Upstream Advantage in Surface Water Systems: Measuring Inequality Using a Water-Graph Approach*

Anton Liutin
University of Wisconsin–Madison

March 26, 2026

Abstract

Surface water flows sequentially through canal networks, reaching upstream users before downstream users. Tail-end deprivation is well documented in developing countries with weak institutions, but whether formal water rights eliminate this inequality in developed settings remains untested at scale. I construct a "water-graph" mapping upstream–downstream relationships among 54,946 agricultural plots across Arizona's canal networks and test whether network position affects crop outcomes during the unprecedented Colorado River drought (2016–2023). A "donut" identification strategy isolates far upstream competition (>1–10 km) from spatially correlated local shocks. Downstream plots experience significantly lower vegetation health, moisture, and evapotranspiration; the drought amplification of this position penalty roughly doubles to triples across outcomes between the full canal path and the 10 km cutoff (2.0× for NDMI, 2.2× for GCVI, 2.5× for ET). A downstream placebo supports the causal interpretation: downstream capacity coefficients are substantially attenuated relative to their upstream counterparts, consistent with a directional water-competition mechanism rather than spatially correlated shocks. These effects emerge within plots over time, ruling out selection on permanent characteristics. The position–drought gradient is significant only in surface-water-dominant districts—groundwater access eliminates it entirely. Physical delivery infrastructure creates distributional consequences not resolved by prior appropriation rights, pointing to infrastructure targeting and groundwater investment as policy levers in arid regions facing increasing scarcity.

1 Introduction

Irrigation canals are a classic common-pool resource: water drawn by one user reduces what remains for those downstream. Since Ostrom (1990), a large literature has asked whether institutions can prevent overextraction from shared resources, yet a fundamental empirical question remains surprisingly untested at scale—do formal water rights actually eliminate the distributional inequality created by sequential access? In the western United States, the prior appropriation doctrine ("first in time, first in right") establishes legal priority, while water markets enable voluntary reallocation. These institutions, in principle, could resolve distributional tensions: legal seniority determines priority during shortage, and farmers facing constraints can purchase additional allocations from those with surplus. In practice, however, western water law was designed for a more abundant era, water markets remain thin with low trading volumes, and institutional adaptation to unprecedented scarcity has been slow—raising the possibility that formal rights provide less protection than commonly assumed.

The sequential structure of canal delivery creates three mechanisms for distributional inequality: upstream farmers may extract slightly more than their allocation within measurement error

*I thank Paul Castañeda Dower (advisor), Corbett Grainger, Daniel Phaneuf, and Laura Schechter for guidance. This research received no external funding. The author has no financial relationships or conflicts of interest to disclose.

margins, delivery losses compound across the network, and tail-end users bear cumulative shortfalls. If water rights and markets function effectively, irrigation network position should not affect agricultural outcomes conditional on legal priority. However, if market frictions, measurement limitations, and institutional constraints prevent full adjustment, physical position may impose persistent costs on downstream farmers.

This paper tests whether network position affects agricultural outcomes in Arizona's surface water systems, where the Colorado River Basin's driest 22-year period since at least 800 CE (Williams, Cook, and Smerdon 2022) triggered a Tier 1 shortage declaration and a 30% cut to Arizona's allocation. I hypothesize that existing institutions do not fully resolve the distributional consequences of sequential water access: despite holding valid water rights, tail-end farmers systematically receive less water than head-end farmers drawing from the same canal. If this hypothesis is correct, the upstream–downstream gap should widen during drought periods, when water constraints bind most tightly and the cumulative effects of upstream extraction become most consequential.

Prior research has documented that droughts negatively impact agricultural yields (Kuwayama et al. 2019), but has not examined whether these impacts fall equally across farmers or systematically disadvantage those in certain positions within irrigation networks. Mieno et al. (2024) demonstrate that groundwater availability buffers drought impacts, yet similar analysis for surface water systems is lacking. The closest precedent is Cody (2018), who shows that network position matters for water availability in Colorado's San Luis Valley, independent of water rights seniority. However, Cody's analysis is limited to a single irrigation system and does not examine how position effects vary with aggregate water scarcity. A systematic framework for measuring farmers' positions across large-scale canal networks—and testing whether these positions predict agricultural outcomes during drought—has been lacking.

This paper develops such a framework. I construct a "water-graph" that maps upstream–downstream relationships among 54,946 agricultural plots across Arizona's canal irrigation networks. Using canal network data from the USGS National Hydrography Dataset, digital elevation models to infer flow directions, and plot boundaries from USDA Crop Sequence Boundaries, I calculate each plot's network position—a normalized measure of how far downstream the plot sits within its canal system—capturing structural water access independent of legal rights, soil quality, or market proximity. To separate genuine water competition from spatially correlated local shocks, I employ a "donut" identification strategy that uses only far upstream variation in water demand while controlling for nearby fallowing patterns. I combine this network structure with satellite-derived crop outcomes (vegetation indices, evapotranspiration) and drought indices (SPEI) spanning 2016–2023 to test whether downstream position causally affects agricultural productivity.

Four main findings emerge. First, downstream plots experience significantly worse crop outcomes in the cross-section—lower moisture (NDMI), chlorophyll vegetation (GCVI), and evapotranspiration (ET)—even after controlling for water rights seniority, soil characteristics, groundwater access, and market proximity. Second, the donut specification reveals that this position penalty strengthens monotonically as near-distance variation is excluded: the drought×position interaction roughly doubles across all three outcomes between the full canal path and the 10 km cutoff, and a downstream position placebo is largely absorbed once upstream position is controlled for—consistent with directional water competition rather than spatially correlated shocks. Third, heterogeneity analysis identifies the key moderator: the position–drought gradient is significant only in surface-water-dominant districts, while groundwater access eliminates it entirely. Within surface-water districts, those with junior priority allocations face steeper drought amplification than senior-rights districts, and formal shortage protocols appear to concentrate rather than distribute shortfalls. Fourth, institutional buffers partially attenuate the gradient—irrigation district membership and groundwater share reduce the remaining position penalty, pointing to specific policy levers.

These findings contribute to the literature on common-pool resource governance by demonstrating that formal property rights—a central institutional response to commons dilemmas—do not fully resolve distributional inequality when the resource flows sequentially through physical infrastructure. The results also speak to climate adaptation policy: as drought intensifies in the Colorado River Basin and other arid systems worldwide, the position penalty will bind more frequently. The heterogeneity results point to actionable levers: volumetric water pricing nearly fully offsets the drought × position gradient in the most affected districts (junior-rights, surface-water-

dominant)—likely because per-unit pricing both requires turnout-level metering infrastructure and creates marginal-cost incentives against over-extraction. Groundwater investment eliminates the penalty entirely in GW-accessible districts, and irrigation district membership attenuates it through collective governance. Shortage protocols in their current form may need redesign to avoid concentrating shortfalls on downstream users.

The remainder of the paper proceeds as follows. Section 2 reviews related literature. Section 3 presents a conceptual model formalizing the three mechanisms through which upstream position confers advantage. Section 4 describes the data. Section 5 details the water-graph methodology for measuring canal network position. Section 6 presents the identification strategy, including the donut approach and downstream placebo. Section 7 reports results, and Section 8 concludes.

2 Literature Review

2.1 Common-Pool Resources and Tail-End Deprivation

Canal irrigation systems exhibit the two defining features of common-pool resources: rivalry in consumption (water diverted by one user reduces supply for others) and difficulty of exclusion (monitoring individual extraction along kilometers of open canal is costly). A foundational literature, anchored by Ostrom (1990), examines how communities design institutions to govern such resources—rotation schedules, monitoring arrangements, graduated sanctions—and under what conditions collective action succeeds or fails.

Much of the empirical evidence on irrigation governance comes from developing countries where formal rights are absent or weakly enforced. Wade (1988) documents that villages at the tail end of canal systems in South India develop stronger collective institutions precisely because they face the greatest water deprivation—a pattern consistent with necessity driving cooperation. Murgai et al. (2002) show that transaction costs localize mutual insurance networks along canal systems in Pakistan: farmers trade water with nearby neighbors but not distant ones, limiting the scope of risk-sharing. Manero et al. (2019) confirm the head–tail yield gap empirically in Tanzanian smallholder schemes, finding that distance along secondary canals—not just proximity to the main intake—drives the income penalty.

Recent work reveals that institutional reforms do not automatically resolve these inequalities. Jacoby and Mansuri (2020) show that politically influential upstream landowners in Pakistan's Indus Basin lobby canal administrators to capture disproportionate water shares, steepening the head-to-tail decline in land values. Jacoby, Mansuri, and Fatima (2021) find that transferring canal management to local farmer organizations actually *increased* water theft and redistributed wealth from downstream to upstream users—decentralization worsened the very inequality it was designed to address. Not all decentralization fails: Das and Dutta (2026) show that elected canal management bodies (Pani Panchayats) in India reduced spatial misallocation, with distant farmers' ability to complain to local representatives as the key mechanism—suggesting that democratic accountability, not decentralization per se, determines whether reforms succeed. Together, these studies establish that position-based inequality is not merely a symptom of weak institutions—it can persist, and even deepen, through the very reforms designed to address it.

A natural question follows: does tail-end deprivation persist when formal property rights and legal enforcement replace community-based governance? Under the standard Coasian logic, clearly defined rights and low transaction costs should allow efficient reallocation regardless of initial position. This paper tests whether that logic holds in practice.

2.2 Water Distribution, Rights, and Crop Outcomes

Burness and Quirk (1979) formalize the efficiency properties of prior appropriation, showing that appropriative rights can achieve efficient allocation under idealized conditions—stochastic river flows, clearly defined priorities, and costless enforcement. Rafey (2023) demonstrates that market-based water reallocation on river networks can generate substantial gains: trading in Australia's Murray-Darling Basin increased output by 4–6%, with the largest gains during drought. In principle, western U.S. water allocation should achieve similar efficiency: senior rights holders receive priority access during shortage, and markets allow reallocation to highest-value uses.

A growing empirical literature questions whether this efficiency is realized. Kuwayama et al. (2019) estimate significant yield losses from drought conditions but treat impacts as aggregate phenomena, without distinguishing how losses distribute across farmers by network position. Mieno et al. (2024) advance this work by showing that groundwater availability buffers drought impacts, suggesting physical water constraints—not just financial constraints—drive the drought penalty. However, their analysis focuses on groundwater systems where no inherent extraction ordering exists. Within the prior appropriation system, Li, Xu, and Zhu (2019) show that shortage risk redistributes sharply by seniority: junior rights holders in Idaho absorb the bulk of drought risk while senior holders are effectively shielded. Cobourn et al. (2022) confirm that junior rights holders respond by fallowing more land and shifting to drought-resilient crops—documenting how institutional position shapes cropping decisions under scarcity.

The closest precedent to this paper in developed-country surface water is Cody (2018), who examines network position and water rights in Colorado's San Luis Valley. Cody demonstrates that a farmer's physical location within the irrigation system matters independent of rights priority—upstream farmers receive water before downstream farmers regardless of legal seniority. However, Cody's analysis focuses on a single irrigation system at the district level, and establishes only that position correlates with water access without distinguishing causal infrastructure effects from selection. Most closely related, Harris (2024) uses a difference-in-differences design to study Arizona farmers' responses to a 35% reduction in Colorado River deliveries via the Central Arizona Project, finding that farmers fallow land but substitute toward groundwater and water-intensive crops rather than conserving. Her analysis examines aggregate responses to policy-induced supply cuts; I study how these cuts are distributed *within* canal networks—asking not just how much total water is lost, but who loses it.

2.3 Contribution

This paper contributes to the literature in three ways. First, I develop a water-graph methodology that maps upstream–downstream relationships among 54,946 agricultural plots across an entire state's canal infrastructure, using publicly available data (NHDPlus, USDA CSB, DEM) that makes the approach replicable in other surface water systems—an order of magnitude larger than prior work limited to single systems or district-level analysis.

