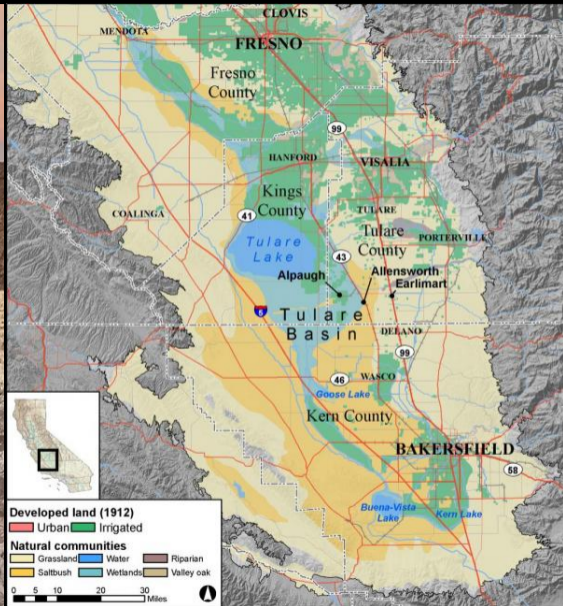


Surface Water Scarcity and Agriculture in the Long Run

Nick Hagerty

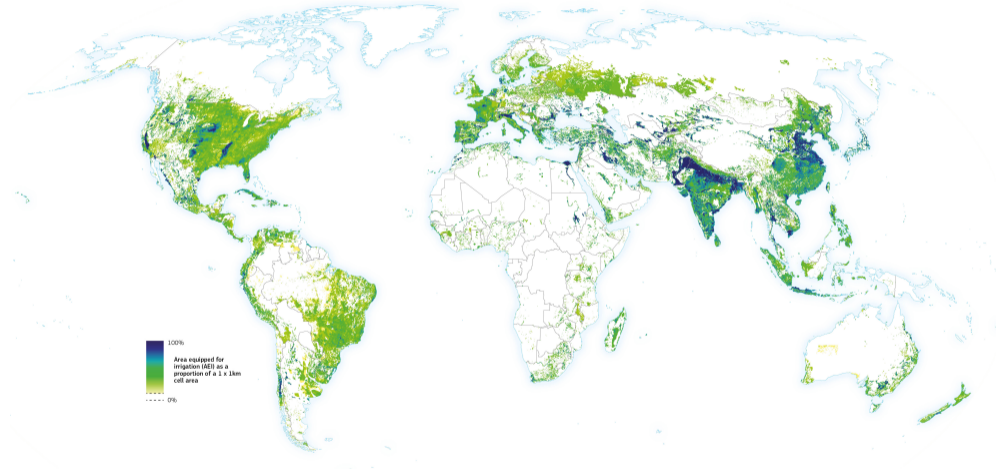
Montana State University

April 2026



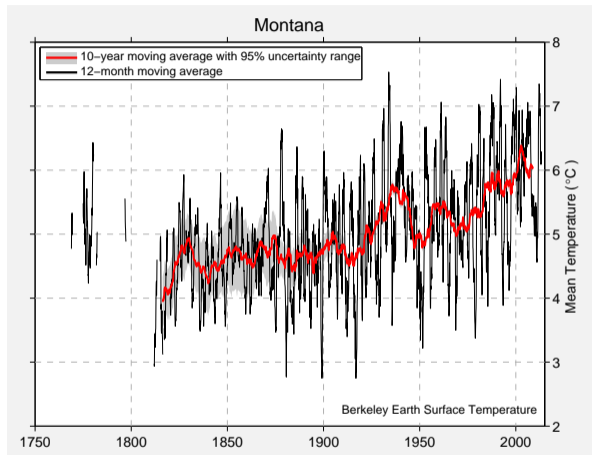
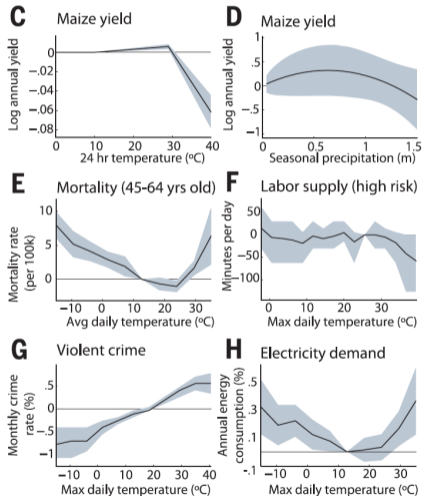
Economic costs of surface water scarcity are uncertain

42% of global crop output relies on irrigation – not just rainfall...



...but so far we don't have enough empirical evidence.

Long-run impacts of environmental change are hard to quantify



Will short-run impacts persist in the long run?

Or will we reduce the costs through adaptation?

Research Question #1


How does surface water affect long-run agricultural production?

Estimate long-run effects using spatial variation from institutional history in California.

- Farm water supplies vary by water district.
- Compare farms across district boundaries using a geographic regression discontinuity.

Long-run effects can inform:

- Water and agricultural policy decisions today.
- Effects of secular changes in future water supplies.



Research Question #2

How do the effects of water scarcity change when farmers can adapt?

- Estimate short-run effects using year-to-year fluctuations in water supplies.
- Compare short-run effects to long-run effects in the same setting.

Effects of adaptation: Difference between short-run and long-run effects.

- Short run? Limited range of ability to respond.
- Long run? Invest in knowledge & equipment to grow new crops.

