factors. In our preferred specification, we estimate larger elasticities than in Table 4: -1.27 for groundwater price changes induced by electricity price changes, and -1.51 for groundwater price changes induced by kWh-to-AF changes. These estimates are not statistically different from each other, which signals that farmers respond similarly to both types of pumping cost changes—as Neoclassical theory would predict, if farmers are indeed rationally optimizing over groundwater as an agricultural input.

Appendix D presents additional robustness checks: sensitivities to how we construct the $\widehat{\frac{\text{kWh}}{\text{AF}}}_{it}$ conversion factor (Appendix Table D5); interacting month-of-sample fixed effects with geographic fixed effects (Appendix Table D7); including weather controls (Appendix Table D8); and sensitivities by assignment to CLUs (Appendix Table D9). We find large, statistically significant elasticities that are quantitatively similar to our preferred point estimate from Column (2) of Table 4 across specifications. As with the electricity regressions described in Section 5.1, we find that our groundwater results are sensitive to the inclusion of observations with zero groundwater pumping: excluding unit-months with zero groundwater extraction attenuates our elasticity to -0.30 (Appendix Table D11). This suggests that an important method of adjustment to high groundwater prices is halting pumping,

31. It is not surprising that Equations (3) and (5) yield similar point estimates, since Q^{water} and P^{water} are multiplicative transformations of Q^{elec} and P^{elec} , and both specifications use the same two-stage least squares model.

which is again consistent with fallowing or crop switching as a mechanism. We examine this possibility in Section 6 below.³²

As with our elasticity estimates for electricity, our groundwater elasticity estimates are quite large relative to the existing literature. Recent studies have also exploited variation in energy prices but have yielded far smaller magnitudes: Hendricks and Peterson (2012) find an elasticity of -0.10 , and Pfeiffer and Lin (2014) find an elasticity of -0.27 (both for agricultural groundwater in Kansas). Bruno and Jessoe (2018) estimate demand elasticities of -0.17 to -0.22 within the Coachella Valley of California, which is a unique setting where groundwater extraction is directly priced.³³

6 Economics of groundwater demand response

We now explore possible mechanisms underlying the large elasticities we find in Section 5. When we observe a farmer consume less electricity at a pump with a dedicated electricity meter, four broad mechanisms could explain this decrease:

1. The farmer applies less irrigation water, without changing crop types.
2. The pump's kWh-to-AF groundwater production function changes.
3. The farmer switches to an alternate water source.
4. The farmer switches crops or fallows the land.

Applying less water, without changing crops We should expect this first mechanism to manifest in the short run, as a response to high-frequency cost shocks. To test for this, we exploit additional richness in our PGE electricity data. Beyond monthly billing data, we also

32. Below, Table 6 further corroborates this intensive-vs.-extensive margin result at the annual level.

33. Prior work on urban water demand, a setting in which researchers also observe prices and quantities for water, has found similar elasticities, ranging from -0.10 to -0.76 (Nataraj (2011); Ito (2013); Baerenklau, Schwabe, and Dinar (2014); Wichman (2014); Buck et al. (2016); Wichman, Taylor, and Haefen (2016); Hagerty (2018)). Previous studies have also estimated farmers' elasticity of demand for surface water, most notably Hagerty (2018), who finds an elasticity of -0.23 for surface water in California agriculture. While estimates of surface water demand are often as large as -0.80 in specific locations (Schoengold, Sunding, and Moreno (2014); Hagerty (2018)), we find agricultural groundwater demand to be even more elastic on average.

observe 336 million observations of hourly electricity consumption for 10,930 service points in our estimation sample. All of these farmers are on time-varying electricity tariffs: during summer months, their average marginal electricity price nearly doubles during 12–6pm peak periods. Figure 5 plots these average marginal prices against the density of their hourly electricity consumption during summer months. We observe no pronounced decrease in consumption in response to these large swings in marginal price. This suggests that short-run adjustments to irrigating existing crops are not the primary mechanism driving our elasticity estimates.³⁴This is also consistent with our informal conversations with California growers, who tend to adhere to crop-specific irrigation targets.³⁵

Changing the groundwater production function By contrast, changes in the kWh-to-AF function should manifest over longer time scales, as a result of pump depreciation, maintenance, and upgrades. To test whether pumping technology changes are driving our elasticity estimates, or whether measurement error in $\widehat{\frac{\text{kWh}}{\text{AF}}}_{it}$ from infrequent pump tests are biasing our results, we re-estimate Equation (5) using only observations that are close in time to a pump test that we observe. Table 5 presents the results, which are quite stable even when we restrict the sample to the 30 percent of observations within 12 months of a pump test. This implies that unobserved changes in kWh-to-AF are unlikely to be impacting our results.³⁶

6.1 A stylized model of farmer decision-making

This leaves two candidate mechanisms: switching water sources and switching crops. Here, we develop a stylized model to characterize the economics underlying both switching margins. Let i index farms, or atomistic pieces of cropland with area A_i . Farm i makes a discrete choice to plant crop k from a set of potential crops \mathcal{K} , with an outside option of fallowing

34. In ongoing work, we are extending our econometric analysis to formally estimate short-run demand elasticities at the daily and hourly levels.

35. Accelerated orchard management would be another manifestation of this mechanism, which we hope to explore further in future work. Farmers may respond to pumping cost shocks by shifting forward the timing of when they replace old trees with young saplings, which require less water as they mature.

36. Another possibility is that unobserved changes in other irrigation capital impact the share of pumped water than ultimately reaches crops. While we lack data on irrigation capital “downstream” from pumps, it is reassuring that we find similar elasticities using a significantly shortened time period for estimation.

($k = 0$). The choice of crop determines *ex ante* expected yields $Y_i(k)$, water required for irrigation $W_i(k)$, and non-irrigation costs $F_i(k)$. Yields, water, and other costs vary cross-sectionally within crops, due to heterogeneous climate, soil quality, irrigation capital, etc.³⁷ All farms are price-takers in the output market, facing common crop prices p^k . Farm i chooses the crop k that maximizes its profits:

$$\max_{k \in \mathcal{K}} \pi_i(k) = p^k Y_i(k) - C_i(W_i(k)) - F_i(k) \quad (6)$$

Irrigation costs $C_i(W)$ are farm-specific and weakly convex in W , since farmers may irrigate using surface water allocated by their water district, pumped groundwater, and/or water purchased on the open market. Water districts typically have the lowest cost per acre-foot, but limited allocations may not be sufficient to meet irrigation needs, especially in drought years. This may force farmers to pump their own groundwater, with costs per acre-foot that are (typically) higher and may increase convexly in the quantity extracted. Open market water prices tend to be much higher than both district-allocated surface water and pumped groundwater, due to high costs of physically moving water. However, district allocations and pumping costs are sufficiently heterogeneous that some California farmers rely on water markets for irrigation (Hagerty (2018)).

The upper left panel of Figure 6 depicts a hypothetical irrigation price schedule for a farm that irrigates with both surface water (allocated by its irrigation district) and pumped groundwater. This representative farm i has chosen crop k^0 , and the shaded region under the price function depicts its total cost of irrigating, $C_i(W_i(k^0))$. If the farm experiences a pumping cost shock due to either an electricity price increase or a groundwater depth increase, the groundwater piece of its price schedule will shift up. The upper right panel of Figure 6 illustrates how such a pumping cost shock would increase farm i 's cost of irrigating crop k^0 by the shaded area $\Delta C_i(W_i(k^0))$.

Our econometric results show that such a pumping cost shock (holding the rest of the price schedule constant) causes average groundwater consumption to decrease. The bottom

37. For simplicity, we abstract away input re-optimization within crops. This focuses our model solely on the planting margin, while also aligning it with standard crop budgeting calculations farmers typically use to make such *ex ante* planting decisions. This stylized static model also ignores the obvious state-dependence inherent in choosing perennial (or annual) crops.

panels of Figure 6 illustrate how this consumption decrease could come through either crop switching or water source substitution. In the lower left panel, farm i switches to a less water-intensive crop k^1 , thereby reducing both its groundwater consumption and its total water consumption. In the lower right panel, a larger pumping cost shock causes farm i to switch to the open market backstop; while it continues to consume $W_i(k^0)$ acre-feet of water, new water purchases now crowd out its groundwater consumption. As long as district-allocated surface water is inframarginal for groundwater users (Hagerty (2019)), water source substitution is only likely to occur at extremely high prices. Hagerty (2018) reports the distribution of prices for 671 California water transactions, with a mean price of \$221 per acre-foot. Since the PGE data almost never imply pumping costs above \$200 per acre-foot, and since California water markets are relatively thin, crop switching (and fallowing) appears more likely as a mechanism for groundwater demand response.

6.2 Empirical tests of crop switching

6.2.1 Linear estimates

We take the model in Section 6.1 to data using several empirical tests. First, we collapse our monthly panel dataset to the service-point-by-year level, and estimate annual elasticities. While monthly data enable us to include more granular fixed effects to non-parametrically control for potential time-varying confounders, annual data more closely align with the timescale of a farmer’s cropping choice. We estimate annual elasticities using modified versions of Equations (3) and (5):

$$\sinh^{-1}(Q_{iy}^{\text{elec}}) = \beta \log(P_{iy}^{\text{elec}}) + \gamma_i + \delta_y + \varepsilon_{iy} \quad (7)$$

and

$$\sinh^{-1}(Q_{iy}^{\text{water}}) = \beta \log(P_{iy}^{\text{water}}) + \gamma_i + \delta_y + \varepsilon_{iy} \quad (8)$$

where the unit of observation is now the service point-(i)-by-year-(y). As in Sections 4.1–4.2, we instrument for the price of electricity (groundwater) with the default within-category price of electricity. Columns (1)–(2) of Table 6 present the main results, where we estimate

an annual demand elasticity of -0.99 for electricity and -0.93 for groundwater. Since these annual results are similar to our monthly estimates, they suggest that farmers are not simply arbitraging within-year fluctuations in groundwater prices by switching to surface water.³⁸

To further investigate crop switching as a mechanism, we estimate pumping changes on the intensive (reducing the amount of non-zero pumping) vs. extensive (stopping pumping entirely) margins. First, we estimate Equation (7) for the subset of service points with non-zero groundwater use in each sample year. Next, we estimate a semi-elasticity, by replacing the dependent variable in Equation (7) with $1[Q_{iy}^{\text{water}}]$, a binary indicator for whether a unit consumes groundwater in a given year. Table 6 reports these results. In Column (3), we find that the intensive margin elasticity (-0.22) is substantially lower than the average elasticity.³⁹

By contrast, Column (4) reports a semi-elasticity of -0.04 for the extensive margin: a 10-percent increase in effective groundwater price increases the probability that a farmer stops pumping entirely by 0.4 percentage points—perhaps via fallowing. Comparing across Columns (2)–(4), we see further evidence that our elasticities are driven by units moving in and out of pumping, which is consistent with the crop-switching mechanism.

6.2.2 Discrete choice estimates

Building on this evidence, we estimate a discrete crop choice model to directly measure the causal effect of groundwater costs on cropping decisions.⁴⁰ We begin by re-writing the farm’s crop choice optimization problem, from Equation (6), as:

$$\max_{k \in \mathcal{K}} \pi_{iy}(k) = \beta_k p_{iy}^{\text{water}} + \gamma_{cyk} + \varepsilon_{iyk} \quad (9)$$

38. Decomposing the variance in monthly P_{it}^{water} into between-unit, within-unit-within-year, and within-unit-across-year components, we find that 68 percent of the price variation is between units. Of the remaining within-unit variation, 84 percent is within-year, and only 16 percent is across-year. Nevertheless, our monthly and annual elasticity estimates are quite similar, suggesting that the former are mostly driven by within-unit-across-year variation in groundwater price. This provides further evidence against (i) farmers applying less water to existing crops and (ii) short-run substitutions between surface water and groundwater.

39. This is consistent with Appendix Table D10, where we compare the inverse hyperbolic sine (IHS) and log transformations at the monthly level. We find substantially larger elasticities with the IHS (which admits zeros) than with the log (which does not).

40. Appendix C provides a step-by-step derivation of the following model.

where farm i in county c in year y chooses among four possible crop types k : annuals, fruit and nut perennials, other perennials, and fallow. Within these groups, crops exhibit roughly similar cost structures. We decompose profits for farm i into two components: the cost of groundwater and all other costs and revenues. The total cost of groundwater is a linear function of the price of groundwater, $\beta_k p_{iy}^{water}$, which we allow to vary by crop type. We represent all remaining costs and revenue with the set of parameters γ_{cyk} , which gives the average annual profits—excluding groundwater costs—from growing crop type k in county c in year y . We estimate the choice of crop type using a multinomial probit model. As in our previous regressions, the groundwater price is potentially endogenous, so we instrument with the unit’s within-category default electricity price to isolate exogenous variation in the price of groundwater.