Second, I introduce a "donut" identification strategy that separates genuine water competition from spatially correlated local shocks by measuring only far upstream capacity while controlling for near-distance fallowing patterns. The monotonic strengthening of effects with distance, combined with a downstream placebo test, provides evidence for a directional water-competition mechanism.

Third, the heterogeneity analysis links district-level institutional characteristics—water source dependence, pricing structure, rights seniority, groundwater access—to the magnitude of the position–drought gradient, identifying which institutional features attenuate or amplify distributional inequality rather than merely documenting that position matters.

3 Conceptual Model

This section develops a conceptual framework to formalize how physical position, water scarcity, and institutional design interact to create systematic disadvantages for downstream users. In Arizona, farmers access surface water through a hierarchical system of canals managed by irrigation districts. Each farmer holds water rights specifying annual volumetric limits and seniority priority.[1] Irrigation districts distribute water through rotational delivery schedules and, during shortage, pro-rata reductions proportional to each farmer's allocation.

[1] Under Arizona's prior appropriation doctrine, water rights are ranked by historical seniority ("first in time, first in right"). During shortage, senior rights holders receive their full allocation before junior holders receive any water. See Arizona Department of Water Resources and Central Arizona Project 2023.

3.1 Two-Farmer Illustration

Consider a hypothetical sub-canal delivering water W_0 to two farmers: Farmer (1) holds senior rights, Farmer (2) holds junior rights. The two panels below vary only in the farmers' physical positions along the canal.

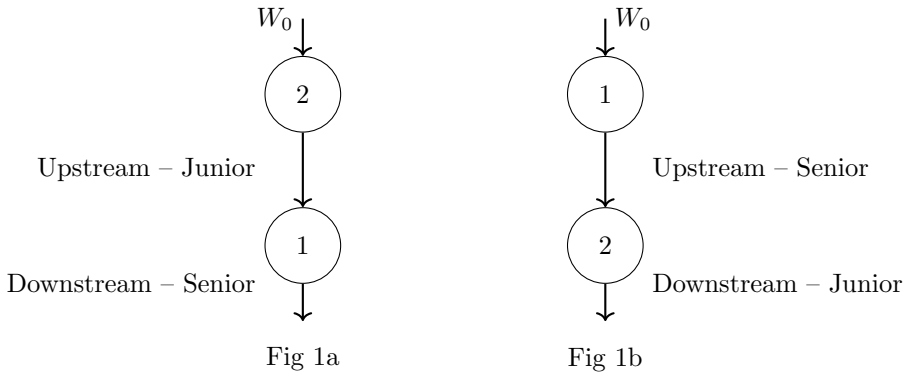


Figure 1: Water allocation under two configurations of seniority and position

In Fig. 1a, the junior farmer sits upstream. If water flow becomes insufficient, the downstream senior can invoke a "call on the river"—a legal mechanism requiring upstream juniors to cease diversion until senior rights are satisfied.[2] The junior farmer therefore faces a binding legal constraint on extraction.

In Fig. 1b, the positions are reversed: the senior farmer sits upstream, the junior downstream. The junior farmer has neither physical control over the flow nor legal priority to demand curtailment. If the upstream senior extracts their full allocation—or slightly more within measurement error—the downstream junior receives only the residual. This configuration, where junior rights and downstream position coincide, produces the sharpest distributional inequality.

If water delivery were perfectly measured and enforcement costless, both configurations would yield identical outcomes: each farmer receives exactly their allocation. In practice, three mechanisms cause the downstream junior to systematically receive less.

3.2 Three Mechanisms

Mechanism 1: Priority and Sequential Access

Head-end farmers extract water *first*. In a canal serving N plots sequentially, each upstream diversion reduces the residual supply available downstream. The precision of extraction monitoring varies substantially across Arizona's irrigation districts. Some Bureau of Reclamation project districts (e.g., Yuma, Wellton-Mohawk) meter water at individual farm turnouts using orifices, propeller meters, or weirs, with reported accuracy of $\pm 3\%$ per device (Stringam, Sauer, and Pugh 2003). Other districts—particularly smaller systems and some Salt River Project sub-districts—allocate water on a per-acre basis without volumetric measurement at individual turnouts.[3] Where metering exists, small per-farmer extraction errors compound along the network: if each of $N-1$ upstream farmers extracts ϵ above their allocation within the undetectable margin, the tail-end farmer faces a cumulative shortfall of $(N-1)\epsilon$. Where metering is absent, the scope for unmonitored over-extraction is larger still. Additionally, delivery losses from evaporation and

[2] A "call on the river" is rooted in the prior appropriation doctrine: senior rights holders can demand that upstream junior appropriators curtail diversions until the senior allocation is met (The Nature Conservancy Arizona and Friends of the Verde River 2020; Feldman 2004).

[3] Arizona has no statewide requirement for metering individual surface water diversions. While Active Management Areas impose strict groundwater pumping regulations and encourage surface water use to offset aquifer depletion, the metering of surface water deliveries to individual farms is determined at the irrigation district level, creating substantial heterogeneity in monitoring capacity across the state.

seepage are not measurement errors—they compound deterministically along the canal regardless of metering infrastructure, further reducing residual supply at the tail.

Prediction: $\frac{\partial Y}{\partial \text{Position}} < 0$—crop health declines with distance from the canal headgate, even after controlling for water rights seniority.

Mechanism 2: Scarcity Amplification

During drought, total water supply falls but upstream extraction behavior does not adjust proportionally. When shortfalls are anticipated *before* delivery begins, the prior appropriation system operates as designed: junior rights holders are curtailed first, regardless of position. However, when shortfalls are realized *during* delivery—due to unexpected supply reductions, infrastructure failures, or emergency redistribution—both seniority and physical position matter. Upstream farmers can maintain their extraction rate while downstream farmers absorb the residual shock. As the marginal value of water rises during drought, incentives to extract within measurement margins intensify, and the cumulative shortfall at the tail grows.

Consider a delivery scheduled to provide W_0 over 10 hours. If an unexpected shock reduces delivery to 9 hours, upstream farmers extract at their planned rate for the full duration while downstream farmers receive whatever remains in the final hour. In Fig. 1b, the downstream junior lacks both flow control and legal priority, bearing the full burden of the supply shock.

Prediction: $\frac{\partial^2 Y}{\partial \text{Position} \cdot \partial \text{Drought}} < 0$—drought *amplifies* the position penalty.

Mechanism 3: Institutional Mediation

If institutions were complete, the position penalty documented in Mechanisms 1 and 2 would vanish. Two institutional channels could, in principle, eliminate it. First, water markets would allow downstream farmers to purchase supplemental allocations from upstream users who value the water less—the standard Coasian solution. Second, regulatory enforcement would prevent upstream over-extraction by detecting and penalizing diversions beyond legal allocations.

In practice, both channels face substantial frictions in Arizona. Water markets remain thin: trading volumes are low, cross-district transfers occur at the district rather than individual level, and conservation-based trading programs often offer compensation below farmers' opportunity costs (Harris 2024). Compounding this, Arizona water rights historically operated under "use it or lose it" provisions—farmers who did not use their full allocation for a specified period risked permanent forfeiture of their rights, creating a perverse incentive to extract the full entitlement even when the marginal value of water was low (Brocious 2021). Although a 2021 legislative reform allows voluntary conservation without forfeiture risk for up to 10 years, this reform covers only the final three years of the sample period, and participation requires formal agreements with the Arizona Department of Water Resources—leaving the underlying friction largely intact during the period I study. Enforcement is limited by the same measurement errors described in Mechanism 1—if individual over-extraction falls within the precision margin of canal flow meters, it is undetectable even with active monitoring. These frictions allow the physical position advantage to persist: upstream farmers retain their extraction benefits because downstream farmers cannot easily bid water away from them, regulators cannot detect cumulative small diversions, and the institutional rules themselves discourage the conservation that would free water for tail-end users.

Given that markets and enforcement do not fully resolve the position penalty, its magnitude should vary with district-level institutional features that either buffer the shortfall or reduce the incentive for over-extraction:

- **Groundwater access** provides an *alternative supply* that bypasses the canal entirely. Farmers with wells can substitute groundwater when surface deliveries fall short, breaking the dependence on canal position. This does not fix the canal distribution problem—it circumvents it. Districts with extensive groundwater infrastructure should exhibit a weaker position penalty.

- **Volumetric pricing** aligns *incentives*. When water is priced per unit volume, each unit of over-extraction has a direct cost, reducing the incentive to divert beyond allocation. Volumetric pricing also counteracts "use it or lose it" behavior by making excess extraction expensive

rather than free. Flat-rate or acreage-based pricing provides no marginal-cost signal and leaves the over-extraction incentive intact.

- **Water rights seniority** provides *legal recourse*. Senior rights holders can invoke curtailment of upstream juniors during shortage—the "call on the river" mechanism described in the two-farmer illustration. In districts where most users hold senior rights, the priority mechanism operates as designed and position should matter less. In junior-rights districts, downstream farmers lack both physical flow control and legal priority, leaving them fully exposed to upstream extraction behavior.
- **Irrigation district membership** provides *collective governance*. Plots within managed irrigation districts benefit from coordinated delivery schedules, infrastructure maintenance, and institutional oversight that individual farmers outside district boundaries cannot access. District membership should attenuate the position penalty by providing a layer of management between individual extraction decisions and aggregate water distribution.

Prediction: The position–drought gradient varies across districts as a function of groundwater access, pricing structure, rights seniority, and institutional governance—with the strongest effects in surface-water-dominant, junior-rights districts lacking volumetric pricing and district-level management.

3.3 Testable Implications

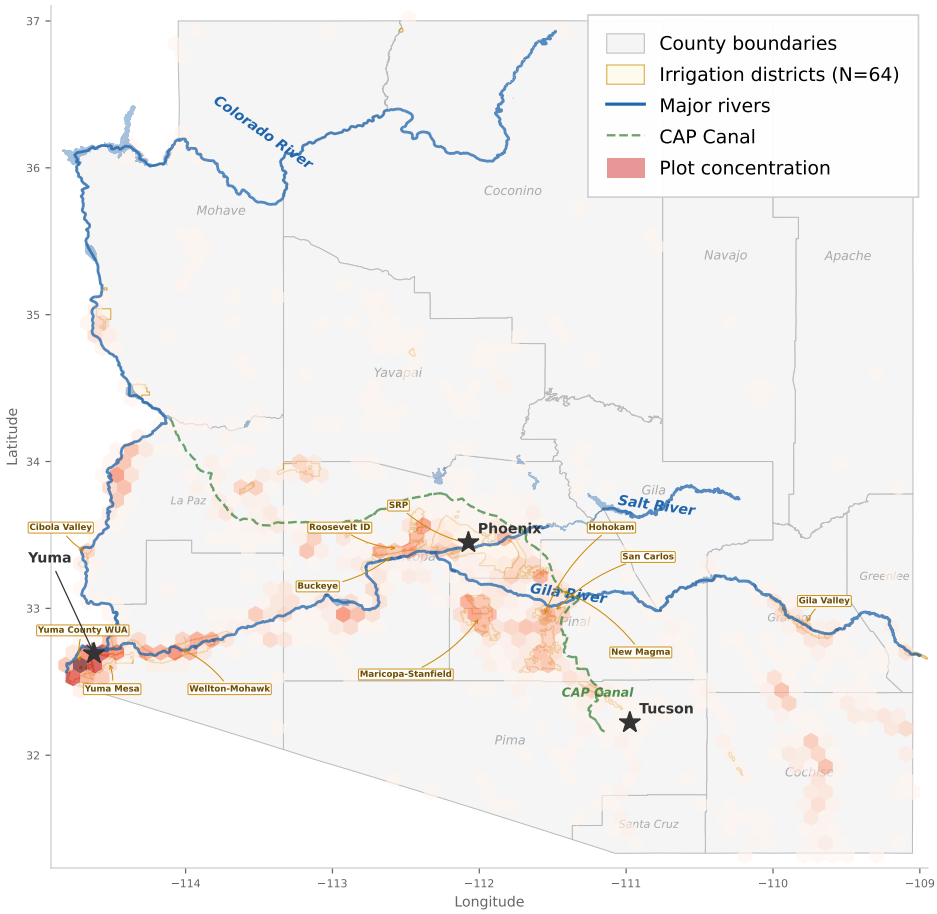
The framework generates three testable predictions that map directly to the empirical analysis:

1. **Level effect**: Downstream plots exhibit lower crop health outcomes (NDMI, GCVI, ET) than upstream plots within the same canal system, conditional on water rights, soil quality, and market access.
2. **Drought amplification**: The position penalty widens during drought years, when water constraints bind most tightly and cumulative upstream extraction becomes most consequential.
3. **Institutional heterogeneity**: The position–drought gradient is attenuated in districts with groundwater access, volumetric pricing, senior water rights, or managed irrigation district governance, and amplified in surface-water-dominant districts with junior allocations lacking these institutional buffers.