Results: Water scarcity matters, and output losses persist

- 1. Surface water scarcity has negative long-run effects on agriculture.**
 - Crop area, crop revenue, and land values are lower in places with less water.
- 2. Land-use responses suggest farmers adapt in some ways.**
 - Differing short-run/long-run responses imply constraints relax over time.
- 3. But adaptation does not offset the fall in agricultural output.**
 - Revenue is *more* sensitive to water scarcity in the long run than in the short run.

Surface water supplies are projected to decline over the next few decades...

⇒ **likely reducing crop area, agricultural output, and welfare.**

Contributions to the literature

Spatial discontinuities: new method for long-run effects of environmental variables

- Long-run effects usually difficult to estimate due to omitted variables concerns (Mendelsohn, Nordhaus, Shaw 1994; Deschênes and Greenstone 2007).
- **Spatial discontinuities approach the ideal experiment.** (Other approaches: Dell et al. 2012; Hornbeck 2012; Burke & Emerick 2016; Barreca et al. 2016; Heutel et al. 2018; Carleton et al. 2022).

Updating the value and effects of water in agriculture

- Groundwater (Hornbeck & Keskin 2014; Sekhri 2014; Ryan & Sudarshan 2021; Burlig et al. 2021).
- Extensive margin (Hansen et al. 2011; Ji & Cobourn 2018; Jones et al. 2019).
- Regression adjustment (Mendelsohn & Dinar 2003; Schlenker, Hanemann, Fisher 2007; Olen, Wu, Langpap 2015; Edwards & Smith 2018).
- Short-run variation (Xu et al. 2014; Manning et al. 2017; Khan et al. 2017).

Dynamics of crop supply adjustment

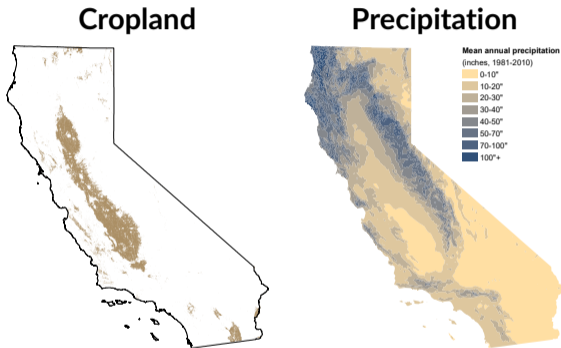
- More adjustment in the long run (Nerlove) or short run (Hendricks et al. 2014)?
- Here: Adjustment to a factor quantity shock is similar or greater in the long run.

Roadmap

1. Introduction
2. **Background**
3. Data
4. Long-Run Effects
5. Adaptation
6. Policy and Conclusion

Agriculture relies on surface water irrigation

Crops need water, but in California it doesn't come from rain.



Farmers instead apply water from:

- **Surface water (rivers & canals): 61%.**
- **Groundwater pumping: 39%.**

Surface water availability determined by water district

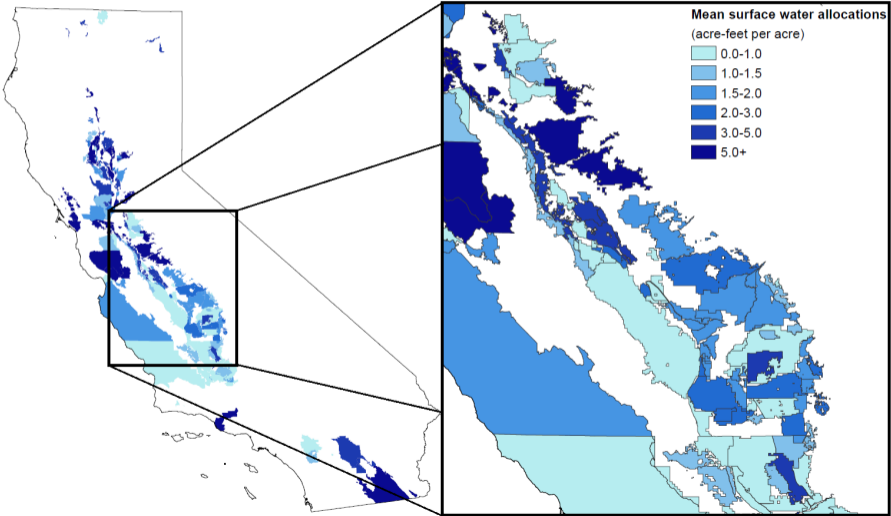
Water districts:

- Local organizations established by groups of farmers, 1860s-1950s.
- Hold long-term entitlements to surface water from:
 - Rivers & streams (permanent rights).
 - Canals run by federal & state governments (long-term contracts).
- Distribute to farmers within their service area, dividing evenly by land area.

⇒ **Spatial distribution of water to farmland has changed little in over 40 years.**

Large differences in average water supplies across districts

Per-acre surface water entitlements by water district



Water supplies fluctuate from year to year

$$Allocation_{dt} = (Maximum\ Entitlement)_d \times (Allocation\ Percentage)_{dt}$$

- **Allocation percentages** are set each year by government agencies
 - Based only on: weather conditions, reservoir levels, environmental regulations.
 - Set separately for 14 different types of entitlements.
 - Announced prior to planting season.
 - [See variation](#)
- Water supplies can differ from allocations
 - Transactions, reservoir storage, supplemental programs.