In this discrete choice model, we continue to use data aggregated to the annual level. In order to assign crops to units, we match each service point to a USDA CLU based on its latitude and longitude, and use CLU as the cross-sectional unit of analysis. As described in Section 3, we use the Cropland Data Layer to assign a crop type to each 30-by-30 meter pixel in each CLU. We aggregate these pixels to the CLU level by assigning each CLU a crop type based on the modal type across pixels within the CLU, for CLUs where the modal crop type covers more than 50 percent of the CLU.⁴¹

Table 7 shows the marginal effects of groundwater price on crop choice and the semi-elasticities of crop choice with respect to groundwater price that are implied by our discrete choice estimates, averaged over all units and reported in percentage points. On average, when the price of groundwater increases by 10 percent, a farmer increases the proportion of land in fruit and nut perennials by 1.0 percentage points and increases fallowed land by 0.4 percentage points, and reduces land in annuals by 0.9 percentage points and other perennials by 0.5 percentage points. This semi-elasticity for fallowing is quite close to the semi-elasticity on the extensive margin of pumping in Table 6 (-0.04). Importantly, because we estimate a static discrete choice model in a setting where dynamics are inherently important, these semi-elasticity estimates likely understate the true land-use response (Scott (2013)).⁴²

41. In this draft, we drop CLUs where the modal crop type covers less than 50 percent of the CLU.

42. In ongoing work, we are adapting the model developed by Scott (2013) to our setting, in order to account for the dynamics inherent in crop choice.

6.2.3 Counterfactual groundwater taxes

We use these discrete choice estimates to simulate the effects of a uniform Pigouvian groundwater tax on cropping decisions. Because the precise size of the groundwater extraction externality is unknown, we consider a range of possible taxes from \$5 per acre-foot to \$15 per acre-foot, which roughly corresponds to a 12–37 percent increase in the average price of groundwater.⁴³ For each tax level, we add the tax to the observed groundwater price faced by each farmer and use the estimated model parameters to calculate the counterfactual choice probability for each crop type.

Table 8 reports the estimated total acreage for each crop type within our sample when farmers face different counterfactual groundwater taxes.⁴⁴ In line with the semi-elasticity estimates, as the groundwater tax increases, farmers shift acreage into fruit and nut perennials and into fallowing, and they decrease acreage in annuals and other perennials. We find that a relatively large share of total acreage is reallocated to different crop types in response to relatively moderate groundwater taxes. For a \$10 per acre-foot tax, more than 24,000 acres—nearly 8 percent of total cropland in our sample—are reallocated to a different crop type. Figure 7 plots how each CLU in our sample would respond to a \$10 groundwater tax.

If we conducted this counterfactual exercise using a Pigouvian tax equal to true (unknown) marginal external costs at the socially efficient level of groundwater extraction, we could interpret the amount of land reallocated due to the tax as the amount of land that is currently misallocated due to unpriced pumping externalities. If the true externality were \$10 per acre-foot, our results would imply that 8 percent of cropland in our 314,884-acre sample is currently misallocated.⁴⁵ Under the assumption that our sample is representative of all 24.3 million acres of California cropland, this result implies that nearly two million acres of cropland are currently growing a socially suboptimal crop type due to groundwater being underpriced. However, to our knowledge there is scant existing evidence on the true

43. The average annual groundwater price for CLUs in our sample is \$41.00 per acre-foot.

44. There are 314,884 acres that we can confidently match to APEP units, which we use for estimation.

45. We are limited to CLUs in PGE service territory, with PGE service points that match to an APEP pump test. In ongoing work, we are extending the sample to also include agricultural groundwater pumpers in Southern California Edison service territory.

external costs from groundwater extraction in California; we encourage future research on this topic.

Alternatively, we use California policy to benchmark a counterfactual groundwater tax. Under SGMA, California has opted for broad quantity-based groundwater sustainability targets. When the regulation goes into effect, major groundwater basins in California are expected to require groundwater pumping restrictions of 20–50 percent (Bruno (2019)). As shown in the bottom row of Table 8, our groundwater demand elasticity estimates imply that a \$10 per acre-foot tax would yield a 27 percent reduction in groundwater extraction. This is at the lower end of the range of expected curtailment, lending credence to our use of a \$10 tax as our central counterfactual scenario. Regardless of the policy mechanism or stringency, restrictions on California farmers’ groundwater use will likely lead to substantial changes in land use.

7 Conclusion

In this paper, we estimate how a key sector—agriculture—is likely to respond to environmental regulation which will increase the costs of an essential input—groundwater—in the setting of California, one of the most productive farming regions in the world. While California accounts for 18 percent of total U.S. crop value, its farmers are heavily dependent on groundwater-based irrigation. As a result of overuse, water levels in California’s underground aquifers have fallen substantially, with 21 of the state’s groundwater basins now deemed “critically overdrafted”.

In an effort to prevent aquifer collapse, California policymakers have recently passed the state’s first comprehensive groundwater legislation, the Sustainable Groundwater Management Act (SGMA), which will require large reductions in groundwater use. Regardless of the exact implementation approach, these regulations will raise the costs of groundwater for California’s farmers.

In order to understand how farmers will respond to increases in water costs, we must overcome measurement and identification challenges: groundwater is typically neither priced nor measured, and groundwater costs are typically not randomly assigned. We leverage the

insight that electricity is the key marginal input into the groundwater production function. We use a novel restricted-access dataset on farmers' electricity consumption and groundwater pump efficiencies, combined with government measurements of groundwater depths and satellite-derived land use designations for the universe of farmers in the Pacific Gas and Electric (PGE) utility service territory, which covers the majority of the farmland in the state. We leverage exogenous variation in electricity tariffs over time to estimate farmers' price elasticity of demand for electricity, and find a surprising large elasticity estimate of -1.17 . We then use the physics of groundwater pumping to compute groundwater costs and groundwater quantities for each pump in our sample, and we estimate the price elasticity of groundwater demand to be -1.12 . These elasticities are much larger than previous estimates in the electricity and groundwater literatures.

We then explore the mechanisms underlying farmers' groundwater demand response. We find evidence consistent with crop switching as the primary mechanism driving our estimated demand elasticities. First, farmers do not appear to respond to large within-day switches in price, meaning that changing water use on existing crops is unlikely to explain our results. Second, we find similar elasticities when we restrict our sample to months around pump tests, making it unlikely that pumping capital upgrades are driving our estimates. Third, because surface water tends to be substantially cheaper than groundwater, and because we estimate similar elasticities at the monthly and annual level, within-year surface water substitution is not likely to be the primary mechanism. Finally, we find a large semi-elasticity on the extensive margin of groundwater use, showing that a key means of adjustment is to cease pumping altogether—consistent with crop switching or fallowing.

We build on this evidence by using a discrete choice model to estimate the impact of increasing groundwater costs on crop choice. We find that higher groundwater costs cause farmers to increase acreage in fruit and nut perennials and increase fallowing, and to decrease acreage in annuals and other perennials. We simulate a counterfactual groundwater tax to estimate the impacts of potential groundwater pricing policies on land use in California. We find that a moderate \$10 per acre-foot tax—approximately the price increase our elasticity estimates imply would be required to meet California's lower-bound target of 20 percent curtailment—would lead farmers to reallocate nearly 8 percent of land to a different crop

type. If the planned curtailments under SGMA reflect the true externality from groundwater extraction, extrapolating our estimates from our sample to the rest of the state suggests that nearly 2 million acres of California cropland may be misallocated due to unpriced groundwater.

References

- Alberini, Anna, and Massimo Filippini. 2011. "Response of residential electricity demand to price: The effect of measurement error." *Energy Economics* 33:889–895.
- Badiani, Reena, and Katrina Jessoe. 2015. "The impact of electricity subsidies on groundwater extraction and agricultural production." Working paper.
- Baerenklau, Kenneth A., Kurt A. Schwabe, and Ariel Dinar. 2014. "The residential water demand effect of increasing block rate water budgets." *Land Economics* 90 (4): 683–699.
- Bellemare, Marc F., and Casey J. Wichman. 2020. "Elasticities and the inverse hyperbolic sine transformation." *Oxford Bulletin of Economics and Statistics* 82 (1): 50–61.
- Blonz, Joshua A. 2016. "Making the best of the second-best: Welfare consequences of time-varying electricity prices." Energy Institute at Haas Working Paper 275.
- Bruno, Ellen. 2019. *Groundwater Policy and California Agriculture*. Technical report. University of California Cooperative Extension.
- Bruno, Ellen, and Katrina Jessoe. 2018. "Water markets and climate change adaptation: Micro-level evidence on agricultural water demand." Mimeo.
- Buck, Steven, Maximilian Auffhammer, Stephen Hamilton, and David Sunding. 2016. "Measuring welfare losses from urban water supply disruptions." *Journal of the Association of Environmental and Resource Economists* 3 (3): 743–778.
- 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 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.
- California Energy Commission. 2005. "California's water-energy relationship." CEC Report 700-2005-011-SF.


- Deryugina, Tatyana, Alexander MacKay, and Julian Reif. 2018. “The long-run dynamics of electricity demand: Evidence from municipal aggregation.” Working paper.
- Famiglietti, J. S. 2014. “The global groundwater crisis.” *Nature Climate Change* 4 (11): 945.
- 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, Catherine Wolfram, C. Anna Spurlock, Annika Todd, Patrick Baylis, and Peter Cappers. 2018. “Default effects and follow-on behavior: Evidence from an electricity pricing program.” Working paper.
- Hagerty, Nick. 2018. “Liquid constrained in California: Estimating the potential gains from water markets.” Working paper.
- . 2019. “The scope for climate adaptation: Evidence from water scarcity in irrigated agriculture.” Working paper.
- Hanack, 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.
- Hewitt, J.A., and W.M. Hanemann. 1995. “A discrete/continuous choice approach to residential water demand under block rate pricing.” *Land Economics* 71 (2): 173–192.
- Ito, Koichiro. 2013. “How do consumers respond to nonlinear pricing? Evidence from household water demand.” Working paper.
- . 2014. “Do consumers respond to marginal or average price? Evidence from nonlinear electricity pricing.” *American Economic Review* 104 (2): 537–563.
- Ito, Koichiro, Takanori Ida, and Makoto Tanaka. 2018. “Moral suasion and economic incentives: Field experimental evidence from electricity demand.” *American Economic Journal: Economic Policy* 10 (1): 240–267.
- Jessoe, Katrina, and David Rapson. 2014. “Knowledge is (less) power: Experimental evidence from residential electricity use.” *American Economic Review* 104 (4): 1417–1438.
- . 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.
- 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.

- Nataraj, Shanthi. 2011. “Does marginal price matter? A regression discontinuity approach to estimating water demand.” *Journal of Environmental Economics and Management* 61:198–212.
- Olmstead, S.M., Michael W. Hanemann, and R.N. Stavins. 2007. “Water demand under alternative price structures.” *Journal of Environmental Economics and Management* 54 (2): 181–198.
- Paul, Anthony, Erica Myers, and Karen Palmer. 2009. “A partial adjustment model of U.S. electricity demand by region, season, and sector.” Resources for the Future Working Paper DP 08-50.
- 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.
- Reiss, Peter C., and Matthew W. White. 2005. “Household electricity demand, revisited.” *Review of Economic Studies* 72:853–883.
- Renwick, E.M., and R.D. Green. 2000. “Do residential water demand side management policies measure up? An analysis of eight California water agencies.” *Journal of Environmental Economics and Management* 40 (1): 37–55.
- Sawyers, Gary. 2007. “A primer on California water rights.” UC Davis Cooperative Extension working paper.
- 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, 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.
- Schoengold, Katrina, David L. Sunding, and Georgina Moreno. 2014. “Price elasticity reconsidered: Panel estimation of an agricultural water demand function.” *Water Resources Research* 42.
- Scott, Paul T. 2013. “Dynamic discrete choice estimation of agricultural land use.” Working paper.
- Siebert, S., J. Burke, J. M. Faures, K. Frenken, J. Hoogeveen, P. Döll, and F. T. Portmann. 2010. “Groundwater use for irrigation – a global inventory.” *Hydrology and Earth System Sciences* 14:1863–1880.
- USDA. 2020. “Crop values: 2019 summary.” United States Department of Agriculture report.

