

NBER WORKING PAPER SERIES

MEASURING WATER MISALLOCATION IN CALIFORNIA

Will Rafey

Working Paper 35176

<http://www.nber.org/papers/w35176>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue

Cambridge, MA 02138

May 2026

This work was conducted with support from the UCLA Hellman Society of Fellows, the Council on Research of the Academic Senate of the Los Angeles Division of the University of California, and the Property and Environment Research Center Lone Mountain Fellowship. This work used computational and storage services associated with the Hoffman2 Cluster which is operated by the UCLA Office of Advanced Research Computing's Research Technology Group. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2026 by Will Rafey. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Measuring Water Misallocation in California

Will Rafey

NBER Working Paper No. 35176

May 2026

JEL No. D23, D61, L1, L2, Q1, Q24, Q25

ABSTRACT

This paper proposes and applies new methods to value water rights and assess misallocation across competing uses in California, the world's fourth largest economy. The empirical strategy combines detailed microdata on farms, evapotranspiration, historical water rights, and the hydrological flow network in order to isolate sources of inefficiency within the hydrological network, assess distributional implications of water access under current property rights, and evaluate alternative mechanisms for water reallocation.

Will Rafey

University of California, Los Angeles

Department of Economics

and NBER

rafey@ucla.edu

An online appendix is available at <http://www.nber.org/data-appendix/w35176>

A link to most current version on author's site is available at https://rafey.io/files/blue_farm.pdf

1 Introduction

Methods to value existing water property rights are critical to design and assess a range of increasingly important environmental and economic policies. This paper proposes and applies a unified approach to value water property rights in settings where they are rarely or never traded. The approach combines insights from recent work on agricultural production, equilibrium land use, and factor misallocation in a setting where water use can be inferred over time from detailed remote sensing data on crop choices, orchard ages, yields, prices, and local evapotranspiration data on weather and soil conditions used to manage irrigation in real time.

The analysis is framed in the context of California, which is intrinsically economically important, notorious for its legacy water rights, and a place with modern, high-resolution data newly collected over the last decade in response to growing concerns with water scarcity. Most of the state’s water is owned and used by sophisticated irrigators who plant a diverse collection of crops, but rarely trade water rights. On one hand, many express concern with water misallocation in California given that the earliest, or most “senior,” appropriative water rights—many of which date to the 19th century—have priority over later claims when water is scarce. On the other hand, California water infrastructure supports the largest agricultural sector and population of any state in the United States.¹ This paper is an attempt to think about water misallocation in this setting, which is especially important given new hydrological and climatic challenges that alter water scarcity and abundance across space and create new environmental concerns that may require reallocating existing sources of water to solve.

I focus on water users above and below the most important component of the water conveyance system in California, the Sacramento–San Joaquin River Delta, through which water must flow in order to reach the southern parts of the state. In particular, I use very detailed high-dimensional data, combined with a model of agricultural production and irrigation scheduling for nearly forty distinct crops—including annual crops like wheat, rice, and hay that must be replanted each year, and perennial crops such as almonds, oranges, and pistachios grown on trees that live for decades—to make some statements about the productivity of water use at a very fine level of resolution, and then think about what that

¹The service area of just one of the two largest surface water infrastructure projects in California contains what would be the eighth-largest economy in the world (DWR, 2023).

might or might not imply about misallocation.

To take advantage of this new high-resolution data, the paper develops a microeconomic model useful to value water rights and misallocation building on two key ingredients not present in prior work. First, I follow the insight of [Burness and Quirk \(1979\)](#) to imagine water rights as random variables that give owners access to certain volumes of water in certain states of the world. As I show in the data, water rights on paper differ significantly in terms of actual water available for diversion, both for water rights with the same face value and within the same water right over time.^{2,3} In the model, water rationing will arise differentially across the network and water right priority tiers, through the interaction of (a) equilibrium diversion decisions by owners of water rights and (b) exogenous idiosyncratic and aggregate shocks to water abundance and physical crop water requirements. Up to water conveyance and delivery costs, the optimal water allocation can be obtained when all users face a shadow value of water that corresponds to the marginal use value, which can be (and often is!) zero or even negative. Testing for misallocation in this incomplete model is challenging even in a fully-connected network without flow constraints: with sufficient heterogeneity across water users or water rights, unobserved planting or other adjustment costs can rationalize any water allocation.

Second, I embed a tractable set of [Rust \(1987\)](#) regenerative optimal stopping problems within the model of water rights, in order to model investments in capital varieties that require uninterrupted water inputs to survive. Water rights are typically perpetual, granting owners access to a sequence of annual water endowments over an infinite horizon, and the reliability of these rights are critical for high-value investment in cities, orchards, and vineyards that rely on continued water access across diverse conditions. I can take this to the data because, like [Rust \(1987\)](#) but unlike any prior work in California of which I am aware, I observe the full age structure of all orchards in California, as well as the panel of replanting decisions. This aspect of the analysis—complementary investment based on the

²The cutoffs that determine priority across competing water rights—the “priority date” or first year of the claim—differ considerably across users, with more senior water rights receiving approximately similar allocations in drought and non-drought years, but junior rights being severely curtailed. With constraints on water trade, this implies differential surface water rationing across the network and within tranches. Senior water rights are disproportionately owned by irrigators rather than cities, but many irrigators also rely on junior rights claims.

³Where feasible, farms also complement surface water with groundwater by incurring [Burlig et al. \(2024\)](#) pumping costs as functions of aquifer depth and extraction volume.

infinite sequence of annual water endowments—allows me to account for some of the adjustment costs that otherwise hinder the analysis of misallocation. That is, it allows me to microfound and estimate what are, in the simplest version of the model, unspecified planting costs that can rationalize any allocation of water. These complementary investments are also crucial to accommodate the fact that my estimates of water used for irrigation in the Central Valley exhibit high degrees of persistence, despite immense cyclical and stochastic components of surface water supply. Even during extreme drought, water volumes used for perennial crops, as well as many annual crops replanted each year, exhibit little variation.

In the model, farmers decide each year how to allocate their water endowment across annual and perennial crops. Conditional on allocating water to some perennial crops, such as trees in orchards or vines in vineyards, farmers then decide whether to maintain or cut down the trees or vines. As in Rust (1987), the productivity of the unit of capital investment depreciates over time—tree yield declines with age—and the farm faces an optimal stopping problem of when to replant, thereby renewing the asset, switching to a different variety of long-lived perennial crop, switching the status of the water to grow annual crops, or trading the water allocation to another user.

Understanding misallocation across water rights in this model, then, starts from each water right’s history of planting decisions, which, together with local growing conditions, determine its sequence of annual productivities. In this setting, there exists a wide range of cross-sectional dispersion in marginal products that need not correspond to any real inefficiency, for the same reason as in models of capital misallocation under uncertainty and adjustment costs (Asker *et al.*, 2014). Planting decisions are characterized by location of use, tree type, and tree age—e.g., some water will be embodied in newly-planted orchards that have yet to bear fruit, some in productive orchards, some in vineyards, and some used flexibly to grow annual crops—and determine the water’s annual productivity, as well as the adjustment costs related to its reallocation. For example, orchards typically remain productive for ten to thirty years, with some crops (e.g., pistachios) having no yield declines for more than one hundred years.

