UC Berkeley
CUDARE Working Papers

Title
Anticipatory Effects of Regulation in Open Access

Permalink
https://escholarship.org/uc/item/58n467v5

Authors
Bruno, Ellen Marie
Hagerty, Nick

Publication Date
2023-08-18

Data Availability
The data associated with this publication are available upon request.

Anticipatory Effects of Regulation in Open Access[*]

Ellen M. Bruno[†] and Nick Hagerty[‡]

August 18, 2023

Abstract

We study the regulation of open-access resources under long implementation horizons. Our theoretical model clarifies when and how future regulation creates either an anticipatory decline or perverse incentives to accelerate extraction (a "Green Paradox"). Then, we evaluate the early effects of a major groundwater regulation in California that does not yet bind. We assemble new data and compare within pairs of neighboring agencies that face varying restrictions on extraction. Differences in future regulation do not affect measures of water-intensive investments or groundwater extraction today. The absence of anticipatory response in either direction can be explained by high private discount rates.

[†]Department of Agricultural and Resource Economics, University of California, Berkeley. Email: ebruno@berkeley.edu

[‡]Department of Agricultural Economics and Economics, Montana State University. Email: nicholas.hagerty@montana.edu

1   Introduction

Policies in natural resource management often face a long implementation horizon. While a gas tax increase may occur with only weeks or months of notice to drivers, carbon targets may be set decades in advance. A long runway can allow for a smooth transition, or alternatively introduce perverse incentives as people race to extract or consume the resource before the regulation binds. Knowing when anticipatory responses are likely to help or hinder the implementation of a policy goal is crucial for understanding both optimal policy design and how policies interact with their political economy context.

In the case of non-renewable resources like fossil fuels, these perverse incentives are known as the "Green Paradox": extraction restrictions in the future reduce scarcity rents today, leading extractors to substitute toward the present and undermining the original policy goals (Sinn, 2008). However, the conditions under which a Green Paradox might occur for renewable or common-pool resources like fisheries and groundwater remain unclear.

This paper studies the anticipatory effects of regulation in the context of groundwater resources. First, we develop a theoretical model that formalizes the conditions under which future regulation gives rise to anticipatory effects in either direction. We show that a Green Paradox can occur for groundwater, but not in open access – when an aquifer is shared among many extractors, each extractor already lacks incentive to save for the future, leaving no opportunity to profitably increase extraction in response to impending regulation. However, when farmers are allowed to make water-intensive capital investments (such as planting orchards or drilling new wells), it is possible for future regulation to decrease extraction now, smoothing the regulatory transition. The net effect of future regulation on extraction in the presence of investment opportunities then becomes an empirical question.

Using this theoretical lens, we empirically evaluate the ongoing effects of California's Sustainable Groundwater Management Act of 2014 (SGMA), arguably the largest-ever regime shift in groundwater management policy in the United States. SGMA provides a useful empirical setting because its decentralized structure gives rise to rich policy variation across the state. Hundreds of new groundwater management agencies are charged with halting groundwater depletion within their jurisdictions by the year 2040. (Previously, most groundwater use in California was not governed by binding regulations.) Areas with greater overdraft at baseline must impose greater future reductions in groundwater extraction to achieve sustainability.[1]

_____

[1]Overdraft refers to the difference between groundwater extraction and recharge through percolation

We use this variation to test how SGMA has affected groundwater extraction and water-intensive agricultural investments to date. Our research design compares changes over time for cropland in neighboring groundwater basins. Within each pair of neighbors, one basin is subject to greater future pumping reductions than the other, yet other factors such as crop suitability or prior groundwater development are similar. We consider only cropland within a relatively close distance of the boundaries and pool all cases of neighboring basins, forming a stacked-pair differences-in-differences design.

We consider two types of capital investment: new plantings of perennial crops (such as orchards or vineyards) and construction of new groundwater wells for irrigation. These investments are both observable (through satellite data products and regulatory reports) and the most relevant to groundwater policy in the context of California, where essentially all cropland is irrigated, farmers produce a diverse mix of annual and perennial crops, and groundwater constitutes a significant portion of the water supply.[2] As for extraction, groundwater pumping is generally unmonitored throughout California and therefore unobserved. As a close proxy, we form an index of water use by combining remote sensing land-use data with scientific estimates of water use by crop.[3]

Measuring future extraction restrictions is not straightforward, due to high scientific uncertainty and lack of agreement over the volume of reductions that will be necessary in the future to halt further depletion in each basin. For anticipatory responses, what matters is extractors' own beliefs, but these are not directly observable. Instead, we assemble measures of overdraft volume and planned future reductions as stated in Groundwater Sustainability Plans (GSPs) submitted by each local groundwater agency to the state. These plans were the product of lengthy public participation and negotiation processes with local stakeholders, so they are likely the best information extractors have about their own future restrictions. Still, it is possible that numbers in GSPs are strategically underestimated and that extractors are aware. We therefore obtain a third estimate by running one of the main hydrological models commonly used for water resource planning in Califor-

---

and lateral flow. Overdraft mechanically results in a decline in groundwater levels, referred to as depletion.)

[2]California's top three crops by revenue and acreage – almonds, grapes, and pistachios – are all permanent crops that feature large upfront investments (high initial capital costs plus several unproductive early years) and long productive lives of 20 to 40 years. California's Central Valley has undergone a major expansion of perennial fruit and nut tree crops over the past couple of decades, with implications for water demand (Mall and Herman, 2019). In fact, since SGMA passed in 2014, acreage in perennial crops has increased by nearly 50%. Similarly, new well construction has shown no evidence of slowing after the passage of SGMA. Agricultural capital investments are likely to be influenced by information on future water supply (Lobell and Field, 2011; Arellano-Gonzalez and Moore, 2020), and more significant changes are expected in areas facing greater restrictions under SGMA.

[3]We control for the other principal source of irrigation water, surface water deliveries, and find it not to affect the results.

nia. This statewide model avoids the risk of manipulation, but the GSPs may incorporate more detailed knowledge of local hydrological systems, plus they represent the officially stated intentions of the relevant regulatory agencies. Because no single measure is clearly superior to the others, we average across all three measures to extract a common signal, and explore robustness to using each measure alone.

Our results show that neither investments (new perennial crops and new well construction) nor groundwater extraction (as proxied by our index of water use) have changed as result of SGMA. All three outcomes followed very similar patterns across neighboring basins that face greater and lesser future pumping restrictions, both before and after SGMA passed and began to be implemented. Confidence intervals are tight, and results are robust to alternative sample definitions, treatment variables, and specifications.

To interpret the empirical results, one tempting explanation might be that the transition-smoothing and Green Paradox effects operate in opposite directions and cancel each other out. But our theoretical model shows that we can rule out this scenario, since *both* investment and extraction have not changed post-SGMA. Instead, the null effects imply that either (1) high private discount rates shrink all anticipatory motives, or (2) groundwater users' beliefs of future regulatory stringency are much lower than implied by law and the best available science.

Our paper makes three main contributions. First, our theoretical model extends a branch of literature that evaluates the anticipatory effects of regulation to include open-access conditions. Stemming from the seminal paper by Hotelling (1931), a rich theoretical literature exists explaining how preemptive resource extraction is altered by policies and other factors over time in the presence of well-defined property rights (Sinn, 1982; Cairns, 2014). This literature considers the endogeneity of total extraction (Heal, 1976), the role of imperfect substitutes (Di Maria et al., 2012) and backstop technologies, and spatial leakage. Second, we expand the model to allow users to invest in water-intensive production technology to more broadly characterize the means with which actors can increase extraction. This allows the model to depict a range of possible outcomes and enables us to characterize scenarios as a function of setting-specific parameters.

Finally, we add to a scant empirical literature that seeks to test the Green Paradox in real-world settings with the first application to groundwater (McDermott et al., 2019; Van der Ploeg and Withagen, 2020). Our study sheds light on how farmers are responding in anticipation of the policy, and contributes new empirical evidence on the preemptive effects to environmental policies. Empirical studies of the Green Paradox have focused on climate and fossil fuel policy (Di Maria et al., 2014; Lemoine, 2017; Jensen et al., 2020), land development in response to the Endangered Species Act (List et al., 2006),

and fisheries (McDermott et al., 2019), with mixed results. In the groundwater context, we find no evidence that perverse preemptive behavior is undermining the policy goal, yet we also do not find evidence that farmers are making early adjustments to meet the regulatory targets.

Many of the world's most productive agricultural regions are experiencing significant declines in groundwater levels and storage (Wada et al., 2010). In addition to the externality issues that arise from the open-access nature of the resource, groundwater also plays a key role in adaptation to climate change because it serves as a buffer to surface water scarcity and variability, reducing drought impacts and weather risk (Tsur and Graham-Tomasi, 1991; Hornbeck and Keskin, 2014). Despite the urgency of groundwater issues, regulation remains rare.[4] California's SGMA has been hailed as a landmark change – a potential model for groundwater management worldwide – and is arguably the biggest statewide regulatory shift in U.S. groundwater history. But its long implementation horizon calls into question if and when intended agricultural adjustments will actually occur.
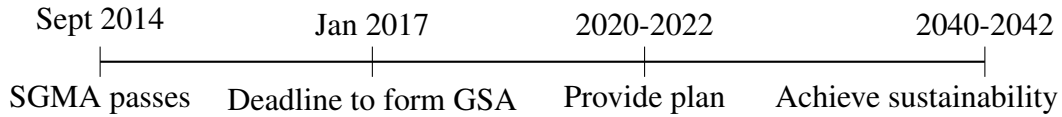
2 Background

Groundwater reserves in California's Central Valley have been declining over the last several decades, raising fears about the long-term availability of the resource. Groundwater serves as a critical buffer during periods of surface water scarcity, with average use increasing from 40 to 80% of the water supply during drought years. The passage of California's Sustainable Groundwater Management Act (SGMA) in 2014 provides an excellent opportunity to study the anticipatory effects of a natural resource policy facing a long implementation horizon.

California's SGMA provides a statewide framework for local agencies to manage groundwater and bring their basins into balance. It requires Groundwater Sustainability Agencies (GSAs) in overdrafted basins throughout California to first form, and then reach and maintain long-term stable groundwater levels. Local agencies are given the authority and flexibility to manage the resource however they see fit, as long as their approach is documented in a "Groundwater Sustainability Plan" (GSP) outlined and approved by the state. The timeline to do so is determined by a state-designated level of priority. All GSPs for high- and medium-priority basins were required to be adopted by January 31, 2022. The subset of GSAs managing groundwater in high- and medium-priority basins subject

---

[4]Examples of groundwater management do exist but are often at local levels and limited to small areas, such as command-and-control policies in parts of Kansas (Drysdale and Hendricks, 2018), price controls in parts of Colorado (Smith et al., 2017) and California (Bruno and Jessoe, 2021), or well drilling moratoria.

4

to critical conditions of overdraft had to adopt a GSP two years earlier, by January 31, 2020. Once they adopt, their plans to reach sustainability by 2040 or 2042 formally go into effect.

Sept 2014	Jan 2017	2020-2022	2040-2042
SGMA passes	Deadline to form GSA	Provide plan	Achieve sustainability

A basin is a geographic region that contains substantial groundwater resources and delineates the area within which groundwater can easily flow laterally. California's Department of Water Resources (DWR) further divides basins into subbasins. Basins are defined according to hydrogeological boundaries, but subbasins are defined for administrative convenience and have no hydrogeological meaning. SGMA required the formation of at least one GSA within each groundwater *subbasin* in basins subject to SGMA.[5] Subbasins could be governed by more than one GSA as long as all area is covered and certain accounting (e.g., water budget) and monitoring was coordinated across the subbasin.[6] Our empirical analysis therefore focuses on comparisons between neighboring GSAs in adjacent subbasins, rather than between GSAs within the same subbasin (because plans might be coordinated in unobserved ways, biasing the results), or between subbasins in different basins (because subbasins within a basin are more similar at baseline).

SGMA created substantial variation in regulatory stringency, since basins with more overdraft must adopt greater pumping restrictions in order to achieve sustainability. There were 111 GSAs determined to be of high and medium priority under SGMA, together covering the majority of agricultural land and accounting for over 95% of the groundwater pumping in the state. Figure 1 shows a map of all groundwater basins in California and distinguishes which are designated as critically overdrafted and subject to a slightly shorter implementation horizon.

GSAs were required by law to conduct stakeholder engagement and outreach via public meetings and public notices with periods of open comment, likely reducing information barriers and increasing the salience of SGMA to landowners. SGMA also created a role for the California Water Resources Control Board to take over management of a given subbasin if local authorities fail to take adequate measures toward achieving sustainability. Bruno et al. (2023b) argue that this role for the state as a backstop reduces the

---

[5]Basin and subbasin boundaries were defined prior to SGMA. DWR's Bulletin 118 describes California's 515 groundwater basins and subbasins.

[6]Each GSA is contained entirely within a single subbasin, but each subbasin may contain multiple GSAs. In many cases, when there exists multiple GSAs that have formed within one subbasin, they have coordinated management under one plan.

likelihood that GSPs lack teeth or enforcement.

Understanding how sustainability is defined and implemented under the law is important for interpreting what it means for farmers' beliefs about their future water availability. Sustainability under SGMA is formally defined by the use and management of groundwater in a manner that can be maintained without causing "undesirable results" in regards to six key indicators. The six indicators include (1) chronic lowering of groundwater levels (depletion of supply), (2) reduction of groundwater storage, (3) seawater intrusion, (4) degraded water quality, (5) land subsidence, and (6) depletion of interconnected surface water. Avoidance of these six features to a "significant and unreasonable" degree constitutes a sustainable outcome. Plans are reviewed by the state for comprehensiveness and sufficiency. Inadequate plans are returned for revisions. Failure to comply results in the state coming in as the backstop and taking over control.

Despite the legal complexity, all six "undesirable results" are closely related both physically and in regulatory plans. Achieving sustainability under SGMA is typically discussed in terms of correcting overdraft, which is relevant for all basins and correlated with each of the sustainability indicators. It is a well-understood metric that can be modeled hydrologically. We take the task of GSAs to be to limit extraction in order to end overdraft.

# 3 A Model of Groundwater Extraction in Anticipation of Regulation

We first set up a general model of groundwater extraction in the absence of regulation. We then introduce regulation in a future period to demonstrate when future regulation induces the Green Paradox. In a third step, we introduce opportunities for investment.

3.1 No regulation

We assume N identical users share an aquifer. Each user i chooses a quantity of groundwater, y_{it}, to extract in each period t to maximize the present value of profits (or net benefits) into the indefinite future. Users each obtain benefits from groundwater, $B(y_{it})$, that are increasing and concave in quantity. They also incur per-unit extraction (pumping) costs, $c(x_{it})$, that are decreasing in the user-specific resource stock, x_{it}. Benefits and costs are discounted at an interest rate $r > 0$.