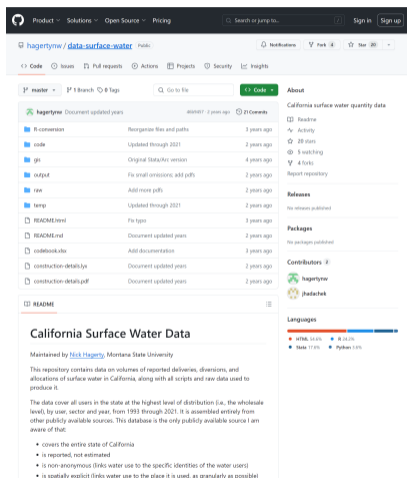
Roadmap

1. Introduction
2. Background
3. **Data**
4. Long-Run Effects
5. Adaptation
6. Policy and Conclusion

New data on surface water supplies

First compilation of all wholesale-level surface water in California.

- Diversions, deliveries, contract maximums, allocation percentages: 1993-2021.
- From archived reports and online databases:
 - Central Valley Project, Lower Colorado (US BOR).
 - State Water Project (California DWR).
 - Appropriative & riparian rights holders (SWRCB).
- Geographic boundaries of districts.



The screenshot shows the GitHub repository page for 'hagertym/data-surface-water'. The repository is maintained by Nick Hagerty, Montana State University. It contains data on volumes of reported deliveries, diversions, and allocations of surface water in California from 1993 to 2021. The repository is organized into folders for 'n-conversion', 'code', 'gis', 'output', 'raw', 'temp', 'README.html', 'README.md', 'codebook.xlsx', 'construction-details.lyx', and 'construction-details.pdf'. The README file is expanded, showing the following text:

California Surface Water Data

Maintained by [Nick Hagerty](#), Montana State University

This repository contains data on volumes of reported deliveries, diversions, and allocations of surface water in California, along with all scripts and raw data used to produce it.

The data cover all users in the state at the highest level of distribution (i.e., the wholesale level), by user, sector and year, from 1993 through 2021. It is assembled entirely from other publicly available sources. This database is the *only* publicly available source I am aware of that:

- covers the entire state of California
- is reported, not estimated.
- is non-anonymous (links water use to the specific identities of the water users)
- is spatially explicit (links water use to the place it is used, as granularly as possible)

The right sidebar of the repository page shows statistics: 29 stars, 5 watching, 4 forks, and 0 releases published. It also lists contributors (hagertym, jhadachuk) and languages (HTML 54.1%, R 24.2%, Markdown 17.6%, Python 3.9%).

Field-level outcome data: Land use and crop choice

Land use and crop choice via satellite data products:

- Statewide Crop Mapping (DWR/Land IQ)
- Cropland Data Layer (USDA NASS)
- All land in California for {2014, 2016, 2018-21}.
- Aggregate to fields (using CLU and PLSS).

Outcome variables: Binary indicators for categories of land use and crop choice.


Crop revenue: Imputed to fields using California County Agricultural Commissioners' Reports.



Field-level outcome data: Property values

Property records for all parcels in agricultural and non-urban parts of California

- Purchased from ParcelQuest
- Price and date of last 3 recorded sales




Field-level data: Covariates

- **Weather:** Gridded data from PRISM via Schlenker & Roberts (2009)
 - Temperature, precipitation, degree days, vapor pressure deficit.
- **Soil characteristics:** Gridded data from the National Cooperative Soil Survey (gSSURGO).
- **Groundwater depth and quality:** Interpolated from observation well readings
 - State Water Boards, Department of Water Resources, U.S. Geological Survey
- **Elevation:** Terrain tiles from USGS/Mapzen/AWS.
- **Historical vegetation:** Central Valley Historic Mapping Project, Base Map Pre-1900 (CSU Chico/USFWS/USBOR).

Roadmap

1. Introduction
2. Background
3. Data
4. **Long-Run Effects**
5. Adaptation
6. Policy and Conclusion

Geographic RD exploits boundaries between neighboring districts



Geographic RD exploits boundaries between neighboring districts

Multi-Cutoff Geographic RD (Cattaneo, Idrobo, Titiunik 2023).

- Pool data across 225 pairs of neighboring districts.
- Normalize by distance to boundary.

Estimate mean difference in outcomes using only data near the boundary:

$$\tau = \underbrace{\lim_{Distance \downarrow 0} \mathbb{E}[Y_i | Distance_i = 0]}_{\text{More Average Water Allocations}} - \underbrace{\lim_{Distance \uparrow 0} \mathbb{E}[Y_i | Distance_i = 0]}_{\text{Fewer Average Water Allocations}}$$

1. **How near?** Bandwidth chosen to be MSE-optimal for each outcome.
2. **Estimation:** Local linear regression on each side, using a triangular kernel.
3. **Inference:** Robust confidence intervals are corrected for misspecification bias.

Geographic RD exploits boundaries between neighboring districts

Scale binary treatment effect to get interpretable units

- Reduced form: Effect of being the relatively water-rich district of each pair.
- First stage: Average cross-boundary difference in mean water supply.
- Ratio (IV or Fuzzy RD): Effect of a continuous increase in mean water supplies.

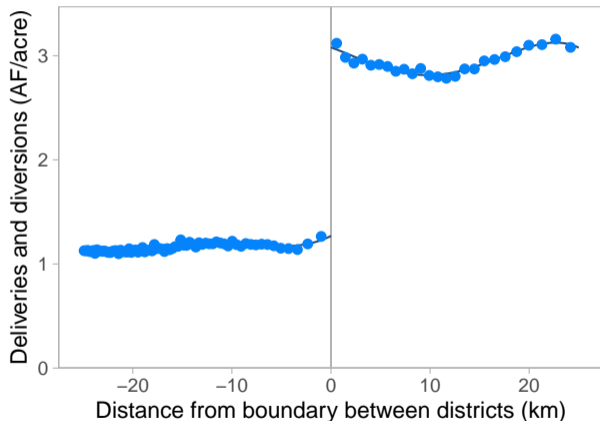
Spatial covariates: Fixed effects for boundary pair b and boundary segments s improve precision & remove bias from geographic imbalance (Keele & Titiunik 2015).

$$Y_{idbst} = \beta \overline{\ln(\text{WaterSupply})}_d + f(\text{Distance}_{idbs}) + \alpha_{bs} + \varepsilon_{idbst}$$

Spatial correlation: Cluster by district d (for both bandwidth selection and inference).