4 Data Description

This study combines multiple geospatial and administrative data sources to construct a water-graph mapping upstream–downstream relationships among 54,946 agricultural plots across Arizona's canal irrigation networks. Plot boundaries are derived from USDA Crop Sequence Boundaries, which define stable field geometries from the full 2008–2023 satellite record—the same plots are observed in every year, forming a balanced panel. The resulting dataset contains 439,576 plot-year observations over 2016–2023, of which approximately 75% represent actively cropped plots and 25% fallow. Canal infrastructure and plot–canal associations are also stable across this period: Arizona's irrigation network did not undergo major construction or realignment during 2016–2023, so the water-graph structure is time-invariant.

Figure 2 displays the geographic setting. Agricultural plots in Arizona concentrate along three major surface water corridors: the Colorado River and its diversion canals in the western border counties (Yuma, La Paz), the Salt–Gila river system feeding the Phoenix metropolitan area (Maricopa, Pinal), and the upper Gila Valley in the east (Graham). The Central Arizona Project (CAP) aqueduct connects the Colorado River to central Arizona, supplying several irrigation districts along its 540 km route. Plot density closely tracks the canal network—the vast majority of irrigated agriculture occurs within a few kilometers of a canal segment, confirming that surface water infrastructure, rather than rainfall, determines the spatial distribution of farming in this arid environment.



54,946 Agricultural Plots (2016–2023)

	SW-dominant (GW<50%)	GW-dominant (GW≥50%)	Total
Inside districts	17,249 (69%)	7,674 (31%)	24,923
Outside districts	13,909 (46%)	16,114 (54%)	30,023

Figure 2: Study area: Arizona agricultural plots and canal infrastructure. Plot locations (dots) cluster along the Colorado River corridor (west), the Salt–Gila system (central), and the upper Gila Valley (east). Major irrigation districts and the Central Arizona Project aqueduct are shown.

Primary Geospatial Data

- **Irrigation Canal Network**: The USGS National Hydrography Dataset Plus (NHDPlus) provides canal geometries, flow attributes, and unique identifiers (NHDPlusID). After filtering ephemeral streams (fcode $\neq$ 46007), the dataset contains approximately 547,000 canal segments across Arizona.
- **Agricultural Plot Boundaries**: The USDA Crop Sequence Boundaries (CSB) dataset provides satellite-derived field boundaries spanning 2016–2023, with each plot identified by a unique CSBID. I restrict to plots in Arizona (STATEFP = '04') across all 15 counties.
- **Digital Elevation Model (DEM)**: High-resolution elevation rasters determine water flow directions within the canal network by calculating elevation gradients along canal segments.
- **Crop Type Classification**: The USDA CropScape Cropland Data Layer (CDL) provides annual satellite-based crop type classification at 30 m resolution for 2016–2023. CDL is used for three purposes: constructing crop-specific fixed effects (crop $\times$ year), classifying plots as active or fallow in each year (which drives the time-varying upstream competition measure), and identifying crop rotation patterns for the fallow type decomposition.

Crop Outcomes

I measure agricultural productivity and water stress using three satellite-derived outcome variables, each computed as a seasonal maximum or sum over the growing season (May–October) for each plot-year. All indices are extracted via Google Earth Engine from imagery with less than 20% cloud cover and are log-transformed using $\ln(1 + x)$ before analysis.

- **Normalized Difference Moisture Index ($\mathbf{NDMI_{max}}$)**: The seasonal maximum of NDMI $= (B8 - B11)/(B8 + B11)$ from Sentinel-2 harmonized surface reflectance (COPERNICUS/S2_SR_HARMONIZED) at 10–20 m resolution. NDMI captures canopy moisture content and serves as the primary water stress proxy. I also compute the within-season standard deviation (NDMI_{sd}) as a measure of intra-annual water stress volatility.
- **Green Chlorophyll Vegetation Index ($\mathbf{GCVI_{max}}$)**: The seasonal maximum of GCVI $= (B8/B3) - 1$ from Sentinel-2. GCVI is sensitive to chlorophyll content and serves as a yield and vigor proxy. Satellite-derived vegetation indices have been validated as reliable yield proxies against ground-truth harvest data (Lobell et al. 2020).
- **Evapotranspiration ($\mathbf{ET_{ensemble}}$)**: The seasonal sum (May–October) of monthly ET from the OpenET Ensemble product (v2.0), which combines multiple energy-balance models (SIMS, eeMETRIC, and others) at 30 m Landsat-scale resolution. I compute the plot-level mean ET at 30 m scale. ET captures actual crop water consumption and complements the stress-based vegetation indices. I additionally extract ET from individual SIMS and eeMETRIC models and construct a "water gap" measure ($\text{ET}_{sims} - \text{ET}_{eemetric}$) as a robustness check.

For the bare soil season (January–March), I extract three Sentinel-2 indices as time-varying soil controls: a salinity index ($\sqrt{B2 \times B4}$, median), a clay ratio ($B11/B12$, median), and a bare soil index (BSI, median). Crop type classifications come from the USDA CropScape Cropland Data Layer at 30 m resolution, enabling crop-specific fixed effects and rotation pattern analysis.

Climatic and Drought Data

- **Standardized Precipitation Evapotranspiration Index (SPEI)**: To disentangle irrigation infrastructure effects from meteorological variability, I incorporate SPEI, which accounts for both precipitation and evaporative demand—critical in arid agriculture where temperature-driven evaporation is a first-order constraint. I extract two timescales from the SPEIbase global gridded dataset (Vicente-Serrano, Beguería, and López-Moreno 2010) at approximately 0.5° resolution, sampling at each plot's centroid for the growing season (May–August):

- **SPEI-3M**: A 3-month index representing immediate meteorological conditions and soil moisture pulses.
- **SPEI-24M**: A 24-month index proxying for reservoir storage and long-term water reliability.

For the regression analysis, I construct binary drought indicators equal to one when the respective SPEI value falls below -0.6, capturing the onset of abnormally dry conditions on the standard SPEI classification scale. This threshold yields five drought years (2018, 2020–2023) and three non-drought years (2016, 2017, 2019) in the sample, providing sufficient variation in both categories for estimation. Notably, the sample period captures the intensification of the ongoing Colorado River Basin megadrought, with four consecutive drought years from 2020 onward. As a robustness check, I also estimate specifications using the continuous SPEI measure to exploit finer variation in drought intensity across years.

Institutional and Infrastructure Data

- **Water Rights Database**: Surface water rights data from the Arizona Department of Water Resources "Statement of Claimant Locator" database provides ownership, priority dates, and locations. These records represent Place of Use (POU) designations—the locations where water may legally be applied for beneficial use. Priority dates enable calculation of weighted seniority metrics for upstream and downstream neighbors.[4]

  A significant limitation is that water rights are recorded as point locations rather than polygons. This creates matching challenges: a single point cannot precisely define which plots fall under a given right, and the service area may extend beyond the recorded coordinates. My water rights controls therefore rely on distance-based matching from plot centroids to the nearest recorded right location, introducing measurement error that likely attenuates estimates of how legal priority affects crop outcomes. Despite this limitation, the priority date information allows me to test whether network position effects persist after controlling for water rights seniority—a key prediction distinguishing physical water constraints from legal allocation mechanisms.

- **Irrigation District Boundaries**: Arizona's surface water governance operates through a complex hierarchy of institutions, including Irrigation Districts, the Central Arizona Project (CAP), the Salt River Project, and Agricultural Improvement Districts—each with distinct water allocation rules and management practices. I obtain Irrigation District boundary shapefiles from the Arizona Department of Water Resources (ADWR) and spatially assign each agricultural plot to the nearest district. This generates an irrigation district identifier, a binary indicator for whether the plot falls within district boundaries, and distance to the nearest district for plots located near but outside boundaries.

  Importantly, irrigation district boundaries partially address the limitations of the point-based water rights data described above. Irrigation districts are typically major surface water rights holders that deliver water to member farms within their service areas. Thus, a plot's location within district boundaries serves as a reliable indicator of surface water access, even when individual water rights records are incomplete or imprecisely geolocated. By controlling for irrigation district membership, I capture much of the variation in legal water access that the point-based rights data may miss, allowing for cleaner identification of the physical network position effects that are the focus of this study.

- **Market Access (OpenStreetMap)**: I calculate Euclidean distance from each plot centroid to nearest urban centers, highways, and rural roads to control for market access and transportation costs. Beyond their direct economic relevance, these distance measures also serve as proxies for water rights seniority. Under the prior appropriation doctrine, early settlers who established farms first obtained senior water rights. Historically, these first farmers

[4] Plots with no matched water right are assigned a priority year of 2025 (the current date), treating them as having the lowest possible seniority. This affects approximately 20% of observations (Table 1).

tended to locate near existing roads and emerging urban centers, where market access and infrastructure made agriculture most viable. As a result, plots closer to roads and cities are more likely to hold senior water rights, even when priority dates are missing or imprecisely recorded in the water rights database. Including these distance controls thus helps absorb variation in legal water priority that might otherwise confound estimates of physical network position effects.

Groundwater Rights and Well Data

- **Grandfathered Groundwater Rights**: I utilize the Arizona Grandfathered Groundwater Rights database to identify historical plot boundaries with established groundwater extraction rights. These grandfathered rights, which predate Arizona's 1980 Groundwater Management Act, provide legal access to groundwater and represent an important alternative water source during surface water shortages.

- **Well Registry**: The Arizona Groundwater Site Inventory provides data on registered wells including GPS coordinates and installation dates. I calculate `WELL_COUNT`—the number of wells within grandfathered rights boundaries associated with each plot—as the primary measure of groundwater access capacity. This count variable captures the intensity of groundwater infrastructure available to supplement surface water deliveries during drought periods.

Soil Characteristics

- **Soil Quality (SSURGO)**: The USDA NRCS Soil Survey Geographic Database provides time-invariant soil properties including clay content (water retention), salinity (EC), and soil organic carbon. These controls address soil heterogeneity that could confound estimates of water delivery efficiency.

Data Processing and Variable Construction

The final analytical dataset contains 54,946 unique agricultural plots observed annually over 2016–2023 (439,576 plot-year observations). Of these, 50,970 plots have canal network positions assigned through the water-graph algorithm described in the next section. The water-graph construction proceeds through several stages: (1) canal flow direction inference using DEM elevation regression; (2) canal graph construction with NHDPlus flow table corrections; (3) parallel canal merging and T-junction splitting; (4) plot–canal linking using elevation-adjusted cost functions; (5) plot neighbor detection via raster flood-fill; and (6) position computation using cumulative canal distance and signed-hop BFS from canal-graph roots. Full details are provided in Section 5.

Table 1 presents summary statistics separately for active and fallow plots.