The resource stock in each period is equal to the resource stock in the previous period

Figure 1: Critical Overdraft Designation of California Groundwater Basins

Critically Overdrafted
Not Critically Overdrafted
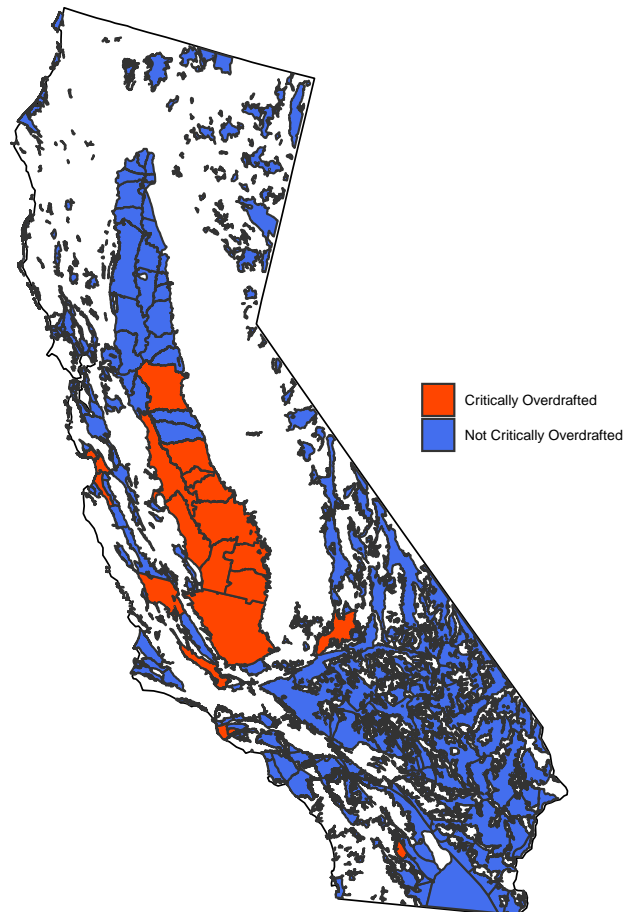
Note: The figure highlights which groundwater basins were designated as critically overdrafted. Our study focuses on groundwater agencies in the Central Valley, which is where the majority of basins subject to SGMA are concentrated.

minus the mean of extraction quantities across all N users, plus natural recharge g. This is a "bathtub" model of groundwater: the water level equalizes across the aquifer between each period, such that each user's extraction affects the resource stock for all users in equal proportion. Together, each user's private extraction problem is:

$$\max_{\{y_{it}, x_{it}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} (1+r)^{-t} \left[B(y_{it}) - c(x_{it})y_{it} \right] \tag{1}$$

$$\text{s.t.} \quad x_{i,t+1} = x_{it} + g - \tfrac{1}{N}\sum_{j=0}^{N} y_{jt} \quad \forall t \geq 0.$$

First-order conditions reveal an expression for resource rents and an Euler equation (for proofs see Appendix Section A.1):

$$\underbrace{B'(y_{it})}_{\text{marginal benefits}} = \underbrace{c(x_{it})}_{\text{marginal cost of pumping}} + \underbrace{\frac{1}{N}(1+r)^{t}\mu_{it}}_{\text{share of scarcity value}} . \tag{2}$$

$$\underbrace{B'(y_{it}) - c(x_{it})}_{\text{marginal net benefits now}} = \underbrace{(1+r)^{-1}\left[B'(y_{i,t+1}) - c(x_{i,t+1})\right]}_{\text{marginal net benefits next period}} + \underbrace{(1+r)^{-1}\frac{1}{N}[-c'(x_{i,t+1})]y_{i,t+1}}_{\text{marginal effect on own pumping costs next period}}. \tag{3}$$

Efficient extraction under complete property rights. Consider the case in which each user's extraction affects only their own stock, $N = 1$. This could represent an isolated aquifer with very low hydraulic conductivity. Here, the user fully internalizes the effect of depletion on their own future extraction costs. They restrain themselves in each period – instead of extracting until current marginal benefits equal the price of extraction, they stop sooner and leave more for future periods. Equation 2 says that marginal benefits equal the per-unit extraction cost plus the full scarcity value μ_{it} (a.k.a. resource rent or marginal user cost). In this case, the problem is equivalent to the social planner's problem for an aquifer with any value of N.

Overextraction in open access. Next, consider the limiting case as $N \to \infty$, representing a large aquifer with many users and high hydraulic conductivity. As N grows, each user's extraction affects their own stock by less and less. In the limit, each user's share of the scarcity value becomes 0, so Equation 2 simplifies to $B'(y_{it}) = c(x_{it})$: Marginal benefits equal marginal costs in each period.

This equation implicitly defines the extraction quantity y_{it}, since there is only one solution for any value of the resource stock x_{it}. Users extract every unit for which the benefits exceed the extraction costs. They do not consider how their extraction today affects future costs, since their own extraction affects the resource stock by a vanishingly small amount. The result is overextraction as compared with social planner's solution. In Equation 3, the third term also becomes 0, so the equation says that the present value of marginal net benefits is equalized across periods.

3.2   Future regulation induces Green Paradox, except in open access

Next, we consider how extraction responds to future regulation of groundwater extraction. We model the regulation as taking the form of quantity limits on extraction. We consider two periods of interest: Period 0 is unregulated, and period 1 is regulated. In period 0, users are aware of the future regulation but choose extraction quantities freely. In period 1, we assume that regulation is a binding constraint: $y_{i1} = \bar{y}$, $\forall i$.[7]

To close the model, we assume that extraction enters a steady state in period 2, such that $x_{it} = x_{i2}$ and therefore $y_{it} = g$ for all $t \geq 2$. This step provides a continuation value of the resource past our two periods of interest; without it, users would mine everything in period 0. Specifying this continuation value as a steady state, rather than some other behavior, is the key that transforms the model into a finite-horizon problem and allows us to obtain analytical solutions. Imposing it in period 2 is an approximation to the asymptotic approach that would occur in an infinite-horizon model: Assuming quantity limits are higher (less stringent) than natural recharge, resource stocks would fall until eventually the regulation no longer binds and extraction declines toward the steady-state value.[8]

Each period can be viewed as lasting many years. In our setting, period 0 represents the time between the passage of SGMA and its implementation, period 1 represents the time following SGMA implementation during which groundwater levels would fall more quickly absent SGMA, and period 2 represents the distant future in which groundwater levels finally stabilize regardless of regulation. Including more periods in the model would

---

[7]Modeling regulation as a tax (a per-unit pumping fee) would exhibit similar dynamics, but we are unable to obtain easily interpretable analytical expressions for that scenario. The reason is that a tax leaves period-1 extraction as an additional free parameter, which increases algebraic complexity. More groundwater basins are planning to comply with SGMA using quantity restrictions than pumping fees (Bruno, Hagerty, and Wardle 2022).

[8]If quantity limits are lower (more stringent) than natural recharge, then resource stocks would rise until they reach a maximum value and a new steady state begins, but this also does not change the qualitative results. Requiring the steady state to begin in period 2 (as opposed to later) is important for obtaining closed-form expressions but not for our qualitative results. Simulations that allow a smoother approach over more periods obtain the same directional results.

allow us to obtain more nuanced approach paths, first to the regulation and later to the steady state. The qualitative results would not change, but we would lose the closed-form analytical insights, because it would need to be simulated.

Finally, we parameterize the marginal cost function as $c(x) = \gamma - psx$, where p is the price of energy and s is the reciprocal of aquifer storativity ($p, s > 0$). This parameterization is based on laws of physics; it is a reasonable approximation for many aquifers and most accurate for those with high hydraulic conductivity (the "bathtub" model). It can also be viewed as a second-order approximation to the cost function.

These assumptions pin down all arguments to the maximand in equation 1 except for three: $\{y_{i0}, x_{i1}, x_{i2}\}$. How does regulatory stringency affect extraction in period 0, before the regulation takes effect?

Proposition 1 (Green Paradox for groundwater extraction). *Extraction decreases with future extraction limits (i.e., increases with future regulatory stringency):*

$$\frac{dy_{i0}}{d\bar{y}} = \frac{ps}{(1+r)NB''(y_{i0})} < 0. \tag{4}$$

Proof. See Appendix Section A.2. □

Extraction before the regulation is implemented decreases with future extraction limits (i.e., increases with future regulatory stringency). Announcing future regulation lowers the benefits that users will be able to obtain from the resource in the future, so it becomes relatively valuable to extract more of the resource before the regulation is implemented. The regulation makes a bigger difference (i) the more expensive is energy p, (ii) the smaller the storativity of the aquifer s^{-1}, (iii) the lower the per-period interest rate r (i.e., the shorter the length of time before the regulation is implemented), (iv) the fewer people that share the aquifer N, and (v) the steeper the slope of marginal benefits.

Corollary 1 (No Green Paradox in open access). *When N is large, extraction is unaffected by future extraction limits:* $\lim_{N\to\infty} dy_{i0}/d\bar{y} = 0$.

Future regulation must affect resource rents in order to change extraction decisions, and in open access there are no rents. A Green Paradox can occur for mineral resources because when users enjoy property rights, they are already taking potential future benefits into account and restraining their extraction relative to a static analysis. In open access, users are already extracting every unit of groundwater for which marginal benefits are less than marginal costs of extraction, so there is nowhere to go.

3.3 Investment opportunities allow an early decline in extraction

Most of the realistic ways that farmers might increase their groundwater extraction do not simply involve applying more water to the same crops, holding everything else constant. Instead, they involve investment decisions that pay off over long periods of time.

To capture this possibility, we allow users an opportunity to invest in a water-intensive production technology. The investment requires an initial cost of K but then delivers greater marginal benefits for any amount of extraction. This setup naturally describes investment in a perennial crop, and it shares basic features with the decision of investment in well construction.[9]

To obtain closed-form expressions, we use a second-order approximation to the benefit function and assume that the investment increases marginal benefits by a constant: $B_0(y) = ay - \frac{1}{2}by^2$ and $B_I(y) = (a+\beta)y - \frac{1}{2}by^2$, where $a, b, \beta > 0$.

Users face a two-stage problem. First, a user chooses whether to make the investment, by comparing the present value of profits with and without the investment. Second, the user chooses extraction quantities to maximize profits, as before, given the investment decision. The problem is:

$$\text{Invest if: } \sum_{t=0}^{\infty}(1+r)^{-t}\left[B_I(y_{it}^I) - c(x_{it}^I)y_{it}^I\right] - K \geq \sum_{t=0}^{\infty}(1+r)^{-t}\left[B_0(y_{it}^0) - c(x_{it}^0)y_{it}^0\right]$$

where y_{it}^I, x_{it}^I, y_{it}^0, and x_{it}^0 are the solutions to the extraction problem in section 3.2, with and without investment.

We study three questions: (1) how investment affects current extraction, (2) how future regulation affects investment, and (3) how future regulation affects extraction as the result of investment.

Lemma 1 (Effect of investment on extraction)**.** *Investment increases extraction in period 0:*

$$y_{i0}^I - y_{i0}^0 = \beta/b > 0. \tag{5}$$

Proof. See Appendix Section A.3. □

Investment increases period-0 extraction simply because it increases the marginal benefits from extraction in period 0. Extracting more in period 0 does increase extraction

---

[9]Well construction is also an up-front investment that pays off over time, with payoffs increasing in extraction. We omit an explicit model of the well construction decision because it would require allowing the marginal cost function either to depend on y_{it} (reflecting a cone of depression within each period) or to be non-convex in x_{it} (reflecting cost discontinuities as wells go dry and must be replaced). Either modification would preclude closed-form solutions for our expressions of interest.

costs in the future, but the marginal cost of this scarcity value is flat in period-0 quantity extracted and does not depend on the investment.

Next, to study how regulation in period 1 affects investment, we define the return on investment Θ, the net present value of the investment excluding the initial cost K_i. We assume the initial cost is a continuous random variable that follows a cumulative density function F_K; the probability density function is f_K. A user invests if $\Theta \geq K_i$, so the greater the return on investment, the more likely a user is to invest.

Proposition 2 (Effect of future regulation on investment and resulting extraction). *Future extraction limits may either increase or decrease both investment and extraction as the result of investment:*

$$\frac{dI_i}{d\bar{y}} = f_K(\Theta)\beta(1+r)^{-1}\left[1 - \frac{ps}{bN}\right]. \tag{6}$$

while extraction as the result of investment is $(y_{i0}^I - y_{i0}^0)(dI_i/d\bar{y})$. When $bN < ps$, a decrease in extraction limits (i.e., an increase in regulatory stringency) raises the probability of investment (equivalently, the share of users who invest), as well as extraction as the result of investment ($dI_i/d\bar{y} < 0$). It lowers investment and resulting extraction when $bN > ps$ ($dI_i/d\bar{y} > 0$), and it has no effect when $bN = ps$ ($dI_i/d\bar{y} = 0$).

Proof. For Equation 6, see Appendix Section A.4. The extension to extraction quantities follows immediately by combining this equation with Lemma 1. □

Equation 6 says that the effect of future extraction limits on investment depends on the benefits and the costs of the additional allowed extraction in period 1. The first term is the benefits of this extraction: It represents the direct effect of regulation, the relatively greater marginal benefits under investment. When more extraction is allowed, the marginal benefits of that extraction are greater with investment, so the investment is more attractive. (More stringent extraction limits reduce the available marginal benefits of investment, so investment is less attractive.)

The second term is the costs of this extraction. A marginal rise in allowed extraction increases total extraction costs by the marginal cost of extraction. Because the investment leads the user to increase period-0 extraction, reducing the stock in period 1, those marginal extraction costs increase, making the investment less attractive. (This second term is the increase in period-0 extraction caused by the investment, multiplied by the marginal increase in period-1 extraction costs caused by the reduced stocks.)

As for extraction, the results show that future extraction limits affect period-0 extraction not just directly, as in Proposition 1, but also through the channel of investment.

12

If future regulation makes investment more attractive, then period-0 extraction increases, because we know from Lemma 1 that investment increases extraction. If future regulation makes investment less attractive, the forgone investment would have increased period-0 extraction, so period-0 extraction decreases as the result of the investment opportunity.

3.4 Net effects of regulation are theoretically ambiguous

With the opportunity for investment, we have multiple simultaneous effects. As a result of stricter future regulation, groundwater users will increase current extraction, exhibiting a Green Paradox conditional on investment (Proposition 1). At the same time, they may decrease investment in water-intensive production technologies, reducing extraction in anticipation of the regulation – or alternatively increase it (Proposition 2). Considering all these effects, we can summarize how future regulation affects extraction overall.