First stage: Boundary effect on surface water supplies

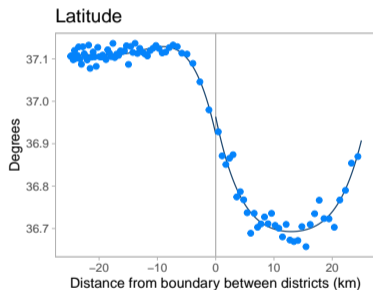
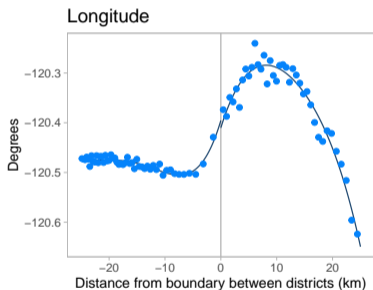
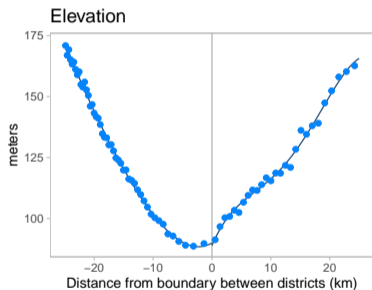
Surface water supplies



Deliveries and diversions (AF/acre)


Optimal bandwidth	3.53 km
RD point estimate	1.80 AF/acre
Robust CI	(1.46, 2.13)
F-statistic	113.2
Pair \times Segment FE	Yes
Eff. Observations	141,565
Local water mean	1.9 AF/acre

Geographic characteristics are continuous at boundaries




Climate variables are continuous at boundaries


Mean temperature (Apr–Sep)




Precipitation (Apr–Sep)




Precipitation (Oct–Mar)




Growing degree days (8–29 C, Apr–Se)



Harmful degree days (>29 C, Apr–Sep)




Mean vapor pressure deficit (Apr–Sep)




Soil characteristics are continuous at boundaries


Soil organic carbon (0–150 mm)




Available water storage (0–150 mm)




Root zone depth




Flooding: occasional or more frequent



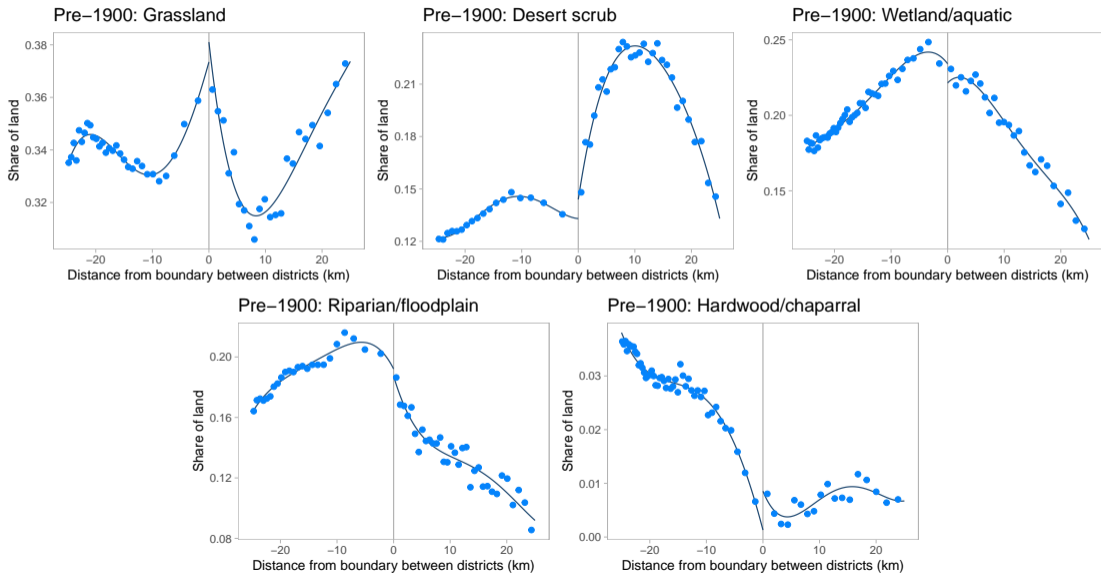
Runoff class: medium to very high



Well-drained



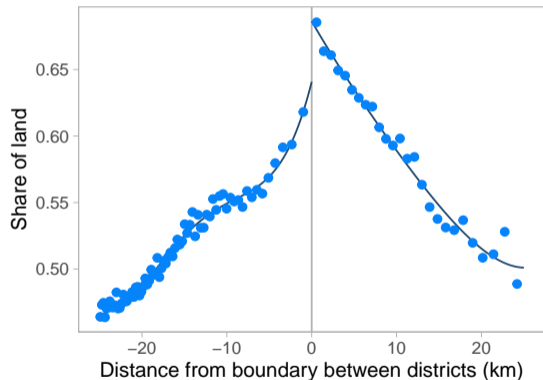
Historical vegetation was continuous at boundaries



Long-run water scarcity reduces crop area

Less land is cropped in places that have less water on average.