- Wichman, Casey J. 2014. "Perceived price in residential water demand: Evidence from a natural experiment." *Journal of Economic Behavior and Organization* 107:308–323.
- Wichman, Casey J., Laura O. Taylor, and Roger H. von Haefen. 2016. "Conservation policies: Who responds to price and who responds to prescription?" *Journal of Environmental Economics and Management* 79:114–134.
- Wolak, Frank A. 2011. "Do residential customers respond to hourly prices? Evidence from a dynamic pricing experiment." *American Economic Review* 101 (3): 83–87.

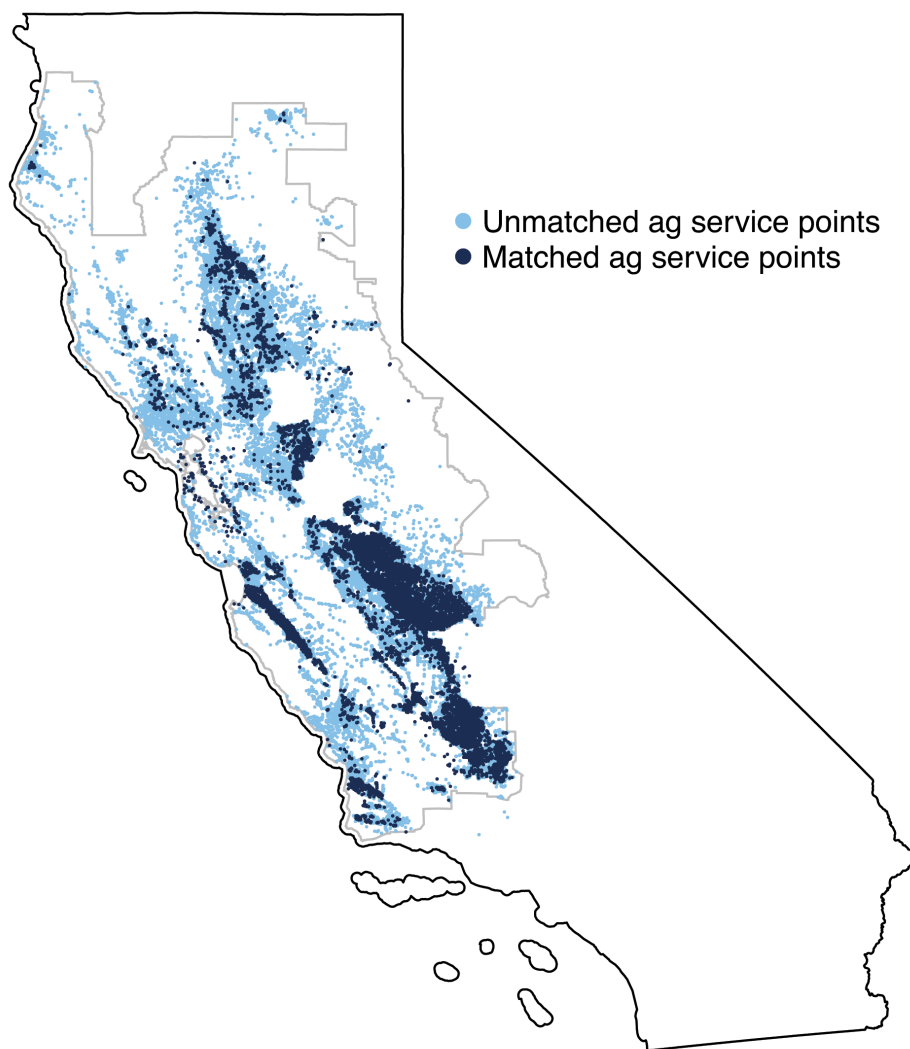
Figures and Tables

Figure 1: Change in cropped area in selected crops, 2010–2017



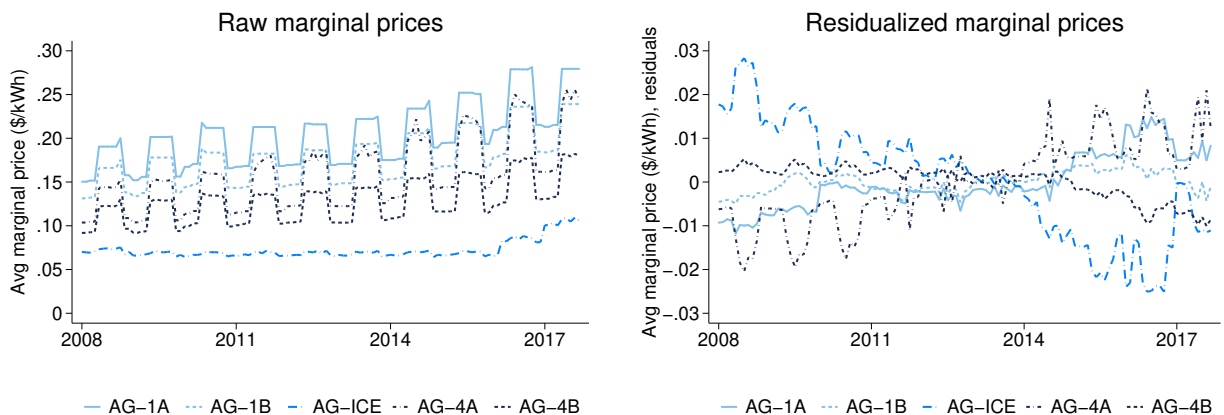
Notes: This figure plots the change in cropped area for nine water-intensive California crops from 2010 (light shaded) to 2017 (dark shaded): before and after the historic droughts of the mid-2010s. Crops in dark blue are relatively lower value per acre, while crops in light blue are relatively high value per acre. Acres planted measured using the USDA's Cropland Data Layer.

Figure 2: PGE agricultural customers




Notes: This figure maps the locations of all agricultural service points served by PGE. Dark blue dots indicate the 11,851 service point that we can match directly to an APEP pump test. Light blue dots indicate unmatched agricultural service points. The light grey outline is the geographic boundary of PGE's service territory.

Figure 3: Average marginal electricity prices




Notes: This figure plots time series of monthly average marginal electricity prices (\$/kWh) for PGE’s five default agricultural tariffs. The left panel plots raw average marginal prices for each month in our estimation sample, taking unweighted averages across all hours. The right panel plots residuals of these same five time series, after partialing out tariff \times month-of-year fixed effects and month-of-sample fixed effects (aligning with the fixed effects we use in estimation). AG-1A and AG-1B are non-time-varying rates (i.e. constant marginal price for all hours within a month), whereas AG-4A and AG-4B are time-varying rates (i.e. higher marginal prices during peak hours and weekdays). AG-1A and AG-4A are for small pumps (< 35 hp), whereas AG-1B and AG-4B are for large pumps (≥ 35 hp). AG-ICE is a time-varying rate for customers with auxiliary internal combustion engines. Marginal prices are systematically higher during summer months (May–October). Our identifying variation comes (a) strict restrictions that segment customers into categories; (b) the fact that the residualized default prices do not move in parallel; and (c) PGE’s smart meter rollout, which exogenously shifted many customers from the AG-1A/1B default tariffs to the AG-4A/4B default tariffs with lower marginal prices.

Figure 4: Histogram of pump horsepower




Notes: This is a histogram of measured horsepower for all 21,851 tests in our APEP pump test dataset. We observe no bunching on either side of the 35 hp cutoff that determines whether PGE classifies pumps as small or large. Bunching would be a sign that farmers optimize against PGE’s tariff schedules when making pump investment decisions.

Figure 5: Hourly electricity prices and consumption



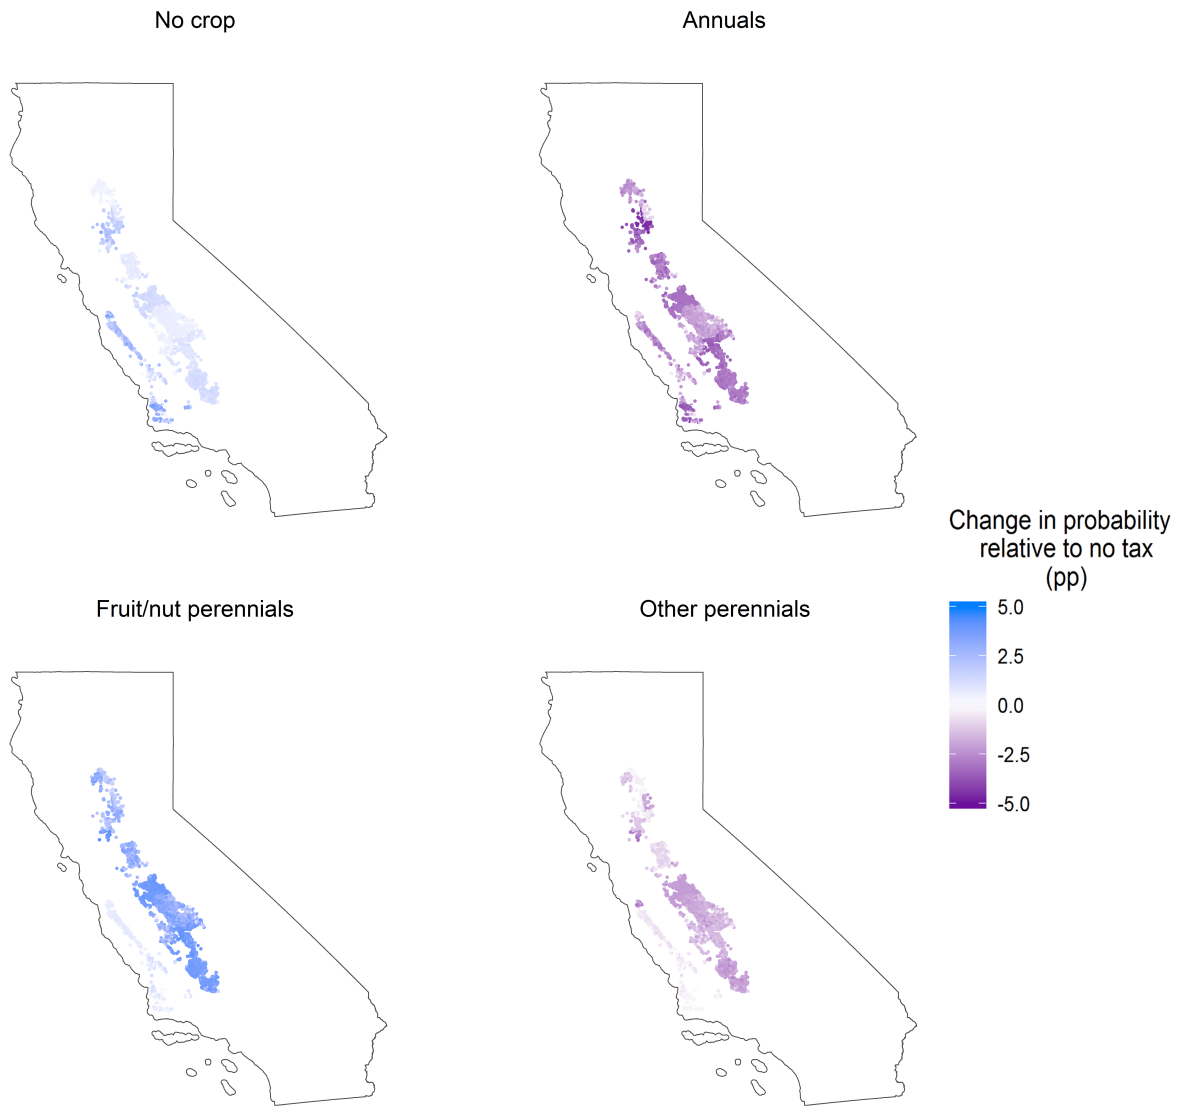
Notes: This figure plots a histogram of electricity consumption at each hour of the day against average marginal price at each hour of the day for summer months. While average marginal prices rise substantially during the afternoon, consumption does not fall in response.

Figure 6: Modeling farm i 's water costs and groundwater demand response



Notes: This figure presents a stylized water price schedule for a representative farm i . The price schedule is nonlinear and comprises water from up to three sources: (i) a low-cost allocation of surface water from farm i 's irrigation district; (ii) medium-cost groundwater pumping, with costs that rise gradually in own extraction; and (iii) a high-cost backstop of open market water transactions, for which we assume farm i is a price taker. For crop k^0 requiring $W_i(k^0)$ acre/feet of water, farm i 's irrigation costs $C_i(W_i(k^0))$ are represented by the shaded region in the top-left panel. If farm i experiences a pumping cost shock (due to either an electricity price increase or a groundwater depth increase), its groundwater costs shift up and its total irrigation costs increase by the shaded region $\Delta C_i(W_i(k^0))$ in the top-right panel. The bottom panels illustrate two ways that this pumping cost shock increase translate to a reduction in farm i 's groundwater consumption. First, the farmer may respond to this pumping cost shock by switching from crop k^0 to a less water intensive crop k^1 , as in the bottom-left panel. Second, for a large enough cost shock, the farmer may continue to grow crop k^0 , but substitute away from groundwater using open market water purchases.