Proposition 3 (Net effect of future regulation on investment and extraction). *The effect of future regulation on current extraction, in total through all channels, is:*

$$\frac{dy_{i0}}{d\bar{y}} = (1+r)^{-1}\left(1 + f_K(\Theta)\frac{\beta^2}{b}\right)\left[\xi^{-1} - \frac{ps}{bN}\right] \tag{7}$$

and the directional effects of tightening future extraction limits on current investment and extraction depend on the following conditions:

Condition	Investment	Net Extraction	
$bN < ps$	Rises $\left(dI/d\bar{y} < 0\right)$		[Green Paradox]
$bN = ps$	No effect $\left(dI/d\bar{y} = 0\right)$	Rises $\left(dy_{i0}/d\bar{y} < 0\right)$	
$ps < bN < ps\xi$			[Mixed results]
$bN = ps\xi$	Falls $\left(dI/d\bar{y} > 0\right)$	No effect $\left(dy_{i0}/d\bar{y} = 0\right)$	
$bN > ps\xi$		Falls $\left(dy_{i0}/d\bar{y} > 0\right)$	[Early decline]

where $\xi := \left(f_K(\Theta)\frac{\beta^2}{b}\right)^{-1} + 1$.

Proof. Results for investment are restated from Proposition 2. For extraction, we write current extraction as a function of the regulation through both direct and indirect channels: $y_{i0} = y_{i0}(\bar{y}, I_i(\bar{y}))$. Totally differentiating with respect to $\bar{y}$ gives

$$\frac{dy_{i0}}{d\bar{y}} = \frac{\partial y_{i0}}{\partial \bar{y}} + \left(y_{i0}^I - y_{i0}^0\right)\frac{dI_i}{d\bar{y}}.$$

The first term is given in Proposition 1 (i.e., $\partial y_{i0}/\partial\bar{y}$ is $dy_{i0}/d\bar{y}$ conditional on the investment decision) and the second is given in Proposition 2. The remaining algebra is given

13

in Appendix Section A.5. □

The results show three main regimes (plus two edge cases):

1. With few users or flat marginal benefits (low bN), investment can actually exacerbate the Green Paradox. The regulation increases overall extraction in period 0, both directly, conditional on investment decisions (Proposition 1), and indirectly, through increased investment in the technology.

2. In open access or with steep marginal benefits (large bN), regulation reduces investment and extraction as a result, and this effect outweighs any Green Paradox tendency to increase extraction conditional on investment.

3. In between, there is an intermediate range of values for which regulation reduces investment while also increasing extraction. Extraction falls because of reduced investment, but not by enough to outweigh the Green Paradox increase conditional on investment.

Of course, these conditions do not guarantee that the effects are large; a high value of the discount rate r can make Equation 7 arbitrarily small.

4   Data and Descriptive Statistics

To take our theory to data, we assemble measures of groundwater extraction and water-intensive investment for all agricultural land in California subject to SGMA. For investment, the outcomes we can observe are the construction of agricultural wells (from well completion reports) and the conversion of land to perennial crops such as orchards and vineyards (from a satellite-based land use data product). For extraction, we form an index of water use by combining the same satellite data on land use with scientific estimates of water use by crop. We also assemble several estimates of expected future groundwater regulations (for the treatment variable) and surface water deliveries (for an important control variable).

Summary statistics are reported in Table 1. The full sample consists of yearly observations during 1993-2022 of all land within GSAs subject to SGMA (i.e., designated as medium or high priority). Each observation represents a quarter-quarter section (about 40 acres) in the Public Land Survey System.[10] The paired sample consists of observations

---

[10]We aggregate spatial variables in this way in order to reduce noise and computation time without losing much information. Hagerty (2021) shows that this division consistently keeps together common units of land use.

from the full sample that fall within 15 km of the boundary between a pair of neighboring groundwater subbasins, with all such subbasin pairs stacked into one dataset. We motivate this sample in Section 5.1.

4.1   Future Extraction Reductions Under SGMA

Our ideal treatment variable would capture farmers' beliefs of the degree to which they will be required to reduce their groundwater pumping in order to achieve the basin's sustainability goals. Because true beliefs are unobservable, we proxy for it by assembling three different measures of the likely future reductions in extraction that will be required in each GSA.

Our first measure, which we refer to as "modeled overdraft," comes from the 1.0 version of the Fine Grid California Central Valley Groundwater-Surface Water Simulation Model (C2VSim), developed by DWR. C2VSim is one of three major statewide hydrological models widely used in water resource planning in California, and the only one that is publicly available. We run C2VSim using default parameters and obtain estimates of the yearly volumetric change in groundwater storage for each year of the 25-year period preceding SGMA (1992-2015).[11] We aggregate gridded values to GSAs by summing over all model grid cells whose centroid falls within each GSA boundary, and take an average across the years of this historical period.

Our other two measures are assembled from Groundwater Sustainability Plans (GSPs) submitted by GSAs to the state. GSPs are multi-thousand-page reports that estimate and report overdraft as well as current and future pumping. One measure, which we refer to as "reported overdraft," is the volume of annual overdraft reported directly in the executive summary of each GSP.[12] The other, which we refer to as "projected reduction," is the difference in annual groundwater extraction between "current" and "future" water budgets. Projected reduction can differ from reported overdraft because many GSPs also project changes in groundwater supply.

For each of these three variables, we divide the GSA-level volumes by the area of undeveloped land in the GSA to obtain a per-acre measure of estimated future pumping reductions for agriculture. By doing so, we assume that future reductions in extraction will be borne exclusively by the agricultural sector and not by municipal users. This is a reasonable approximation, since agriculture is responsible for the vast majority of ground-

---

[11]Change in storage and overdraft are conceptually very similar; however one incorporates lateral flow. Overdraft tells us the difference between pumping (out) and recharge (in), net of lateral flows.

[12]Each plan contains several water budgets that are based on different subsets of historical data. The plans state their preferred water budget and corresponding preferred overdraft estimate, which we use.

Table 1: Descriptive Statistics

	Observations	Mean	Std. Dev.
Full Sample			
Future reductions, mean of 3 measures (AF/acre)	22137750	0.078	0.13
Projected reduction, from GSPs (AF/acre)	22137750	0.054	0.15
Reported overdraft, from GSPs (AF/acre)	22137750	0.085	0.18
Modeled overdraft, from C2VSim (AF/acre)	22137750	0.094	0.17
Crop water intensity (AF/acre)	12217269	2.2	2.1
New perennials planted (share of land)	11408667	0.0067	0.18
New wells per square mile	24447150	0.03	2.6
Stock of perennials planted (share of land)	12223572	0.14	0.35
Stock of wells per square mile	24447150	0.95	32
Surface water deliveries (AF/acre)	24447150	1.3	2.3
Paired Sample			
Future reductions, mean of 3 measures (AF/acre)	13924988	0.12	0.16
Projected reduction, from GSPs (AF/acre)	13924988	0.086	0.19
Reported overdraft, from GSPs (AF/acre)	13924988	0.13	0.22
Modeled overdraft, from C2VSim (AF/acre)	13924988	0.15	0.19
Crop water intensity (AF/acre)	7202578	2.9	1.8
New perennials planted (share of land)	6722406	0.011	0.23
New wells per square mile	13924988	0.043	3.8
Stock of perennials planted (share of land)	7202578	0.24	0.43
Stock of wells per square mile	13924988	1.3	46
Surface water deliveries (AF/acre)	13924988	1.4	2.2

Notes: This table reports units, observations, means, and standard deviations (SD) for the full and paired samples. The full sample includes all land within GSAs as yearly observations of quarter-quarter sections. The paired sample is the subset of observations within 15 km of the boundary between pairs of neighboring groundwater subbasins, with all such pairs stacked into one dataset. The paired sample excludes land within 1.14 km (i.e., 1 mile $\times \frac{\sqrt{(2)}}{2}$) of the boundary to avoid classifying wells to the wrong side of the border; well construction data is rounded to the nearest mile for anonymity. Water is measured in acre-feet (AF). Dataset runs 1993-2022; crop water intensity and perennials have fewer observations because they are derived from remote sensing data available 2007-2021. New perennials planted is the first difference of the stock of perennials planted, so it is not calculable for 2007. Measures of future pumping reductions are inherently cross-sectional but repeated for each year of the panel.

water extraction (in many GSAs, the extent of overdraft alone exceeds total municipal use) and the value of water tends to be much higher in residential and industrial uses. We also assume that pumping reductions will be divided evenly across all agricultural land in the GSA. In the absence of more specific regulatory plans, this is a reasonable assumption because of strong pre-existing allocation norms; surface water districts in California almost always allocate reductions in irrigation water equally across cropland area (Hagerty, 2021). We also censor negative values at zero. Negative values mean that a GSA has room to extract more groundwater each year without suffering overdraft. Because our focus is on future reductions in extraction, we only care about the extent of overdraft, not the extent of resource under-utilization.

Our final treatment variable averages across the per-acre versions of these three proxies and is shown in Figure 2. In some cases, multiple contiguous GSAs joined together to collaboratively develop one GSP; we combine and treat them as one unit in our analysis. We also exclude GSAs that exclusively or primarily cover cities. The subset of groundwater basins that reside in the Central Valley form the basis of our full-sample analysis and consist of both critically and non-critically overdrafted basins. The estimated reduction in groundwater extraction under SGMA ranges from 0 to 1.1 acre-feet per acre (AF/acre)[13] and averages 0.12 AF/acre. For context, California crops like fruits, vegetables, and nuts can use 1.5 to 4 AF of water per year depending on the crop.

4.2 Land Use

Our land use data consist of annual information on crops grown in the state at a 30-meter grid resolution spanning 2007-2021. We use the USDA's Cropland Data Layer, which is a remotely sensed data product of 119 distinct land-use classifications. We aggregate pixels to fields (quarter-quarter sections as described above). We classify land use into six categories: annual crops, perennial crops, fallowed/idled land, grassland, nature, and developed space. Figure 3 plots trends in these land use categories over time.

Throughout our sample, we observe a trend of annual acreage declining and perennial acreage increasing. This trend is visible in years prior to the passing of the SGMA legislation. The drop in annuals appears to have leveled off in the initial years after the announcement of the policy before continuing a downward trend in recent years. Perennial acreage has steadily increased since 2010, roughly doubling over a 10 year period with no visible changes in the trajectory in the years after SGMA. In fact, perennial crops have increased nearly 50% since SGMA passed in 2014.

---

[13] An acre-foot is the volume of water that would cover one acre of land to a depth of one foot.

Figure 2: Spatial Variation in Regulatory Stringency
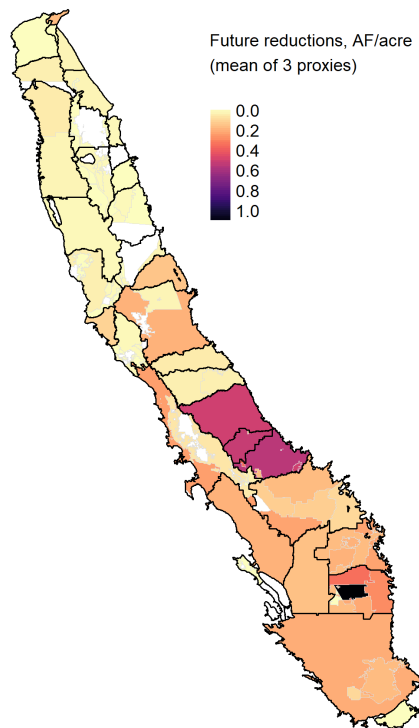
Note: The map shows average expected reduction in groundwater pumping required under SGMA in acre-feet per acre (AF/acre) for basins in the Central Valley. This average reduction is estimated by averaging across the three treatment variables: reported overdraft, projected reduction, and modeled overdraft.

Figure 3: Land Use, 2007-2021
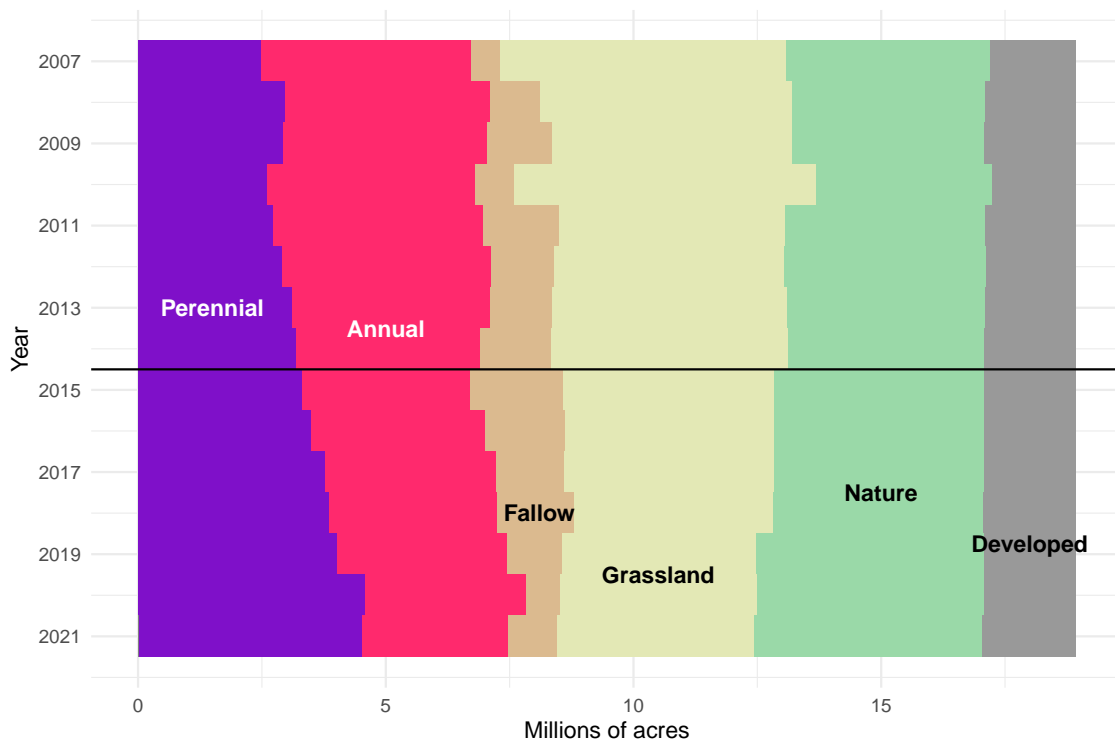
Note: Data come from USDA's Cropland Data Layer. The horizontal line marks the passage of SGMA.

For our outcome variable of new perennial plantings, we take the first difference of a binary indicator for whether each field is planted with perennials in a given year. This first difference may amplify noise from classification error, so we attempt to reduce error by applying a data correction procedure that leverages the panel structure of the data.[14] Overall, about 1% of fields are newly planted with perennials in each year of our sample.

4.3 Agricultural Well Construction

A secondary outcome variable is new agricultural well construction, which is another long-term investment decision that may be affected by expected future water supply. We use the Well Completion Reports from the Department of Water Resources, which represent the universe of agricultural wells drilled in California. The data run through 2022 and extend back many decades, but we use data beginning in 1993 for congruence with our other variables. The dataset includes information on each well's location, drilled depth, and intended use.