Network Position Variables

The water-graph algorithm (Section 5) represents the canal irrigation network as a directed graph of agricultural plots, where an edge from plot A to plot B indicates that water flows past A before reaching B. Three key structural concepts organize this graph:

- A **canal block** is the primary spatial unit of analysis. The canal graph contains 21,998 independent canal networks (weakly connected components). Of these, 1,060 contain at least one agricultural plot. For compact networks whose plots fit within a 50×50 km bounding box (1,054 of 1,060), the entire network forms a single block. For the 6 largest networks spanning multiple irrigation districts, I apply DBSCAN spatial clustering (eps = 5 km) on plot locations to separate distinct agricultural clusters, yielding 1,162 total blocks after a post-clustering connectivity check ensures each block's canal subgraph is connected.

- A **weakly connected component (WCC)** is a maximal set of plots connected by canal paths in either direction within a block—effectively, a single independent irrigation subnetwork. Some blocks contain multiple WCCs because intermediate canals with no associated

Table 1: Summary Statistics: Active vs. Fallow Plots

		Active Plots			**Fallow Plots**		
	TV	Mean	SD	N	Mean	SD	N
Crop Outcomes							
NDMI (max)	✓	0.253	0.156	328,006	0.022	0.138	107,155
GCVI (max)	✓	1.78	0.629	328,006	0.930	0.493	107,155
ET Ensemble (log)	✓	6.29	0.552	328,006	5.26	0.809	107,155
Network Position & Upstream Competition							
Network Position (0–1)		0.507	0.258	305,807	0.512	0.280	101,953
log(1 + upstream active >1km)	✓	3.45	1.67	253,120	2.74	1.76	81,368
log(1 + downstream active >1km)	✓	3.47	1.67	253,799	2.58	1.73	81,225
Canal Network Size (plots)		4400.5	4556.0	305,807	2736.6	3762.9	101,953
Drought							
Drought_{24M} (binary)	✓	0.141	0.348	331,334	0.000	0.003	108,242
SPEI 24-Month	✓	-0.108	0.512	328,304	-0.092	0.601	107,352
SPEI 3-Month	✓	-0.303	0.545	328,304	-0.355	0.514	107,352
Fallow & Upstream Activity							
Stress fallow (non-rotation)	✓	0.000	0.000	331,334	0.279	0.448	108,186
Rotation fallow	✓	0.000	0.000	331,334	0.260	0.439	108,186
Upstream non-rot. fallow share	✓	0.051	0.099	294,373	0.114	0.178	95,459
Upstream rot. fallow share	✓	0.036	0.091	294,373	0.146	0.238	95,459
Plot Characteristics							
Plot Size (acres)		21.39	36.12	331,334	14.58	40.89	108,242
Well Count		1.48	4.51	331,334	1.59	4.84	108,242
Distance to Canal (m)		726.2	964.6	331,334	1011.3	1320.3	108,242
Slope (%)		1.33	1.75	327,450	1.67	2.50	106,430
Elevation Δ to Intake (m)		-0.192	4.83	331,334	0.400	7.14	108,242
Pumping (dummy)		0.411	0.492	331,334	0.500	0.500	108,242
Soil & Water Access							
Soil Salinity (EC)		6.81	0.499	327,478	7.25	0.427	106,864
Clay Content		0.849	0.116	327,478	0.776	0.071	106,864
GW Share (HUC12)		0.436	0.310	323,974	0.521	0.302	101,466

Notes: 439,576 plot-year observations (54,946 unique plots, 2016–2023; 50,970 with canal network positions). Active $\approx$ 75%, Fallow $\approx$ 25%. ET Ensemble is log-transformed. Network position normalized 0–1 within each weakly connected component.

plots create gaps in the plot graph. There are 3,362 WCCs in Arizona, ranging from 2 to over 12,000 plots. Position is normalized within each WCC, since plots in different WCCs share no water.

- **Breadth-first search (BFS)** is the graph traversal algorithm used to compute each plot's position within its WCC. Starting from the root plot (the most upstream plot on the highest-elevation source canal), BFS assigns each plot a signed hop count: +1 per step downstream, −1 per step upstream along secondary inlets. This hop count is then normalized 0–1 within each WCC.

The algorithm produces two complementary position measures. The *static* measure, `signed_hop_pos`, is the BFS-based normalized position (0 = headgate, 1 = tail), fixed across years. The *time-varying* measure captures upstream competitive pressure: for each plot-year, I count active upstream plots (those not classified as fallow in the CDL) beyond a Euclidean distance threshold d and compute

$$\text{log_pos_far}_d = \ln\big(1 + \text{count of active upstream plots with Euclidean centroid distance} > d\big).$$

I construct this measure at thresholds $d \in \{0, 1, 2, 3, 5, 7.5, 10\}$ km. The "donut" identification strategy uses the far measure ($d > 0$) as the treatment variable while controlling for near-distance variation.

The WCC identifier (`wcc_id`) serves as the primary fixed-effect grouping for position normalization.

Upstream Fallow Shares

To implement the donut strategy, I decompose upstream fallowing into near and far components at each distance threshold. For each plot-year, I compute the share of upstream plots that are fallow, separately for two fallow types:

- **Non-rotation fallow**: Plots classified as fallow in the CDL with a predicted rotation probability below 0.3 (from a machine learning classifier trained on crop sequence patterns; details in Appendix A). These plots are likely fallowed due to water constraints rather than planned crop rotation.
- **Rotation fallow**: Plots classified as fallow with a predicted rotation probability above 0.7, representing planned agronomic rotation.

For a given distance threshold d, fallow shares are computed separately for *near* (within Euclidean distance $\leq d$) and *far* (beyond d) upstream plots, and separately for non-rotation and rotation fallow types. In the donut specification (Section 6.3), only the *near* spatial fallow shares enter as controls—they absorb spatially correlated local shocks (shared weather, pest pressure, market conditions) within the donut hole. The far upstream variation is captured directly by the treatment variable log_pos_far_d, which counts active upstream plots beyond d; adding far fallow shares as separate controls would double-count this variation.

Irrigation District Institutional Characteristics

The heterogeneity analysis requires district-level variables capturing institutional features that may moderate the position–drought gradient. I collect these for 64 Arizona irrigation districts from official district records, Arizona Department of Water Resources publications, and publicly available governance documents.[5] Districts with fewer than 50 plots are excluded from per-district analysis, leaving 49 districts. The key institutional variables are:

[5]District-level data were compiled from multiple sources including ADWR records, district annual reports, CAP allocation documents, and supplementary web-based research. The complete dataset and source documentation are available upon request.

- **Water source composition**: I compute the groundwater share of irrigation withdrawals (`gw_share`) from the USGS National Water Availability Assessment (NWAA) Data Companion, which provides monthly HUC12-level estimates of irrigation water withdrawals disaggregated by groundwater (`irrwdgw_mgd`) and surface water (`irrwdsw_mgd`) sources. Each plot inherits the groundwater share of its containing HUC12 watershed. For the heterogeneity analysis, I aggregate this plot-level measure to the district level by computing the mean `gw_share` across all plots within each irrigation district. I classify districts as *GW-dominant* ($\geq$70% groundwater), *SW-dominant* ($\leq$30% groundwater), or *mixed* (30–70% groundwater). After applying the 50-plot minimum, 49 districts enter the heterogeneity analysis.
- **Pricing structure**: Whether the district charges volumetric rates per acre-foot (`price_per_unit`), flat annual fees (`price_flat_fee`), per-acre assessments (`price_per_acre`), tiered/progressive rates (`has_tiered_rates`), or conservation penalties for over-use. Pricing categories are not mutually exclusive; many districts combine per-acre assessments with volumetric charges or tiered surcharges. Volumetric pricing creates marginal-cost incentives that should reduce upstream over-extraction.
- **Shortage and governance**: Whether the district has a formal shortage curtailment protocol (`has_shortage_protocol`, $N = 18$), uses rotational delivery schedules (`has_rotation_schedule`, $N = 17$), and whether the district falls within an Active Management Area subject to state groundwater regulation (`within_AMA`, $N = 35$).
- **Water rights seniority**: Colorado River priority ranking from 1 (most senior) to 7 (most junior), with mean 5.1 across districts. I classify districts as *senior* (rank ≤ 3) or *junior* (rank > 3). I also identify districts vulnerable to Colorado River shortage declarations—those with junior priority or dependence on Central Arizona Project allocations, which are curtailed first under the 2007 Interim Guidelines.

Table 2 summarizes these institutional characteristics across the GW-dominant, Mixed, and SW-dominant subsamples, as well as the senior/junior rights split within SW-dominant districts.

5 Measuring Position in the Water Network: The Water-Graph Approach

5.1 Overview

The water-graph is a directed graph in which nodes represent agricultural plots and edges represent water flow relationships: an edge from plot A to plot B indicates that water passes A's canal intake before reaching B. The algorithm takes three public datasets as inputs—NHDPlus canal geometries, USDA Crop Sequence Boundaries, and digital elevation models—and produces a normalized position measure $\in [0, 1]$ for each plot, where 0 denotes the canal headgate and 1 the tail end.

The construction proceeds through four stages: (i) building a directed canal graph from NHDPlus flowlines, (ii) linking each plot to its nearest canal segment, (iii) assembling a directed plot graph from canal-structure and spatial-neighbor edges, and (iv) computing positions via signed-hop BFS from canal-graph roots. The final graph covers 54,946 plots linked to approximately 160,000 canal graph nodes connected by 201,000 directed edges, organized into 3,362 weakly connected components across 1,162 canal blocks. Position is based on canal network topology, not Euclidean distance: two plots 100 m apart may have very different positions if they sit on different canal branches.

5.2 Canal Graph Construction

I extracted canal and ditch segments from the NHDPlus flowline layer, filtering by feature code to retain canals, ditches, and aqueducts while excluding ephemeral streams (fcode $\neq$ 46007), yielding approximately 547,000 segments. For each segment, I inferred flow direction by regressing elevation along the segment's coordinates using high-resolution DEM rasters. A statistically significant

Table 2: Irrigation District Institutional Characteristics

	Full Sample			*SW-Dependent Subsample*	
	GW-dom (≥70%)	Mixed (30–70%)	SW-dom (≤30%)	Senior (rank≤3)	Junior (rank>3)
N districts	16	22	11	10	23
N plots	16,545	17,858	19,609	18,394	19,073
Panel A: Water Source					
Pct groundwater (mean)	82.4	53.7	13.2	20.9	48.6
(SD)	(8.5)	(10.8)	(3.9)	(19.2)	(16.4)
Panel B: Pricing Structure					
Volumetric (N districts)	1	7	1	1	7
Per-acre assessment (N)	13	11	7	8	10
Flat annual fee (N)	0	3	6	6	3
Tiered/progressive (N)	3	2	2	2	2
Panel C: Governance					
Shortage protocol (N)	3	11	3	3	11
Rotation schedule (N)	4	3	1	2	2
Within AMA (N)	7	10	0	2	8
Panel D: Water Rights					
Mean priority rank	5.0	4.9	2.9	2.1	5.2
(SD)	(2.3)	(1.4)	(1.9)	(1.0)	(1.2)
CR-shortage vulnerable (N)	2	7	7	5	9

Notes: Groundwater share computed as mean of plot-level `gw_share` (USGS NWAA) within each district. GW-dominant: ≥70%; Mixed: 30–70%; SW-dominant: ≤30%. Senior: Colorado River priority rank ≤ 3; Junior: rank > 3. Right panel restricted to SW-dominant and Mixed districts (≤70% GW). Pricing categories are not mutually exclusive. Districts with fewer than 50 plots excluded. CR-shortage vulnerable = subject to Colorado River Tier 1 shortage reductions.

negative slope ($R^2 > 0.6$) indicates the segment was digitized in the direction of flow; a positive slope triggers reversal. Where the NHDPlus Value Added Attributes (VAA) flow table provided ground-truth flow connections between segments, I used these to correct DEM-based inferences.

I connected segments into a directed graph G_{canal} by spatial proximity: if the downstream endpoint of segment A falls within 100 m of the upstream endpoint of segment B, I added an edge $A \to B$. Two additional processing steps address infrastructure complexity. First, I merged parallel canal segments—common in NHDPlus where a single physical canal is represented by multiple overlapping polylines—into unified nodes, reducing redundancy without losing connectivity. Second, I split canals at T-junctions where lateral canals connect mid-segment, ensuring that the graph correctly routes water through branching points rather than treating the entire parent segment as a single node.