Figure 7: Crop choice changes in response to a \$10 groundwater tax



Notes: This figure plots the estimated counterfactual changes in crop choice resulting from a \$10 groundwater tax, relative to no tax. Each dot is the centroid of a CLU in our sample. To generate these estimates, we first compute the average probability each CLU farms each crop type across all sample years in the no-tax baseline. We then compute similar probabilities with the groundwater price increased by \$10 per acre-foot. Finally, we subtract the no-tax baseline from the \$10 tax counterfactual to compute the change in probability of having each land type. The top left panel shows the change in following. The top right panel shows the change in annual crops. The bottom left panel shows the change in fruit and nut perennials, and the bottom right shows changes in other perennials.

Table 1: Summary statistics – Electricity data

	All Ag Customers	Matched to Pumps
Service point-month observations	9,991,458	1,168,511
Unique service points (SPs)	108,172	11,849
SPs that switch tariff categories	44,414	2,844
SPs that switch categories (pumping capital)	3,454	561
SPs that switch categories (smart meters)	43,045	2,553
Share of SP-months on time-varying tariffs	0.702	0.886
Share of SP-months on peak-day tariffs	0.295	0.152
Monthly electricity consumption (kWh)	6080.9 (39783.1)	12055.6 (25075.3)
Monthly electricity consumption (kWh), summer	8249.6 (45660.8)	17589.0 (29818.9)
Monthly electricity consumption (kWh), winter	3849.7 (32498.8)	6362.8 (17232.7)
Average marginal electricity price (\$/kWh)	0.148 (0.050)	0.113 (0.042)
Average marginal electricity price (\$/kWh), summer	0.171 (0.051)	0.130 (0.044)
Average marginal electricity price (\$/kWh), winter	0.126 (0.037)	0.096 (0.032)
Average monthly bill (\$, non-zero bills)	936.66 (4662.71)	1814.15 (3285.26)
Average monthly bill (\$, non-zero bills), summer	1398.90 (5847.34)	2821.16 (4020.99)
Average monthly bill (\$, non-zero bills), winter	456.17 (2888.86)	764.99 (1742.21)

Notes: The left column reports summary statistics for the universe of agricultural electricity customers in PGE service territory, from 2008–2017. The right column includes the subset of agricultural customers that we successfully match to a groundwater pump in the APEP pump test dataset—i.e., our main estimation sample. “Pumping capital” denotes tariff category switches driven by shifts between small pumps (< 35 hp) and large pumps (\geq 35 hp), or adding/removing an auxiliary internal combustion engine. Most tariff category switches were driven by PGE’s smart meter rollout. Time-varying tariffs (i.e. all except 1A and 1B) have higher marginal prices during peak demand hours. Peak-day tariffs (i.e. 4A, 4D, 4C, 4F, 5C, 5F) have very high marginal prices during peak hours on the 14 highest-demand summer days. Monthly bills include both volumetric (\$/kWh) and fixed charges (\$/kW, \$/hp, and \$/day). Summer months are May–October. Standard deviations of sample means in parentheses.

Table 2: Summary statistics – Pump tests and groundwater consumption

	Matched to Pumps
Service point-month observations	1,168,511
Unique service points (SPs)	11,849
Matched APEP points per SP	1.67 (1.73)
Operating pump efficiency (%)	54.46 (11.52)
kWh per AF conversion factor (APEP measured)	430.35 (254.23)
kWh per AF conversion factor (constructed)	321.12 (185.81)
Monthly groundwater consumption (AF)	46.8 (133.1)
Monthly groundwater consumption (AF), summer	69.7 (155.8)
Monthly groundwater consumption (AF), winter	23.2 (99.4)
Average marginal groundwater price (\$/AF)	40.35 (30.08)
Average marginal groundwater price (\$/AF), summer	43.27 (32.38)
Average marginal groundwater price (\$/AF), winter	37.34 (27.19)

Notes: These summary stats are from the merged panel of groundwater prices and quantities, which combines electricity data, pump test data, and groundwater data. We observe 3.45 unique APEP pump tests for the average matched service point, although 37 percent of service points match to only a single APEP test. Our constructed kWh per AF conversion factor (i.e. $\widehat{\text{kWh}/\text{AF}}_{it}$) uses monthly groundwater rasters to capture changes in (measured) kWh per AF over time, and estimation error compresses the right tail of distribution of measured kWh per AF. Monthly groundwater consumption divides electricity consumption (kWh) by $\widehat{\text{kWh}/\text{AF}}_{it}$. Groundwater prices multiply marginal electricity prices (\$/kWh) by $\widehat{\text{kWh}/\text{AF}}_{it}$. Summer months are May–October. Standard deviations of sample means in parentheses.

Table 3: Estimated demand elasticities – Electricity

	(1)	(2)	(3)	(4)	(5)	(6)
	OLS	IV	IV	IV	IV	IV
$\log(P_{it}^{\text{elec}})$	-1.31*** (0.11)	-1.58*** (0.17)	-1.17*** (0.16)	-0.95*** (0.14)	-1.18*** (0.21)	-0.76*** (0.17)
Instrument(s):						
Default $\log(P_{it}^{\text{elec}})$		Yes	Yes	Yes		Yes
Default $\log(P_{it}^{\text{elec}})$, lagged					Yes	
Fixed effects:						
Unit \times month-of-year	Yes	Yes	Yes	Yes	Yes	Yes
Month-of-sample	Yes	Yes	Yes	Yes	Yes	Yes
Unit \times physical capital			Yes	Yes	Yes	Yes
Water basin \times year				Yes		
Water district \times year				Yes		
Unit-specific linear time trends						Yes
Service point units	11,173	11,173	11,173	11,167	10,922	11,173
Months	117	117	117	117	105	117
Observations	1.05M	1.05M	1.05M	1.05M	0.91M	1.05M
First stage F -statistic		4136	7382	7542	757	4776

Notes: Each regression estimates Equation (3) at the service point by month level, where the dependent variable is the inverse hyperbolic sine transformation of electricity consumed by service point i in month t . We estimate IV specifications via two-stage least squares, instrumenting with either unit i 's within-category default logged electricity price in month t or the 6- and 12-month lags of this variable. "Physical capital" is a categorical variable for (i) small pumps, (ii) large pumps, and (iii) internal combustion engines, and unit \times physical capital fixed effects control for shifts in tariff category triggered by the installation of new pumping equipment. Water basin \times year fixed effects control for broad geographic trends in groundwater depth. Water district \times year fixed effects control for annual variation in surface water allocations; we include a common "no-water-district" dummy for units not assigned to a water district, to avoid dropping them from the regression. All regressions drop solar NEM customers, customers with bad geocodes, and months with irregular electricity bills (e.g. first/last bills, bills longer/shorter than 1 month, overlapping bills for a single account). Standard errors (in parentheses) are two-way clustered by service point and by month-of-sample. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 4: Estimated demand elasticities – Groundwater

	(1)	(2)	(3)	(4)	(5)	(6)
	OLS	IV	IV	IV	IV	IV
$\log(P_{it}^{\text{water}})$	-0.88*** (0.07)	-1.12*** (0.15)	-1.16*** (0.17)	-1.12*** (0.15)	-0.90*** (0.14)	-1.14*** (0.21)
Instrument(s):						
Default $\log(P_{it}^{\text{elec}})$		Yes	Yes	Yes	Yes	
Default $\log(P_{it}^{\text{elec}})$, lagged						Yes
Fixed effects:						
Unit \times month-of-year	Yes	Yes	Yes	Yes	Yes	Yes
Month-of-sample	Yes	Yes	Yes	Yes	Yes	Yes
Unit \times physical capital	Yes	Yes	Yes	Yes	Yes	Yes
Water basin \times year					Yes	
Water district \times year					Yes	
Groundwater time step	Month	Month	Month	Quarter	Month	Month
Only basins with > 1000 SPs			Yes			
Service point units	10,155	10,155	9,337	10,155	10,149	9,922
Months	117	117	117	117	117	105
Observations	0.93M	0.93M	0.85M	0.93M	0.93M	0.82M
First stage F -statistic		3021	2791	3198	4562	477

Notes: Each regression estimates Equation (5) at the service point by month level, where the dependent variable is the inverse hyperbolic sine transformation of groundwater consumed by service point i in month t . We estimate IV specifications via two-stage least squares, and Columns (2)–(5) instrument for P_{it}^{water} with unit i 's within-category default logged electricity price. Column (6) instruments with the 6- and 12- month lags of this variable. “Physical capital” is a categorical variable for (i) small pumps, (ii) large pumps, and (iii) internal combustion engines, and unit \times physical capital fixed effects control for shifts in tariff category triggered by the installation of new pumping equipment. Water basin \times year fixed effects control for broad geographic trends in groundwater depth. Water district \times year fixed effects control for annual variation in surface water allocations; we include a common “no-water-district” dummy for units not assigned to a water district, to avoid dropping them from the regression. Column (3) restricts the sample to only the three most common water basins (San Joaquin Valley, Sacramento Valley, and Salinas Valley), each of which contains over 1000 unique SPs in our estimation sample. Column (4) uses a quarterly panel of groundwater depths to construct both Q_{it}^{water} and P_{it}^{water} , rather than a monthly panel. All regressions drop solar NEM customers, customers with bad geocodes, months with irregular electricity bills (e.g. first/last bills, bills longer/shorter than 1 month, overlapping bills for a single account), and pumps with implausible test measurements. Standard errors (in parentheses) are two-way clustered by service point and by month-of-sample. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 5: Sensitivity to recent pump tests – Groundwater

	(1)	(2)	(3)	(4)	(5)
	IV	IV	IV	IV	IV
$\log(P_{it}^{\text{water}})$	-1.10*** (0.16)	-1.00*** (0.17)	-0.92*** (0.18)	-1.01*** (0.20)	-0.90*** (0.23)
Months away from pump test:	60	48	36	24	12
IV: Default $\log(P_{it}^{\text{elec}})$	Yes	Yes	Yes	Yes	Yes
Fixed effects:					
Unit \times month-of-year	Yes	Yes	Yes	Yes	Yes
Month-of-sample	Yes	Yes	Yes	Yes	Yes
Unit \times physical capital	Yes	Yes	Yes	Yes	Yes
Groundwater time step	Month	Month	Month	Month	Month
Service point units	10,144	10,129	10,110	10,054	9,826
Months	117	117	117	105	93
Observations	0.82M	0.74M	0.63M	0.48M	0.28M
First stage F -statistic	2803	2557	2098	1517	902

Notes: Each regression replicates our preferred specification in Column (2) from Table 4, while restricting the sample to units with pump tests within m months of sample month t . For example, a March 2013 observation for unit i is only included in Column (3) if we observe a pump test for unit i between March 2010 and March 2016. These regressions reveal that unobserved changes in pump specifications are unlikely to be systematically biasing our groundwater elasticity estimates. Stated differently, the mechanism underlying our estimates is unlikely to be unobserved changes to farmers' irrigation capital. See notes under Table 4 for further detail. Standard errors (in parentheses) are two-way clustered by service point and by month-of-sample. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 6: Annual demand elasticities – Intensive vs. extensive margin

	Electricity	Groundwater		
	Overall elasticity (1)	Overall elasticity (2)	Intensive margin (3)	Extensive margin (4)
$\log(P_{iy}^{\text{elec}})$	-0.99*** (0.25)			
$\log(P_{iy}^{\text{water}})$		-0.93*** (0.24)	-0.22** (0.09)	-0.04*** (0.01)
Outcome:				
$\sinh^{-1}(Q_{iy})$	Yes	Yes	Yes	
$1[Q_{iy} > 0]$				Yes
Sample restriction:				
$Q_{iy} > 0$ in all years			Yes	
Instrument:				
Default $\log(P_{iy}^{\text{elec}})$	Yes	Yes	Yes	Yes
Fixed effects:				
Unit \times physical capital	Yes	Yes	Yes	Yes
Water basin \times year	Yes	Yes	Yes	Yes
Water district \times year	Yes	Yes	Yes	Yes
Service point units	10,113	9,058	4,270	9,058
County \times years	270	268	236	268
Observations	58.2K	51.8K	25.0K	51.8K
First stage F -statistic	2165	2425	1092	2425