Because the data source reports only where wells were constructed, not where they were not, we form a consistent sample frame by joining well observations to the farm fields we defined above for land use observations. Many (but not all) well locations are anonymized by rounding to the nearest node in a one-square-mile grid. This means that some of our fields have an implausible number of wells while most others have none, but this is not a problem because all analysis smooths over fields within each basin. The most concerning type of measurement error would be misclassification of a well into the wrong subbasin. We eliminate this error in the paired sample by excluding fields that may be misclassified: those within 1 mile $\times \frac{\sqrt{2}}{2} = 1.14$ km of the boundary.

Our final variable is the number of new wells per year per square mile, which we construct by dividing the number of new wells in a field by the square mileage of the field. In all analysis, we weight by land area of field observations, to ensure estimates are geographically representative and do not depend on the method of aggregation. The mean number of new wells per year in our full sample is 0.03 per square mile. Taking a cumulative sum of all new wells through the observed year for each field, the mean number of total wells is 0.95 per square mile.

---

[14]Perennial crops by definition must exist for more than one year, so for each field, we examine sliding five-year windows. If the land use code is identical in years 1, 2, 4, and 5, but different in year 3 – and either the year-3 value is a perennial crop and the surrounding years are not, or vice versa – we correct the year-3 value to be the same as the surrounding years.

4.4 Water Use

To proxy for groundwater extraction, we create an index of crop water use intensity, which estimates the volume of irrigation water used at each field. To form this index, we combine the land use data above with estimates of crop-specific water use by fine geographic regions and year, provided by DWR and described further in Hagerty (2021). We join each field to the region that contains it and impute the water use estimate for the crop observed at that field.

Our goal is to measure groundwater use, but total water use includes both surface water and groundwater. This is not a problem for our difference-in-differences analysis so long as any post-SGMA changes in surface water quantities are uncorrelated with expected future reductions in groundwater extraction under SGMA. In case this is not true, we also collect data on surface water supplies from Hagerty (2021). This dataset includes annual volumes of surface water deliveries from the Central Valley Project (CVP), State Water Project (SWP), and Lower Colorado operations, and estimated diversions on the basis of surface water rights, spanning 1993-2022. On average, surface water use amounts to 1.3 AF/acre, or 62% of total applied water.

5 Empirical Approach

5.1 Research Design: Paired Difference-in-Difference

To measure the effects of future reductions in groundwater extraction, we leverage the fact that SGMA has created substantial variation in future regulatory stringency across geography in California. Our basic approach is to compare outcomes across different GSAs that are subject to greater or lesser future pumping reductions. However, a simple analysis that pools together all GSAs into a single comparison raises immediate problems. Regions facing greater reductions under SGMA are very different from regions facing fewer (or no) reductions. Figure 1 illustrates this well: basins deemed to be in "critical overdraft" largely reside in the southern half of the Central Valley, where weather and growing conditions are quite different from regions in the northern half.

A difference-in-difference analysis would help by subtracting out baseline trends, but even this relies on a parallel trends assumption that is difficult to justify from institutional knowledge. Farms in the southern Central Valley have been planting perennials and depleting groundwater at a faster rate than those further north, so the post-SGMA trajectory of northern GSAs is unlikely to be a plausible counterfactual for that of southern GSAs.

We illustrate the challenges of a full-sample analysis in Figure 4, which plots our three outcome variables over time by critical overdraft status, a coarse binary classification that correlates with future pumping reductions. Not only are critically overdrafted basins quite different from others – for example, they grow considerably more water-intensive crops – they also fail to exhibit non-parallel trends in the pre-treatment period (prior to 2014).

Instead, we use a paired difference-in-difference approach. Rather than comparing each GSA to all other GSAs in the state, we identify the impacts of impending groundwater regulation by comparing each GSA to other neighboring GSAs, before and after the restrictions became known. Neighboring GSAs often still have variation in expected future groundwater restrictions, yet they are more similar to each other in other ways. A paired approach relaxes the parallel trends assumption – it is more plausible to think that neighboring GSAs would have similar responses to common shocks absent SGMA, than to think the same of all GSAs in the state.

An alternative research design in this setting might be a geographic regression discontinuity (RD) at the boundaries between neighboring GSAs. The main reason we prefer a difference-in-difference approach is that many boundaries of GSAs coincide with boundaries of water districts, which supply surface water and have been shown to introduce discontinuous effects on land use and agricultural production (Hagerty, 2021). If we estimated a geographic RD in the post-SGMA period, it would likely include bias from these other borders. Instead, the difference-in-difference design accounts for this bias, by allowing us to ask how much the spatial difference across GSA boundaries changed post-SGMA relative to the pre-SGMA period.

Still, our analysis includes some elements of an RD design to deal with another important concern: Future restrictions on groundwater extraction are determined not randomly but by amount of overdraft. Regions with greater overdraft tend to have lower groundwater levels, so they are likely to respond to economic shocks differently than would regions with less overdraft. However, GSA boundaries represent only administrative boundaries, not hydrological boundaries, so underground groundwater levels equalize across GSA boundaries. Two neighboring GSAs *on average* might have very different groundwater levels (and therefore face different future restrictions), but close to the boundary between them, groundwater levels (and therefore the cost of extraction) will be nearly identical.[15]

---

[15]This will not be true if the GSA boundaries are drawn to coincide with physical barriers that restrict underground flow. This is why, as mentioned in Section 2, we do not use GSA comparisons across boundaries of basins, which are defined by hydrogeological features. We use only comparisons across boundaries of subbasins, which are defined for administrative convenience and have no hydrogeological meaning. The exception is that in the Central Valley we combine the Sacramento Valley and San Joaquin Valley basins, which are connected underground but defined separately because of their surface hydrology.

Figure 4: Investment and Extraction Outcomes in Full Sample

A. New Perennial Plantings

B. New Well Construction

C. Water Use Intensity
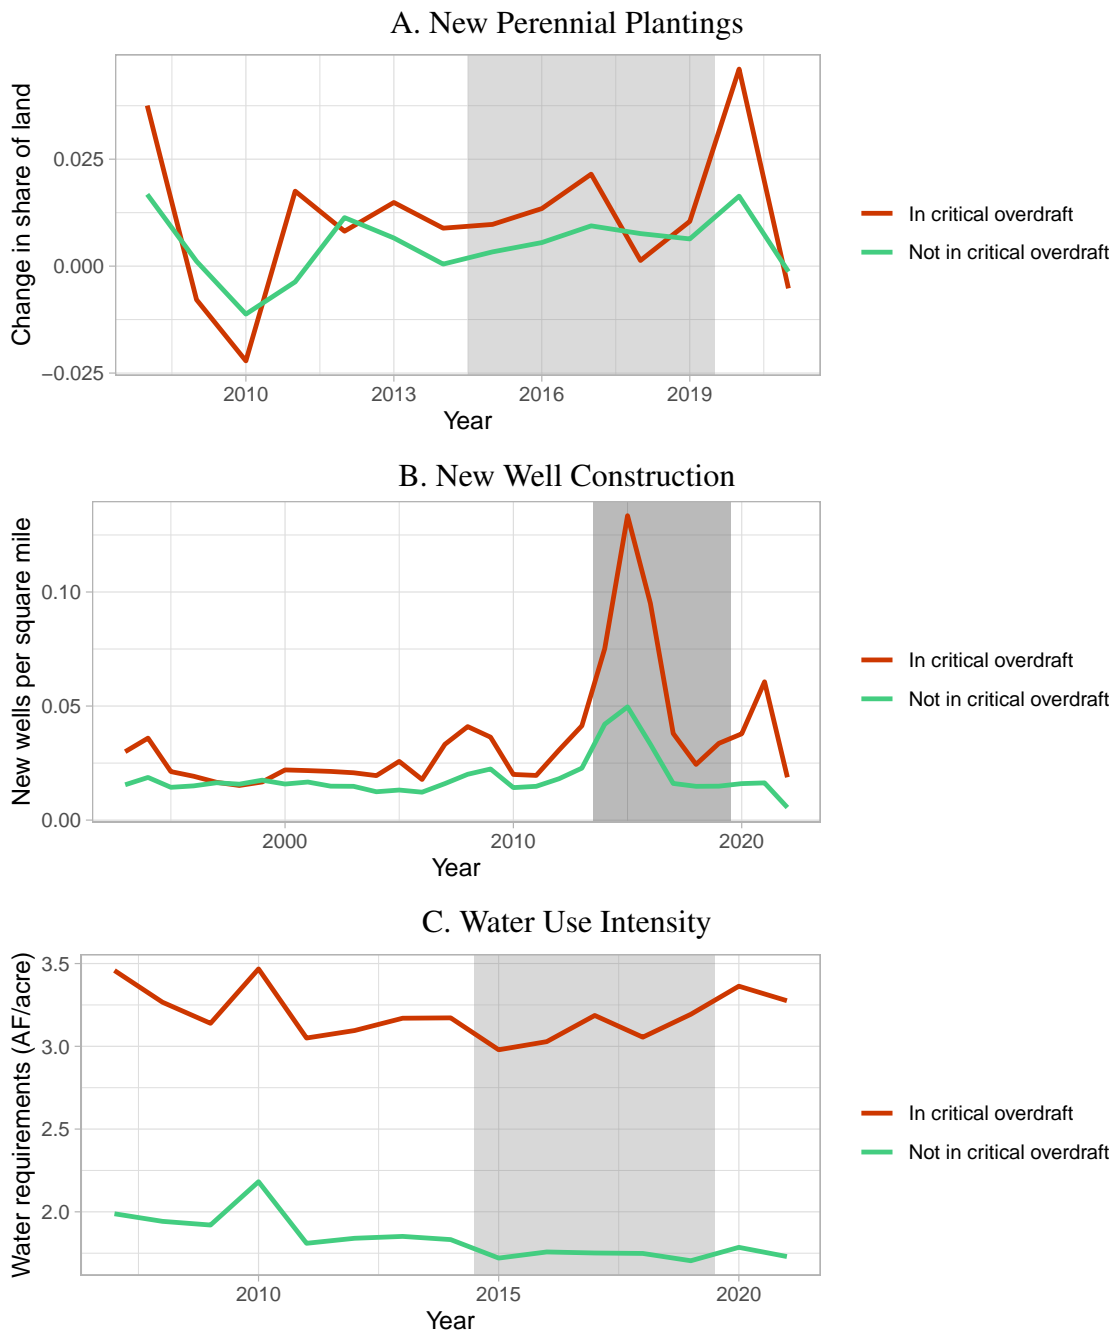
Note: Graphs plot the three outcome variables by critical overdraft designation using the full sample – all agricultural land in GSAs affected by SGMA. Years shaded in gray denote the time between passage of SGMA and release of GSPs; the "pre-treatment" period is before the gray period and the "post-treatment" period is after it.

As a result, areas immediately around a GSA boundary have different values of the treatment variable (future restrictions change discretely at the boundary and are likely to apply equally throughout each GSA) but share more similar environmental conditions than areas further away from the boundary. So in the spirit of an RD design, our main specification controls for distance to the boundary and applies triangular kernel weights that put greater weight on areas closer to the boundary, although results turn out to be insensitive to these choices.

To form the paired sample for our main analysis, we find all pairs of contiguous groundwater subbasins in California, restrict the full sample to observations that fall within 15 km of the boundary between each pair of neighboring subbasins, and then stack observations from all such pairs into one dataset. The paired sample is therefore both restricted and repeated; many observations appear multiple times as part of distinct pair comparisons. The radius of 15 km is chosen to be large enough to include a substantial mass of observations on both sides of the boundary while small enough to ensure they are similar; we show that results are insensitive to this specific choice.

5.2   Timing of Treatment

To select time periods for the before-after comparison, we want to isolate periods that are completely unaffected by SGMA, and those during which the future pumping restrictions are clear. The pre-treatment period is reasonably straightforward, since SGMA passed in September 2014. For the outcome variables of new perennial plantings and water use intensity, we consider 2014 to be the last pre-treatment year. These two variables are derived from observations of land use, which would not have responded late in the calendar year, since planting decisions are made in early spring. For well construction, we consider 2013 to be the last pre-treatment year, since wells are drilled at discrete times, so decisions during 2014 could have been affected by the legislative process in that year.

We define the post-treatment period as only starting in 2020. We exclude the intervening years of 2015-19 from both pre- and post-treatment periods and consider them to be a "coordination" or "middle" period. The reason is that the post-treatment period should consist of a time during which we can be confident that farmers have changed their beliefs about the future availability of water under SGMA. The years immediately after SGMA do not fit this description: The deadline for GSAs to form was June 30, 2017, so before then, farmers did not even know what GSA they would be in. It was not until 2018 that sustainability plans were drafted and public hearings held. However, after this point,

GSAs undertook significant community outreach and engagement.[16] By the time each GSA published a Notice of Planned Adoption of their sustainability plans – late 2019 for almost all GSAs in our dataset – it is likely that landowners successfully updated their beliefs about changes to future pumping.

Because the timing of our treatment variable is simultaneous across all units, we avoid many of the problems identified in the recent literature on difference-in-differences (Baker et al., 2021).

5.3   Regression Model

To build intuition for our main regression specification, consider a simple scenario in which two GSAs g in neighboring subbasins differ in expected future pumping reductions. The treatment variable T_g takes a value of 1 for the GSA facing greater cutbacks and 0 for the other. The timing variables Mid_t and $Post_t$ equal 1 in the coordination period (after SGMA was announced in 2014 but before GSPs were finalized in 2019) and the post-treatment period (after farmers have had a chance to update their beliefs about future pumping restrictions), respectively. If we regress an outcome Y_{igt} for field i on these variables and their interactions:

$$Y_{igt} = \gamma T_g + \lambda_1 Mid_t + \sigma_1 (T_g \times Mid_t) + \lambda_2 Post_t + \sigma_2 (T_g \times Post_t) + \varepsilon_{igt} \qquad (8)$$

the coefficient on $T_g \times Post_t$ captures the additional effect of being in the GSA with greater future pumping restrictions (relative to its neighbor) in the post-treatment period (relative to the pre-treatment period, excluding the coordination period).

Our main specification stacks together all 73 pairs of neighboring subbasins by using the paired sample. It pools the coefficient of interest β across pairs p:

$$Y_{igpst} = \gamma T_{gp} + \delta (T_{gp} \times Mid_t) + \beta (T_{gp} \times Post_t) + \alpha_{pst} + \omega' X_{igpst} + \varepsilon_{igpst}. \qquad (9)$$

As described above, the baseline sample is restricted to observations within 15 km of the boundary of each subbasin pair. The variable α_{pst} represents year $\times$ subbasin pair $\times$ boundary-segment fixed effects. These fixed effects control for time-invariant subbasin pair characteristics as well as annual shocks shared by GSAs on both sides of the subbasin boundary. We split each boundary pair into 5-km pieces we call boundary segments s to

---

[16]Community outreach and engagement were codified into the law under SGMA. In fact, GSAs were required to record their public outreach efforts. Stakeholder engagement included the dissemination of resources regarding SGMA implementation and several public comment hearings at the local level.

ensure the regression is comparing observations that are near each other in both perpendicular and parallel dimensions. The fixed effects thus ensure our coefficient of interest is identified by comparing fields only directly across a subbasin boundary from each other.