Crops planted




Land area cropped (share)

Optimal bandwidth	3.65 km
RD-IV estimate	0.027 (per AF/acre)
Robust CI	(0.005, 0.049)
<i>p</i> -value	0.02
Pair \times Segment FE	Yes
Eff. Observations	143,866
Local water mean	1.9 AF/acre


Translation: 10% less water \Rightarrow 0.5 p.p. less crop area.

Instead, land is fallowed or kept in natural vegetation.


Fallowed/idle land




Natural vegetation



Grassland


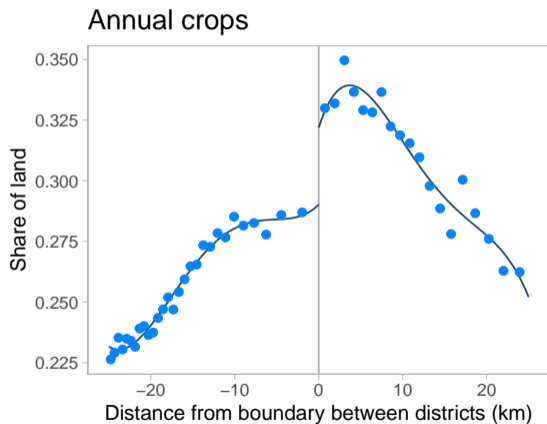


Development



Response comes from... annuals, not perennials


Perennials are insensitive to long-run water supplies.




Response comes from... low-value crops

High-value crops are insensitive to long-run water supplies.

Low-revenue crops




High-revenue crops



Gross output? Somewhat reduced by long-run water scarcity

Land with less water earns less revenue... but not by a lot




Imputed revenue (Poisson regression)

Optimal bandwidth	4.63 km
RD-IV estimate	0.037 per AF/acre
Standard error	(0.020)
Pair \times Segment FE	Yes
Eff. Observations	200,616
Local outcome mean	\$2,595/acre

Translation: 10% less water \Rightarrow 0.7% less revenue. (1 AF produces \$99 in revenue.)

Land values? Lower with long-run water scarcity

Property values are lower in districts with less water.



Log property sales price (2020\$)

Optimal bandwidth	4.84 km
RD-IV estimate	0.095 per AF/acre
Robust CI	(0.016, 0.174)
<i>p</i> -value	0.02
Pair × Segment FE	Yes
Eff. Observations	49,923
Local outcome mean	\$13,900/acre

Translation: 10% less water \Rightarrow 2% lower price. (MWTP for 1 AF/year is \$1,300.)

Results are robust to many checks

Long-run results are robust to: [Table 1](#) [Table 2](#)

- Controls for pre-treatment variables
- Alternative bandwidths
- “Donut hole” regressions
- Alternative definitions of crop revenue
- No controls for latitude and longitude
- No border segments (only border pair effects)
- Smaller border segments (2 km instead of 5 km)
- Rectangular kernel (instead of triangular)
- OLS estimation
- Controlling for variance of water supplies


Groundwater substitution is likely an important response margin

Why isn't output more sensitive?

- No data on extraction, but water tables have fallen faster in surface water-scarce districts.

Groundwater feedback is likely to intensify the effects of surface water.

- Further depletion reduces ability to substitute.



Roadmap

1. Introduction
2. Background
3. Data
4. Long-Run Effects
5. **Adaptation**
6. Policy and Conclusion

Research Question #2

How do the effects of water scarcity change when farmers can adapt?

- Estimate short-run effects using year-to-year fluctuations in water supplies.
- Compare short-run effects to long-run effects in the same setting.

Effects of adaptation: Difference between short-run and long-run effects.

- Short run? Limited range of ability to respond.
- Long run? Invest in knowledge & equipment to grow new crops.

Short run: Two-way fixed effects regression

Exploit year-to-year fluctuations in water supplies driven by weather.

- Regress outcomes (field i , district d , year t) on yearly water supply:

$$Y_{idt} = \alpha_{id} + \gamma \ln(\text{WaterSupply})_{dt} + \lambda_t + \varepsilon_{idt}$$


- Fixed effects for field (α_{id}) and year (λ_t).
- Instrument $\ln(\text{WaterSupply})_{dt}$ with $\ln(\text{WaterAllocation})_{dt}$.

Identification assumptions

First stage

Short-run water scarcity reduces crop area

Crop area is lower in years with less surface water.



IV estimate:

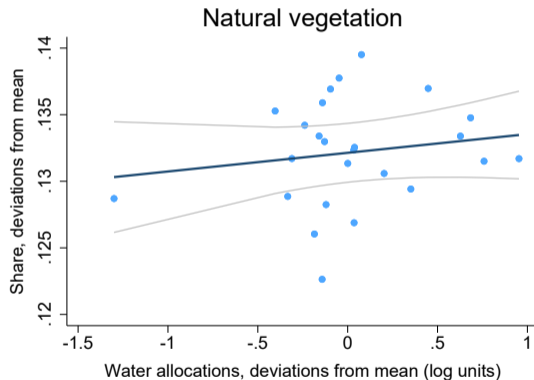
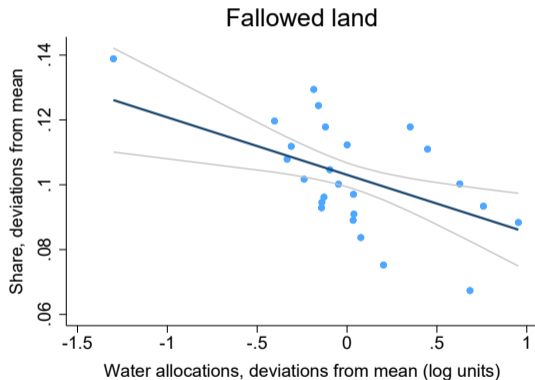
(p.p./% Δ)