A key step in combining the two ingredients introduced above—state-contingent Burness and Quirk (1979) water rights and state-dependent Rust (1987) optimal investment—into a tractable empirical model is to distinguish between the perma-

nent and transient components of a water right. In general, variable water rights make the optimal crop allocation nonseparable across fields, I cannot directly apply the canonical [Scott \(2013\)](#) approach to simplify the general equilibrium dynamics of land use by studying a continuum of independent field-level choices. However, when water, not land, is the binding constraint, I can analyze a more tractable problem in “water space,” where units of water rather than land are assigned to different crops, making the problem separable even when water constraints bind. As these water-level decisions can be made independently from one another, I can leverage the computationally tractable [Berry \(1994\)](#) discrete choice apparatus and its extensions (e.g., [Gowrisankaran and Rysman, 2012](#)) to value water at an extremely fine level of spatial and temporal resolution, allowing for unobserved heterogeneity in planting costs that can help explain the diverse allocations of water to crops. Unlike existing revealed-preference models of land use, which will rationalize inefficient uses of water through differences in farmer tastes for crops, this paper’s model will not directly foreclose questions of water misallocation because the assumption used to identify planting costs is that, given the various water rights that they own, a farm will grow the most valuable crops.

Four main empirical findings flow from interpreting the data through the lens of the model and its estimates. First, using data on field-level planting decisions, crop evapotranspiration, agricultural yields and prices, the model allows me to estimate irrigation volumes and annual marginal products of water, or “water productivities,” for every field in California from 2014–2022. I find that estimated marginal products of water exhibit significant dispersion within and across regions, as well as remarkable persistence over time as mentioned earlier, despite large hydrological variability, with greater dispersion and lower average water values during drought. As emphasized above, this dispersion in annual water values cannot be interpreted directly as evidence of misallocation due to unobserved differences in planting costs.

Second, combining the water productivity estimates with data on the hydrological flow network, I document a clear gradient in estimated water values above and below the most critical chokepoint in the network, the Sacramento-San Joaquin Delta, with persistently higher values below the constraint, where trade constraints are likely to bind. I show that this gradient mirrors gaps in willingness-to-pay to extract groundwater obtained from well depths, as well as

the typical price gradient among the few annual water allocation trades from 1987–2024, which further support the hypothesis that Delta flow constraints lead to water misallocation. This finding corroborates some of the longstanding policy concerns that have led to extensive debate in California over the construction of new conveyance infrastructure in the Delta. Those conversations typically focus on the opportunity to reallocate water from agricultural users in the Sacramento Valley to large urban suppliers in southern California; these results indicate possible gains from trade even within agricultural uses, which is important given that most of California’s water is used for agriculture and many marginal values of urban water use are not clearly greater than perennial irrigation.⁴

Third, combining the water productivity estimates with administrative data on all water rights on paper, I find that water values correlate with the seniority of water rights; watersheds endowed with more senior water rights exhibit less dispersion in marginal values and higher values overall. These findings provide support for theories of complementarity in capital investments and more reliable senior water rights (Burness and Quirk, 1979) as well as more general models of directed technical change (Acemoglu, 2002) where more abundant factor supply leads to innovation that raises that factor’s productivity.⁵ However, the finding that more senior water owners are more productive on average contravenes the common perception (and prediction of some models of moral hazard) of more wasteful water use by senior water owners (“use-it-or-lose-it”).

Fourth, combining water productivity estimates with the full age distribution of tree varieties, I find an endogenous partition of water rights, where most (>90%) of the water used by orchards is “locked in” in a given year, in the sense that it would typically not make economic sense to reallocate to other uses. This follows directly from the observed orchard demographics and an assumption—strongly supported by the data and the implied economic costs of the counterfactual—that farmers almost never cut down orchards in the first decade of their productive life. The resulting endogenous distribution of water-

⁴About 80% of California’s water is used for agriculture. While this paper focuses on irrigation, similar arguments here apply to water that sustains human populations, which also entails large fixed investments in housing stock that require reliable water supply to operate. A major open empirical question is the extent to which water scarcity, versus other zoning and land use restrictions that create barriers to new development, inhibits urban development and expansion (see, e.g., Edelstein, 2025 for progress in this direction).

⁵Two countervailing forces are at work: abundance raises the value of such investments, but also lowers equilibrium factor prices (Acemoglu, 2002). If water is not priced at its opportunity cost, however, there will be no countervailing equilibrium price channel.

augmenting capital varieties—reminiscent of the Atkeson and Kehoe (1999) microfoundation for low short-run and large long-run energy price elasticities in macroeconomics—provides an explanation for why aggregate water demand appears conspicuously inelastic to fluctuations in aggregate water surface supplies by factors of two or more. It also has implications for long-run water contracts in the design of water markets—because very little of California’s water can be productively reallocated within a year, partial reforms to liberalize annual trade without also allowing longer-run trades seem unlikely to deliver substantial gains.⁶

Beyond the work cited above, this paper makes two primary contributions to the existing literature. One, the paper’s general empirical model enables the valuation of water rights by combining the rich detail of agricultural production models with some of the information implied from revealed preference about crop choices and replanting decisions within fields over time. Modern agricultural production models can make use of rich data and do not rely on revealed preference, which allow them to deliver large differences across water values (e.g., D’Odorico *et al.*, 2020; Medellin-Azuara *et al.*, 2022; Boser *et al.*, 2022). However, these models typically rule out unobservable heterogeneity that correlates with water use, rather than using observed decisions to learn about unobservable heterogeneity. In contrast, modern equilibrium models of land use (Scott, 2013; Burlig *et al.*, 2024), when identified with valid instruments, can recover certain forms of unobservable heterogeneity from aggregate data in a consistent way, for the same reason as in other product markets (Berry *et al.*, 1995). However, the source of that advantage—to rely on revealed preference and quasi-experimental price shifters to recover valuations—can limit the ability of these models to test for arbitrary misallocation. The key idea here to combine the two approaches without foreclosing the analysis of misallocation is that, holding fixed a given water right of a certain type in a given location, the crop choices made for that water right can reveal information about the relative planting costs across different potential crops, as well as the costs of renewal and replanting for perennial crops like almonds, pistachios, or wine grapes—even when that water could have a much greater value elsewhere.

Two, the paper’s substantive empirical work contributes the new findings on the value of water in California discussed above: how value varies with flow con-

⁶Under California law, trading water rights for more than one year or water rights with pre-1914 seniority status involve several additional legal frictions than annual trades.

straints in the surface water network, characteristics of historical water property rights, and over time through equilibrium investment decisions. Each of these aspects of water interact to determine the value of water reallocation,⁷ a recurring question in this literature—whether places that so rarely trade water on markets, like California (Hanak *et al.*, 2021; Hagerty, 2023), stand to gain by introducing more advanced water markets. Valuing water rights across competing uses is central to answering this question, but has been inhibited by data issues with extant water contracts and prices,⁸ except in rare cases like Australia (Rafey, 2023). This paper confronts the challenge of valuing water without trades, by far the most common situation worldwide. It follows in the line of supply-side industrial organization papers like Borenstein *et al.* (2002), Syverson (2004), and Asker *et al.* (2019) to start with details of the production process—for example, the physics of electricity, oil, and water flow at the point of generation, extraction, or diversion—and try to learn from the resulting distribution of economic activity, tailoring the analysis as close as possible to the observed set of producers. This allows the study of rich distributions of marginal water productivities over time and across the hydrological network and varied hydrological conditions. Working with water use data also lets me account for conjunctive use of surface and groundwater, an important limitation to previous research in California.