5.3 Plot–Canal Linking

I matched each of the 54,946 agricultural plots to its nearest canal segment using an elevation-adjusted cost function:

$$\text{Cost}_{ij} = \text{Distance}_{ij} + \max\big(0,\ \alpha_{pump} \times (\text{Elev}_{plot_i} - \text{Elev}_{canal_j})\big) \tag{1}$$

where i indexes plots, j indexes canal segments, and $\alpha_{pump} = 30$ penalizes assignments requiring uphill pumping—each meter of elevation difference where the plot sits above the canal adds 30 meters of equivalent distance cost. When the canal is at equal or higher elevation (gravity-fed delivery), the cost reduces to Euclidean distance alone. In practice, 43% of plots sit above their nearest canal (median elevation difference of 1.1 m among these plots), generating a median penalty of approximately 34 m—modest relative to the median plot–canal distance of 398 m. However, for plots 5 m above a canal, the penalty reaches 150 m (38% of median distance), sufficient to redirect the assignment to a more elevation-compatible canal nearby. This function thus balances physical proximity against elevation feasibility: a plot may bypass a nearer canal at lower elevation in favor of a slightly more distant canal at a favorable gradient. The assignment identifies both the matched canal segment and the specific *intake point*—the location along the canal geometry closest to the plot centroid, projected onto the canal linestring. For plots sharing the same canal segment, I ordered them by projecting their intake points along the canal's flow direction.

5.4 Plot Graph and Position Computation

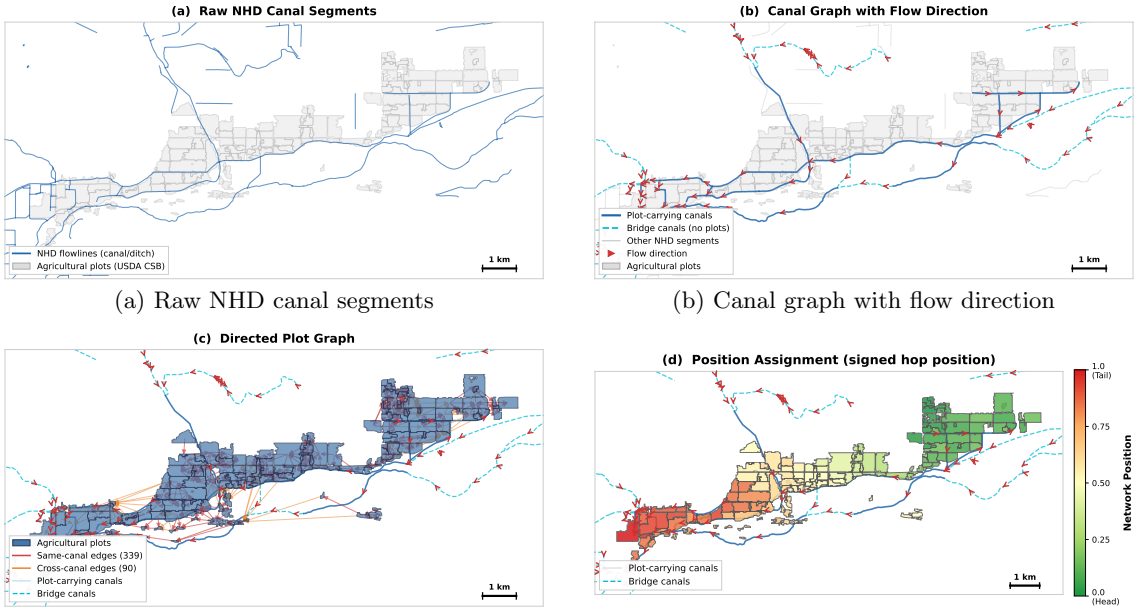
I constructed the directed plot graph G_{plot} from two types of edges. *Canal-structure edges* connect plots whose canal segments are linked in G_{canal}: if a directed path exists from $canal_A$ to $canal_B$, all plots on A are upstream of all plots on B. For plots on the same canal, their ordering follows the intake-point projection along flow direction. *Spatial-neighbor edges* connect plots identified as geographic neighbors via raster flood-fill at approximately 11 m resolution, providing local connectivity where canal paths are indirect.

To compute positions within each WCC, I identified a *root plot*—the most upstream plot in the network—using the following procedure: among all source canals (in-degree = 0 in G_{canal}), I selected the canal whose associated plots have the highest mean elevation, then chose the in-degree-0 plot on that canal with the most graph descendants. From this root, I performed a signed-hop breadth-first search: each step downstream increments the hop count by $+1$, while steps upstream along secondary inlets decrement by -1. The resulting signed hop counts are normalized to $[0, 1]$ within each WCC by min-max scaling.

Finally, I grouped WCCs into 1,162 canal blocks using DBSCAN spatial clustering (eps = 5 km, min_samples = 3) of plot locations, as described in Section 4. Figure 3 illustrates the three-step construction: canal network with intake points, directed plot graph, and BFS-derived positions.

5.5 From Graph to Regression: Two Position Measures

The water-graph produces a single static network structure, but the empirical strategy requires two complementary position measures that serve different identification purposes.



(a) Raw NHD canal segments

(b) Canal graph with flow direction

(c) Directed plot graph

(d) Position assignment ($0 = \text{head} \rightarrow 1 = \text{tail}$)

Figure 3: Water-graph construction for WCC 20_17 (68 plots, 20 canals). (a) Raw canal/ditch segments from NHDPlus. (b) Canal graph: WCC canals (blue) with DEM-inferred flow direction arrows (red). (c) Directed plot graph: canal-structure edges (red) and spatial-neighbor edges (orange). (d) Position via signed-hop BFS from root (★), normalized 0–1 within WCC.

The *time-invariant* measure, `hops_relative_rc`, is each plot's signed-hop BFS position normalized to $[0, 1]$ within its WCC. This captures *where* a plot sits on the canal—its structural location. Because it is fixed over time, it is absorbed by plot fixed effects in the panel specifications and can only be estimated in the cross-section, where identification comes from comparing plots at different positions within the same canal-year cell.

The *time-varying* measure, $\text{log_pos_far}_d = \ln(1 + N_{it,>d}^{\text{active, upstream}})$, counts the number of active upstream plots beyond Euclidean distance d in year t. This measure changes year to year as plots enter or exit fallow, creating within-plot variation that *survives* plot fixed effects. It captures *how much upstream competition* a plot faces in a given year—not where it sits, but how crowded the canal is above it. This is the treatment variable in the donut specifications.

The distinction matters for interpretation. A negative coefficient on `hops_relative` says: "plots further from the headgate have worse outcomes." A negative coefficient on log_pos_far_d says: "when more upstream plots are actively irrigating, downstream outcomes worsen"—a direct test of the water competition mechanism. The donut restriction ($d > 0$) ensures that this upstream activity variation is far enough away to be plausibly exogenous to the downstream plot's local conditions.

5.6 Validation

I validated the water-graph through several approaches. Visual inspection of plot–canal associations overlaid on satellite imagery confirmed that assignments match observable field connections. Inferred flow directions agreed with NHDPlus VAA ground truth where available, and a manual audit of 100 random canals showed 100% agreement between the `flowdir` column and the start-point coordinates. A direction-consistency reconciliation step aligns canal graph edges with elevation-based flow directions, using the graph structure as ground truth where available and falling back on spatial elevation gradients otherwise. After this reconciliation, residual direction disagreements are concentrated in NHDPlus flow table errors and T-junction ambiguities at inter-segment junctions; same-canal ordering—which determines within-canal plot positions—is unaffected.

Key limitations include: (i) the graph captures static infrastructure and does not reflect seasonal canal closures or operational changes; (ii) accuracy depends on NHDPlus completeness and DEM precision, which vary regionally; (iii) the model assumes gravity-fed delivery and does not account

for all control gates, check structures, or complex pumping systems; and (iv) 171 plots in Yuma County grid-loop irrigation structures receive valid positions but are flagged (`in_cycle = True`) and excluded in robustness checks. Despite these limitations, the water-graph provides the first systematic, reproducible framework for measuring canal network position at the scale of an entire state.

6 Identification Strategy

6.1 Cross-Section Specification

I begin with a cross-sectional regression that establishes whether network position predicts crop outcomes and whether drought amplifies the gradient:

$$Y_{irt} = \alpha_1 \cdot \text{Position}_{ir} + \alpha_2 \cdot \left(\text{Position}_{ir} \times \text{Drought}_{rt}\right) + \beta X_{ir} + \gamma_{\text{crop},t} + \delta_{d,c,t} + \epsilon_{irt} \tag{2}$$

where i indexes plots, r indexes canal networks, and t indexes years. Y_{irt} is a satellite-derived crop outcome (NDMI, GCVI, or ET). Position_{ir} = `hops_relative` $\in [0, 1]$ is the time-invariant normalized network position (0 = headgate, 1 = tail). Drought_{rt} is a binary indicator equal to one when the upstream 24-month SPEI falls below -0.6. X_{ir} includes 42 plot-level controls: plot size, slope, well count, SPEI drought indices, soil characteristics (salinity, clay, organic carbon, bulk density, erosion factors, pH), groundwater access (IGFR distance, pumping dummy, pump stress index), water rights seniority (nearest priority year), market access (distances to city, highway, rural roads), network topology (in-degree, out-degree, head/tail dummies), protected area distances, and canal type interactions.[6] $\gamma_{\text{crop},t}$ denotes crop $\times$ year fixed effects. $\delta_{d,c,t}$ denotes irrigation district $\times$ WCC $\times$ year fixed effects, which absorb all time-varying institutional heterogeneity at the district-canal level, including district-specific drought responses.[7] Standard errors are clustered two-way by crop-year and plot (CSBID).

I predict $\alpha_1 < 0$ (downstream plots have worse outcomes) and $\alpha_2 < 0$ (drought amplifies the position penalty). However, the interaction faces a power limitation in this specification: the district $\times$ WCC $\times$ year fixed effects absorb most of the drought variation needed to identify α_2, since within a given canal network and year, drought conditions are nearly uniform across plots. As shown in Section 7.1, the position level effect (α_1) is strongly significant but the interaction (α_2) is small and insignificant in the cross-section. The panel specification with the donut strategy (Section 6.3) provides the necessary within-plot temporal variation to identify the drought amplification mechanism.

A separate limitation is that Position_{ir} is time-invariant: it is absorbed by plot fixed effects and therefore cannot be estimated in a panel specification. The cross-section cannot rule out selection on unobserved time-invariant plot characteristics correlated with canal position.

6.2 The Endogeneity Problem: Fallow Contagion

Moving to a panel specification with plot fixed effects requires a time-varying position measure. The natural candidate is the log count of active upstream plots, which changes year to year as plots enter or exit fallow. But this creates an endogeneity problem. When a nearby upstream plot goes fallow, two things happen simultaneously: (1) the downstream plot's measured upstream competition decreases—its "position" improves; and (2) the same local shock that caused the neighbor to fallow (drought, pest pressure, market conditions) likely affects the downstream plot as well. The observed correlation between upstream activity and downstream outcomes confounds genuine water competition with shared local shocks.

This concern is not hypothetical. Figure 4 shows that fallow status is spatially correlated far above random chance—correlation persists beyond 30 km during drought years. Fallowing is driven

[6] The full control list is reported in the table notes.

[7] The FE variable `fe_irr_wcc_year` interacts irrigation district ID, WCC identifier, and year.

by shared local conditions, not independent decisions, making nearby fallow changes endogenous to the downstream plot's environment.