We control for surface water supplies in case the treatment variable happens to be correlated with any post-SGMA shocks to surface water quantities. Our covariates X_{igpst} include surface water supplies in both the same year and the previous year, to capture the recent past of any decisions that affect the outcome variables, since investment and extraction decisions are made throughout the year.[17] We also include interactions of these surface water variables with a full set of year indicators, to flexibly allow the effects of surface water across GSAs to vary separately for each year in the data.

In the spirit of an RD design, we also control for perpendicular distance to the subbasin boundary, and interact this distance with T_{gp} to estimate separate terms on each side of the boundary. Observations are weighted both by field acres (to obtain estimates that are representative of land area) and by a triangular kernel in distance to the boundary (following Cattaneo et al. (2019)). Standard errors are clustered by the unit of treatment – GSA, or sets of GSAs that submit a joint GSP – to account for both serial and spatial correlation.

To show effects over time, we also deploy an event study framework that estimates separate effects for each year of our data:

$$Y_{igpst} = \gamma T_{gp} + \sum_{t \neq 2014} \theta_t T_{gp} + \alpha_{pst} + \omega' X_{igpst} + \varepsilon_{igpt}. \tag{10}$$

relative to an excluded year of 2014 (for new perennial plantings and water use intensity) or 2013 (for well construction).

In our baseline specification for both event studies and the pooled regression, we use a simple binary indicator for the treatment variable T_{gp}. It measures the effect of being in a subbasin that faces greater future pumping restrictions than its neighbor, on average across all pairs of neighboring subbasins. This effect tells us about the direction of response, but to interpret it quantitatively, we also need to know the average difference in future pumping reductions between subbasin pairs in the paired sample: 0.12 AF/acre. We also estimate an alternative specification that uses the raw estimated value of future pumping reductions as a continuous treatment variable.

Identification in our setting requires that in the absence of the sustainability mandate, differences in outcomes between the treated and counterfactual comparison groups would

---

[17]Bruno et al. (2023a) show that well construction does not respond to surface water supplies more than one year later.

have remained constant over time. We lean on the panel of pre-treatment data to test for differences in outcomes between treated and control units in years prior to SGMA. The failure to identify a difference in the pre-treatment years provides evidence to support the assumption that in the absence of the policy, treatment and comparison groups would have trended similarly.

Despite empirical evidence for the absence of pre-trends, it could still be the case that GSAs subject to greater pumping restrictions would have trended differently after 2014. The post-treatment years in our sample mark a tumultuous time for California farmers, many of whom produce goods for international buyers and suffered losses from retaliatory tariffs, port congestion, and continuing supply chain issues. While many of these shocks may have differential effects on growers of different crop types, they are unlikely to be correlated with GSA-level variation in overdraft within neighboring subbasin pairs.

6   Results

We present results for three outcome variables: new perennial plantings and new well construction as measures of investment, and crop water intensity as a proxy for groundwater extraction.

6.1   New Perennial Planting

To measure whether future regulation leads farmers to increase or decrease their rate of investment in water-intensive capital, we first consider the rate of new plantings of perennial crops, such as fruit and nut orchards and vineyards.

To start, we assess trends in new perennials in the pre-treatment period. Figure 5 shows that in the paired sample – unlike in the full sample – new perennial plantings tracked each other very closely prior to 2014. Not only do they appear to move in parallel, they also closely match in levels. Since the "fewer" and "more" groups behave so similarly prior to SGMA, it increases confidence that they would have also behaved similarly afterward without SGMA – and that the parallel trends assumption is much more plausible in the paired sample than the full sample (Figure 4).

Next, we examine how new perennial plantings changed in the post-treatment period, after SGMA passed and future pumping reductions became clearer. The answer appears obvious from the time series plot in Figure 5: there was no change. GSAs facing fewer vs. more future pumping reductions continued to plant perennials at the same rate as each other in the post-treatment period just as much as in the pre-treatment period. One concern

Figure 5: New Perennial Plantings by Treatment Status, Paired Sample
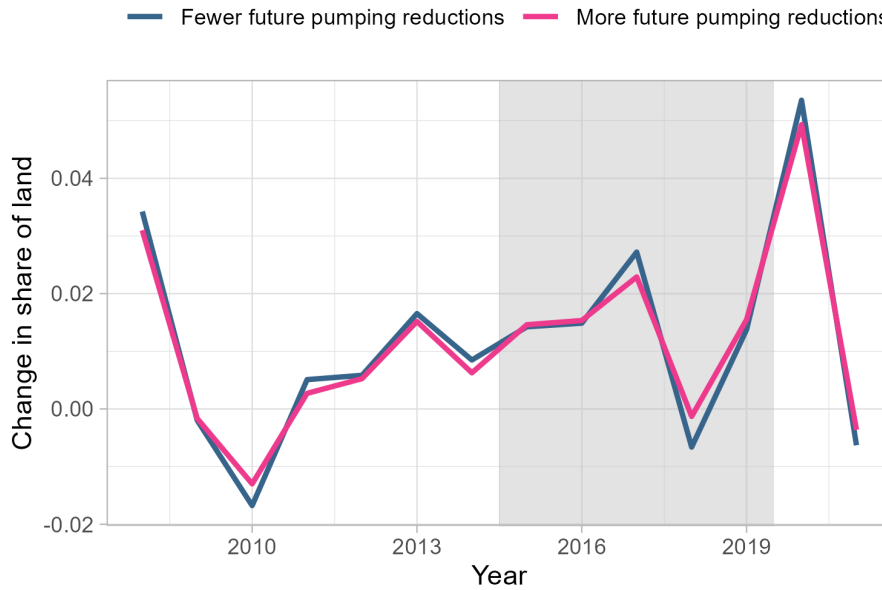
Note: Figure plots the annual change in the share of fields planted in perennial crops in the paired sample, which stacks all neighboring subbasins and includes only observations within 15 km of their boundary. "More" and "Fewer" are within these pairs, relative to neighbors in the same year. Gray shading indicates the "coordination" period between when SGMA was passed and when local sustainability plans were published. Means weighted by area.

Figure 6: Effect of Greater Future Reductions on New Perennial Plantings
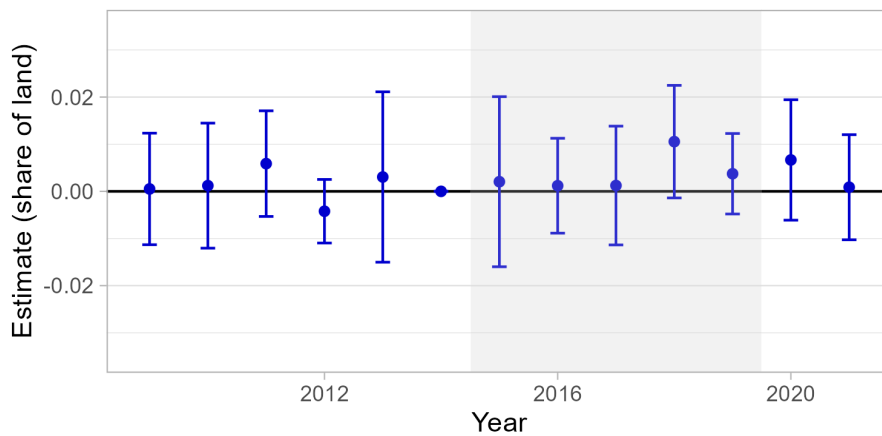
Note: Figure plots year-specific coefficients from the estimation of Equation 10. Each coefficient represents the difference in new perennial plantings between GSAs facing more or fewer future pumping restrictions (within each pair of neighboring subbasins) in that year, minus the same difference in 2014, the last year of planting decisions before SGMA became law. Estimates also adjust for surface water supplies and distance to the boundary and are weighted by area and a triangular kernel in distance to boundary. Sample is limited to fields within 15 km of the boundary. Vertical bars denote 95% confidence intervals. Standard errors clustered by GSA.

might be whether the "more" and "fewer" groups really do have meaningful differences in the treatment variable. But despite the similarity in the outcome variable, the average difference in future pumping reductions is 0.12 AF/acre – roughly the same as the average *value* of future reductions across the paired sample as a whole.

To confirm this apparent result, we proceed to showing results from a formal event study: the effects over time estimated from equation 10. Figure 6 plots the year-specific average effect of being in the GSA with "more reductions" between each neighboring pair, relative to 2014, the last pre-treatment year. This figure plots the same data as Figure 5, but it shows differences between the two groups in each year net of their 2014 difference, controls for surface water supplies and distance to subbasin boundary, and adds confidence intervals. Since farmland near the boundary is very similar other than the change in expected future pumping restrictions, we can interpret any differences in new perennial plantings relative to 2014 as the effect of being in a GSA with greater future regulation.

In each of the five years preceding the passage of SGMA, we fail to reject that the difference in average new perennial plantings across all subbasin pairs is statistically different from that in 2014, again lending confidence to the identifying assumption. However, the effects of greater future pumping reductions in each of the two years in the post-treatment period similarly show no statistical difference in new perennial plantings relative to 2014. These results suggest that farmers are not making anticipatory adjustments in new perennial plantings as a result of SGMA in these early years.

6.2 New Well Construction

We next turn to a second measure of water-intensive capital investments: new construction of irrigation wells. In Figure 7, we report trends in new wells constructed per square mile over time by treatment status in the paired sample, focusing only on farms within 15km of the boundary between agencies. For this outcome variable, we can lean on a longer panel of pre-treatment data to investigate the parallel trends assumption. We again see that subbasins that face more and fewer future pumping reductions under SGMA closely tracked each other in the pre-treatment period – as well as in the post-treatment period.

Turning to the event study, Figure 8 plots coefficient estimates from the estimation of equation 10 with new well construction as the outcome variable. With few exceptions, we cannot reject that the differences in new well construction between subbasin pairs are significantly different from that in 2013. The estimated effects in the coordination years of 2014-2019 and the post-treatment years of 2020-2022 similarly show no statistical differ-

Figure 7: New Well Construction by Treatment Status, Paired Sample
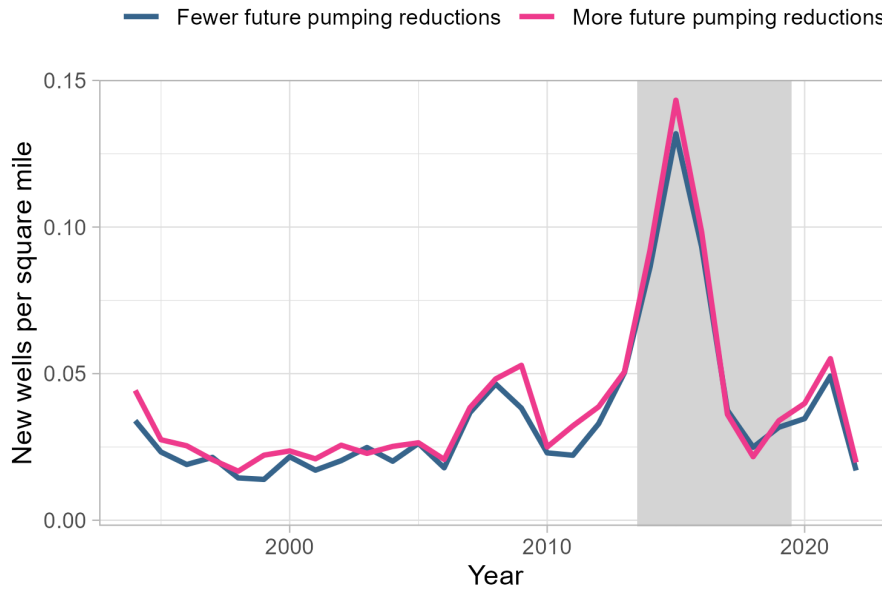
Note: Figure plots the mean annual count of new wells constructed per unit area in the paired sample, which stacks all neighboring subbasins and includes only observations within 15 km of their boundary. "More" and "Fewer" are within these pairs, relative to neighbors in the same year. Gray shading indicates the "coordination" period between when SGMA was passed and when local sustainability plans were published. Means weighted by area.

Figure 8: Effect of Greater Future Reductions on New Well Construction
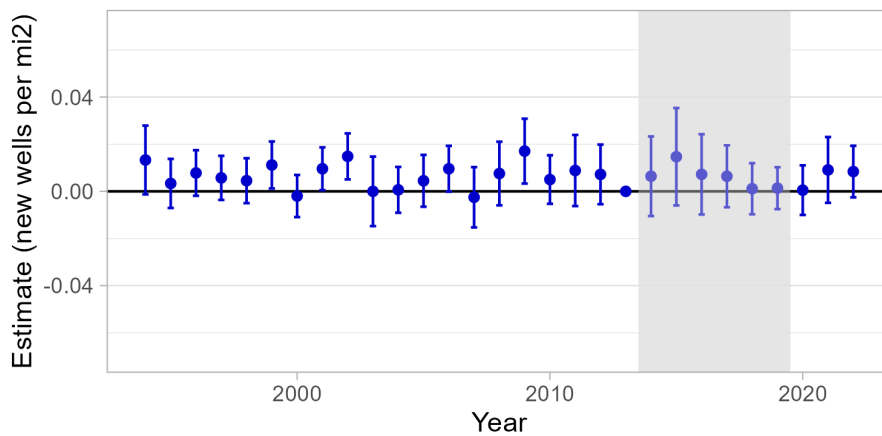
Note: Figure plots year-specific coefficients from the estimation of Equation 10. Each coefficient represents the difference in new well construction between GSAs facing more or fewer future pumping restrictions (within each pair of neighboring subbasins) in that year, minus the same difference in 2013, the last full year before SGMA became law. Estimates also adjust for surface water supplies and distance to the boundary and are weighted by area and a triangular kernel in distance to boundary. Sample is limited to fields within 15 km of the boundary. Vertical bars denote 95% confidence intervals. Standard errors clustered by GSA.

30

ence in new well construction relative to 2014, suggesting that farmers are not responding to greater future pumping restrictions by investing in new irrigation wells.

6.3 Water Use Intensity

Finally, to measure whether future regulation leads farmers to increase or decrease their rate of groundwater extraction before the regulation binds, we turn to our index of crop water use intensity. We again first plot changes in water use intensity between basins facings greater and fewer future pumping reductions in Figure 9 and then show coefficient estimates from the estimation of equation 10, expressed relative to 2014, in Figure 10. Figure 9 suggests that basins trended similarly prior to the announcement of SGMA, with regions that were facing more pumping restrictions on average having higher water requirements. Figure 10 shows formally that there were no statistically significant differences in crop water use intensity before or after the announcement of the regulation. We fail to find evidence that farmers are altering water use in anticipation of future groundwater restrictions.