In addition to the primary contributions, this builds on a large body of work using extensive satellite imagery to characterize economic activity on the surface of the Earth and address measurement issues with conventional economic data.⁹ Much of this study’s empirical progress is enabled by observables beyond land cover; here, differentiated capital varieties implied by tree planting dates and local water requirements implied by daily evapotranspiration. Finally, a broader

⁷Some recent work on California water includes Gartrell *et al.* (2017), Hanak *et al.* (2021), Ayres *et al.* (2021), Zeff *et al.* (2021), Medellin-Azuara *et al.* (2022), Hagerty (2023, 2022), Burlig *et al.* (2024), and Leonard *et al.* (2025), as well as recent PhD dissertations analyzing groundwater regulations and labor on small farms (Sum, 2024), urban-agricultural water price gaps (Ferguson, 2024), and water impact fees and housing development (Edelstein, 2025).

⁸See Rafey (2023, pp. 433–434) for a discussion. Water is rarely traded, the terms of these contracts are often unobserved, and water rates are typically not market prices, but rather fees set formulaically by utilities or irrigation districts to recover fixed costs.

⁹On evapotranspiration, see D’Odorico *et al.* (2020), Wong *et al.* (2021), Boser *et al.* (2022), Wong *et al.* (2025), and Leonard *et al.* (2025). On groundwater subsidence, see Carleton *et al.* (2025). Beyond water, the satellite data renaissance has enabled progress on several other problems in the economics of land use and the environment, such as deforestation (Burgess *et al.*, 2012; Souza-Rodrigues, 2019; Hsiao, 2026; Balboni *et al.*, forthcoming), agriculture (Scott, 2013; Costinot *et al.*, 2016), and wetland conservation (Aronoff and Rafey, 2023).

literature documents significant dispersion in marginal products across various factors, like capital, labor, land, and oil; how we interpret that dispersion is obviously a challenge and this is one paper in that line of thinking.¹⁰

The rest of the paper starts with some background on California water institutions from an economic perspective in Section 2. I then introduce the model in Section 3, describe what I do with data in Section 4, and finally provide some results and some discussion of the limitations of interpreting the estimates as allocative inefficiency in Section 5.

2 Institutional details

This section discusses where California’s water comes from, where we move it and how, and who uses it. This context will then allow us to think about overlapping water rights and some of the constraints on water reallocation.

2.1 Water sources and supply network

The water in California comes from two places. About 200 million acre-feet in an average year falls from the sky, or about four to five times as much water as the entire state uses annually. Most of this water is not directly usable, but some stays in the mountains, turns into snow, and then slowly melts down into valleys through massive rivers. The triangles in Figure A1A depict the dams built in California to capture this surface water, modulate its variability, and deliver it to where it is most economically useful. This surface water then travels throughout California, flowing through the rivers plotted in Figure A1C, and the network of canals, primarily operated by the State Water and the Central Valley Projects, depicted in Figure A1D, which move water from Northern California and, critically, as shown in detail in the inset with more precision, through the Sacramento–San Joaquin Delta.

Surface water exhibits large cyclical and stochastic components, as illustrated by the useable river inflows in Figure 1. Annual volumes from 1980–2021 for the Sacramento and San Joaquin River Basins, given existing conveyance infrastructure, range from less than 10 million acre-feet to more than 30 million acre-feet across years.

¹⁰On misallocation and its discontents, see, e.g., Hsieh and Klenow (2009), Midrigan and Xu (2014), Adamopoulos and Restuccia (2014), Asker *et al.* (2014, 2019), Baqaee and Farhi (2020).

The natural surface water flow network also serves as the origin of California’s groundwater. Much of the water that lands on California seeps into the earth over many thousands of years, forming the gigantic aquifers depicted in Figure A1B. The dark blue aquifer in the Central Valley is the state’s largest source of groundwater, though aquifers also serve as important water sources in the Central Coast, where there are a lot of vineyards, and in southern California, especially for municipal users in desert areas.

2.2 Water users

The water that travels through the hydrological network has three main competing users in California: irrigators who use water to grow crops; people who drink water to not die; and ecosystems, which rely on natural flows to survive and thrive.¹¹ The water rights defined for these different uses are primarily owned by irrigators, and in particular, as I show below, many of the irrigation water rights were established prior to the complete settlement of California. Figure A2 reports a map of agricultural production in California in 2020. It shows just two colors; purple corresponds to perennial crops such as almonds, pistachios, oranges, and grapes, and orange corresponds to annual crops like wheat, rice, and pasture.

Agricultural users comprise about 70–80% of California water use, or 30–35 million acre-feet of water, in a given year. Most of the agriculture occurs in the Central Valley.¹² In contrast, Figure A2, Panel B, shows where the people are. Although California is the largest state by population in the US, it is sparsely populated outside of Los Angeles and the Bay Area; in the aggregate, municipal uses account for only about 15% of water in California. Finally, water flows serve critical ecological functions. Water throughout the network delivers value to the species that live in the rivers and water bodies; I discuss some of these values in Section 2.4 below.

¹¹California also generates a lot of hydroelectricity, but hydroelectricity is typically a “non-consumptive” use that does not compete with agricultural or municipal uses.

¹²While the main focus of this study is the Central Valley, agriculture also occurs in the Central Coast as well as the Imperial Irrigation District in the southeast corner of the state that uses senior Colorado River rights, as seen in Figure A2, Panel A.

2.3 Water rights, contracts, and rationing

The immense cyclical and stochastic components of water supply and flow create several difficulties with specifying well-defined water property rights. Most water rights in the western United States emerged based on historical claims by original water users. The fundamental problem is that these “appropriative” rights are typically defined at the place of use (“point of diversion”), and not the point of origin. This creates significant problems if (and only if) multiple competing uses conflict with one another, because they then have to be reconciled. In an ideal world, if we know exactly where all the water is coming from, for example, from a single dam, then we just define property rights over the water in the dam. One still needs to charge users along the conveyance network for transmission costs and solve the optimal routing problem, but there’s no incompleteness because we have fully specified the rights over the water.

But historically, surface water rights are often defined in ways that depart from this ideal case, because water comes from so many diverse places and water abundance varies across so many states of the world. Water property rights evolved in California, and the western United States more broadly—and, similarly in places like Australia—by granting property rights to initial claimants. Many claims in California originated when people started moving to the state in the Gold Rush of 1848–1855. Figure 2, which depicts cumulative, reported face values of current water rights for irrigation in the Central Valley by decade claimed from 1800–2020, shows nearly 20 million acre-feet of water claims prior to the 1914 Water Commission Act. These records contain measurement issues, but show evidence of large legacy rights that date prior to California becoming part of the US,¹³ as well as water claims that grow as people start moving to California to mine gold with water-reliant methods like sluice boxes and hydraulic mining. Figure 2 also shows large jumps in cumulative rights in 1940 and 1960, the completion decades of the state’s two largest infrastructure conveyance projects, which significantly expanded available surface water supply. By the 1980s, all sources of surface water for the Central Valley have been fully appropriated.