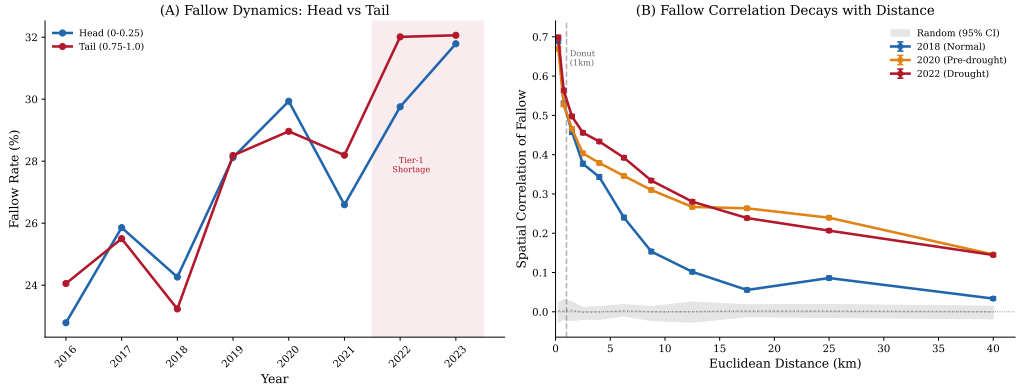


Figure 4: Fallow dynamics and spatial correlation. (A) Mean fallow rate by position tercile (head/mid/tail), 2016–2023. (B) Spatial autocorrelation of fallow status vs. distance: correlation persists far above random (gray band) and extends beyond 30 km during drought years.

Not all fallow is equally problematic. I distinguish two types using a machine learning classifier trained on 8-year crop sequence patterns (details in Appendix A): *non-rotation fallow* (predicted rotation probability < 0.3), which reflects water constraints or economic stress, and *rotation fallow* (probability > 0.7), which reflects planned agronomic rotation. Only stress-driven fallow poses an endogeneity threat; rotation fallow is agronomically predetermined and largely orthogonal to current-year water conditions. The spatial fallow controls in the donut specification separate these two channels.

6.3 The Donut Solution

The donut identification strategy resolves the fallow contagion problem by restricting the treatment variable to *far* upstream variation while controlling for *near* spatial fallow activity. The intuition is illustrated in Figure 5: a nearby plot going fallow is endogenous (shared local shock), but a far upstream plot going fallow provides clean variation—far enough that local conditions are independent, but connected via the canal network so that its fallowing reduces upstream water demand and frees water for downstream users.

The specification is:

$$Y_{it} = \beta_1 \cdot \text{log_pos_far}_{d,it} + \beta_2 \cdot \left(\text{log_pos_far}_{d,it} \times \text{Drought}_t\right) + \theta_1 \cdot \text{SpatialFallow}^{\text{nr}}_{d,it} + \theta_2 \cdot \text{SpatialFallow}^{\text{rot}}_{d,it} + \beta X_{it} + \alpha_i + \gamma_{\text{crop},t} + \epsilon_i \tag{3}$$

where $\text{log_pos_far}_{d,it} = \ln(1 + N^{\text{active, upstream}}_{it,>d})$ counts active upstream plots beyond Euclidean distance d in year t. This is the treatment variable—it captures far upstream water demand that is plausibly exogenous to the downstream plot's local conditions. $\text{SpatialFallow}^{\text{nr}}_{d,it}$ and $\text{SpatialFallow}^{\text{rot}}_{d,it}$ are the shares of all plots within Euclidean radius d that are non-rotation fallow and rotation fallow, respectively. These spatial fallow shares are computed over all plots within radius d regardless of canal-network direction, because the local shocks they absorb—weather, pests, market conditions—propagate spatially, not directionally along the canal. α_i denotes plot fixed effects, absorbing all time-invariant characteristics (soil, elevation, baseline productivity, farmer type). $\gamma_{\text{crop},t}$ denotes crop $\times$ year fixed effects. I estimate the specification at distance thresholds $d \in \{0, 1, 2, \ldots, 10\}$ km in 1 km increments.

The treatment variable log_pos_far_d encodes far upstream activity directly—it counts active plots beyond d. I therefore do not include separate far fallow share controls, which would double-count the same variation. The spatial fallow shares serve a distinct purpose: absorbing local shocks within radius d that could contaminate the treatment variable but are not part of the directional water competition mechanism.

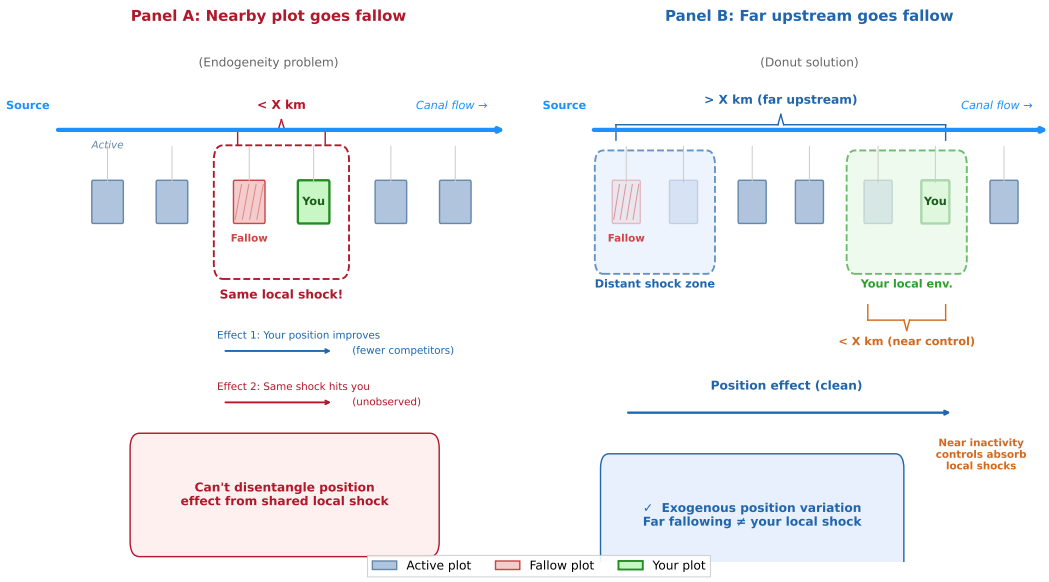


Figure 5: The "donut" identification strategy. A nearby plot going fallow (left) creates endogeneity—the same local shock affects both plots. A far upstream plot going fallow (right) provides clean variation: far enough that local conditions are independent, but connected via the canal so its fallowing reduces upstream water demand.

Three features of this design support a causal interpretation. First, plot fixed effects absorb all permanent plot characteristics. Second, spatial fallow shares absorb time-varying local shocks within Euclidean radius d—the region excluded from the treatment variable precisely because it is contaminated by shared local conditions. Third, what remains in log_pos_far_d is variation in upstream water demand driven by fallowing decisions of plots far enough to face independent local conditions, but connected through the canal network.

If identification is valid, the drought × position coefficient should strengthen monotonically as d increases, because larger donuts exclude more contaminated near-distance variation. If coefficients instead weaken or become unstable with distance, the effect is likely spurious local correlation rather than water competition.

6.4 Downstream Placebo

I conduct two validation tests to confirm that the donut specification captures directional water competition rather than symmetric spatial shocks.

Fallow shares symmetry test. If fallow spatial correlation reflects directional water competition (upstream extraction reducing downstream supply), upstream fallow should predict a plot's own fallow status more strongly than downstream fallow. If instead it reflects symmetric local shocks (shared weather or market conditions), upstream and downstream fallow shares should have similar predictive power. I find that upstream and downstream non-rotation fallow coefficients are of comparable magnitude, confirming that fallow contagion is driven by local shocks, not directional water competition. This validates the need for the donut: near fallow variation is contaminated by symmetric local shocks and must be controlled for, not used as the treatment variable.

Downstream position placebo. I construct a mirror-image downstream position measure—the log count of active *downstream* plots beyond distance d—and include both upstream and downstream position × drought interactions in the same regression. If upstream captures genuine water competition, the upstream coefficient should survive while downstream attenuates, since a plot's downstream neighbors cannot affect its water supply. If both coefficients survive equally, the pattern reflects symmetric spatial shocks. At the 10 km cutoff, the upstream GCVI drought

interaction (-0.054^{***}) is roughly four times larger in magnitude than the downstream coefficient (-0.013, insignificant). In the horse race, upstream coefficients strengthen monotonically with distance while downstream coefficients flatten toward zero—consistent with directional water competition.

Together, the monotonic donut pattern, the fallow shares symmetry test, and the downstream position placebo support the interpretation that the donut specification isolates genuine upstream water competition from spatially correlated local shocks.

7 Results

7.1 Cross-Section: Position Effects

Table 3 presents cross-sectional estimates of the position–drought–crop health relationship. The position level effects are significant across all three outcomes (NDMI -0.007^{**}, GCVI -0.039^{***}, ET -0.026^{**}): moving from the canal headgate to the tail reduces crop health within the same irrigation district, canal network, crop type, and year. In economic magnitude, the full head-to-tail GCVI reduction of 0.039 corresponds to 6.2% of the active-plot standard deviation (0.629); a one-standard-deviation increase in position (0.26) implies a 1.6% SD reduction—modest but persistent across every year in the panel. However, the drought interaction terms are small and statistically insignificant across all three outcomes. The interaction lacks power in the cross-section because the irrigation district × WCC × year fixed effects absorb most of the drought variation needed to identify it—within a given canal network and year, drought conditions are nearly uniform across plots. The panel specification with plot fixed effects and the donut strategy (Section 7.2) provides the necessary within-plot temporal variation to identify the drought amplification mechanism.

Table 3: Network Position × Drought and Crop Health (Cross-Section)

	NDMI (1)	GCVI (2)	ET (3)
hops_relative_rc	**−0.007****	**−0.039*****	**−0.026****
	(0.003)	(0.011)	(0.010)
× Drought_{24M}	−0.001	−0.013	−0.009
	(0.005)	(0.018)	(0.016)
Drought_{24M}	0.000	0.007	0.001
	(0.004)	(0.016)	(0.015)
FE: Crop × Year + Irr. Dist. × WCC × Year	✓	✓	✓
Observations	297,019	297,019	297,019
R^2	0.439	0.484	0.499

Notes: Active plots only. Position = `hops_relative_rc` $\in [0, 1]$. $\text{Drought}_{24M} = 1$ if upstream SPEI-24 < -0.6. FE: Crop×Year + Irrigation District×WCC×Year. 42 controls. Two-way clustered SE (Crop×Year + CSBID). $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

7.2 Donut Results: Distance Cutoff Analysis

Table 4 presents the donut panel estimates at six selected distance thresholds (Full path, >1, >2, >3, >5, and >10 km). Two patterns emerge. First, the position *level* effect generally strengthens as the donut radius increases, though the pattern is not strictly monotonic at every threshold. At the full path ($d = 0$), the level coefficient is modest and often insignificant after plot fixed effects absorb cross-sectional position variation; as d increases, the treatment variable captures progressively cleaner far upstream variation, and the coefficients on GCVI and ET grow larger and more significant at intermediate thresholds.

Second, the drought × position interaction strengthens with distance. At the full path, the GCVI interaction is -0.028^{***} and the ET interaction is -0.019^{**}; by $d > 10$ km, these roughly

double to -0.061^{***} and -0.047^{***} respectively (ratios of 2.2× and 2.5×; NDMI shows a similar 2.0× pattern). To anchor the economic magnitude: the within-plot standard deviation of $\text{log_pos_far}_{5\text{km}}$ is 0.229 (after demeaning by plot fixed effects), reflecting year-to-year variation in upstream activity driven by fallowing decisions. At the $d > 5$ km threshold, a one-within-SD increase in far upstream activity during drought reduces GCVI by $0.037 \times 0.229 = 0.008$, or 2.3% of the within-plot GCVI standard deviation (0.368)—a modest effect per unit of treatment variation, but one that reflects conservative identification: plot fixed effects absorb the large cross-sectional position differences (6.2% of SD head-to-tail in Table 3), and the donut identifies only from residual year-to-year fluctuations in upstream competition. That these small temporal perturbations produce statistically significant effects suggests the underlying water competition mechanism is strong. The two specifications are complementary: the cross-section provides the full magnitude of the position penalty, while the panel with the donut provides the cleanly identified causal estimate. The overall strengthening pattern—despite a slight dip at the 1 km threshold for GCVI—is the core evidence for the water competition mechanism: if the effects were driven by spatially correlated local shocks, they would weaken (not strengthen) as contaminated near-distance variation is excluded.