6.4 Pooled Regressions

To quantify our results, we report estimates of equation 9 in Table 2. These regressions pool together years in the pre- and post-treatment years, providing an overall average difference-in-difference estimate. They potentially improve statistical power over any single year's estimate in the event study.

Looking at the coefficient of interest in the top row, estimates for all variables are small, with standard errors that cannot reject a zero effect. For new perennial plantings, the point estimate is 0.3 percentage points per year, which is relatively small compared with the sample mean value of new perennial plantings (1.1 percentage points per year). Recall that the average difference in future reductions between neighbors represented by the "More Reductions" treatment variable is about the same as the sample average value of future reductions, so we can interpret its effect as the effect of SGMA overall without further scaling.

Estimates for other two outcome variables are considerably more precise. For new well construction, we can reject an anticipatory response in either direction of 0.006 per square mile per year. This value is small compared with the sample mean of 0.043 per square mile per year. For water use intensity, we can reject an anticipatory increase of 0.06 AF/acre or an anticipatory decrease of 0.02 AF/acre per year, again small compared with the sample mean value of 2.9 AF/acre.

Figure 9: Water Use Intensity by Treatment Status, Paired Sample
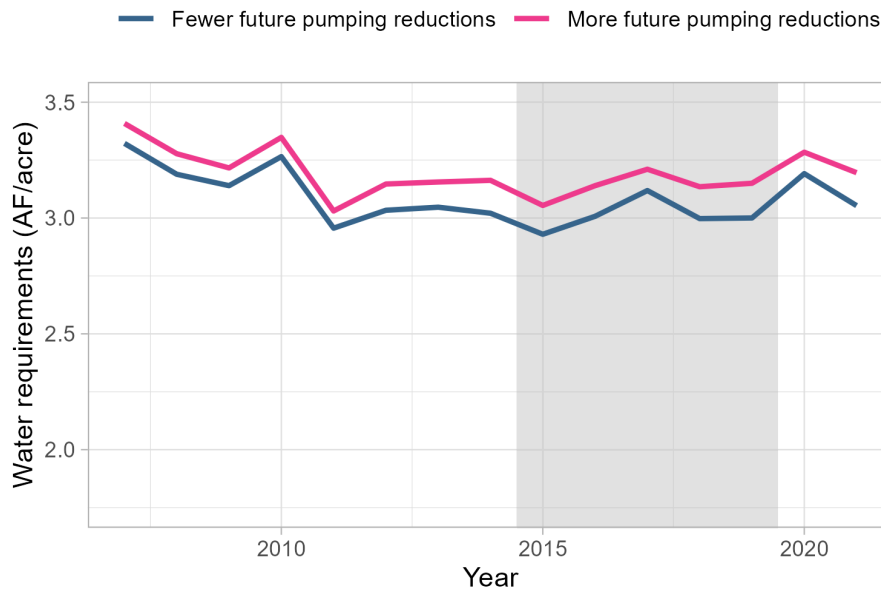
Note: Figure plots mean water-use intensity in the paired sample, which stacks all neighboring subbasins and includes only observations within 15 km of their boundary. "More" and "Fewer" are within these pairs, relative to neighbors in the same year. Water-use intensity is estimated by combining remote sensing land use data with scientific estimates of crop-specific water use. Gray shading indicates the "coordination" period between when SGMA was passed and when local sustainability plans were published. Means weighted by area.

Figure 10: Effect of Greater Future Reductions on Water Use Intensity
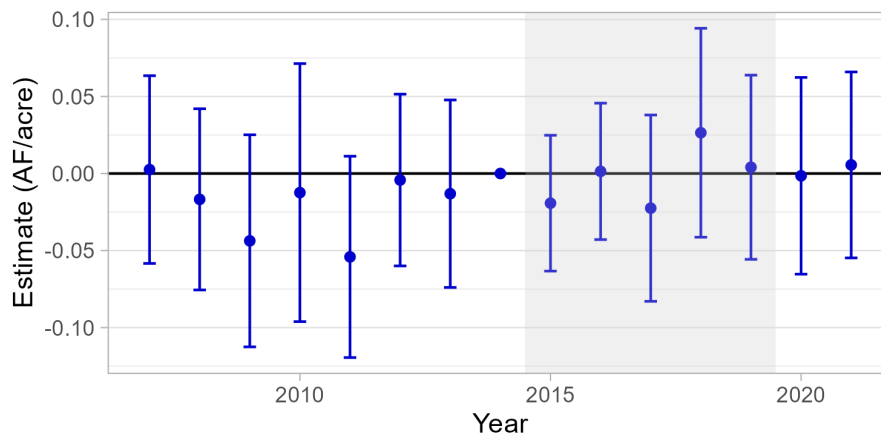
Note: Figure plots year-specific coefficients from the estimation of Equation 10. Each coefficient represents the difference in water-use intensity between GSAs facing more or fewer future pumping restrictions (within each pair of neighboring subbasins) in that year, minus the same difference in 2014, the last year of planting decisions before SGMA became law. Estimates also adjust for surface water supplies and distance to the boundary and are weighted by area and a triangular kernel in distance to boundary. Sample is limited to fields within 15 km of the boundary. Vertical bars denote 95% confidence intervals. Standard errors clustered by GSA.

Table 2: Paired Difference-in-Difference Regression Estimates

	New Perennial Plantings (share)	New Well Construction (per sq. mile)	Water Use Intensity (AF/acre)
	(1)	(2)	(3)
More Reductions × Post	0.003	0.000	0.020
	(0.002)	(0.003)	(0.027)
More Reductions × Middle	0.003	0.000	0.016
	(0.002)	(0.003)	(0.018)
More Reductions	0.000	0.005	0.134
	(0.001)	(0.003)	(0.078)
Distance to boundary	✓	✓	✓
Distance to boundary × More Reductions	✓	✓	✓
Year-Subbasin Pair-Boundary Segment FE	✓	✓	✓
Year FE × Surface water supplies	✓	✓	✓
Year FE × Lagged surface water	✓	✓	✓
Observations	6,242,234	13,924,988	7,202,578
Clusters	104	104	104

Notes: Table reports regression estimates of Equation 9 in the paired sample, which includes all observations within 15 km of the boundary between pairs of neighboring subbasins, with all such pairs stacked into one dataset. Observations are fields or units of land, most commonly quarter quarter sections, about 40 acres, from the Public Land Survey System, per year. "More Reductions" is a binary indicator for whether the field lies in the subbasin with greater expected future pumping reductions under SGMA than its neighbor, within each pair. "Post" is a binary indicator for the post-treatment period after future reductions under SGMA became clearer (2020-22); "Middle" is a binary indicator for the coordination period (2014-19 for wells and 2015-19 for the other variables) after SGMA passed. Water use intensity is an index constructed from remotely sensed land use data and scientific estimates of crop-specific water use. Well construction is drawn from required state reports. Perennial crops are observed from remotely sensed land use data. Data begin in 1993 for wells and 2007 for the other outcome variables. Observations are weighted by area and a triangular kernel in distance to boundary. Standard errors (in parentheses) are clustered by GSA.

6.5 Robustness

Before concluding that future groundwater restrictions under SGMA are not altering behavior around extraction or investment, we want to ensure that our null results are not driven by some of the specific choices we made in our data processing and regression analysis. We explore the sensitivity of our results to different controls, ways of measuring treatment status, and to alternative sample definitions. Figure 11 plots estimates of the overall difference-in-difference coefficient from equation 9, for each of our three outcome variables, for a range of modifications to the baseline specification.

The first row presents our base (preferred) specification, which corresponds to the estimates in Table 2. Recall that our baseline treatment variable is a binary indicator for whether or not a field is within a GSA facing greater future restrictions than its neighbor. In the second row, we instead use the continuous measure of expected future pumping restrictions, in units of AF/acre, as described in Section 4.1. Coefficient estimates here are in different units; they give the change in outcomes due to a one AF/acre increase in overdraft in the years following the announcement of the policy. Quantitatively, they cannot be directly compared to the base specification, so the fact that the confidence intervals are wider does not mean that the estimates are less precise. Directionally, they tell a similar story: we do not see evidence that future pumping reductions affect present decisions.

Next, we narrow our analysis to a comparison that is a priori more likely to respond more strongly to groundwater restrictions: areas outside of surface water districts. These regions are solely dependent on groundwater, so a given reduction in pumping constitutes a greater share of their total water use. They also may be asked to shoulder a greater share of the pumping reductions within a given GSA, since they may be responsible for a greater share of groundwater extraction in the past and present. Still, in row 3, we restrict our sample to only areas outside of water districts, but find similar results across all three outcome variables.

We next vary the bandwidth used to construct the paired sample, which restricted our sample to observations within 15 km of the boundary between neighboring subbasins. Larger bandwidths allow us to include more data and improve precision, but smaller bandwidths can reduce concerns about omitted variables. In the next four rows, we report estimates from constructing the paired sample using four alternative bandwidths: 5, 10, 25, and 50 km.[18] Marginally significant positive results are seen in new perennial plant-

---

[18]Although optimal bandwidths can be calculated in a basic RD setting, it is not straightforward to do so while incorporating a pre/post difference, spatial correlation, a multidimensional cutoff, and pooling across subbasin comparisons.

Figure 11: Robustness of Treatment Effects
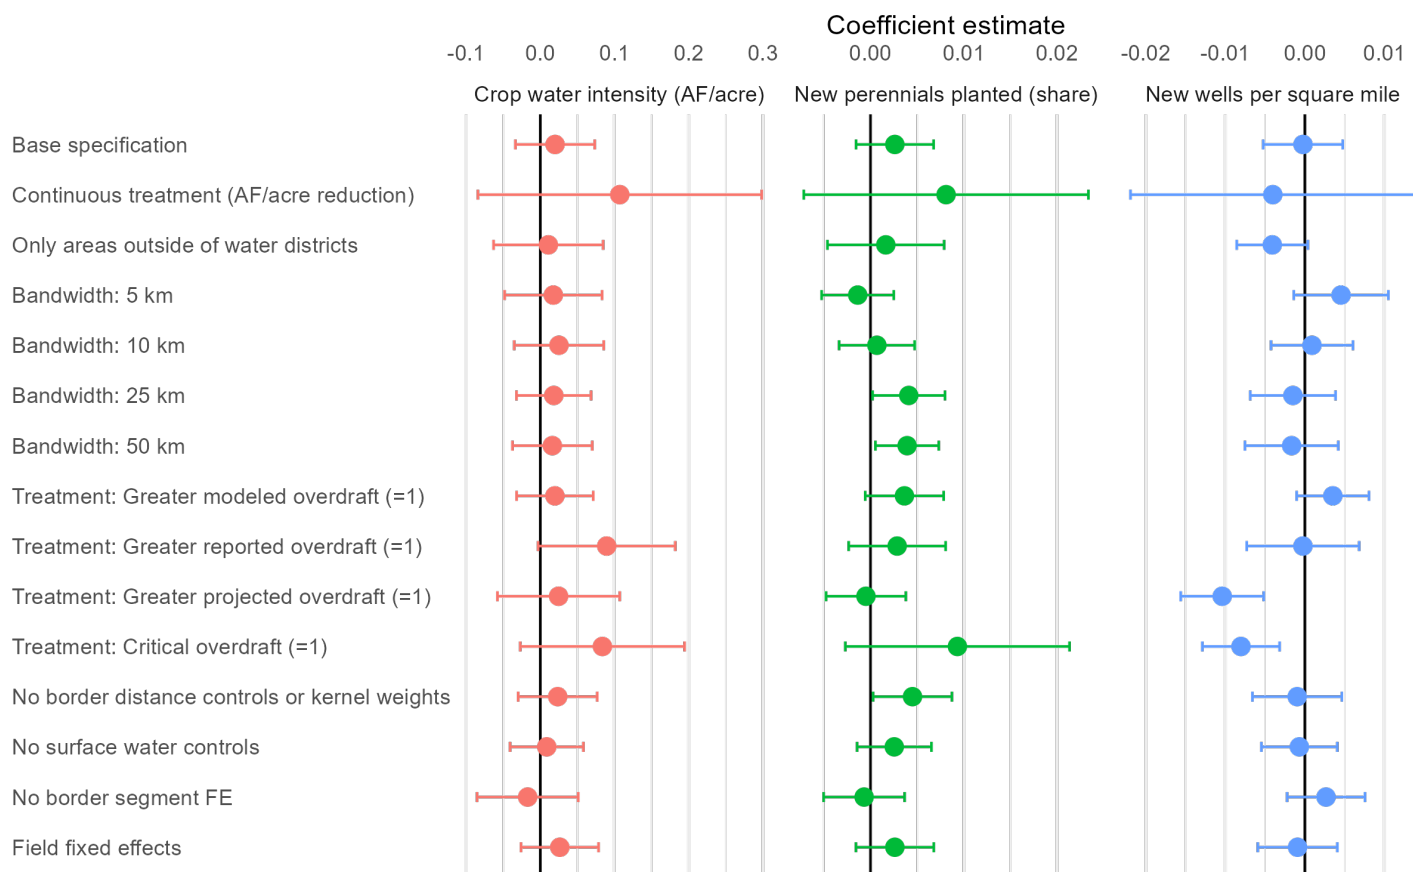
Note: Figure reports difference-in-difference estimates from equation 9, pooled across years in the post-treatment and pre-treatment periods. Each row presents results from a different regression specification for all three outcome variables. Horizontal bars denote 95% confidence intervals.

ings for larger bandwidths, but they are negative for smaller bandwidths. Overall, we find that varying the bandwidth used to construct the sample does not change the conclusions of null results across outcome variables.

We next check to see if our results are sensitive to the choice of treatment variable. Recall that our preferred treatment variable was derived from an average across three proxy variables: modeled overdraft, reported overdraft, and projected overdraft. Rows 8-10 show results with alternative treatment variables that instead use each of these proxies individually. While the pooled treatment effect on new well construction appears to vary by the choice of proxy, results for other outcomes variables are stable to this choice. The story is similar with one additional variation on our binary treatment variable, which considers basins that are deemed by the state to be in conditions of critical overdraft. While two estimates here are statistically significant, further investigation (not shown here) reveals that they in turn fail to survive minor specification changes and do not show patterns of heterogeneity that align with theory. We also note that after conducting many null hypothesis tests, we should expect a few to be statistically significant; otherwise our confidence intervals would be too wide.

A final set of robustness checks relates to our choice of control variables and the inclusion or exclusion of various fixed effects. We test the sensitivity of our baseline results to the exclusion of (a) border distance control and kernel weights, (b) surface water controls, and (c) border segment fixed effects, and (d) to the inclusion of field fixed effects. Across alternative specifications, results are consistent: statistically and economically insignificant effects in the post-SGMA period, estimated with similar magnitudes and precision to results in the main table.

7   Discussion

The precisely estimated zero effects of future pumping restrictions on new perennial plantings, new well construction, and crop water intensity suggest that the policy is not yet altering extraction or investment in water-intensive production technologies. Across the board, we find that null effects are robust across specifications for all outcome variables, with no detectable heterogeneity. No consistent pattern emerges across this large swath of alternative specifications.