0.048***

(0.019)

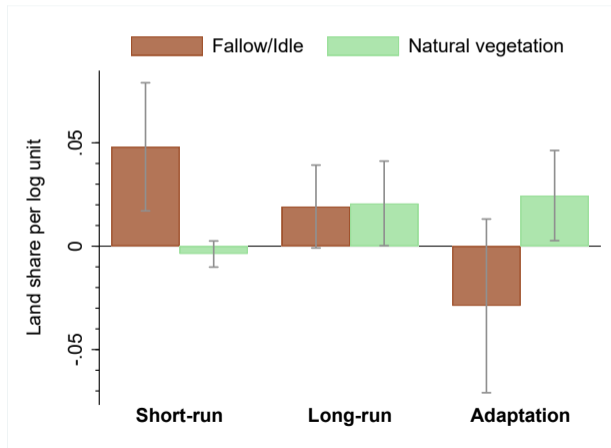
Following accounts for the short-run response

In the short run, land taken out of production is held fallow.



Comparison reveals effects of long-run adaptation

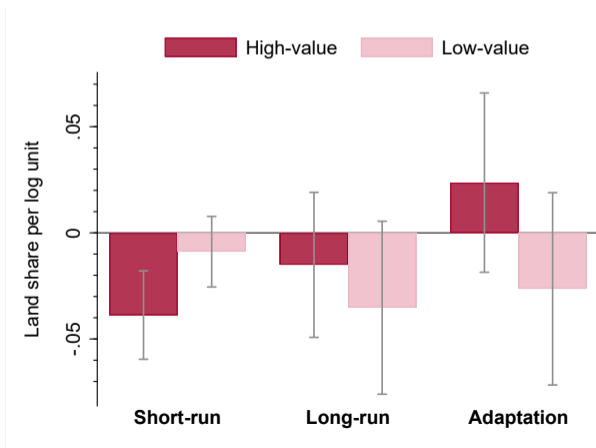
As the short run turns into the long run... farmers permanently retire fallowed land and keep it in natural vegetation.



Crop choice responds differently in short run vs. long run

Crops forgone in the short run are high-value (vs. low-value in the long run).

- On net, farmers shift toward higher-value crops.
- (More cotton/grapes, less alfalfa/rice/corn.)



Other results

Shift toward lower-water crops (weaker evidence)

- More wheat & grains.
- Fewer fruit & nut orchards.
- [Graph](#)

Perennial crops are unaffected by long-run adaptation

- Only annuals respond to water scarcity in either SR or LR.
- [Graph](#)

Differential responses evince long-run adaptation

Land-use choices respond differently in the short run vs. long run.


We can infer:

- Farmers face constraints in the short run that relax in the long run.
- This long-run adaptation is profit-maximizing.

But what does adaptation mean for overall agricultural output?

Short-run water scarcity reduces agricultural output

When farmers get less water, they earn less crop revenue.



Poisson IV elasticity:
0.049**
(0.023)

Adaptation exacerbates output losses

Water scarcity reduces revenue by more in the long run than in the short run.

Crop area

	Semi-elasticity	Percent
Short-run	4.8 p.p.	100%
Long-run	5.1 p.p.	106%
Adaptation	+0.3 p.p.	+6%

Crop revenue (imputed)

	Elasticity	Percent
Short-run	0.049	100%
Long-run	0.072	147%
Adaptation	+0.023	+47%

Results are robust to many checks

Short-run results are robust to: [Table](#)


- Controls for local weather conditions and their lags
- County-specific time trends
- Alternative definitions of crop revenue
- Linear functional forms

Roadmap

1. Introduction
2. Background
3. Data
4. Long-Run Effects
5. Adaptation
6. **Policy and Conclusion**

Policy reforms could reduce revenue losses

Treatment effect heterogeneity implies potential gains from water reallocation.



Extensions

- Can cross-sectional regression replicate the RD results? [Go](#)
- Long-run effects on short-run sensitivity [Go](#)

In summary

I estimate the long-run effects of surface water scarcity in irrigated agriculture

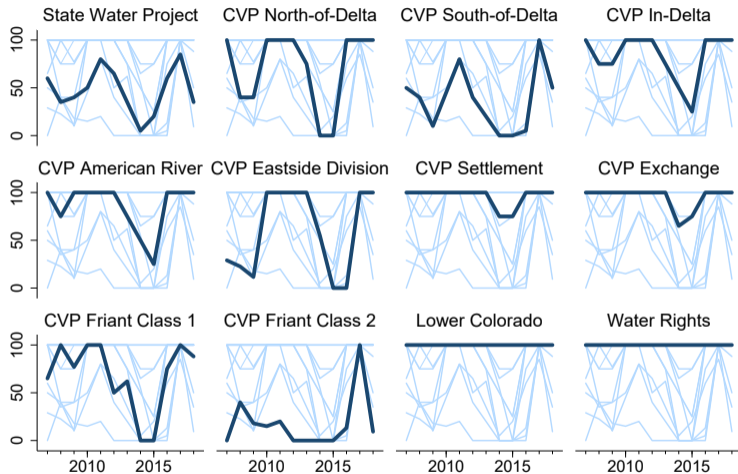
- And how adaptation responses change these effects over time.
- 1. Surface water scarcity is harmful in the long run.** (Water scarcity reduces crop area, crop value, and land values.)
 - 2. Land-use responses suggest farmers adapt.** (Crop choices respond differently in the short run and long run.)
 - 3. Adaptation does not mitigate output losses.** (Revenue losses persist and grow.)