These appropriative surface water rights are defined by two attributes. First, a “face value” amount, denominated in acre-feet per year, which equals the maxi-

¹³One can see similar patterns in land parcel data. There are still ranches exactly the size of Spanish land grants, because many of these land rights were given as concessions to landowners when the United States annexed California from Mexico in 1848 (Gates, 1991, pp. 3–13).

mum amount of water that the user can take in a given year. Second, a “priority” year or tranche, which determines how water rationing occurs in states of the world in which the water right conflicts with other rights. Rationing can arise due to (i) shocks that lead to lower river inflows and aggregate water availability and (ii) excess demand by other water rightsholders with greater or equal seniority. In other words, an owner of an appropriative water right cannot take water if there is no divertable water in the river; even when there is water in the river, they may not take water when other, overlapping water rights have legal priority. All of this implies that rationing can occur in some, but not all, states of the world, as reported in Table 1. This fact is also partly why climate change raises new risks for water rights, because only in states of the world where supply constraints bind is it important to sort out these competing claims.¹⁴

In contrast to appropriative surface rights, property rights to groundwater are tied to the land. While there have been some regulations that have started to change the way people drill for groundwater in California, broadly, the landowner’s right in California extends to subsurface drilling rights.¹⁵ Figure A4, Panel A, shows meaningful groundwater extraction capacity relative to appropriative rights, with about 20–25 million acre-feet of capacity of groundwater wells by 2022.¹⁶ Figure A4, Panel B, which stacks cumulative well capacity together with appropriative water rights, shows that farmers start to drill for groundwater as surface water becomes fully appropriated.

An important aspect of groundwater, unlike most surface water, is its non-trivial marginal extraction cost. The cost of extracting water from a well can be calculated in proportion to drilling depth (“lift-height”) (Timmins, 2002). Table A9 calculates these costs using data on all the wells that have been drilled, showing meaningful differences across watersheds within the Central Valley, as well as differences of about a factor of two in average lift heights and implied pumping

¹⁴A separate but related concern with appropriative rights defined by historical use, distinct from inefficient rationing across users, is that these rights can create perverse “use it or lose it” incentives. In practice, the threshold to show beneficial use is very low. And, in part to encourage people to transfer water rights, and to avoid the kind of overuse incentives under a use-it-or-lose-it regime, since the 1970s, it’s been very clear in California law that, in most cases, the owner retains full ownership rights even if they do not use the water.

¹⁵This subsurface water is, essentially, open-access, limited by rising extraction costs as aquifer levels fall. The California Sustainable Groundwater Management Act, passed in 2014, aimed to limit groundwater extraction in critically overdrafted regions; how its regional objectives will be enforced in practice remains unclear more than a decade after the law’s passage.

¹⁶In California, drilling a well requires a permit (“well report”), so there are about a million of these well reports, about 100,000 of which are in the Central Valley.

costs above and below the Delta.

2.4 Trade constraints and frictions

Even a very distorted initial allocation should not give rise to allocative inefficiency if water endowments are freely and costlessly tradeable prior to use. Water, however, is rarely traded (Libecap, 2011). Constraints on water trade among owners of water rights in California—as in many other places—arise for several reasons. Some constraints reflect environmental priorities. Figure A2, Panel C illustrates some of the environmental assets in California. The purple polygon outlines the legal boundary of the Sacramento–San Joaquin Delta; many endangered species that inhabit the Delta are federally regulated under the Endangered Species Act of 1973. The Delta is also ecologically critical to prevent saltwater intrusion because it drains into the Pacific Ocean through the San Francisco Bay. California hydrology has been so altered from its natural state that the saline ocean water could overrun the Delta without minimum outflows from the Delta to protect California from turning into a wasteland (Tolkien, 1965, p. 287). These environmental priorities restrict how water can be moved, stipulating certain seasonal outflow requirements to maintain species and deliver environmental value in the Delta. Figure A5 plots variation in the water used for environmental requirements over 1980–2022.

Infrastructural constraints on water transport also limit water transfers. The Central Valley, where most of California’s agriculture occurs, is largely divided between the Sacramento Basin and the San Joaquin Basin. Water flow through the Delta to southern California must be lifted at the bottom of the Delta before it can flow into the San Joaquin River Valley.¹⁷ This is the primary flow constraint that precludes greater reallocation of water across these two basins. While significant volumes of water flow through the Delta each year (about 4–8 million acre-feet)—the source of much of the water used in the southern component of the network—there are still considerable water volumes consumed in the North that may not be able to cross the Delta.¹⁸

¹⁷Only two pumping plants transport water out of the Delta, the (federal) Bill Jones and (state) Harvey Banks Pumping Plants. The proposed Bethany Reservoir Pumping Plant, which would convey about 400,000 acre-feet more water to flow through the Delta to southern users, is planned for completion by 2040 at an estimated cost of about \$20 billion.

¹⁸Another, more general, friction to trade is the prospect of legal challenges to water owners who propose to transfer some or all of their water property rights to another user. For example,

In part due to these constraints, annual trade in water allocation rights in California—as is well-known in this literature (Hanak *et al.*, 2021)—appears to usually be less than one percent of water use and never much more, even during severe drought.¹⁹ Despite this sparsity, there are a few transactions. I did my best to pull as much water trading data as I could for these two regions to construct price distributions for each region. The comparisons corroborate estimates of Regnacq *et al.* (2016) and Hagerty (2023) of large transaction costs across the Delta. The ratio of regional prices, reported in the rightmost column in Table 2, ranges between 1.2 and 2.3 from the third to seventh deciles. Corroborating these differences, Figure A6 shows, in 2023, two lines. The blue line shows prices south of the Delta (SOD), water that travels through the Delta to the San Joaquin River Valley. The red line shows prices of water transacted above the flow constraint, north of the Delta. One can see a large wedge, a factor of three or four, even out to August, when people start wanting water to plant. Prices are very flat along the red line, north of the Delta; the SOD price exhibits much larger fluctuations. The lower panel of Figure A6 shows a similar pattern in 2024. Taken together, these prices would suggest a wide Delta price gradient if one takes the data seriously—approximately a factor of two, even towards the end of the growing season. This further motivates my analysis below of how water productivity and water use differ across these two regions.

3 Model

This section introduces an empirical model of water rights and irrigated agricultural production to clarify some of the economic mechanisms at work in the

imagine that someone who takes water from one location on a network instead declines to take that water so that someone else in the network can take more. In principle, that reallocation would affect other flows throughout the network. So, anyone between the two counterparties can litigate. Under the California Environmental Quality Act of 1970, there also needs to be environmental review to make sure that all the species whose existences and locations and habitats have been predicated on the existing distribution of water are not adversely affected. This has, in practice, turned out to be an easy way to obtain standing to file lawsuits and obtain injunctions against proposed water transfers.

¹⁹Figure A8 reports California water trades from 1985 to 2022. The units are in acre-feet, and range in the top panel from zero to one-and-a-half million acre-feet. The numbers on the top of the bars are the number of trades. For example, there were 17 trades in 1993 and 16 trades in 1994. This is a very sparse market. The lower panel of Figure A8 reports the same trade volumes alongside the reported water volumes used (after mid-2010s legislation introduced mandatory reporting requirements), to emphasize that trade volumes are less than one percent of water.

institutions described above, and to motivate the use of data to explore the possibilities of water misallocation in the subsequent section.