To interpret these empirical results, we turn back to the theoretical model. Our theoretical model showed how both investment and net extraction (the effect of future regulation on current extraction through all channels) changed with beliefs about future water supply. Under certain conditions, countervailing Green Paradox and early-decline effects might

cancel out, leaving zero effects on net. But in fact, we can rule out this possibility, because we are able to look at effects on both extraction and investment. The conditions for zero effects are different for different outcomes ($bN = ps$ for investment and $bN = ps\xi$ for extraction), and they cannot be true simultaneously. This leaves us with the remaining explanation: that a high value of farmers' private discount rate deflates away considerations that are at least 10 to 20 years down the road, leaving the effects small.

Another potential explanation for null effects could be that farmers' true beliefs about future regulation are smaller than what our measures are capturing. This would manifest in our model as an attenuation bias from underestimating the change in $\bar{y}$. This could be due to either lack of salience – perhaps landowners lack information – or they may have low confidence in the enforcement of the regulation. If farmers perceive the future restrictions to be small or unenforceable, then any effects on current extraction may be too small to detect. There is no one clear way of knowing what the future regulation will be. But given that the state operates as a backstop for non-compliant GSAs, public outreach was codified into the law, and SGMA has dominated local news headlines about water since its passing, we consider this a less likely explanation.

8 Conclusion

This paper studies whether producers respond to future groundwater regulation by changing groundwater extraction or investing in long-term agricultural capital like planting perennial crops and constructing new irrigation wells. Our theoretical model shows formally that a Green Paradox can occur for groundwater, but that it is unlikely in conditions of open access. Allowing for investment opportunities like adopting water-intensive production technology – a main mechanism for farmers to increase groundwater use – complicates the story and allows for the possibility of an anticipatory decline in extraction. Our model generates testable scenarios that we take to data on California's agricultural groundwater.

Empirically, we evaluate the early effects of California's Sustainable Groundwater Management Act of 2014, a sweeping groundwater regulation that is affecting over 95% of the agricultural groundwater pumping in the state. The regulation is particularly remarkable given the fact that groundwater use was largely open access prior to its passing. The policy required groundwater agencies to establish sustainable pumping criteria and develop plans for how to achieve that over the next two decades. The decentralized nature of the mandate led to large variation in expected future pumping restrictions across the state, creating a policy experiment to study questions about anticipatory behavior.

Our analysis uses spatial land use data for all agricultural parcels subject to the legislation and estimates how groundwater extraction and farmland investments responded to changes in future pumping access. Although investments in perennial crops have increased by nearly 50% since SGMA passed at the end of 2014, we find that this boom occurred despite, not because of, the policy. Likewise, when comparing within pairs of neighboring subbasins that face greater and lesser future pumping restrictions, we find no evidence of changes in water use intensity or new well construction in basins facing greater future pumping restrictions. Our theoretical model suggests that the most likely explanation for our findings – that SGMA is not yet altering extraction or investment – is a high private discount rate that diminishes the importance of future regulation and shrinks both types of anticipatory motives.

References

Arellano-Gonzalez, J. and Moore, F. C. (2020). Intertemporal arbitrage of water and long-term agricultural investments: Drought, groundwater banking, and perennial cropping decisions in California. American Journal of Agricultural Economics, 102(5):1368–1382.

Baker, A., Larcker, D. F., and Wang, C. C. (2021). How much should we trust staggered difference-in-differences estimates? Working Paper SSRN 3794018.

Bruno, E., Hadachek, J., Hagerty, N., and Jessoe, K. (2023a). External costs of climate change adaptation: Agricultural wells and access to drinking water. Working Paper.

Bruno, E. M., Hagerty, N., and Wardle, A. R. (2023b). The political economy of groundwater management: Descriptive evidence from California. In Libecap, G. and Dinar, A., editors, American Agriculture, Water Resources, and Climate Change. University of Chicago Press.

Bruno, E. M. and Jessoe, K. (2021). Missing markets: Evidence on agricultural groundwater demand from volumetric pricing. Journal of Public Economics, 196:104374.

Cairns, R. D. (2014). The green paradox of the economics of exhaustible resources. Energy Policy, 65:78–85.

Cattaneo, M. D., Idrobo, N., and Titiunik, R. (2019). A practical introduction to regression discontinuity designs: Foundations. Cambridge University Press.

Di Maria, C., Lange, I., and Van der Werf, E. (2014). Should we be worried about the green paradox? Announcement effects of the acid rain program. European Economic Review, 69:143–162.

Di Maria, C., Smulders, S., and Van der Werf, E. (2012). Absolute abundance and relative scarcity: Environmental policy with implementation lags. Ecological Economics, 74:104–119.

Drysdale, K. M. and Hendricks, N. P. (2018). Adaptation to an irrigation water restriction imposed through local governance. Journal of Environmental Economics and Management, 91:150–165.

Hagerty, N. (2021). Adaptation to surface water scarcity in irrigated agriculture. Working paper.

Heal, G. (1976). The relationship between price and extraction cost for a resource with a backstop technology. The Bell Journal of Economics, pages 371–378.

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

Hotelling, H. (1931). The economics of exhaustible resources. Journal of Political Economy, 39(2):137–175.

Jensen, S., Mohlin, K., Pittel, K., and Sterner, T. (2020). An introduction to the green paradox: The unintended consequences of climate policies. Review of Environmental Economics and Policy.

Lemoine, D. (2017). Green expectations: Current effects of anticipated carbon pricing. Review of Economics and Statistics, 99(3):499–513.

List, J. A., Margolis, M., and Osgood, D. (2006). Is the endangered species act endangering species? NBER Working Paper 12777.

Lobell, D. B. and Field, C. B. (2011). California perennial crops in a changing climate. Climatic Change, 109(1):317–333.

Mall, N. K. and Herman, J. D. (2019). Water shortage risks from perennial crop expansion in California's Central Valley. Environmental Research Letters, 14(10):104014.

McDermott, G. R., Meng, K. C., McDonald, G. G., and Costello, C. J. (2019). The blue paradox: Preemptive overfishing in marine reserves. Proceedings of the National Academy of Sciences, 116(12):5319–5325.

Sinn, H.-W. (1982). Taxation, growth, and resource extraction: A general equilibrium approach. European Economic Review, 19(2):357–386.

Sinn, H.-W. (2008). Public policies against global warming: A supply side approach. International tax and public finance, 15:360–394.

Smith, S. M., Andersson, K., Cody, K. C., Cox, M., and Ficklin, D. (2017). Responding to a groundwater crisis: The effects of self-imposed economic incentives. Journal of the Association of Environmental and Resource Economists, 4(4):985–1023.

Tsur, Y. and Graham-Tomasi, T. (1991). The buffer value of groundwater with stochastic surface water supplies. Journal of Environmental Economics and Management, 21(3):201–224.

Van der Ploeg, F. and Withagen, C. (2020). Global warming and the green paradox: A review of adverse effects of climate policies. Review of Environmental Economics and Policy.

Wada, Y., Van Beek, L. P., Van Kempen, C. M., Reckman, J. W., Vasak, S., and Bierkens, M. F. (2010). Global depletion of groundwater resources. Geophysical Research Letters, 37(20).

A Appendix: Proofs

A.1 Proof of Equations 2 and 3

The Lagrangian of this problem is:

$$\mathscr{L} = \sum_{t=0}^{\infty}(1+r)^{-t}\left[B(y_{it}) - c(x_{it})y_{it}\right] + \sum_{t=0}^{\infty}\mu_{it}\left[x_{it} + g - \frac{1}{N}\sum_{j=0}^{N}y_{jt} - x_{i,t+1}\right].$$

Its first-order conditions are:

$$
\begin{aligned}
y_{it} &: \quad (1+r)^{-t}\left[B'(y_{it}) - c(x_{it})\right] - \tfrac{1}{N}\mu_{it} = 0 \quad \forall i,t \\
x_{it} &: \quad -(1+r)^{-t}c'(x_{it})y_{it} + \mu_{it} - \mu_{i,t-1} = 0 \quad \forall i,t > 0 \\
\mu_{it} &: \quad x_{it} + g - \tfrac{1}{N}\sum_{j=0}^{N}y_{jt} - x_{i,t+1} = 0 \quad \forall i,t.
\end{aligned}
$$

Rearranging the first condition reveals Equation 2. To obtain Equation 3, we can substitute the first-order conditions for y_{it} and $y_{i,t-1}$ into the one for x_{it} and rearrange:

$$
\begin{aligned}
-(1+r)^{-t}c'(x_{it})y_{it} &= \mu_{i,t-1} - \mu_{it} \\
-(1+r)^{-t}c'(x_{it})y_{it} &= (1+r)^{-(t-1)}\left[B'(y_{i,t-1}) - c(x_{i,t-1})\right]N - (1+r)^{-t}\left[B'(y_{it}) - c(x_{it})\right]N \\
-\frac{1}{N}c'(x_{it})y_{it} &= (1+r)\left[B'(y_{i,t-1}) - c(x_{i,t-1})\right] - B'(y_{it}) + c(x_{it}) \\
(1+r)\left[B'(y_{i,t-1}) - c(x_{i,t-1})\right] &= B'(y_{it}) - c(x_{it}) + \frac{1}{N}(-c'(x_{it}))y_{it} \\
B'(y_{i,t-1}) - c(x_{i,t-1}) &= (1+r)^{-1}\left[B'(y_{it}) - c(x_{it})\right] + (1+r)^{-1}\frac{1}{N}(-c'(x_{it}))y_{it} \\
B'(y_{it}) - c(x_{it}) &= (1+r)^{-1}\left[B'(y_{i,t+1}) - c(x_{i,t+1})\right] + (1+r)^{-1}\frac{1}{N}[-c'(x_{i,t+1})]y_{i,t+1}.
\end{aligned}
$$

A.2 Proof of Proposition 1

Starting with the Lagrangian from before, expanding sums, and substituting the assumptions $y_{i1} = \bar{y}$ and $x_{it} = x_{i2}$ and $y_{it} = g$ for all $t \geq 2$:

$$
\begin{aligned}
\mathscr{L} &= \sum_{t=0}^{\infty}(1+r)^{-t}\left[B(y_{it}) - c(x_{it})y_{it}\right] + \sum_{t=0}^{\infty}\mu_{it}\left[x_{it} + g - \frac{1}{N}\sum_{j=0}^{N}y_{jt} - x_{i,t+1}\right] \\
&= B(y_{i0}) - c(x_{i0})y_{i0} + (1+r)^{-1}\left[B(y_{i1}) - c(x_{i1})y_{i1}\right] + \sum_{t=2}^{\infty}(1+r)^{-t}\left[B(g) - c(x_{i2})g\right] + \\
&\quad \mu_{i0}\left[x_{i0} + g - \frac{1}{N}\sum_{j=0}^{N}y_{j0} - x_{i1}\right] + \mu_{i1}\left[x_{i1} + g - \bar{y} - x_{i2}\right] + \sum_{t=2}^{\infty}\mu_{it}\left[x_{i2} - x_{i2}\right] \\
&= B(y_{i0}) - c(x_{i0})y_{i0} + (1+r)^{-1}\left[B(\bar{y}) - c(x_{i1})\bar{y}\right] + (1+r)^{-1}\frac{1}{r}\left[B(g) - c(x_{i2})g\right] + \\
&\quad \mu_{i0}\left[x_{i0} + g - \frac{1}{N}\sum_{j=0}^{N}y_{j0} - x_{i1}\right] + \mu_{i1}\left[x_{i1} + g - \bar{y} - x_{i2}\right].
\end{aligned}
$$

The third equality uses the fact that $\sum_{t=1}^{\infty}(1+r)^{-t} = r^{-1}$ and therefore $\sum_{t=2}^{\infty}(1+r)^{-t} = r^{-1}(1+r)^{-1}$, through either substitution or a change of variables.

The first-order conditions of this new Lagrangian are:

$$
\begin{aligned}
y_{i0} : \quad 0 &= B'(y_{i0}) - c(x_{i0}) - \frac{1}{N}\mu_{i0} \\
x_{i1} : \quad 0 &= -(1+r)^{-1}c'(x_{i1})\bar{y} - \mu_{i0} + \mu_{i1} \\
x_{i2} : \quad 0 &= -(1+r)^{-1}\frac{1}{r}c'(x_{i2})g - \mu_{i1}
\end{aligned}
$$

and the Euler equation is:

$$
\begin{aligned}
\mu_{i0} &= \mu_{i1} - (1+r)^{-1}c'(x_{i1})\bar{y} \\
N\left[B'(y_{i0}) - c(x_{i0})\right] &= -(1+r)^{-1}\frac{1}{r}c'(x_{i2})g - (1+r)^{-1}c'(x_{i1})\bar{y} \\
B'(y_{i0}) - c(x_{i0}) &= -(1+r)^{-1}\frac{1}{N}\left[\frac{1}{r}c'(x_{i2})g + c'(x_{i1})\bar{y}\right] \\
B'(y_{i0}) - \gamma + psx_{i0} &= -(1+r)^{-1}\frac{1}{N}\left[\frac{1}{r}(-ps)g + (-ps)\bar{y}\right] \\
B'(y_{i0}) - \gamma + psx_{i0} &= (1+r)^{-1}\frac{1}{N}ps\left[\frac{1}{r}g + \bar{y}\right].
\end{aligned}
$$

Using the Implicit Function Theorem:

$$
\begin{aligned}
G &:= B'(y_{i0}) - c(x_{i0}) - (1+r)^{-1}\frac{1}{N}ps\left[\frac{1}{r}g + \bar{y}\right] = 0 \\
\frac{\partial G}{\partial y_{i0}} &= B''(y_{i0}) \\
\frac{\partial G}{\partial \bar{y}} &= -(1+r)^{-1}\frac{1}{N}ps \\
\frac{dy_{i0}}{d\bar{y}} &= -\frac{\partial G/\partial \bar{y}}{\partial G/\partial y_{i0}} = \frac{ps}{(1+r)NB''(y_{i0})}.
\end{aligned}
$$

$B(y)$ is concave, so $B''(y)$ is negative, and $\{p,s,r,N\}$ are all positive, so this derivative is always negative.

A.3   Proof of Lemma 1

Starting with the Euler equation above and taking other users' investment decisions as given, we substitute in the benefit function parameterization for each investment choice:

$$
\begin{aligned}
a - by_{i0}^0 - \gamma + psx_{i0} &= \frac{1}{N}(1+r)^{-1}ps\left[\frac{1}{r}g + \bar{y}\right] \\
a + \beta - by_{i0}^I - \gamma + psx_{i0} &= \frac{1}{N}(1+r)^{-1}ps\left[\frac{1}{r}g + \bar{y}\right]
\end{aligned}
$$

Substituting these equations to find $y_{i0}^I - y_{i0}^0$:

$$
\begin{aligned}
-by_{i0}^0 &= \beta - by_{i0}^I \\
by_{i0}^I - by_{i0}^0 &= \beta \\
y_{i0}^I - y_{i0}^0 &= \beta/b.
\end{aligned}
$$

This expression is always positive, since both β and b are positive.