Absent new investments or policy changes...

projected future declines in surface water supplies

are likely to reduce the land area, output, and value of agriculture.

Rich temporal variation in water allocation



Panel A. Summary statistics

	Observations	Mean	Std. Dev.	Minimum	Maximum	Within-district S.D.	Between-district S.D.
Year	3,815,232	2012.5	3.5	2007	2018		
Field area (acres)	3,815,232	34.4	11.9	0.1	79.9		
<i>Water supplies</i>							
Deliveries (acre-feet/acre/year)	3,815,232	2.65	2.14	0.01	10.33	0.33	2.10
Allocations (acre-feet/acre/year)	3,815,232	2.79	2.41	0.01	11.15	0.38	2.79
Ln (Deliveries)	3,815,232	0.615	0.938	-4.490	2.335	0.562	0.669
Ln (Allocations)	3,815,232	0.522	1.238	-4.627	2.411	0.298	0.660
Allocation percentage	3,815,232	0.790	0.294	0	1		
<i>Land use</i>							
Cropped	3,815,232	0.606	0.489	0	1		
Fallow	3,815,232	0.103	0.304	0	1		
Grassland	3,815,232	0.159	0.365	0	1		
Natural vegetation	3,815,232	0.132	0.339	0	1		
<i>Crop outcomes</i>							
Water needs (acre-feet/acre/year)	3,815,229	2.18	2.04	0.00	8.00		
Revenue (2009\$/acre/year)	3,460,126	2,032	3,143	0	74,586		
arcsinh (Water needs)	3,815,229	1.16	1.00	0.00	2.78		
arcsinh (Revenue)	3,460,126	5.40	3.80	0.00	11.91		

Panel B. Crop characteristics by category

Crop category	Share	Mean water needs per acre (acre-feet)	Mean revenue per acre (2009\$)
<i>Perennial, long-term crops</i>			
Almonds, pistachios	13.5%	4.82	5,081
Grapes	6.0%	3.88	7,735
Citrus, other subtropical fruit	2.4%	3.83	6,683
Other tree fruits, nuts	3.8%	4.03	4,811
<i>Annual & short-term crops</i>			
Alfalfa	7.7%	4.97	1,303
Grains	7.0%	1.31	622
Rice	5.9%	2.89	1,594
Corn	4.3%	2.61	868
Cotton	3.7%	3.62	1,736
Tomatoes	2.5%	2.59	4,171
Safflower	0.6%	1.99	453
Onions, garlic	0.5%	3.03	5,943
Melons, squash, cucumbers	0.3%	2.41	5,844
Sugar beets	0.2%	3.62	2,187
Dry beans	0.2%	2.12	1,215
Potatoes	0.2%	1.72	7,825
Pasture, grass	0.1%	7.58	125
Other vegetables, berries	1.0%	1.94	18,128
Other field crops	0.7%	2.67	1,036
<i>Not crops</i>			
Grassland (unirrigated rangeland)	15.9%	0.00	11
Fallow	10.3%	0.00	0
Natural vegetation	13.2%	0.00	0

Long run: Identification assumptions

Assumptions for causal effect of districts:

- (*Continuity*) Pre-treatment factors vary smoothly at district borders.
- (*Non-manipulation*) No transport of surface water across district borders.
 - Robust to excluding a “donut hole” of observations very close to the border.

Assumptions to interpret as causal effect of water supplies:

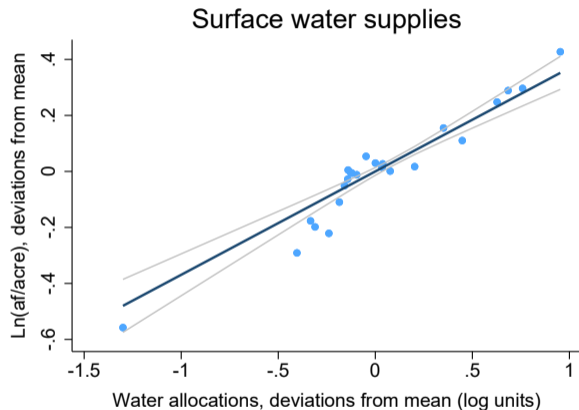
- (*Exclusion*) Districts affect outcomes only through water supplies.
- (*Relevance*) Water supplies are different between neighbors.

Short run: Panel specification

Identification assumptions:

1. Allocations are as good as random (conditional on mean allocations and statewide water availability).
 - Robust to including county-specific linear time trends.
2. (*Exclusion*) Allocations affect agriculture only through actual supplies.
 - Robust to controls for local precipitation and other weather variables.
3. (*Relevance*) Allocations do affect actual supplies.

Short run: First stage



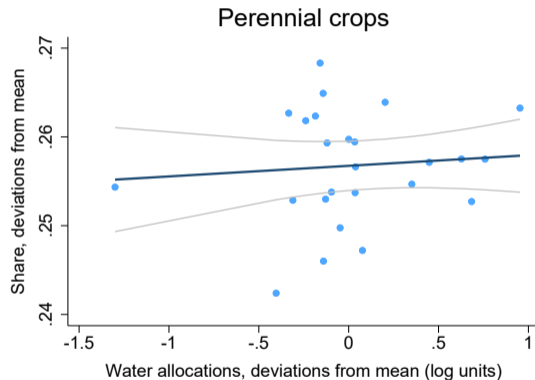
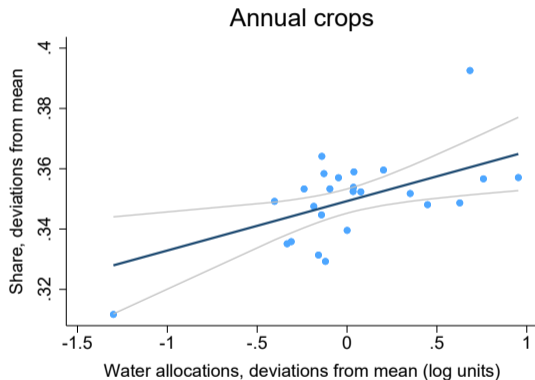
	Ln (Supply)
Ln (Allocations)	0.357*** (0.032)
F-statistic	124.6
Field fixed effects	Yes
Year effects	Yes
Observations	3,815,232
Clusters	2,172

* $p < .1$, ** $p < .05$, *** $p < .01$.