I specify the model in two ways. First, I outline a general model of irrigation and water rights for which I can estimate annual marginal products of water under a relatively small set of assumptions that leave water rights, the water market, and planting costs unspecified. Water misallocation is not identified from the distribution of water productivities without further structure, such as assumptions that unobserved planting costs are the same across certain places in the network. Second, I specialize the general model to a fully-specified model of equilibrium water use, which provides a microfoundation for heterogeneous planting costs under the assumptions that permanent and transient water rights can be distinguished and that a water right owner allocates water efficiently within each reliability tier of their water right. The planting costs and implied values of water rights can then be identified in the fully-specified model with micro panel data on fields within a water right and instruments for crop prices.

3.1 Irrigation technology

Irrigators, the main users of water, indexed by i , produce various crops indexed by c with fields (units of water) indexed by $\iota \in i$. Agricultural calendars operate on an annual basis, with annual production in each year t .²⁰ Total annual crop production, Q_{ict} , is a function of the land that the farmer uses and the irrigation that they apply to the field. On field ι in year t , each crop c has a potential yield per acre, $A_{\iota ct}$. Annual production for irrigator i , crop c , in year t , is then given by

$$Q_{ict} = \sum_{\iota \in i} A_{\iota ct} \min \{K_{\iota ct}, W_{\iota ct}\} \quad (1)$$

where output Q_{ict} is measured in tons, yield $A_{\iota ct}$ in tons per acre, land $K_{\iota ct}$ in acres, and effective annual field-level irrigation in acre-feet per year is calculated as

$$W_{\iota ct} = \min_{\tau} \left\{ \frac{w_{\iota ct}(\tau)}{\zeta_{\iota ct}(\tau)} \right\}_{\tau=0}^{365}, \quad (2)$$

using the field-level irrigation application rate $w_{\iota ct}(\tau)$ in acre-feet/day. The key feature of (1) is to allow the value of daily irrigation $w_{\iota ct}(\tau)$ in production to

²⁰Some crops are grown more than once a year. While the empirical work below abstracts from multi-cropping, to incorporate it, just let c index each potential multi-crop combination.

depend on each field’s inherent and crop-specific water efficiency, $\zeta_{\iota ct}(\tau)$, which corresponds to how much water the crop needs in acre-feet per acre, in theory, on day τ of that year and field ι . A crop’s water demand varies over the growing season and with climate conditions. For example, $\zeta_{\iota ct}(\tau)$ will be zero in times when the crop needs no water, and $\zeta_{\iota ct}(\tau)$ will be larger as the crop becomes thirstier.

The irrigation scheduling algorithm allows water efficiency to differ (i) across crops at the same location ι and year t , (ii) across locations $\iota \in i$ for the same crop and year, and (iii) across years t within a location ι growing the same crop c . The reason $\zeta_{\iota ct} : [0, 365] \rightarrow \mathbb{R}_+$ will vary with ι , c , and t is that location influences the climate conditions that determine the plant’s water demand through local soil characteristics, daily weather conditions, and the relevant growing season. For example, a crop will typically need more water on hotter, windier, or sunnier days, and more water when fully grown than when still a seedling.

Equations (1)–(2) make no restrictions on how a farm chooses which crop to plant. To make progress, however, once a crop is planted on a given field, I will always restrict how farms irrigate that crop within the year. Specifically, the key assumption that I make in order to learn about the water used in California from the crop choices that I see is that irrigation scheduling for each crop on field ι over its growing season is optimal. These assumptions are stated on endogenous objects, but straightforward to state in terms of primitives:

A1. No overwatering over the growing season (nonnegative marginal irrigation costs or declining yield), i.e., $w_{\iota ct}(\tau)/K_{\iota ct} \leq \zeta_{\iota ct}(\tau)$ for all τ .

A2. No deficit irrigation over the growing season, i.e., $K_{\iota ct} \leq \min_{\tau} \frac{w_{\iota ct}(\tau)}{\zeta_{\iota ct}(\tau)}$.

Assumption **A1** rules out farms that irrigate more than the crop needs, which can be microfounded with any nonzero marginal cost of irrigation, or a yield function where overwatering lowers yield. Similarly, **A2** rules out under-irrigation. While there is immense interest in deficit irrigation—conceptually, it would be great to get more for less—agronomy experiments have largely found that deficit irrigation is not a viable strategy in most settings because under-watering severely compromises crop yield.²¹ No-deficit irrigation will be violated by pasture planting with multiple cuts in a single year (e.g., alfalfa), where deficit irrigation

²¹One can think about the global production function in water as quadratic, a parabola where yield declines very fast on either side of the optimal amount of water that the plant needs.

corresponds to fewer plantings within the year.²² Regardless of their microfoundations, A1 and A2 imply that, at the field level, water demand coincides with optimal irrigation requirements. This lets me use (1)–(2) to take field-level crop choices K_{ict} and crop-location-day irrigation efficiencies to characterize optimal irrigation on day τ as $w_{ict}^*(\tau) = K_{ict}\zeta_{ict}(\tau)$. I then identify total water use—the annual volume of the water right i used in equilibrium—from data on land use and crop evapotranspiration each day of year t , via

$$W_{it} = \sum_{\iota,c} K_{\iota c} \cdot \int_0^{365} \zeta_{\iota c}(\tau) d\tau. \quad (3)$$

Note that while A1 and A2 clearly restrict input demand within a year on a given field, they impose no restrictions on crop choice—which occurs prior to within-year irrigation—or the assignment of water across water rights.

3.2 Sketch of the farm’s problem

Given the production technology in (1), the farm will solve a problem that I first specify at a general level that I do not directly take to the data. Suppose that an irrigator has some amount of land, $\bar{K}_i$, and water rights, $\widetilde{W}_{it}$, which they can assign to different crops c at t , given some general planting cost function $\Gamma_{it}(K)$ that depends on the full vector of crop choices. The land allocation solves

$$\max_{\{K_{\iota c}\}} \sum_c P_{ict} Q_{ict} - \Gamma_{it}(K) \quad (4)$$

such that $\sum_{\iota,c} K_{\iota c} \leq \bar{K}_i$ and $W_{it} \leq \widetilde{W}_{it}$. That is, the farm’s crop choices maximize profits given their productivities embedded in Q_{ict} , the prices that they obtain for each crop, the planting and other costs Γ_{it} , and the land and water that they own. If, for example, the cost function Γ_{it} is linear, the two constraints in (4) form a convex polytope that one explores to find the optimal crop choice.

Equation (4) is intended to emphasize how the value of water rights will, in general, depend on two critical economic forces in addition to the marginal product of water analyzed in Section 3.1. First, the vicissitudes of water scarcity or abundance, given that a water right $\widetilde{W}_{it}$ is a random variable. Second, what I refer to as “planting costs,” written in (4) as some arbitrary function Γ_{it} of

²²However, one can proceed (as I do) with data or a model of the number of alfalfa cuts.

the full vector of crop choices. In the fully-specified dynamic setting of Section 3.3, this Γ_{it} reflects underlying technological primitives, the farm's (endogenous) past decisions, and the implications of current decisions for (endogenous) future decisions and payoffs. Here, all of these details are represented implicitly by an unknown function that varies with i and t .