Notes: Each regression estimates Equation (7) or Equation (8) at the service point by year level. Column (1) reports results for electricity consumption, and Columns (2)–(4) report results for groundwater consumption. Columns (1) and (2) report annual demand elasticities for electricity and water, respectively. These results are analogous to the monthly demand elasticities reported in Column (4) of Table 3 and in Column (5) of Table 4, respectively. Column (3) reports an analogous demand elasticity for the subset of service points that consume water in every year of our sample. Column (4) reports the semi-elasticity for the extensive margin by replacing the outcome variable with a binary indicator water consumption. We estimate these regressions using two-stage least squares, instrumenting with unit i 's within-category default logged electricity price in year y . “Physical capital” is a categorical variable for (i) small pumps, (ii) large pumps, and (iii) internal combustion engines, and unit \times physical capital fixed effects control for shifts in tariff category triggered by the installation of new pumping equipment. Water basin \times year fixed effects control for broad geographic trends in groundwater depth. Water district \times year fixed effects control for annual variation in surface water allocations. All regressions drop solar NEM customers, customers with bad geocodes, years with irregular electricity bills (e.g. first/last bills, bills longer/shorter than 1 month, overlapping bills for a single account), and incomplete years. Groundwater regressions use a monthly time interval to assign rasterized groundwater levels. Standard errors (in parentheses) are two-way clustered by service point and by county-year. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 7: Crop switching estimates – Discrete choice model

	Marginal effect	Semi-elasticity
Annuals	−0.252*** (0.042)	−0.088*** (0.013)
Fruit and nut perennials	0.307*** (0.055)	0.104*** (0.020)
Other perennials	−0.162*** (0.034)	−0.053*** (0.010)
Fallow	0.107*** (0.033)	0.038*** (0.013)
Instrument:		
Default $\log(P_{iy}^{\text{elec}})$		Yes
Fixed effects:		
County \times year \times crop type		Yes
Common land units		7,380
Observations		36,799
First stage χ^2 -statistic		1134

Notes: We estimate the discrete choice model of Equation (9) using IV probit, instrumenting for groundwater price with the default within-category electricity price. We include county-by-year-by-crop-type fixed effects to flexibly estimate profit excluding the cost of groundwater pumping. The left column presents the average marginal effects of groundwater price on crop type choice probabilities. The right column reports the average semi-elasticities for each crop type choice probability with respect to the groundwater price. Each value is calculated for every observation in our analysis, and then we take the mean over all observations to yield these average values. Standard errors (in parentheses) are clustered at the common land unit (CLU). Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 8: Counterfactual groundwater taxes


	No tax	\$5 tax	\$10 tax	\$15 tax
Simulated acreage (thousands of acres)				
Annuals	61.74	57.85	54.09	50.45
Fruit and nut perennials	147.30	151.69	155.93	160.03
Other perennials	34.24	31.94	29.71	27.57
Fallow	71.60	73.40	75.15	76.83
Total reallocation		12.37	24.35	35.92
Total reallocation (percent)		(3.9%)	(7.7%)	(11.4%)
Change in groundwater consumption (percent)		-13.7%	-27.3%	-41.0%

Notes: This table reports the results of adding counterfactual taxes on groundwater to the observed electricity prices in our sample. To simulate the impacts of a groundwater tax, we first calculate the choice probability of each crop type (annuals, fruit/nut perennials, other perennials, and no crop) for each CLU in our sample over our time series. This baseline allocation is represented in the first column, labeled “No tax.” The sample average marginal price is \$41.00 per acre-foot. In the subsequent columns, we take each CLU’s average annual marginal price and add the reported tax level to it. We then calculate choice probabilities for this counterfactual groundwater price. The first four rows correspond to the four crop types in our analysis, and the table displays the total acreage in our sample that we predict would be cropped in each crop type under each of the tax levels. The fifth row reports the total acreage of cropland that is reallocated to a different crop type due to the groundwater tax, as compared to no tax. The sixth row displays the total percent change in land use for each tax level, as compared to no tax. These reallocations estimates are based on the 314,884 acres of agricultural land matched to our sample. The final row reports the estimated change in groundwater consumption, using our groundwater elasticity estimate of -1.12 , for each tax level.

Appendix A PGE electricity prices

As described in Section 3 of the main text, PGE offers 23 distinct agricultural tariffs, which fall into 5 categories, depending on a farmer’s meter type (conventional vs. smart) and pumping capital (small, large, or previously powered by an internal combustion engine). Here, we present additional details on these rates. Table A1 describes each rate in detail, including a description of the eligibility category, the broad pricing schedule on each rate, and the share of customers on each rate within our sample. Figure A1 shows a time series of each rate over our sample. This is analogous to Figure 3 in the main text, but shows all rates in addition to the “default” within-category rate. All rates that are the same color in Figure A1 belong to the same category; the default rate for each category is bolded. The left panel shows the raw rate time series. The right panel shows residualized rates, after partialling out tariff \times month-of-year fixed effects and month-of-sample fixed effects.

Figure A1: Average marginal electricity prices (all rates)



Notes: This figure plots times series of monthly average marginal electricity prices (\$/kWh) for all of PGE agricultural tariffs. These 23 tariffs are divided into 5 mutually-exclusive categories, based on the type of electricity meter on a farm (smart vs. conventional) and a farm’s pumping capital (small, large, or previously internal combustion engine). All rates belonging to the same category are the same color. The five “default” rates, which we also show in main text Figure 3 are bolded. The left panel plots raw average marginal prices for each month in our estimation sample, taking unweighted averages across all hours. The right panel plots residuals of these same five time series, after partialling out tariff \times month-of-year fixed effects and month-of-sample fixed effects (aligning with the fixed effects we use in estimation). AG-1A and AG-1B are non-time-varying rates (i.e. constant marginal price for all hours within a month), whereas AG-4A and AG-4B are time-varying rates (i.e. higher marginal prices during peak hours and weekdays). AG-1A and AG-4A are for small pumps (< 35 hp), whereas AG-1B and AG-4B are for large pumps (≥ 35 hp). AG-ICE is a time-varying rate for customers with auxiliary internal combustion engines. Marginal prices are systematically higher during summer months (May–October). Our identifying variation comes (a) strict restrictions that segment customers into categories; (b) the fact that the residualized default prices do not move in parallel; and (c) PGE’s smart meter rollout, which exogenously shifted many customers from the AG-1A/1B default tariffs to the AG-4A/4B default tariffs with lower marginal prices.

Table A1: PGE agricultural tariffs

Category	Tariff	Description	Percent
Small pumps, conventional meters single motor < 35 hp, or multiple motors summing to < 15 hp	1A	High price per kWh (not time-varying), fixed charge per hp connected	3.0
Large pumps, conventional meters single motor \geq 35 hp, or multiple motors summing to \geq 15 hp, or single overloaded motor \geq 15 hp	1B	High price per kWh (not time-varying), fixed charge per max kW consumed	8.1
Small pumps, smart meters single motor < 35 hp, or multiple motors summing to < 15 hp	4A (4D)	High prices per kWh (higher in peak hours), fixed charges per hp connected, very high peak prices on 14 summer Event Days	7.2
	5A (5D)	Lower prices per kWh (peak & offpeak), no Event Day price increases, higher fixed charges per hp	2.7
	RA (RD)	Lower peak prices per kWh, higher off-peak prices per kWh, no Event Day price increases, choice between MTW or WTF peak days	1.2
	VA (VD)	Lower peak prices per kWh, higher off-peak prices per kWh, no Event Day price increases, choice of 3 shorter 4-hour peak periods	0.9
Large pumps, smart meters single motor \geq 35 hp, or multiple motors summing to \geq 15 hp, or single overloaded motor \geq 15 hp	4B (4E)	High prices per kWh (higher in peak hours), fixed charges per max kW consumed	20.1
	5B (5E)	Much lower prices per kWh (peak & offpeak), higher fixed charge per max kW	37.8
	4C (4F)	Slightly lower prices per kWh (peak & offpeak), higher fixed charges per kW shifted to peak, very high peak prices on 14 summer Event Days	2.4
	5C (5F)	Much lower prices per kWh (peak & offpeak), higher fixed charges per kW shifted to peak, very high peak prices on 14 summer Event Days	7.8
	RB (RE)	Higher prices per kWh (peak & off-peak), choice between MTW or WTF peak days, lower fixed charges per max kW (in summer)	1.5
	VB (VE)	Higher prices per kWh (peak & off-peak), choice of 3 shorter 4-hour peak periods, lower fixed charges per max kW (in summer)	0.6
Customers transitioning off internal combustion engines	ICE	Very low price per kWh (high in peak hours), fixed charge per max kW consumed	6.8

Notes: This table provides a rough summary of PGE's 23 electricity tariffs for agricultural customers. The first column lists the 5 disjoint categories of customers, defined (primarily) by physical pumping capital and electricity meters. Effective default tariffs within each group are in bold, and farmers may switch tariffs *within* a category (but not *across* categories). All tariffs have fixed (\$/kW) and volumetric (\$/kWh) prices that vary by summer vs. winter. All time-of-use tariffs (i.e. all but 1A and 1B) also vary between peak (12:00pm–6:00pm on summer weekdays), partial peak (8:30am–9:30pm on weekends), and off-peak periods. DEF tariffs are functionally equivalent to their ABC analogs, and are holdovers for the earliest customers to adopt time-of-use pricing. Actual tariffs are far more complex, and tariff documents are available at <https://www.pge.com/tariffs/index.page>. The right-most column reports the percent of observations in our main estimation sample on each tariff.

Appendix B Groundwater demand estimation

In Section 4.2 of the main text, we present an approach for estimating the price elasticity of demand for groundwater. Here, we use a decomposition approach and estimate the elasticity with respect to electricity prices and water pumping costs separately.

As in Section 4.2, we aim to estimate causal effect of groundwater price on groundwater consumption, and this demand elasticity is linearly approximated by the coefficient β :

$$\log(Q_{it}^{\text{water}}) = \beta \log(P_{it}^{\text{water}}) \quad (3)$$

We construct Q_{it}^{water} and P_{it}^{water} using the estimated conversion factor $\widehat{\frac{\text{kWh}}{\text{AF}}}_{it}$, which has measurement error and is also potentially endogenous.

Hence, the same measurement error and endogeneity is present on both the left-hand side and the right-hand side of Equation (3). We can rewrite this expression decomposing $\widehat{\frac{\text{kWh}}{\text{AF}}}_{it}$ on both sides:

$$\log(Q_{it}^{\text{elec}}) - \log\left(\widehat{\frac{\text{kWh}}{\text{AF}}}_{it}\right) = \beta \left[\log(P_{it}^{\text{elec}}) + \log\left(\widehat{\frac{\text{kWh}}{\text{AF}}}_{it}\right) \right] \quad (4)$$

Rearranging:

$$\log(Q_{it}^{\text{elec}}) = \beta \log(P_{it}^{\text{elec}}) + (\beta + 1) \log\left(\widehat{\frac{\text{kWh}}{\text{AF}}}_{it}\right) \quad (5)$$

This expression is algebraically equivalent to Equation (3), but it isolates the endogenous estimated conversion factor in one right-hand-side variable. We estimate an analogous regression specification:

$$\sinh^{-1}(Q_{it}^{\text{elec}}) = \beta^e \log(P_{it}^{\text{elec}}) + (\beta^w + 1) \log\left(\widehat{\frac{\text{kWh}}{\text{AF}}}_{it}\right) + \gamma_i + \delta_t + \varepsilon_{it} \quad (6)$$

This specification is similar to Equation (3), except that we can now interpret β^e and β^w as the price elasticity of demand for groundwater. We allow this elasticity to vary depending on the source of variation in pumping costs—groundwater depths may be more salient to farmers than electricity prices, or vice versa.¹ As in the electricity regressions, as well as those in the main text, we purge electricity price endogeneity by instrumenting P_{it}^{elec} with within-category default prices.

To identify β^w , we must overcome three potential sources of bias. First, farmers may choose to alter their pumping technologies in order to change $\widehat{\frac{\text{kWh}}{\text{AF}}}_{it}$, and such changes are likely correlated with Q_{it}^{elec} . Second, $\widehat{\frac{\text{kWh}}{\text{AF}}}_{it}$ is a function of unit i 's groundwater depth, which is mechanically linked to Q_{it}^{elec} —when unit i consumes electricity to extract groundwater,

1. A strict Neoclassical interpretation would assume $\beta^e = \beta^w$, as the optimizing farmer should respond to all short-run changes in P_{it}^{water} identically.

its localized groundwater level falls, thereby increasing $\widehat{\frac{\text{kWh}}{\text{AF}}}_{it}$. Third, $\widehat{\frac{\text{kWh}}{\text{AF}}}_{it}$ incorporates measurement error both from interpolating rasterized groundwater depths across space and from interpolating/extrapolating unit i 's APEP measurements across time.