Table 4: Donut Results: Position Effects by Distance Cutoff

	Full	>1km	>2km	>3km	>5km	>10km
Position × *Drought*$_{24M}$						
NDMI	-0.007^{***}	-0.007^{**}	-0.008^{**}	-0.008^{**}	-0.009^{**}	-0.014^{***}
	(0.002)	(0.003)	(0.003)	(0.003)	(0.004)	(0.005)
GCVI	-0.028^{***}	-0.025^{**}	-0.029^{***}	-0.031^{***}	-0.037^{***}	-0.061^{***}
	(0.009)	(0.010)	(0.011)	(0.012)	(0.013)	(0.016)
ET	-0.019^{**}	-0.023^{**}	-0.025^{**}	-0.027^{**}	-0.027^{**}	-0.047^{***}
	(0.008)	(0.010)	(0.010)	(0.011)	(0.013)	(0.017)
Position Level						
NDMI	-0.012^{***}	-0.008^{***}	-0.008^{***}	-0.009^{***}	-0.008^{**}	-0.006
GCVI	-0.026	-0.021^{*}	-0.030^{***}	-0.033^{***}	-0.028^{**}	-0.040^{**}
ET	-0.034^{**}	-0.034^{***}	-0.039^{***}	-0.042^{***}	-0.035^{***}	-0.014
N	295K	246K	198K	162K	117K	66K

Notes: Active plots only. Treatment = ln(1 + active upstream count beyond d). Spatial fallow shares (non-rotation + rotation) within radius d control for local shocks. FE: Crop×Year + Plot. Two-way clustered SE (Crop×Year + CSBID). $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Figure 6 plots the drought × position coefficients continuously across 11 distance thresholds (0–20 km) for GCVI and ET. The monotonic strengthening is visually striking: confidence intervals exclude zero beyond approximately 3 km, and coefficients stabilize at 7–10 km, suggesting that the remaining far upstream variation is free of local-shock contamination.

7.3 Downstream Placebo

The downstream position placebo (Appendix Table A6) confirms the directional interpretation. I include both upstream and downstream log active counts × drought in the same regression. At the 10 km cutoff, the upstream GCVI drought interaction (-0.054^{***}) is roughly four times larger than the downstream coefficient (-0.013, insignificant). Figure 7 shows that upstream coefficients strengthen monotonically with distance while downstream coefficients flatten toward zero—consistent with directional water competition, not symmetric spatial shocks. The fallow shares symmetry test (Appendix Table A5) further validates the donut design: upstream and downstream fallow shares predict own fallow status with comparable magnitude, confirming that fallow contagion reflects local shocks rather than directional water competition.

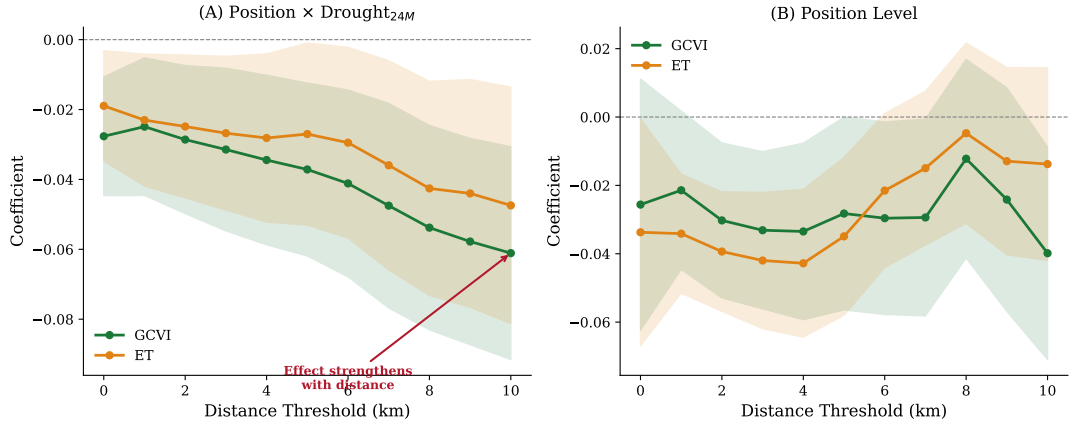


Figure 6: Donut coefficients across distance thresholds. Each point shows the drought × position interaction coefficient at threshold d. Bands show 95% CI. Same specification as Table 4.

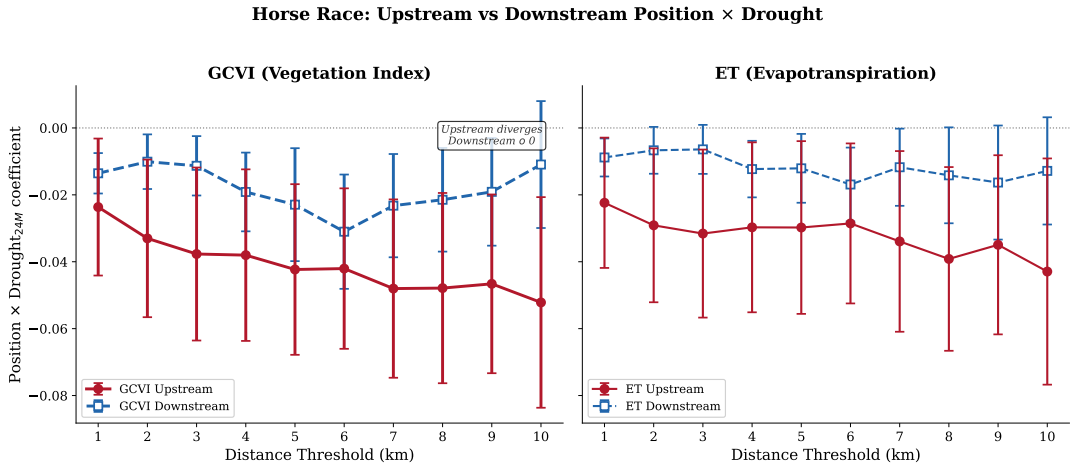


Figure 7: Upstream vs. downstream horse race across distance thresholds. Both upstream and downstream position × drought in the same regression. Upstream strengthens monotonically; downstream flattens toward zero.

7.4 Institutional Heterogeneity

Triple Interactions: Which Institutions Attenuate Inequality?

Table 5 presents triple interaction estimates (Position × Drought × Institutional Variable) that identify which district-level characteristics moderate the position–drought gradient. Continuous institutional variables are standardized so coefficients reflect the effect of a one-standard-deviation change. In the full specification (column 4), two institutional features significantly attenuate the drought × position penalty: irrigation district membership ($+0.050^{***}$) and groundwater share ($+0.022^{***}$). Volumetric water pricing is strongly significant in the bivariate specification ($+0.034^{***}$, column 3) but is absorbed when all institutional controls enter jointly, suggesting it correlates with district membership.

Table 5: Triple Interaction: Position × Drought × Institutional Characteristics

	(1) Physical + GW	(2) Rules	(3) Pricing + Belong	(4) All
Soil salinity	−0.008			−0.008
	(0.007)			(0.006)
Canal tail-end reach	0.012^{**}			0.010^{*}
	(0.006)			(0.005)
GW share	0.012			0.022^{***}
	(0.008)			(0.008)
GW share × AMA				−0.017
				(0.018)
Rotation schedule		−0.008		−0.015
		(0.009)		(0.021)
Within AMA		0.046^{***}		0.036
		(0.015)		(0.030)
Volumetric pricing			0.034^{***}	−0.000
			(0.010)	(0.012)
Within district boundary			0.035^{**}	0.050^{***}
			(0.015)	(0.014)
Base: Pos × Drought	-0.014^{*}	-0.041^{***}	-0.053^{***}	-0.054^{***}
N	247,792	247,681	247,681	247,610
R^2	0.7255	0.7246	0.7247	0.7265

Notes: DV: GCVI. Same donut specification as Table 4 at $d > 1$ km. Plot FE + Crop×Year FE. Institutional variables interacted with Position × Drought. Positive = attenuates the negative drought×position gradient. Two-way clustered SE. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Water Source Split

Table 6 splits the sample by district water source composition. The position–drought gradient is concentrated exclusively in surface-water-dominant districts (GCVI: -0.043^{***}, ET: -0.042^{***}). In groundwater-dominant districts, the interaction is approximately zero—wells eliminate the canal position penalty entirely. Mixed districts show intermediate effects. This finding confirms Mechanism 3 from the conceptual framework: groundwater access breaks the dependence on canal position by providing an alternative water source immune to upstream extraction.

Senior vs. Junior Water Rights

Within surface-water-dominant districts, Table 7 further splits by water rights seniority. Senior-rights districts show small or insignificant position × drought interactions—legal priority protects tail-end users as designed. Junior-rights districts face substantially steeper drought amplification (GCVI: -0.040^{**}, ET: -0.047^{***}). Among junior-rights districts, those subject to formal shortage protocols experience the steepest gradient (GCVI: -0.114^{***}, ET: -0.107^{***}), suggesting that

Table 6: Heterogeneity by Water Source: GW-Dominant, Mixed, SW-Dominant

	GW-dominant		**Mixed**		**SW-dominant**	
	GCVI	ET	GCVI	ET	GCVI	ET
Pos × Drought (base)	+0.009	+0.005	+0.006	-0.003	-0.043***	-0.042***
	(0.009)	(0.006)	(0.005)	(0.003)	(0.015)	(0.015)
Position (level)	-0.042*	-0.018	+0.023*	-0.001	-0.037**	-0.054***
	(0.024)	(0.023)	(0.013)	(0.009)	(0.015)	(0.018)
Institutional triple interactions (Pos × Drought × X):						
Rotation schedule			-0.005	-0.003	+0.023	+0.062***
			(0.008)	(0.006)	(0.027)	(0.021)
Shortage protocol	-0.049***	-0.007	-0.006	-0.004	-0.042**	-0.039**
	(0.017)	(0.010)	(0.009)	(0.007)	(0.018)	(0.019)
Volumetric pricing			+0.003	-0.003	+0.038	+0.029
			(0.009)	(0.005)	(0.030)	(0.029)
Tiered rates	+0.059*	-0.018*	-0.005	-0.003	+0.032**	+0.032**
	(0.031)	(0.010)	(0.008)	(0.007)	(0.016)	(0.016)
Within district	+0.018	-0.009	-0.002	-0.004	+0.066***	+0.069***
	(0.019)	(0.009)	(0.009)	(0.005)	(0.020)	(0.020)
N	37,479	37,479	83,623	83,623	123,826	123,826
Districts	21	21	12	12	14	14

GW-dom: ≥70% GW; Mixed: 30–70%; SW-dom: ≤30% GW. FE: Crop×Year + Plot.
Institutional rows show Pos×Drought×X within each subsample. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$

Notes: GW-dominant: ≥70% groundwater; Mixed: 30–70%; SW-dominant: ≤30% groundwater. Same donut specification as Table 4 at $d > 1$ km. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

these protocols may concentrate shortfalls on downstream users rather than distributing them equitably.

Volumetric Pricing: Metering and Incentives

The volumetric pricing result deserves closer interpretation. In the pooled triple interaction (Table 5), volumetric pricing attenuates the drought × position gradient ($+0.034^{***}$ in the bivariate specification) but is absorbed with full controls. However, the effect is most striking within the subsample where the position penalty is most acute—junior-rights, surface-water-dominant districts—where volumetric pricing is the strongest attenuator (GCVI: $+0.091^{***}$, ET: $+0.088^{***}$ in Table 7), more than offsetting the junior-rights base interaction (GCVI: -0.040^{**}, ET: -0.047^{***}).