The value function $V_{it}(\bar{K}_i, \bar{W})$, defined as the maximand of (4) at $\bar{W}_{it} = \bar{W}$, then reveals the value of water rights through farm i 's water constraint at t . Various comparative statics with respect to attributes of that water right can then be used to trace the value of the water right through the changes in irrigator i 's value at t . In doing so, we obtain the marginal (shadow) value of changing the water right,

$$\frac{\partial V_{it}}{\partial \bar{W}} = \underbrace{\sum_{\iota, c} \frac{A_{\iota ct} P_{\iota ct}}{\int \zeta_{\iota ct}(\tau) d\tau} \frac{\partial \int w_{\iota ct}^*(\tau) d\tau}{\partial \bar{W}}}_{\Delta \text{ marginal products}} - \underbrace{\sum_c \frac{\partial \Gamma_{it}}{\partial K_c} \frac{\partial K_{\iota ct}}{\partial \bar{W}}}_{\Delta \text{ planting costs}}, \quad (5)$$

the sum of two components; first, the change in the annual marginal product of water, which depends on how water is reallocated across crops and each crop's value per unit water; second, the change in marginal planting costs, including dynamic considerations embedded in Γ_{it} . A farm might, for example, reallocate land to crops with much higher annual marginal products under a different $\bar{W}$, but that reallocation could involve new planting costs such that the total value of (5) is not large. Alternatively, switching costs may be so great that the change in water rights does not change the crops grown, so that (5) vanishes.

The reason I write down the derivative in (5) is to emphasize when we can, and when we cannot, learn about the value of water rights from data on the joint distribution of the marginal productivity of a resource and its allocation. To the extent that planting costs are similar across different irrigators, the difference in annual marginal products of water is a valid way to compare the true value of the water right (embodied in the crops that are being grown) in some place i with the value of the water right (embodied in crops) in some other place i' . When planting costs are identical, the joint distribution of the annual marginal products of water and the allocation of land, $\{K_{\iota ct}, \frac{A_{\iota ct} P_{\iota ct}}{\int \zeta_{\iota ct}(\tau) d\tau}\}_{\iota, c, i, t}$, is sufficient to characterize misallocation among hydrologically connected, freely transferable water rights. In contrast, where planting costs differ meaningfully across i and i' , differences in marginal products across i and i' need not be evidence of inefficiency.

For example, where irrigators embody water in different kinds of trees, planting costs depend on the history of a farm’s crop choices. In this case, systematic differences in marginal products of water could be entirely efficient; reflecting differences in optimal investment, not misallocation.²³ The richer model below is designed to account for both possibilities.

3.3 Farm’s problem, fully specified

I now specialize and extend the above model to specify planting costs across fields that reflect the evolving opportunity cost of replanting long-lived orchards of different ages and varied types. This model allows me to (a) more directly capture the varied value of water rights by decomposing planting costs across reliability tiers and sunk investments and (b) obtain some necessary ingredients to study alternative water allocation mechanisms. While this requires some additional structure on the cost function Γ_{it} and equilibrium behavior that determines crop choices, the exercise will remain disciplined by the additional panel micro-data that I have on the age of trees and within-orchard planting and replanting decisions across many fields within the same water right.

Water rights. Water rights are random variables that are realized at the start of each year, $\widetilde{W}_{it} \sim G_i$. For tractability, I distinguish between the permanent (“reliable”) and transient (“unreliable”) component of each water right,

$$\widetilde{W}_{it} = W_i^p + W_{it}^a, \quad (6)$$

with the reliable component of the water right defined as delivering water almost surely, i.e., $G_i(w) = 0$ for all $w < W_i^p$, so that $\mathbb{P}(\widetilde{W}_{it} \geq W_i^p) = 1$. The decomposition in (6) allows me to define values for different units of water in the water right; without loss of generality, I index the units of a water right $\iota \in [0, \widetilde{W}_{it}]$ such that the first $\iota \in [0, W_i^p]$ units are reliable.

In practice, reliable rights can be a senior surface water right that always delivers water, or a surface water right that delivers in fewer than all years combined with a groundwater well that can make up the difference in low surface water years. While $\widetilde{W}_{it}$ captures all available water—both surface and groundwater—

²³This is reminiscent of Asker *et al.* (2014), who show how substantial cross-sectional dispersion in marginal products of capital can arise in models of investment by firms with different productivity shocks over time, even when very little of the dispersion is welfare-relevant in the sense that there exists a feasible reallocation to improve allocative efficiency.

the costs of using that water may differ depending on the source. Variable irrigation costs are accounted for with idiosyncratic “planting costs” defined below that differ by water right $\times$ reliability tier $\times$ year. The source of water in a given year—surface, ground, or a combination—affects idiosyncratic irrigation costs because (i) diverting surface water and pumping groundwater entail quite different costs, (ii) the share of pumping can vary over time and across space with the surface water right’s seniority and abundance of surface water (determining the share of groundwater use), and (iii) the cost of pumping varies with electricity prices, groundwater regulation, and the local well depth (determining the marginal cost of a unit of groundwater).^{24,25}

Flow payoffs. I assume that water, not land, is the scarce factor of production. Each year, a unit ι (acre-foot) of water allocated by farm i to a crop c of age $x \in \{0, 1, 2, \dots\}$ delivers a payoff, denominated in \$/af, of

$$\pi_{it}(c, x) = \xi_{ict} + \sigma_{\epsilon^k}^i \varepsilon_{ct}(\iota), \quad (7)$$

where $\pi_{it}(c, x) = \alpha_{ict}(x)P_{ict} = \frac{A_{ict}(x)}{\int \zeta_{ict}(\tau) d\tau} P_{ict}$ is the revenue (\$/af) from the crop derived from equation (1) and assumptions A1 and A2, given the (random) field-level ι yield-per-unit-water coefficient $\alpha_{ict}(x)$, observed crop prices P_{ict} , irrigator i ’s relative planting costs ξ_{ict} , and the field ι ’s idiosyncratic planting costs $\sigma_{\epsilon^k}^i \varepsilon_{ct}(\iota)$, independently and identically distributed Type 1 Extreme Value (T1EV) at the field-crop-level, scaled by $\sigma_{\epsilon^k}^i \in [0, \infty)$, which determines the extent of heterogeneity across fields ι within the water right i ; this scalar is indexed $k \in \{a, p, r\}$ to differ across the annual, perennial, and renewal choice sets defined below. In addition, I assume that irrigators take output prices, P_{ict} , as given. Aggregate yield shocks and acreage in California may affect these prices, but agricultural commodity markets are assumed to be perfectly competitive.

Interannual choice sets. In equation (7), payoffs on ι depend on the year

²⁴As the earlier discussion of priority and rationing indicated, this random variable could depend on all the equilibrium different decisions of other users in the network. But from the farm’s standpoint as a producer, they care only about G_i , the probability that they have a certain amount of water, $\bar{W}_{it}$.

²⁵A related but distinct issue are the implications of the origin of the water right for the counterfactuals discussed below. The source of the water used in a directed hydrological network affects the interpretation of reallocation; surface water often travels hundreds of miles, while groundwater remains local. While groundwater-to-groundwater reallocation cannot occur across aquifer boundaries, surface water reallocation can affect groundwater use through users of both surface and groundwater (conjunctive users).

t as well as the state of the water right, (c, x) . Dynamics arise because the state of the water unit, (c, x) , can determine the choice set of the farmer. For annual planting, the problem is truly static; age $x = 0$ and the choice set is $c \in \mathcal{C}_i^a$. For new perennial plantings, $c \in \mathcal{C}_i^p$ and the age in the next period becomes $x = 1$. For perennial plantings with age $x > 0$, the choice set becomes either to maintain the orchard, leading to $(c, x + 1)$, or to cut down the tree to return to $(0, 0)$. Consequently, when growing a given perennial crop, the farm's subproblem resembles a Rust (1987) optimal regenerative stopping problem. When selecting a new perennial planting, the farm's problem can be reduced to a static Berry *et al.* (1995) discrete choice over payoffs that correspond to the value functions derived below and random coefficients that correspond to the $\alpha_{ict}(x)$ that I take from field-level crop evapotranspiration data discussed below.