A.4   Proof of Proposition 2

The probability of investment is $I_i = \Pr(K_i \leq \Theta_i) = F_K(\Theta_i)$, and the probability density function is defined as $f_K(\Theta_i) := dF_K(\Theta_i)/d\Theta_i$. Applying the Chain Rule:

$$
\frac{dI_i}{d\bar{y}} = \frac{dF_K(\Theta_i)}{d\bar{y}} = \frac{dF_K(\Theta_i)}{d\Theta_i}\frac{d\Theta_i}{d\bar{y}} = f_K(\Theta_i)\frac{d\Theta_i}{d\bar{y}}.
$$

The remaining task is to find $d\Theta_i/d\bar{y}$.

The return on investment Θ_i is defined as:

$$\Theta_i \; := \; \sum_{t=0}^{\infty}(1+r)^{-t}\Big[B_I(y_{it}^I)-c(x_{it}^I)y_{it}^I\Big]-\sum_{t=0}^{\infty}(1+r)^{-t}\Big[B_0(y_{it}^0)-c(x_{it}^0)y_{it}^0\Big].$$

Substituting in the cost function parameterization and rearranging:

$$
\begin{aligned}
\Theta_i \; &= \; \sum_{t=0}^{\infty}(1+r)^{-t}\Big[B_I(y_{it}^I)-B_0(y_{it}^0)-c(x_{it}^I)y_{it}^I+c(x_{it}^0)y_{it}^0\Big]\\
&= \; \sum_{t=0}^{\infty}(1+r)^{-t}\Big[B_I(y_{it}^I)-B_0(y_{it}^0)-(\gamma-psx_{it}^I)y_{it}^I+(\gamma-psx_{it}^0)y_{it}^0\Big]\\
&= \; \sum_{t=0}^{\infty}(1+r)^{-t}\Big[\big(B_I(y_{it}^I)-B_0(y_{it}^0)\big)-(y_{it}^I-y_{it}^0)\gamma+(x_{it}^I y_{it}^I-x_{it}^0 y_{it}^0)ps\Big].
\end{aligned}
$$

Expanding the sum to 3 periods:

$$
\begin{aligned}
\Theta_i \; = \; & \Big[\big(B_I(y_{i0}^I)-B_0(y_{i0}^0)\big)-(y_{i0}^I-y_{i0}^0)\gamma+(x_{i0}^I y_{i0}^I-x_{i0}^0 y_{i0}^0)ps\Big]+\\
& (1+r)^{-1}\Big[\big(B_I(y_{i1}^I)-B_0(y_{i1}^0)\big)-(y_{i1}^I-y_{i1}^0)\gamma+(x_{i1}^I y_{i1}^I-x_{i1}^0 y_{i1}^0)ps\Big]+\\
& \sum_{t=2}^{\infty}(1+r)^{-t}\Big[\big(B_I(y_{i2}^I)-B_0(y_{i2}^0)\big)-(y_{i2}^I-y_{i2}^0)\gamma+(x_{i2}^I y_{i2}^I-x_{i2}^0 y_{i2}^0)ps\Big].
\end{aligned}
$$

Substituting in $y_{i1}=\bar{y}$ and $y_{it}=g$ for $t\geq 2$:

$$
\begin{aligned}
\Theta_i \; = \; & \Big[\big(B_I(y_{i0}^I)-B_0(y_{i0}^0)\big)-(y_{i0}^I-y_{i0}^0)\gamma+(y_{i0}^I-y_{i0}^0)x_{i0}ps\Big]+\\
& (1+r)^{-1}\Big[\big(B_I(y_{i1}^I)-B_0(y_{i1}^0)\big)-(\bar{y}-\bar{y})\gamma+(x_{i1}^I\bar{y}-x_{i1}^0\bar{y})ps\Big]+\\
& \sum_{t=2}^{\infty}(1+r)^{-t}\Big[\big(B_I(y_{i2}^I)-B_0(y_{i2}^0)\big)-(g-g)\gamma+(x_{i2}^I g-x_{i2}^0 g)ps\Big]\\
= \; & \Big[\big(B_I(y_{i0}^I)-B_0(y_{i0}^0)\big)-(y_{i0}^I-y_{i0}^0)\gamma+(y_{i0}^I-y_{i0}^0)x_{i0}ps\Big]+\\
& (1+r)^{-1}\Big[\big(B_I(y_{i1}^I)-B_0(y_{i1}^0)\big)+(x_{i1}^I-x_{i1}^0)\bar{y}ps\Big]+\\
& \sum_{t=2}^{\infty}(1+r)^{-t}\Big[\big(B_I(y_{i2}^I)-B_0(y_{i2}^0)\big)+(x_{i2}^I-x_{i2}^0)gps\Big].
\end{aligned}
$$

Substituting in the equations of motion, holding constant the extraction choices of other

users:

$$
\begin{aligned}
\Theta_i &= \left[\left(B_I(y_{i0}^I) - B_0(y_{i0}^0) \right) - (y_{i0}^I - y_{i0}^0)\gamma + (y_{i0}^I - y_{i0}^0)x_{i0}ps \right] + \\
&\quad (1+r)^{-1}\left[\left(B_I(y_{i1}^I) - B_0(y_{i1}^0) \right) - \frac{1}{N}(y_{i0}^I - y_{i0}^0)\bar{y}ps \right] + \\
&\quad \sum_{t=2}^{\infty}(1+r)^{-t}\left[\left(B_I(y_{i2}^I) - B_0(y_{i2}^0) \right) - \frac{1}{N}(y_{i0}^I - y_{i0}^0)gps \right] \\
&= \left(B_I(y_{i0}^I) - B_0(y_{i0}^0) \right) + (y_{i0}^I - y_{i0}^0)(x_{i0}ps - \gamma) + \\
&\quad \left(B_I(\bar{y}) - B_0(\bar{y}) \right)(1+r)^{-1} - (y_{i0}^I - y_{i0}^0)\bar{y}\frac{1}{N}ps(1+r)^{-1} + \\
&\quad \left(B_I(g) - B_0(g) \right)r^{-1}(1+r)^{-1} - (y_{i0}^I - y_{i0}^0)\frac{1}{N}gpsr^{-1}(1+r)^{-1}.
\end{aligned}
$$

How does Θ depend on the period-1 quantity limits? Taking the derivative with respect to $\bar{y}$:

$$
\begin{aligned}
\frac{d\Theta_i}{d\bar{y}} &= B_I'(y_{i0}^I)\frac{dy_{i0}^I}{d\bar{y}} - B_0'(y_{i0}^0)\frac{dy_{i0}^0}{d\bar{y}} + \left(\frac{dy_{i0}^I}{d\bar{y}} - \frac{dy_{i0}^0}{d\bar{y}}\right)(x_{i0}ps - \gamma) + \\
&\quad \left(B_I'(\bar{y}) - B_0'(\bar{y}) \right)(1+r)^{-1} - \left(\frac{dy_{i0}^I}{d\bar{y}} - \frac{dy_{i0}^0}{d\bar{y}}\right)\bar{y}\frac{1}{N}ps(1+r)^{-1} + \\
&\quad -(y_{i0}^I - y_{i0}^0)\frac{1}{N}ps(1+r)^{-1} - \left(\frac{dy_{i0}^I}{d\bar{y}} - \frac{dy_{i0}^0}{d\bar{y}}\right)\frac{1}{N}gpsr^{-1}(1+r)^{-1} \\
&= B_I'(y_{i0}^I)\frac{dy_{i0}^I}{d\bar{y}} - B_0'(y_{i0}^0)\frac{dy_{i0}^0}{d\bar{y}} + \\
&\quad \left(\frac{dy_{i0}^I}{d\bar{y}} - \frac{dy_{i0}^0}{d\bar{y}}\right)\left(x_{i0}ps - \gamma - \frac{1}{N}ps(1+r)^{-1}(\bar{y} + gr^{-1})\right) + \\
&\quad \left(B_I'(\bar{y}) - B_0'(\bar{y}) \right)(1+r)^{-1} - (y_{i0}^I - y_{i0}^0)\frac{1}{N}ps(1+r)^{-1}.
\end{aligned}
$$

We know $dy_{i0}/d\bar{y}$ from Proposition 1. Plugging in equation 4:

$$
\begin{aligned}
\frac{d\Theta_i}{d\bar{y}} &= B_I'(y_{i0}^I)\frac{ps}{N(1+r)B_I''(y_{i0}^I)} - B_0'(y_{i0}^0)\frac{ps}{N(1+r)B_0''(y_{i0}^0)} + \\
&\quad (\frac{ps}{N(1+r)B_I''(y_{i0}^I)} - \frac{ps}{N(1+r)B_0''(y_{i0}^0)})(x_{i0}ps - \gamma - \frac{1}{N}ps(1+r)^{-1}(\bar{y} + gr^{-1})) + \\
&\quad (B_I'(\bar{y}) - B_0'(\bar{y}))(1+r)^{-1} - (y_{i0}^I - y_{i0}^0)\frac{1}{N}ps(1+r)^{-1} \\
&= \frac{1}{N}(1+r)^{-1}ps\left[\frac{B_I'(y_{i0}^I)}{B_I''(y_{i0}^I)} - \frac{B_0'(y_{i0}^0)}{B_0''(y_{i0}^0)}\right] + \\
&\quad \frac{1}{N}(1+r)^{-2}ps\left[\frac{1}{B_I''(y_{i0}^I)} - \frac{1}{B_0''(y_{i0}^0)}\right](x_{i0}ps - \gamma - \frac{1}{N}ps(1+r)^{-1}(\bar{y} + gr^{-1})) + \\
&\quad (B_I'(\bar{y}) - B_0'(\bar{y}))(1+r)^{-1} - (y_{i0}^I - y_{i0}^0)\frac{1}{N}ps(1+r)^{-1}.
\end{aligned}
$$

Substituting in the parameterized benefit functions and equation 5:

$$
\begin{aligned}
\frac{d\Theta_i}{d\bar{y}} &= -\frac{1}{N}(1+r)^{-1}ps\left[\frac{a + \beta - by_{i0}^I}{b} - \frac{a - by_{i0}^0}{b}\right] + \\
&\quad (a + \beta - b\bar{y} - a + b\bar{y})(1+r)^{-1} - (y_{i0}^I - y_{i0}^0)\frac{1}{N}ps(1+r)^{-1} \\
&= -\frac{1}{bN}(1+r)^{-1}ps\left[\beta - b(y_{i0}^I - y_{i0}^0)\right] + \beta(1+r)^{-1} - (y_{i0}^I - y_{i0}^0)\frac{1}{N}ps(1+r)^{-1} \\
&= -\frac{1}{bN}(1+r)^{-1}ps\left[\beta - \beta\right] + \beta(1+r)^{-1} - (y_{i0}^I - y_{i0}^0)\frac{1}{N}ps(1+r)^{-1} \\
&= \beta(1+r)^{-1} - (y_{i0}^I - y_{i0}^0)\frac{1}{N}ps(1+r)^{-1} \\
&= \beta(1+r)^{-1} - \beta\frac{1}{bN}(1+r)^{-1}ps \\
&= \beta(1+r)^{-1}\left[1 - \frac{ps}{bN}\right].
\end{aligned}
$$

Finally, we can plug this expression into the equation at the start of this proof:

$$
\frac{dI_i}{d\bar{y}} = f_K(\Theta_i)\beta(1+r)^{-1}\left[1 - \frac{ps}{bN}\right].
$$

A.5 Proof of Proposition 3

From Proposition 1, and substituting in the benefit function parameterization for either investment decision, we have:

$$\frac{\partial y_{i0}}{\partial \bar{y}} = \frac{ps}{(1+r)NB''(y_{i0})} = -(1+r)^{-1}\frac{ps}{bN}.$$

And from Proposition 2, we have:

$$\left(y_{i0}^I - y_{i0}^0\right)\frac{dI_i}{d\bar{y}} = f_K(\Theta)\frac{\beta^2}{b}(1+r)^{-1}\left[1 - \frac{ps}{bN}\right].$$

Totally differentiating $y_{i0}(\bar{y}, I_i(\bar{y}))$ and substituting in the expressions above:

$$
\begin{aligned}
\frac{dy_{i0}}{d\bar{y}} &= \frac{\partial y_{i0}}{\partial \bar{y}} + \left(y_{i0}^I - y_{i0}^0\right)\frac{dI_i}{d\bar{y}} \\
&= -(1+r)^{-1}\frac{ps}{bN} + f_K(\Theta)\frac{\beta^2}{b}(1+r)^{-1}\left[1 - \frac{ps}{bN}\right] \\
&= (1+r)^{-1}\left[f_K(\Theta)\frac{\beta^2}{b} - \left(1 + f_K(\Theta)\frac{\beta^2}{b}\right)\frac{ps}{bN}\right]. \\
&= (1+r)^{-1}\left(1 + f_K(\Theta)\frac{\beta^2}{b}\right)\left[\frac{f_K(\Theta)\frac{\beta^2}{b}}{1 + f_K(\Theta)\frac{\beta^2}{b}} - \frac{ps}{bN}\right].
\end{aligned}
$$

Defining

$$
\begin{aligned}
\xi &:= \left(f_K(\Theta)\frac{\beta^2}{b}\right)^{-1} + 1 \\
&= \frac{1 + f_K(\Theta)\frac{\beta^2}{b}}{f_K(\Theta)\frac{\beta^2}{b}}
\end{aligned}
$$

and substituting it into the expression above:

$$\frac{dy_{i0}}{d\bar{y}} = (1+r)^{-1}\left(1 + f_K(\Theta)\frac{\beta^2}{b}\right)\left[\xi^{-1} - \frac{ps}{bN}\right].$$

Next, we sign the factors in this expression. All of $\{r, f_K, \beta, b\}$ are positive, so $(1+$

$r)^{-1}\left(1+f_K(\Theta)\frac{\beta^2}{b}\right)$ is positive, and we can ignore it:

$$\text{sign}\left(\frac{dy_{i0}}{d\bar{y}}\right) = \text{sign}\left(\xi^{-1} - \frac{ps}{bN}\right).$$

Therefore, $dy_{i0}/d\bar{y} > 0$ when $\xi^{-1} > \frac{ps}{bN}$, or $bN > ps\xi$. Similarly, $dy_{i0}/d\bar{y} < 0$ when $bN < ps\xi$, and $dy_{i0}/d\bar{y} = 0$ when $bN = ps\xi$.

Finally, we can split the range of bN into the three regimes $\{(-\infty, ps), (ps, ps\xi), (ps\xi, \infty)\}$ because $\{f_K, \beta, b, p, s\}$ are all positive, meaning that $\xi > 1$ and therefore $ps < ps\xi$.