This subsample-specific effect likely reflects two complementary channels. First, districts that charge per acre-foot must meter individual deliveries to compute bills—the pricing structure *requires* measurement infrastructure that directly reduces the scope for undetected over-extraction described in Mechanism 1. Districts with flat-rate or per-acre pricing have no operational need for turnout-level metering, leaving the measurement error channel unchecked. Second, the price signal itself creates a marginal-cost incentive: each additional unit of water diverted has a direct cost, discouraging extraction beyond the allocation even when monitoring is imperfect.

These two channels are complementary—metering without pricing detects but does not penalize over-extraction, while pricing without metering creates incentives that cannot be enforced. That volumetric pricing matters precisely where the theory predicts the strongest position penalty—junior-rights districts dependent on surface water during drought—provides additional support for the water competition mechanism and points to a concrete, implementable policy lever.

7.5 Robustness

The main results are robust to several alternative specifications. Alternative drought measures (3-month SPEI, continuous SPEI) produce qualitatively similar patterns (Appendix Table A7). Excluding plots that enter or exit the sample during the panel period does not materially affect

Table 7: Within SW Districts: Senior vs. Junior Rights and Shortage Protocols

	Senior rights		**Junior rights**	
	GCVI	ET	GCVI	ET
Pos × Drought (base)	-0.004	-0.002	-0.040**	-0.047***
	(0.005)	(0.004)	(0.018)	(0.017)
Position (level)	-0.004	-0.021*	-0.048**	-0.091***
	(0.013)	(0.012)	(0.020)	(0.026)
Institutional triple interactions (Pos × Drought × X):				
Rotation schedule	+0.006	+0.009	+0.012	+0.004
	(0.009)	(0.006)	(0.030)	(0.029)
Shortage protocol	-0.004	-0.003	-0.114***	-0.107***
	(0.008)	(0.006)	(0.025)	(0.023)
Volumetric pricing	+0.017	+0.012	+0.091***	+0.088***
	(0.011)	(0.009)	(0.020)	(0.016)
Tiered rates	+0.002	-0.002	+0.004	-0.000
	(0.009)	(0.008)	(0.037)	(0.036)
Within district	+0.003	+0.007	+0.084***	+0.078***
	(0.010)	(0.008)	(0.020)	(0.020)
N	143,621	143,621	63,819	63,819
Districts	15	15	11	11

Senior: priority rank ≤ 3; Junior: rank > 3. SW+Mixed districts only ($\leq$70% GW).
FE: Crop×Year + Plot. Cluster: Crop×Year + CSBID. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$

Notes: Sample restricted to SW-dominant and Mixed districts ($\leq$70% GW). Same donut specification as Table 4 at $d > 1$ km. Senior: priority rank ≤ 3; Junior: rank > 3. Shortage protocol = formal curtailment rules. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

the estimates (Appendix Table A11). The donut specification using area-weighted hop-distance measures instead of simple upstream counts yields similar monotonic strengthening (Appendix Table A9). Canal size interactions confirm that the position penalty is strongest in medium and large canal systems where water competition is most intense (Appendix Table A1). Full robustness tables are available in the appendix.

8 Conclusion

This study asks whether formal water rights eliminate the distributional inequality created by sequential access in canal irrigation networks. Using a water-graph methodology that maps upstream–downstream relationships among 54,946 agricultural plots across Arizona's canal systems, I measure each plot's network position and test whether downstream location causally affects crop outcomes during the Colorado River Basin's ongoing megadrought.

Four main findings emerge. First, downstream plots exhibit significantly lower moisture content (NDMI), chlorophyll vegetation (GCVI), and evapotranspiration (ET) than upstream plots within the same canal system, even after controlling for water rights seniority, soil quality, groundwater access, and 42 additional covariates. Second, a "donut" identification strategy—which isolates far upstream water demand (>1–10 km) while controlling for spatially correlated local shocks—confirms the causal interpretation: the drought × position interaction strengthens monotonically as contaminated near-distance variation is excluded, and a downstream placebo shows that downstream position coefficients are substantially attenuated relative to upstream coefficients. Third, drought roughly doubles the position penalty, concentrating shortfalls on tail-end users during the years when water constraints bind most tightly. Fourth, heterogeneity analysis reveals that the position–drought gradient exists exclusively in surface-water-dominant districts—groundwater access eliminates it entirely—and within surface-water districts, those with junior priority allocations face steeper drought amplification, while formal shortage protocols appear to concentrate rather than distribute shortfalls on downstream users.

These findings contribute to the literature on common-pool resource governance by providing the first large-scale empirical test of whether formal property rights resolve position-based inequality in developed-country irrigation systems. The answer is nuanced: senior water rights do protect tail-end users as designed, but junior-rights districts—which include the most drought-vulnerable communities on the Colorado River—experience persistent and growing distributional inequality that the prior appropriation doctrine does not address. This result connects to the developing-country evidence on tail-end deprivation (Wade 1988; Jacoby and Mansuri 2020; Jacoby, Mansuri, and Fatima 2021) while demonstrating that the problem persists even with formal legal frameworks, consistent with Coasian failures driven by thin water markets and high transaction costs.

The heterogeneity results point to actionable policy levers. Perhaps the most striking finding is that volumetric water pricing nearly fully offsets the drought × position gradient in the most vulnerable districts—junior-rights, surface-water-dominant systems where the penalty is largest. This result likely reflects a dual mechanism: districts charging per acre-foot must install turnout-level metering to compute bills, which directly reduces undetected over-extraction, while the price signal itself creates marginal-cost incentives against excess diversion. That pricing works precisely where the theory predicts the strongest position penalty provides both a test of the water competition mechanism and a concrete policy prescription. Groundwater investment eliminates the penalty entirely in GW-accessible districts by providing an alternative supply immune to canal position. Irrigation district membership attenuates the gradient through collective governance. Shortage protocols in their current form may need redesign: the finding that shortage-exposed junior-rights districts face the steepest gradient suggests that curtailment rules concentrate rather than distribute the burden of scarcity.

Several limitations warrant acknowledgment. The water rights data, recorded as points rather than polygons, introduces measurement error in controlling for legal priority. The water-graph captures static infrastructure and does not account for operational dynamics such as seasonal canal closures or rotational delivery adjustments. The donut strategy relies on Euclidean distance thresholds rather than canal distance for the near/far classification, which may not perfectly correspond to the boundary where local shock correlation decays. Finally, the 8-year panel (2016–2023) captures the intensification of the megadrought but provides limited pre-drought variation for estimating long-run equilibrium effects.

Future research could extend this analysis by examining whether downstream farmers adapt through crop choice or technology adoption, investigating land consolidation as a mechanism for internalizing upstream–downstream externalities, and exploring whether improved water measurement infrastructure reduces the scope for upstream over-extraction. The water-graph methodology is replicable using publicly available data (NHDPlus, USDA CSB, DEM) and could be applied to other surface water systems globally—including the developing-country settings where tail-end deprivation has been documented but rarely quantified at scale.

References

Arizona Department of Water Resources and Central Arizona Project. 2023. *Joint Briefing on Shortage Allocations.* Technical Report. Arizona Department of Water Resources, May. https://www.azwater.gov/sites/default/files/2023.05.05_Joint_ADWR_CAP_Shortage%20Briefing_v12.pdf.

Brocious, Ariana. 2021. *Arizona Changes 'Use It or Lose It' Water Law.* AZPM News, February. https://www.azpm.org/s/84588-arizona-changes-use-it-or-lose-it-water-law/.

Burness, H Stuart, and James P Quirk. 1979. "Appropriative Water Rights and the Efficient Allocation of Resources." *American Economic Review* 69 (1): 25–37.

Cobourn, Kelly M., Xinde Ji, Siân Mooney, and Neil F. Crescenti. 2022. "The Effect of Prior Appropriation Water Rights on Land-Allocation Decisions in Irrigated Agriculture." *American Journal of Agricultural Economics* 104 (3): 947–975.

Cody, Kelsey C. 2018. "Upstream with a shovel or downstream with a water right? Irrigation in a changing climate." *Environmental Science & Policy* 80:62–73.

Das, Sabyasachi, and Souvik Dutta. 2026. "Decentralization and Spatial (Mis-)Allocation of Irrigation Water." *Agricultural Economics* 57 (1): e70074. https://doi.org/10.1111/agec.70074.

Feldman, Eli. 2004. "Death Penalty for Water Thieves." *U. Denv. Water L. Rev.* 8:1.

Harris, Lena. 2024. "Farmer Response to Policy Induced Water Reductions: Evidence from the Colorado River." *Journal of Environmental Economics and Management* 125:102986.

Jacoby, Hanan G., and Ghazala Mansuri. 2020. "Enforcement for Sale: Lobbying for Water in Pakistan's Indus Basin." *Journal of Public Economics* 190:104261.

Jacoby, Hanan G., Ghazala Mansuri, and Freeha Fatima. 2021. "Decentralizing Corruption: Irrigation Reform in Pakistan." *Journal of Public Economics* 202:104497.

Kuwayama, Yusuke, Alexandra Thompson, Richard Bernknopf, Benjamin Zaitchik, and Peter Vail. 2019. "Estimating the impact of drought on agriculture using the US Drought Monitor." *American Journal of Agricultural Economics* 101 (1): 193–210.

Li, Man, Wenchao Xu, and Tingju Zhu. 2019. "Agricultural Water Allocation under Uncertainty: Redistribution of Water Shortage Risk." *American Journal of Agricultural Economics* 101 (1): 134–153.

Lobell, David B, George Azzari, Marshall Burke, Sydney Gourlay, Zhenong Jin, Talip Kilic, and Siobhan Murray. 2020. "Eyes in the sky, boots on the ground: Assessing satellite- and ground-based approaches to crop yield measurement and analysis." *American Journal of Agricultural Economics* 102 (1): 202–219.

Manero, Ana, Sarah Ann Wheeler, Alec Zuo, and Makarius Mdemu. 2019. "Exploring the Head Versus Tail-End Dichotomy on Yield and Farm Incomes in Smallholder Irrigation Schemes in Tanzania." *Water Resources Research* 55 (5): 4322–4342.

Mieno, Taro, Timothy Foster, Shunkei Kakimoto, and Nicholas Brozović. 2024. "Aquifer depletion exacerbates agricultural drought losses in the US High Plains." *Nature Water,* 1–11.

Murgai, Rinku, Paul Winters, Elisabeth Sadoulet, and Alain de Janvry. 2002. "Localized and incomplete mutual insurance." *Journal of Development Economics* 67 (2): 245–274.

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

Rafey, Will. 2023. "Droughts, Deluges, and (River) Diversions: Valuing Market-Based Water Reallocation." *American Economic Review* 113 (2): 430–471.

Stringam, Blair L., Thomas C. Sauer, and C. Albert Pugh. 2003. "Accurate water delivery using a simplified automated farm turnout." *Irrigation and Drainage* 52 (4): 355–361.

The Nature Conservancy Arizona and Friends of the Verde River. 2020. *Understanding Surface Water Rights in the Verde Valley, Arizona.* Accessed: 2025-01-14. https://verderiver.org/wp-content/uploads/2020/10/understandingsurfacewaterrights_lowres.pdf.

Vicente-Serrano, Sergio M, Santiago Beguería, and Juan I López-Moreno. 2010. "A multiscalar drought index sensitive to global warming: the standardized precipitation evapotranspiration index." *Journal of climate* 23 (7): 1696–1718.

Wade, Robert. 1988. *Village Republics: Economic Conditions for Collective Action in South India.* Cambridge University Press.

Williams, A Park, Benjamin I Cook, and Jason E Smerdon. 2022. "Rapid intensification of the emerging southwestern North American megadrought in 2020–2021." *Nature Climate Change* 12 (3): 232–234.

Appendix

Appendix tables referenced in the text (Tables A1, A5, A6, A7, A9, A11) and details of the fallow rotation classifier (Appendix A) are available from the author upon request and will be included in the next version of this paper.