We instrument for $\log(\widehat{\frac{\text{kWh}}{\text{AF}}}_{it})$ using logged groundwater depth averaged across unit i 's full groundwater basin.² This purges potential endogeneity driven by changes in pumping technologies, and eliminates bias induced by measurement error in unit i 's pump specifications in month t . It also breaks the mechanical relationship between $\widehat{\frac{\text{kWh}}{\text{AF}}}_{it}$ and Q_{it}^{elec} , as farm i 's extraction should have a negligible contemporaneous effect on average groundwater levels across the whole basin. Finally, instrumenting with basin-wide average depth mitigates measurement error from having spatially interpolated groundwater measurements into a (potentially overfit) gridded raster.

Table B2 presents our results for estimating farmers' groundwater demand. Each column estimates Equation (6) using our preferred strategy for identifying the elasticity with respect to the electricity price: instrumenting for $\log(P_{it}^{\text{elec}})$ with within-category default prices, and interacting unit fixed effects with indicators for each category of physical pumping capital. Note that we report $\hat{\beta}^e$ and $\hat{\beta}^w$, where the latter subtracts 1 from the regression coefficient on $\log(\widehat{\frac{\text{kWh}}{\text{AF}}}_{it})$. We interpret each coefficient as the elasticity of demand for groundwater with respect to one component of the price of groundwater, holding the other component constant.

In Column (1), we present a quasi-OLS specification: while we instrument for $\log(P_{it}^{\text{elec}})$ with the within-category default electricity price, we do not instrument for $\log(\text{kWh}/\text{AF}_{it})$. In this specification, we recover a somewhat lower elasticity of demand with respect to pumping costs (-0.99) than with respect to electricity prices (-1.21).

Column (2) reports our preferred estimates of $\hat{\beta}^e$ and $\hat{\beta}^w$, where we instrument for $\log(\widehat{\frac{\text{kWh}}{\text{AF}}}_{it})$ with logged groundwater depth in month t averaged across unit i 's groundwater basin. Comparing $\hat{\beta}^w$ in Columns (2) vs. (1), instrumenting with average depth appears to alleviate bias due to measurement error in $\log(\widehat{\frac{\text{kWh}}{\text{AF}}}_{it})$, and our estimate rises to (-1.51) .³ The exclusion restriction requires that unit i 's pumping behavior have no contemporaneous impact on basin-wide average groundwater depths. Such feedback effects between the dependent variable and the instrument would be extremely unlikely for three reasons: (i) unit i is small relative to the geographic footprint of its groundwater basin; (ii) thousands of other pumpers are also extracting from the same basin; (iii) basin-wide average groundwater levels do not instantaneously reequilibrate after extraction at one point in space. Column

2. We instrument with groundwater depth in logs (rather than levels) because logging both sides of Equation (2) implies that $\log(\text{kWh}/\text{AF}_{it})$ is linear in $\log(\text{lift})$, and a percentage change in depth should yield a similar percentage change in lift.

3. We discuss three potential sources of bias in β^w in Section 4.2: (i) endogenous changes to pumping technologies, (ii) the mechanical relationship between extraction and depth at a given location, and (iii) measurement error. Bias from (i) and (ii) appear unlikely, as they should bias our β^w away from zero, rather than towards zero.

(3) restricts the sample to the 3 largest groundwater basins, each of which has over 1,000 units in our estimation sample.⁴

The magnitudes of our $\hat{\beta}^e$ estimates are relatively similar (if slightly larger) than the results in our electricity-only regressions, especially comparing $\hat{\beta}^e = -1.21$ from Column (1) of Table B2 with the analogous estimate ($\hat{\beta} = -1.17$) from Column (3) of Table 3. This is not surprising, since Equation (6) simply adds one regressor to Equation (3). $\hat{\beta}^e$ is quite close to our instrumented $\hat{\beta}^w$ estimate (-1.27 vs. -1.51). This implies that a 1 percent change in the effective price of groundwater has close to the same effect on farmers' pumping behavior, whether that change comes via their marginal electricity price or via their pump's kWh/AF conversion factor. It also suggests that farmers are quite attentive to their true costs of pumping, and that they reoptimize their pumping behavior relatively similarly in response to either type of price variation—as Neoclassical theory would predict.

Columns (4)–(6) report three alternate versions of our preferred estimates in Column (2). First, to account for the inherent tradeoff between spatial density vs. temporal frequency of groundwater measurements, Column (4) re-estimates Equation (6) using groundwater data rasterized at the quarterly (rather than monthly) level. Whereas our preferred monthly rasters are able to capture groundwater measurements at greater temporal frequency, quarterly rasters have greater accuracy in the cross-section by incorporating more distinct measurement sites. The resulting $\hat{\beta}^w$ estimate decreases in magnitude slightly, and comes closer to the $\hat{\beta}^e$ estimate. Column (5) includes water basin by year and water district by year fixed effects, yielding only slightly attenuated point estimates despite eliminating much of the variation in the average depth instrument. In Column (6), we instrument with 6- and 12-month lags of average depth (rather than contemporaneous depth), as it is possible (albeit unlikely) that farmers pump less in months with lower groundwater levels for some reason other than pumping costs. These lagged instruments marginally increase $\hat{\beta}^e$ and substantially increase $\hat{\beta}^w$; however, the small first stage F -statistic indicates a weak instrument, and we interpret these results with caution.

4. These basins are the San Joaquin Valley, the Sacramento Valley, and the Salinas Valley. The number of agricultural groundwater pumpers in each basin is likely much larger, as our estimation sample comprises only the subset of PGE customers that we can confidently match to an APEP-subsidized pump test.

Table B2: Estimated demand elasticities decomposed – Groundwater

	(1)	(2)	(3)	(4)	(5)	(6)
	IV	IV	IV	IV	IV	IV
$\log(P_{it}^{\text{elec}}): \hat{\beta}^e$	-1.21*** (0.17)	-1.27*** (0.17)	-1.27*** (0.18)	-1.23*** (0.17)	-1.05*** (0.16)	-1.47*** (0.19)
$\log\left(\frac{\widehat{\text{kWh}}}{\text{AF}_{it}}\right): \hat{\beta}^w$	-0.99*** (0.10)	-1.51*** (0.32)	-1.47*** (0.38)	-1.25*** (0.33)	-1.14*** (0.28)	-2.09*** (0.56)
Instrument(s):						
Default $\log(P_{it}^{\text{elec}})$	Yes	Yes	Yes	Yes	Yes	Yes
$\log(\text{Avg depth in basin})$		Yes	Yes	Yes	Yes	
$\log(\text{Avg depth in basin}), \text{lagged}$						Yes
Fixed effects:						
Unit \times month-of-year	Yes	Yes	Yes	Yes	Yes	Yes
Month-of-sample	Yes	Yes	Yes	Yes	Yes	Yes
Unit \times physical capital	Yes	Yes	Yes	Yes	Yes	Yes
Water basin \times year					Yes	
Water district \times year					Yes	
Groundwater time step	Month	Month	Month	Quarter	Month	Month
Only basins with > 1000 SPs			Yes			
Service point units	10,155	10,141	9,337	10,141	10,140	10,108
Months	117	117	117	117	117	105
Observations	0.93M	0.87M	0.83M	0.91M	0.87M	0.77M
First stage F -statistic	6935	87	71	49	55	18

Notes: Each regression estimates Equation (6) at the service point by month level, where the dependent variable is the inverse hyperbolic sine transformation of electricity consumed by service point i in month t . We report estimates for $\hat{\beta}^e$ and $\hat{\beta}^w$, where the latter subtracts 1 from the estimated coefficient on $\log(\widehat{\text{kWh}}/\text{AF}_{it})$. We estimate IV specifications via two-stage least squares, and all regressions instrument for P_{it}^{elec} with unit i 's within-category default logged electricity price in month t (consistent with our preferred specification from Table 3). We instrument for $\log(\widehat{\text{kWh}}/\text{AF}_{it})$ with either logged average groundwater depth across unit i 's basin, or the 6- and 12-month lags of this variable. “Physical capital” is a categorical variable for (i) small pumps, (ii) large pumps, and (iii) internal combustion engines, and unit \times physical capital fixed effects control for shifts in tariff category triggered by the installation of new pumping equipment. Water basin \times year fixed effects control for broad geographic trends in groundwater depth. Water district \times year fixed effects control for annual variation in surface water allocations. Column (3) restricts the sample to only the three most common water basins (San Joaquin Valley, Sacramento Valley, and Salinas Valley), each of which contains over 1000 unique SPs in our estimation sample. Column (4) uses a quarterly panel of groundwater depths to construct $\log(\widehat{\text{kWh}}/\text{AF}_{it})$ and the instrument, rather than a monthly panel. All regressions drop solar NEM customers, customers with bad geocodes, months with irregular electricity bills (e.g. first/last bills, bills longer/shorter than 1 month, overlapping bills for a single account), and pumps with implausible test measurements. Standard errors (in parentheses) are two-way clustered by service point and by month-of-sample. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Appendix C Discrete choice modeling framework

To arrive at the discrete choice model we present in Section 6.2.2 of the main text, we begin with a general model of farm profits:

$$\pi_{iy}(k) = r_{iyk} - c_{iyk} + \varepsilon_{iyk}$$

where $\pi_{iy}(k)$ is the profit for farmer i in year y growing crop type k , r_{iyk} is farm revenues, c_{iyk} are farm costs, and ε_{iyk} is an error term. Rewriting revenues as a function of crop prices p_{ky} (common across farmers) and quantity grown q_{iyk} , and decomposing costs into non-water costs and water costs, we can re-write this as:

$$\begin{aligned}\pi_{iy}(k) &= p_{yk}q_{iyk} - c_{iyk}^{\text{non-water}} - c_{iyk}^{\text{water}} + \varepsilon_{iyk} \\ &= p_{yt}q_{iyk} - c_{iyk}^{\text{non-water}} - p_{iy}^{\text{water}}q_{iyk}^{\text{water}} + \varepsilon_{iyk}\end{aligned}$$

Assuming yield and per-unit costs are constant within county c , and denoting acreage as A_{iy} :

$$\begin{aligned}\pi_{iy}(k) &= \alpha_{cyk} + \gamma_{cyk}A_{iy} - p_{iy}^{\text{water}}q_{iyk}^{\text{water}} + \varepsilon_{iyk} \\ &= \alpha_{cyk} + \gamma_{cyk}A_{iy} + \beta_{cyk}A_{iy}p_{iy}^{\text{water}} + \varepsilon_{iyk}\end{aligned}$$

Assuming all costs scale with acres (i.e. $\alpha_{cky} = 0$):

$$\pi_{iy}(k) = \gamma_{cyk} + \beta_{cyk}p_{iy}^{\text{water}} + \varepsilon_{iyk}$$

And, finally, assume that the impact of water prices on profits is time-invariant and location-invariant (i.e. $\beta_{cyk} = \beta_k \forall y, c$):

$$\pi_{iy}(k) = \gamma_{cyy} + \beta_k p_{iy}^{\text{water}} + \varepsilon_{iyk}$$

Farmer i maximizes profits by choosing crop k (identical to Equation (9) in the main text):

$$\max_{k \in \mathcal{K}} \pi_{iy}(k) = \gamma_{cky} + \beta_k p_{iy}^{\text{water}} + \varepsilon_{iyk}$$

We further assume that the error terms are i.i.d. Normal: $\varepsilon_{iyk} \sim \mathcal{N}(0, \sigma)$, and estimate this model using the instrumental variables probit model described in the main text.

Appendix D Additional tables and figures

In this Appendix, we present a variety of sensitivity analyses and robustness checks which build upon the results we present in the main text.

D.1 Instrumenting with modal tariffs

In our main estimates, we instrument for $\log(P_{it}^{\text{elec}})$ and $\log(P_{it}^{\text{water}})$ using the “default” within-category electricity tariff for each of PGE’s 5 rate categories. For the AG-1A, AG-1B, and AG-ICE tariffs, this designation is trivial – each of these rates is a singleton within its category. However, for the small pumps and smart meters category and the large pumps and small meters category, there are 8 and 12 separate tariffs, respectively. We define the default tariff as the rate within each category that has the least complex marginal pricing structure: AG-4A and AG-4B. In Appendix Table D3, we instead present results where we instrument for $\log(P_{it}^{\text{elec}})$ with the *modal* tariff in each category: AG-1A, AG-1B, AG-ICE, AG-4A, and AG-5B. Our preferred specification, shown in Column (2), produces an identical elasticity (-1.17) to our preferred estimate in Table 3 in the main text.