Aggregate uncertainty. An attractive aspect of the Rust (1987) approach to optimal renewal is that stochastic state transitions due to observable events are straightforward to incorporate into the model. Only data and computational limits restrict this. I focus on arguably the most important such random, yet persistent, state of the world for irrigation in California, an indicator for drought. It is plausible—and will turn out to be the case in the resulting estimates—that the replanting/renewal costs that create frictions to adjust reliable water rights committed to orchards differ under drought conditions.

Irrigators hold rational expectations about drought transitions: they observe the current drought state when they decide whether to renew their orchard and when they decide which new perennial crop to grow; they incorporate the stochastic and persistent features of the drought process when they calculate the expected value of these choices; and, consequently, the value of the water right also depends on the drought state. For concision, I suppress the dependence of the payoff and value functions $V_i(\cdot)$ and $\alpha_{i,ct}(\cdot)$ on the drought state s , as well as the conditional expectation for the next year's drought state given the current s implicit in $\bar{V}_i(\cdot)$.

Equilibrium values of water. Farms discount periods with a common factor $\beta < 1$. The value of a unit of water ι committed for x years to a perennial crop c can be constructed recursively via

$$V_i(\iota, c, x) = \max \left\{ \alpha_{ict}(x)P_{ict} - \xi_{icxt} + \sigma_{\epsilon^i}^i \epsilon_{ct}(\iota) + \beta \bar{V}_i(\iota, c, x + 1), \right. \\ \left. -\xi_{i0t} + \sigma_{\epsilon^i}^i \epsilon_{0t}(\iota) + \beta \bar{V}_i(\iota, 0, 0) \right\}, \quad (8)$$

for $x > 0$, with the overline $\bar{V}_i$ denoting the ex-ante value function prior to

drawing the ε shocks for the field ι that period, and

$$V_i(\iota, 0, 0) = \max_{c \in C_i^p} \{ \bar{V}_i(\iota, c, 1) - \xi_{ic1} + \sigma_{\varepsilon^p}^i \varepsilon_{ct}(\iota) \} \quad (9)$$

at the event of replanting.

Two special cases of (8)–(9) illustrate the generality of this empirical model as a way to value water rights. First, as $\sigma_{\varepsilon^k}^i = \{\sigma_{\varepsilon^k}^i\}_k \rightarrow 0$, within-water-right field-crop-level heterogeneity in planting costs vanish. If, in this case, crop prices and planting costs are also time-invariant, there will be a perennial crop c_i^* that always maximizes (9). Further, when $\frac{\partial}{\partial x} \alpha_{\iota ct}(x) < 0$ for all $x \geq a$ for some bounded $a > 0$, there exists a finite stopping time x_{ι}^* at which replacement becomes optimal, so that the value of the water right (9) admits the simple closed form:

$$V_i(\iota, 0, 0) = \frac{1}{1 - \beta^{x_{\iota}^*}} \sum_{s < x_{\iota}^*} \beta^s [\alpha_{\iota c^* t}(s)P - \xi_{ic^*}] - \frac{\beta^{x_{\iota}^*}}{1 - \beta^{x_{\iota}^*}} \xi_{i0}. \quad (10)$$

In this case, the water right's value equals the net-present-discounted profits from growing and harvesting the best crop for its optimal lifespan, then replanting the same crop for a cost ξ_{i0} , ad nauseam, ad infinitum. Second, as $\sigma_{\varepsilon^k}^i \rightarrow \infty$, any pattern of choices in a finite sample within the water right can be explained by the model. In this sense, the model nests the case in which microdata on water use and marginal products cannot reveal anything about misallocation.

Importantly, the additional heterogeneity in planting costs across crop types and the evolution in crop prices, switching costs, and yields do not qualitatively change the nature of the farm's problem from the simpler case of (10) above. The basic economic principle is always the same: a water right tied to a specific location and irrigator is used by that irrigator to grow the most valuable crop or sequence of crops. The additional heterogeneity simply enriches the decision space to better account for data where an otherwise identical (on observables) water right is observed to be used to grow a different sequence of crops over time. Misallocation, in this model, arises from the way in which nontradable water rights are tethered to specific locations and irrigators, not from mistakes made by irrigators in the use of the water they own.

3.4 Identification of the value of water rights

The previous section describes the choices for a single irrigator owning a given water right. Practically, we would like to use data from many irrigators, and within the same irrigator over time, to estimate the structural parameters. Here, I describe additional assumptions within irrigator choices over time that one can use to identify planting costs; these are related to standard assumptions used in the dynamic discrete choice literature (Rust, 1987; Berry, 1994; Gowrisankaran and Rysman, 2012; Berry and Haile, 2024), described here for the model of water rights where water, not land, is the scarce factor of production:

A3. Planting costs from revealed preference. For each water right i , water uses $\iota \in i$ have

- (i) Price instruments (Berry and Haile, 2024). Shifters Z_{ict} of P_{ict} such that $\mathbb{E}[Z_{ict}\xi_{ict}] = 0$ and $\mathbb{P}(\frac{\partial c_{it}(\iota)}{\partial Z_{ict}} \neq 0) > 0$.
- (ii) Common average unobserved characteristics (Berry *et al.*, 1995). Fields within the same water right, $\iota \in i$, have the same average planting costs ξ_{ict} , and the same curvature $\sigma_{\epsilon^k}^i$ across $c \in \mathcal{C}_i^k$.
- (iii) Common beliefs over future states (Rust, 1987) and Gowrisankaran and Rysman (2012) inclusive value sufficiency for the renewal action.

With Assumption A3, the fully-specified model is identified. Under A3(i)–A3(ii), choices within a water right across fields $\iota \in i$ can be used to recover the relative planting costs that rationalize planting decisions at the start of each year—which, by assumption, do not depend on future periods. For annual crops, the probability of water ι allocated to $c \in \mathcal{C}_i^a$ is

$$\mathbb{P}(c_{it}(\iota) = c) = \exp \left\{ \frac{1}{\sigma_{\epsilon^a}^i} (\alpha_{\iota ct}(0)P_{ict} - \xi_{ict}) \right\} / \sum_{c' \in \mathcal{C}_i^a} \exp \left\{ \frac{1}{\sigma_{\epsilon^a}^i} (\alpha_{\iota c't}(0)P_{ic't} - \xi_{ic't}) \right\} \quad (11)$$

for each annual crop $c \in \mathcal{C}_i^a$ in each year t and each water right i .