◀ Back

Perennial crops are maintained in the short run

Annual crops account for most of the full crop area effect.

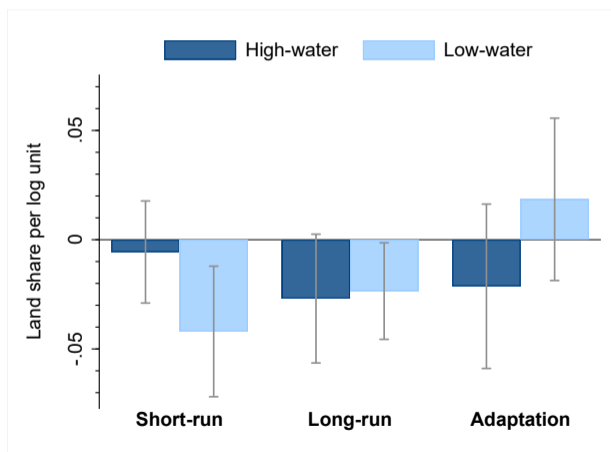


[← Back](#)

2. Farmers adapt in some ways in the long run

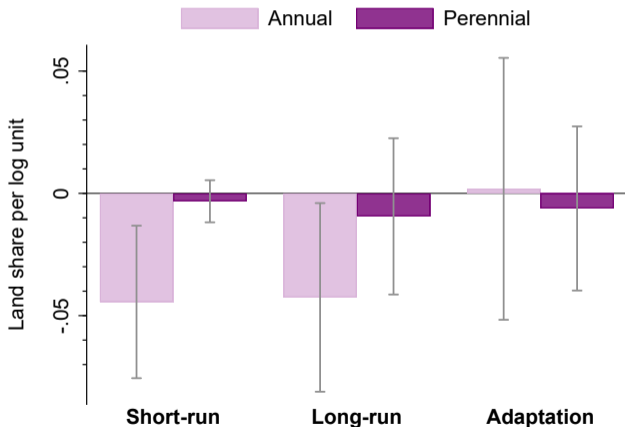
Farmers may also shift toward lower-water crops

- More wheat & grains. Fewer fruit & nut orchards.



Adaptation does not change effect on perennials

Orchards and other perennial crops do not respond, even in the long run.



◀ Back

Short-run results are robust

Specification	Crop area		Crop revenue	
Baseline	0.048	(0.019)	0.362	(0.160)
With county-specific time trends	0.047	(0.018)	0.334	(0.153)
With weather covariates	0.047	(0.020)	0.398	(0.168)
Alternative revenue definitions				
Yields & prices per county & year			0.325	(0.161)
Yields per county & year; prices per year			0.321	(0.157)
Yields per county & year; prices constant			0.362	(0.160)
Yields per county; prices per year			0.285	(0.146)
Yields & prices constant			0.375	(0.144)
Linear outcome			229.7	(95.21)

[← Back](#)

Long-run results are robust (conceptual checks)

Specification	Crop area		Crop revenue	
Baseline	0.050	(0.017)	0.306	(0.137)
With pre-treatment covariates	0.051	(0.018)	0.324	(0.141)
Donut hole (0.57-km radius)	0.058	(0.020)	0.382	(0.152)
Donut hole (1.14-km radius)	0.052	(0.023)	0.293	(0.166)
Alternative revenue definitions				
Yields & prices per county & year			0.288	(0.137)
Yields per county & year; prices per year			0.303	(0.137)
Yields per county & year; prices constant			0.306	(0.137)
Yields per county; prices per year			0.311	(0.131)
Yields & prices constant			0.304	(0.133)
Linear outcome			97.3	(82.6)

Long-run results are robust (RD specification checks)

Specification	Crop area		Crop revenue	
Baseline (10 km)	0.050	(0.017)	0.306	(0.137)
Bandwidth: 25 km	0.064	(0.021)	0.381	(0.148)
Bandwidth: 5 km	0.042	(0.015)	0.256	(0.123)
Rectangular kernel	0.055	(0.018)	0.338	(0.141)
Without latitude & longitude	0.049	(0.017)	0.302	(0.131)
No border segments	0.042	(0.015)	0.249	(0.118)
Smaller border segments (2 km)	0.044	(0.021)	0.239	(0.174)
Stricter definition of neighbors	0.051	(0.018)	0.319	(0.142)
Estimated by OLS (not IV)	0.047	(0.009)	0.337	(0.073)

◀ Back


Can cross-sectional regression replicate the RD results?

Only when combining multiple alternative research designs.

Specification	Mean squared bias (10^{-3})
Cubic control in lat & lon	1.4
Physical covariates	6.6
County fixed effects	55.8
Matched neighbor pair fixed effects	6.6
All of the above	0.1

◀ Back

Long-run scarcity does not change short-run sensitivity



◀ Back