Table D3: Instrumenting with within-category modal tariffs – Electricity

	(1)	(2)	(3)	(4)
	IV	IV	IV	IV
$\log(P_{it}^{\text{elec}})$	-1.53*** (0.15)	-1.17*** (0.16)	-1.00*** (0.15)	-1.19*** (0.20)
Instrument(s):				
Modal $\log(P_{it}^{\text{elec}})$	Yes	Yes	Yes	
Modal $\log(P_{it}^{\text{elec}})$, lagged				Yes
Fixed effects:				
Unit \times month-of-year	Yes	Yes	Yes	Yes
Month-of-sample	Yes	Yes	Yes	Yes
Unit \times physical capital		Yes	Yes	Yes
Water basin \times year			Yes	
Water district \times year			Yes	
Service point units	11,173	11,173	11,142	10,922
Months	117	117	117	105
Observations	1.05M	1.05M	1.04M	0.91M
First stage F -statistic	5796	5006	5202	1043

Notes: This table reestimates Columns (2)–(5) from Table 3, instrumenting with the average marginal price of the modal tariff within each category. These instruments produce very similar results, demonstrating that our main results are not sensitive to our choice of default tariff. See notes under Table 3 for further detail. Standard errors (in parentheses) are two-way clustered by service point and by month-of-sample. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

D.2 Sensitivity to trends in pump characteristics

A potential endogeneity concern in our setting is farmers choosing their pumping capital in order to attain a more favorable electricity tariff. We find no evidence of farmers “bunching” pump characteristics around the 35 hp cutoff in Figure 4, and we include a unit \times physical capital fixed effect in our preferred specifications to help control for this. Ultimately, our identification strategy relies on a parallel-trends type argument, which requires that electricity consumption for farmers in different tariff categories would be trending similarly in the absence of differential rate increases or decreases over time. To provide evidence in support of this assumption, in Appendix Table D4, we interact our month-of-sample fixed effects with bins of three different pump characteristics: horsepower, kW, and operating pump efficiency (OPE). We use 11 bins in horsepower: 1 below PGE’s 35 hp cutoff, and 10 bins for deciles of the horsepower distribution above 35 hp. We similarly use 11 bins in kW: 1 for measured kW for pumps below PGE’s 35 hp cutoff (equivalent to 26.1 kW), and 10 for each decile of kW for pumps above this 35 hp cutoff. Finally, we use 10 bins for OPE. In Columns (1)–(3) of Table D4, we exclude our unit \times physical capital fixed effect, and find larger elasticities

Table D4: Sensitivity to trends in HP, kW, and OPE – Electricity

	(1)	(2)	(3)	(4)	(5)	(6)
	IV	IV	IV	IV	IV	IV
$\log(P_{it}^{\text{elec}})$	-1.56*** (0.16)	-1.54*** (0.16)	-1.60*** (0.17)	-1.15*** (0.16)	-1.13*** (0.16)	-1.17*** (0.16)
Month-of-sample FEs interaction	HP bins	kW bins	OPE bins	HP bins	kW bins	OPE bins
IV: Default $\log(P_{it}^{\text{elec}})$	Yes	Yes	Yes	Yes	Yes	Yes
Fixed effects:						
Unit \times month-of-year	Yes	Yes	Yes	Yes	Yes	Yes
Month-of-sample	Yes	Yes	Yes	Yes	Yes	Yes
Unit \times physical capital				Yes	Yes	Yes
Service point units	11,173	11,173	11,173	11,173	11,173	11,173
Months	117	117	117	117	117	117
Observations	1.05M	1.05M	1.05M	1.05M	1.05M	1.05M
First stage F -statistic	3983	4056	4164	6989	7123	7406

Notes: This table conducts sensitivity analysis on our monthly electricity regressions by interacting month-of-sample fixed effects with bins of pump horsepower, kW, and operating efficiency. Columns (1)–(3) replicate Column (2) from Table 3, while Columns (4)–(6) replicate Column (3) from Table 3. Columns (1) and (4) interact month-of-sample fixed effects with 11 bins of nameplate horsepower: 1 bin below PGE’s 35 hp cutoff, and 10 bins for deciles of the distribution of hp above this cutoff. Columns (2) and (5) interact month-of-sample fixed effects with 11 bins of kW usage, as measured in APEP pump tests: 1 bin below PGE’s 35 hp cutoff (equivalent to 26.1 kW), and 10 bins for deciles of the distribution of measured kW above this cutoff. Columns (3) and (6) interact month-of-sample fixed effects with 10 bins for deciles of operating pump efficiency recorded in APEP pump tests. See notes under Table 3 for further detail. Standard errors (in parentheses) are two-way clustered by service point and by month-of-sample. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

than in our preferred model. In Columns (4)–(6), we include this fixed effect, and recover estimates that are quantitatively similar to our preferred estimate (-1.17 from Table 3), providing reassurance that our results are not being driven by differential sorting into tariff categories over time.

D.3 Sensitivity to $\widehat{\text{kWh}}/\text{AF}$ construction

Because we do not observe groundwater extraction or costs directly, we must construct Q_{it}^{water} and P_{it}^{water} by scaling our electricity data by a conversion factor:

$$\frac{\widehat{\text{kWh}}}{\text{AF}_{it}} = \frac{[\text{Lift (feet)}] \times [\text{Constant}]}{\text{Operating pump efficiency } (\%)_{it}}$$

While operating pump efficiency is a variable in our data, lift is a function of the pump’s drawdown and the static water level. We use rasterized versions of CASGEM well measure-

ments to construct lift. Appendix Table D5 presents sensitivity analyses for our groundwater elasticity estimates using a variety of approaches to construct $\widehat{\frac{\text{kWh}}{\text{AF}}}_{it}$. In Column (1), we instrument *only* with the (log of) the average groundwater depth in each unit's basin. The resulting point estimate, (-1.36) , lies between our preferred electricity-instrument-only estimate of (-1.12) presented in Table 4 and our preferred dual-instrument $\hat{\beta}^w$ estimate of -1.51 presented in Appendix Table B2. In Column (2), we use only the average depth instrument, and assign units $\widehat{\frac{\text{kWh}}{\text{AF}}}_{it}$ directly from an APEP test, rather than attempting to estimate it. In Columns (3)–(6), we use the electricity instrument only. In Column (3), we remove units that do not have reliable measures of drawdown in their APEP test data. In Column (4), we predict drawdown as a function of groundwater depth, rather than retaining a static drawdown measurement from an APEP test. In Column (5), we again predict drawdown, this time using the average basin-wide groundwater level. In Column (6), we restrict our sample to units that have a groundwater depth measurement within 8 miles prior to rasterization. Across all specifications, we find estimates that are quantitatively similar to our central estimate of -1.12 .

Table D5: Sensitivity to $\widehat{\text{kWh/AF}}$ construction – Groundwater

	(1)	(2)	(3)	(4)	(5)	(6)
	IV	IV	IV	IV	IV	IV
$\log(P_{it}^{\text{water}})$	-1.36*** (0.32)	-1.09*** (0.15)	-1.26*** (0.37)	-1.12*** (0.16)	-1.13*** (0.15)	-1.28*** (0.18)
$\widehat{\text{kWh/AF}}_{it}$ criteria:						
Measured, not estimated		Yes				
Drop tests with bad drawdown			Yes			
Time-varying predicted drawdown				Yes	Yes	
Mean groundwater depth					Yes	
Depth measured w/in 8 miles						Yes
Instrument:						
$\log(\text{Avg depth in basin})$	Yes	Yes				
Default $\log(P_{it}^{\text{elec}})$			Yes	Yes	Yes	Yes
Fixed effects:						
Unit \times month-of-year	Yes	Yes	Yes	Yes	Yes	
Month-of-sample	Yes	Yes	Yes	Yes	Yes	
Unit \times physical capital	Yes	Yes	Yes	Yes	Yes	
Groundwater time step	Month	Month	Month	Month	Month	Month
Service point units	10,141	10,155	1,562	10,155	10,155	9,930
Months	117	117	117	117	117	117
Observations	0.87M	0.93M	0.12M	0.93M	0.93M	0.45M
First stage F -statistic	161	5398	645	2824	2420	2382

Notes: Each regression replicates our preferred specification from Column (2) of Table 4, while altering our preferred method of specifying units' kWh/AF conversion factor. Columns (1)–(2) maintain our preferred $\widehat{\text{kWh/AF}}$ definition, but instrument for groundwater price using basin-wide average groundwater depths. This leverages only variation in P^{water} driven by changes in depth. Column (2) directly assigns kWh/AF as measured in an APEP pump test, which yields a P^{water} variable that is independent of changes in groundwater depth. Column (3) removes units without a reliable drawdown measurement from an APEP pump test. Columns (4)–(5) construct $\widehat{\text{kWh/AF}}$ using predicted drawdown as a function of groundwater depth, rather than fixed drawdown within pumps over time. Column (5) also applies basin-wide average depth to construct $\widehat{\text{kWh/AF}}$, rather than using localized measurements from groundwater rasters. Column (6) uses rasterized groundwater measurements, but drop the (roughly half of) observations without a contemporaneous groundwater measurement within 8 miles. See notes under Table 4 for further detail. Standard errors (in parentheses) are two-way clustered by service point and by month-of-sample. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

D.4 Sensitivity to geographic controls

In our preferred specifications in Tables 3 and 4, we include unit-by-month-of-year fixed effects and month-of-sample fixed effects. However, it is possible that confounders that vary both by location and time remain. In Appendix Tables D6 and D7, we add additional geographic controls by interacting our month-of-sample fixed effects with fixed effects for a

Table D6: Sensitivity to geographic controls – Electricity

	(1)	(2)	(3)	(4)	(5)
	IV	IV	IV	IV	IV
$\log(P_{it}^{\text{elec}})$	-1.12*** (0.15)	-1.09*** (0.15)	-1.08*** (0.15)	-1.10*** (0.15)	-1.02*** (0.15)
Month-of-sample FEs interaction	Climate zone	County	Basin	Sub-Basin	Water district
IV: Default $\log(P_{it}^{\text{elec}})$	Yes	Yes	Yes	Yes	Yes
Fixed effects:					
Unit \times month-of-year	Yes	Yes	Yes	Yes	Yes
Month-of-sample	Yes	Yes	Yes	Yes	Yes
Unit \times physical capital	Yes	Yes	Yes	Yes	Yes
Service point units	11,167	11,170	11,159	11,151	11,156
Months	117	117	117	117	117
Observations	1.04M	1.05M	1.04M	1.04M	1.04M
First stage F -statistic	7548	7527	7375	7579	7648

Notes: This table conducts sensitivity analysis on our preferred electricity specification from Column (3) of Table 3, by interacting month-of-sample fixed effects with different geographic variables. California comprises 16 climate zones, and PGE agriculture customers are distributed across 11 distinct climate zones. Sub-basins are administrative sub-divisions of groundwater basins; this estimation sample includes agricultural consumers from 46 unique groundwater basins and 95 unique sub-basins. The sample also includes units assigned to 125 unique water districts; Column (5) includes a separate set of month-of-sample fixed effects for units not assigned to a water district. See notes under Table 3 for further detail. Standard errors (in parentheses) are two-way clustered by service point and by month-of-sample. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

variety of different geographic scales that may be relevant for agricultural production: climate zone, county, groundwater basin, groundwater sub-basin, and water district to allay these concerns. In Appendix Table D6, we present results for electricity. When we interact our month-of-sample fixed effects with these geographic fixed effects, we find very similar estimates to those in Table 3. Including month-of-sample-by-water-district fixed effects attenuates the estimates the most, to -1.02 , though we cannot reject that this is the same as the -1.12 in our preferred specification. Appendix Table D7, presents results for groundwater. When we interact our month-of-sample fixed effects with these geographic fixed effects, we again find very similar estimates to those in Table 4. Once again, including month-of-sample-by-water-district fixed effects attenuates the estimates the most, to -0.95 , though this elasticity is still large and highly statistically significant.

Table D7: Sensitivity to geographic controls – Water

	(1)	(2)	(3)	(4)	(5)
	IV	IV	IV	IV	IV
$\log(P_{it}^{\text{water}})$	-1.08*** (0.15)	-1.05*** (0.15)	-1.05*** (0.15)	-1.05*** (0.15)	-0.95*** (0.14)
Month-of-sample FEs interaction	Climate zone	County	Basin	Sub-Basin	Water district
IV: Default $\log(P_{it}^{\text{elec}})$	Yes	Yes	Yes	Yes	Yes
Fixed effects:					
Unit \times month-of-year	Yes	Yes	Yes	Yes	Yes
Month-of-sample	Yes	Yes	Yes	Yes	Yes
Unit \times physical capital	Yes	Yes	Yes	Yes	Yes
Service point units	10,150	10,152	10,142	10,135	10,136
Months	117	117	117	117	117
Observations	0.93M	0.93M	0.93M	0.93M	0.93M
First stage F -statistic	3827	4379	3573	4511	4239

Notes: This table conducts sensitivity analysis on our preferred water specification from Column (2) of Table 4, by interacting month-of-sample fixed effects with different geographic variables. California comprises 16 climate zones, and PGE agriculture customers are distributed across 11 distinct climate zones. Sub-basins are administrative sub-divisions of groundwater basins; this estimation sample includes agricultural consumers from 46 unique groundwater basins and 95 unique sub-basins. The sample also includes units assigned to 125 unique water districts; Column (5) includes a separate set of month-of-sample fixed effects for units not assigned to a water district. See notes under Table 3 for further detail. Standard errors (in parentheses) are two-way clustered by service point and by month-of-sample. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

D.5 Sensitivity to weather controls

Weather is a key input in the agricultural production process. While we do not include weather controls in our main estimates, because we expect weather to be orthogonal to our within-category electricity tariff instrument, conditional on unit-by-month-of-year and month-of-sample fixed effects. Nevertheless, we present sensitivities to the inclusion of weather controls here. We obtained gridded daily temperature and precipitation data from PRISM, and geographically matched this weather data to our units using CLU centroids. We average daily maximum and minimum temperatures and sum daily precipitation over all days in each month to construct monthly weather controls. Appendix Table D8 presents the results of adding weather controls to our main estimates. In Columns (1) and (4), we add a monthly precipitation control to our electricity and water regressions, respectively. In Columns (2) and (5), we also add average daily minimum and maximum temperature. In Columns (3) and (6), we also add 1-month-lagged temperature and precipitation. Across all specifications, our electricity and groundwater elasticity estimates remain quantitatively similar to our preferred estimates in Table 3 and 4: -1.17 and -1.12 , respectively.

Table D8: Sensitivity to weather controls

	Electricity			Groundwater		
	(1) IV	(2) IV	(3) IV	(4) IV	(5) IV	(6) IV
$\log(P_{it}^{\text{elec}})$	-1.18*** (0.16)	-1.17*** (0.16)	-1.19*** (0.16)	-1.12*** (0.15)	-1.12*** (0.15)	-1.14*** (0.15)
Weather controls:						
Precipitation	Yes	Yes	Yes	Yes	Yes	Yes
Temperature		Yes	Yes		Yes	Yes
Lagged precipitation			Yes			Yes
Lagged temperature			Yes			Yes
IV: Default $\log(P_{it}^{\text{elec}})$	Yes	Yes	Yes	Yes	Yes	Yes
Fixed effects:						
Unit \times month-of-year	Yes	Yes	Yes	Yes	Yes	Yes
Month-of-sample	Yes	Yes	Yes	Yes	Yes	Yes
Unit \times physical capital	Yes	Yes	Yes	Yes	Yes	Yes
Service point units	11,173	11,173	11,172	10,155	10,155	10,154
Months	117	117	116	117	117	116
Observations	1.05M	1.05M	1.04M	0.93M	0.93M	0.93M
First stage F -statistic	7387	7390	7348	3033	3054	2977

Notes: This table adds weather controls to our preferred specifications for electricity (Column (3) of Table 3) and groundwater (Column (2) of Table 4). We assign daily precipitation, maximum temperature, and minimum temperatures to each unit's latitude and longitude, using daily rasters from PRISM. We sum daily precipitation over all days in each month, and average daily maximum and minimum temperatures over all days in each month. Columns (3) and (6) control for 1-month lags in all three variables. See notes under Tables 3 and Table 4 for further details. Standard errors (in parentheses) are two-way clustered by service point and by month-of-sample. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

D.6 Sensitivity to field assignments

Our PGE data are geographically resolved to the service point (SP) level. In order to estimate impacts of electricity and groundwater costs on land use, we must match these SPs to geographic features with agricultural meaning. We use the USDA's Common Land Unit (CLU) as our main agricultural unit of analysis. A CLU is defined as "is the smallest unit of land that has a permanent, contiguous boundary, a common land cover and land management, a common owner and a common producer in agricultural land associated with USDA farm programs. CLU boundaries are delineated from relatively permanent features such as fence lines, roads, and/or waterways."⁵ We use a 2008 CLU shapefile: in the 2008

5. <https://www.fsa.usda.gov/programs-and-services/aerial-photography/imagery-products/common-land-unit-clu/index>

Table D9: Sensitivity to CLU assignments and groupings – Groundwater

	(1)	(2)	(3)	(4)	(5)
	IV	IV	IV	IV	IV
$\log(P_{it}^{\text{water}})$	-1.11*** (0.19)	-1.09*** (0.16)	-1.35*** (0.26)	-0.96*** (0.19)	-1.09*** (0.17)
Sample criteria:					
Inside CLU polygon	Yes				
Drop CLU inconsistencies		Yes			
Pumps per CLU group			1	2+	
Cluster by CLU group					Yes
IV: Default $\log(P_{it}^{\text{elec}})$	Yes	Yes	Yes	Yes	Yes
Fixed effects:					
Unit \times month-of-year	Yes	Yes	Yes	Yes	Yes
Month-of-sample	Yes	Yes	Yes	Yes	Yes
Unit \times physical capital	Yes	Yes	Yes	Yes	Yes
Groundwater time step	Month	Month	Month	Month	Month
Service point units	5,913	9,702	2,677	7,072	4,708
Months	117	117	117	117	117
Observations	0.55M	0.89M	0.25M	0.65M	0.90M
First stage F -statistic	2059	2896	1334	2009	2548

Notes: Each regression replicates our preferred specification in Column (2) from Table 4, while conducting sensitivity on unit-specific assignments to CLU polygons (i.e. fields). Column (1) includes only units with coordinates that are fully inside their assigned CLU polygon. Column (2) drops units with inconsistent, problematic, or internally conflicting CLU assignments. Columns (3)–(5) group CLUs that lie within the same tax parcels. Column (3) includes only units that are the singleton (confirmed) groundwater pump in their CLU group. Column (4) includes units in CLU groups with multiple (confirmed) groundwater pumps. Column (5) two-way clusters by CLU group and by month-of-sample, with 4708 unique CLU groups. See notes under Table 4 for further detail. Standard errors (in parentheses) are two-way clustered by service point and by month-of-sample, except in Column (5). Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Farm Bill, CLUs were deemed to be confidential, and future shapefiles were not made publicly available. In part because PGE’s SPs often lie at the edge of property boundaries (i.e., on roads, and therefore easily accessible), there is the potential for measurement error in the SP–CLU match. In Appendix Table D9, we present sensitivities on this assignment process. In Column (1), we keep only units that lie within a CLU polygon. In Column (2), we drop units with inconsistent CLU assignments. In Column (3), we restrict the sample to SPs whose CLUs do not contain any other APEP pumps. In Column (4), we restrict the sample to SPs whose CLUs contain multiple pumps. In Column (5), we two-way cluster our standard errors by CLU and month-of-sample. We find similar elasticities across all specifications.

D.7 Sensitivity to functional form

In our main specifications in Tables 3 and 4, we use the inverse hyperbolic sine (IHS) transformation for our dependent variables, $\sinh^{-1}(Q_{it}^{\text{elec}})$ and $\sinh^{-1}(Q_{it}^{\text{water}})$, since 14 percent of our monthly observations are zeroes. The IHS behaves much like the log, but admits zeroes (Bellemare and Wichman (2020)). Here, we present sensitivities using different dependent variables: the IHS transformation, excluding zero-valued observations, $\log(Q)$, $\log(1 + Q)$, and $\log(1 + 100Q)$. Appendix Table D10 presents electricity results, and Appendix Table D11 presents groundwater results. Both tables are laid out identically. In Column (1), we replicate our preferred specification from the main text, where we use the IHS transformation with no sample restrictions. In Column (2), we retain the IHS transformation, but only use strictly positive-valued observations, to match the sample used by the standard log. We find that our estimates attenuate strongly, moving from an elasticity of -1.17 to -0.31 (electricity) and -1.12 to -0.30 (water). Reassuringly, we find identical point estimates when we use $\log(Q)$ as the dependent variable in Column (3). In Column (4), we use $\log(1 + Q)$, in order to use the log while retaining the zero-valued observations, and our point estimates rise (a much larger -0.78 for electricity; a smaller -0.36 for water). Finally, Column (5) uses $\log(1 + 100Q)$, and we again recover less attenuated estimates: -1.12 for electricity and -0.71 for water. This suggests that our choice of functional form is not driving our results. However, zeroes are important in this context, because they are likely associated with farmers fallowing their land or switching crops. Finding that our elasticity estimates are driven in part by these mechanisms is consistent with our discussion in Section 6 in the main text, and echoed by our intensive-vs.-extensive margin results in Table 6.

Table D10: Sensitivity to IHS vs. log transformation – Electricity

	(1)	(2)	(3)	(4)	(5)
	IV	IV	IV	IV	IV
$\log(P_{it}^{\text{elec}})$	-1.17*** (0.16)	-0.31*** (0.08)	-0.31*** (0.08)	-0.78*** (0.11)	-1.12*** (0.15)
LHS transformation:	$\sinh^{-1}(Q)$	$\sinh^{-1}(Q)$	$\log(Q)$	$\log(1 + Q)$	$\log(1 + 100Q)$
Sample restriction: $Q_{it} > 0$		Yes	Yes		
IV: Default $\log(P_{it}^{\text{elec}})$	Yes	Yes	Yes	Yes	Yes
Fixed effects:					
Unit \times month-of-year	Yes	Yes	Yes	Yes	Yes
Month-of-sample	Yes	Yes	Yes	Yes	Yes
Unit \times physical capital	Yes	Yes	Yes	Yes	Yes
Service point units	11,173	11,109	11,109	11,173	11,173
Months	117	117	117	117	117
Observations	1.05M	0.89M	0.89M	1.05M	1.05M
First stage F -statistic	7382	6841	6841	7382	7382

Notes: This table conducts sensitivity analysis on the transformation of the dependent variable Q^{elec} . Column (1) reproduces our preferred specification from Column (3) of Table 3 using the inverse hyperbolic sine transformation, where the dependent variable is the inverse hyperbolic sine transformation of electricity consumed by service point i in month t . Column (2) uses the same transformation but removed zeros to align with the natural log transformation in Column (3). Columns (4)–(5) apply the natural log + 1 transformation. Column (5) also scales the dependent variable by 100, which nearly matches our results using the inverse hyperbolic sine transformation. See notes under Table 3 for further detail. Standard errors (in parentheses) are two-way clustered by service point and by month-of-sample. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table D11: Sensitivity to IHS vs. log transformation – Water

	(1)	(2)	(3)	(4)	(5)
	IV	IV	IV	IV	IV
$\log(P_{it}^{\text{water}})$	-1.12*** (0.15)	-0.30*** (0.08)	-0.30*** (0.08)	-0.36*** (0.05)	-0.71*** (0.10)
LHS transformation:	$\sinh^{-1}(Q)$	$\sinh^{-1}(Q)$	$\log(Q)$	$\log(1 + Q)$	$\log(1 + 100Q)$
Sample restriction: $Q_{it} > 0$		Yes	Yes		
IV: Default $\log(P_{it}^{\text{elec}})$	Yes	Yes	Yes	Yes	Yes
Fixed effects:					
Unit \times month-of-year	Yes	Yes	Yes	Yes	Yes
Month-of-sample	Yes	Yes	Yes	Yes	Yes
Unit \times physical capital	Yes	Yes	Yes	Yes	Yes
Service point units	10,155	10,091	10,091	10,155	10,155
Months	117	117	117	117	117
Observations	0.93M	0.79M	0.79M	0.93M	0.93M
First stage F -statistic	3021	2629	2629	3021	3021

Notes: This table conducts sensitivity analysis on the transformation of the dependent variable Q^{water} . Column (1) reproduces our preferred specification from Column (2) of Table 4 using the inverse hyperbolic sine transformation, where the dependent variable is the inverse hyperbolic sine transformation of electricity consumed by service point i in month t . Column (2) uses the same transformation but removed zeros to align with the natural log transformation in Column (3). Columns (4)–(5) apply the natural log + 1 transformation. Column (5) also scales the dependent variable by 100, which nearly matches our results using the inverse hyperbolic sine transformation. See notes under Table 4 for further detail. Standard errors (in parentheses) are two-way clustered by service point and by month-of-sample. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.