Under the common beliefs of A3(iii), choices for $\iota \in i$ reveal the planting costs and yield curvature that rationalize within-orchard decisions over time, because the only variation in choices over time arises through the structural parameters θ , ξ , and σ . Let $\kappa_{it}(\iota, c, x) = 1$ if i chooses to renew ι at (c, x) and normalize $\sigma_{\epsilon^r}^i \equiv 1$; then the odds ratio that characterizes the renewal probability for ι can

be obtained as

$$\ln \frac{1 - \mathbb{P}(\kappa_{it}(\iota, c, x) = 1)}{\mathbb{P}(\kappa_{it}(\iota, c, x) = 1)} = \alpha_{\iota ct}(x)P_{ict} - \xi_{ic_x t} + \beta \bar{V}_i(\iota, c, x+1) - [\beta \bar{V}_i(\iota, 0, 0) - \xi_{i0t}]. \quad (12)$$

For fields $\iota \in i$ that share both common planting costs and beliefs, i.e., A3(ii) and A3(iii), micro choices for each $\iota \in i$ can be used to recover the planting costs that rationalize new perennial planting decisions at the start of the year; that is, for newly planted perennials (ι such that $\kappa_{i,t-1}(\iota, \cdot, \cdot) = 1$) in that year:

$$\mathbb{P}(c_{it}(\iota) = c | \kappa_{i,t-1} = 1) = \exp \left\{ \frac{1}{\sigma_{\epsilon^p}^i} (\bar{V}_i(\iota, c, 1) - \xi_{ic_1 t}) \right\} / \sum_{c' \in \mathcal{C}_i^p} \exp \left\{ \frac{1}{\sigma_{\epsilon^p}^i} (\bar{V}_i(\iota, c', 1) - \xi_{ic'_1 t}) \right\} \quad (13)$$

which is analogous to the previous equation (11) except choice-specific payoffs depend on the value functions derived from the full lifecycle of the dynamic choice, as in Gowrisankaran and Rysman (2012). The inclusive-value-sufficiency component of A3(iii) makes the ex-ante or *de novo* value of the new planting choice a sufficient statistic for the water right owner's continuation value at renewal, so that the renewal problem does not track the full vector of possible future crop-specific payoffs. This simplifies computation below by making the Rust (1987) contraction for the renewal decision separable from the Berry (1994) contraction in the discrete choice over crop choices.

The necessity of—and intuition for—price variation in Assumption A3(i) is the same as in canonical discrete choice models with micro data (Berry and Haile, 2024). The concern is that equilibrium planting and replanting decisions, characterized by (11), (12), and (13), reveal the core identification concern in this model, as well as standard approaches to address the endogeneity. Equilibrium water use across water rights and over time reflects differences in water right reliability G_i , water productivity, $\alpha(x)$, the drought state s , and unobserved planting costs ξ , all of which could be correlated with crop prices P .

Here, price instruments allow moments of the distribution of planting cost heterogeneity to be identified via the responsiveness of crop choices—if large (quasi-)random shifts in a given crop's price at the time of planting induce limited changes in water allocated to that crop by the owner of the water right, then relative planting costs must be quite heterogeneous to explain the small measure of marginal fields; in contrast, larger changes in planting decisions imply less heterogeneous planting costs across fields. Section 4.8 describes the estimation

algorithm in detail after describing the exact data in my empirical work.

3.5 Measuring misallocation

As emphasized, the model delivers a microfoundation to value water rights. To value water right i , which recall is distributed via G_i , let $(c_i(\ell), x_i(\ell))_{\ell \in i}$ denote the state of i 's planting decisions. Then the value of the water right G_i is

$$\mathcal{V}_i(G_i) = \int_0^\infty \bar{V}_i(\ell, c_i(\ell), x_i(\ell)) dG_i(\ell),$$

which can be decomposed into the reliable and unreliable components as

$$\mathcal{V}_i(G_i) = \int_0^{W_i^p} \bar{V}_i(\ell, c_i(\ell), x_i(\ell)) dG_i(\ell) + \int_{W_i^p}^\infty \bar{V}_i(\ell, c_i(\ell), x_i(\ell)) dG_i(\ell), \quad (14)$$

where the first integral integrates only over the distribution of planting decisions, since $1 - G_i(w) = 1$ for all $w \in [0, W_i^p)$, while the second integral integrates over the probability distribution of the remaining (unreliable) component.

Asymmetric incentives for trade. Interestingly, equation (14) shows how the presence of reliable water rights leads short-run incentives for water trade to generally differ from long-run incentives, due to the rigidities associated with water rights invested in perennial crops with nonzero ages. In particular, i 's reservation value, or minimum price at which they may sell a marginal unit of reliable water, equals the lowest opportunity cost of keeping the water within their existing orchard fleet, $\{c_i(\ell), x_i(\ell)\}_\ell$,

$$\partial_- \mathcal{V}_i = \min_\ell \bar{V}_i(c_i(\ell), x_i(\ell), s)$$

whereas their marginal value for a unit of new reliable water, if their land constraint does not bind, equals its value in new production, $\partial_+ \mathcal{V}_i = \bar{V}_i(0, 0, s)$. Given that $\bar{V}_i(c_i(\ell), x_i(\ell), s)$ includes the option to cut down the orchard and renew to obtain $\bar{V}_i(0, 0, s)$, generally this implies that reliable water rights will be less valuable to transfer until the tree capital in which they are embodied is sufficiently depreciated.

Aggregate effects of reallocation. Let the vector $\mathbf{V} = (\mathcal{V}_i)_i$ collect the values of existing water rights across owners i , so that total surplus equals $\sum_i \mathcal{V}_i$. A natural question—and a first step towards exploring the space of counterfac-

tual water allocations—is to ask how a vector of water right changes, $\Delta \ln \widetilde{W} = (\Delta \ln \widetilde{W}_i)_i$, will affect total surplus in the economy. In partial equilibrium, water reallocation always creates social value when water moves from lower to higher marginal value uses. More generally, this need not hold—if water used by some irrigators also creates or diminishes value for others (spillovers), then moving water from a water owner with a lower value to a water owner with a higher value can lower aggregate surplus. We can make progress with a matrix that summarizes complementarities in water use across firms, $\left[\frac{\partial \mathbf{V}}{\partial \ln \widetilde{W}} \right] \equiv \left(\frac{\partial}{\partial \ln \widetilde{W}_j} \mathbf{V}_i \right)_{i,j}$, by collecting the semi-elasticities of the value functions with respect to one’s own and others’ water rights.

One can then fully characterize to second-order the value of a given reallocation. As Baqaee and Sangani (2025, Proposition 5) show, the aggregate value of this water reallocation can be approximated via the quadratic form

$$\Delta \ln \sum_i \mathcal{V}_i = \mathbf{V}'(\Delta \ln \widetilde{W}) + \frac{1}{2}(\Delta \ln \widetilde{W})' \left[\frac{\partial \mathbf{V}}{\partial \ln \widetilde{W}} \right] (\Delta \ln \widetilde{W}), \quad (15)$$

which reflects the correlation of the reallocation vector with the marginal productivity of current water rights, adjusted for off-diagonal spillovers via $\left[\frac{\partial}{\partial \ln \widetilde{W}} \mathbf{V} \right]$.

Summary of identifying assumptions. The two most important unobserved sources of planting cost heterogeneity that I do not want to interpret as inefficiency—and help to rationalize dispersion in marginal products—can be seen in the data I introduce below on (i) the dispersion in marginal products of water across units of water within a given right of a given reliability class and (ii) the dispersion in water productivity over time, through the lifecycle of an orchard, for the reliable component of a given water right.

The model assumes that the water within a given water right (of a given reliability) is allocated efficiently at the time of planting (discrete crop choice), and that the reliable water is committed and re-committed to perennial crops efficiently (optimal regenerative stopping). This implies that every unit of water within a given water right has the same average value within a year, and that units within a given reliable water right have the same counterfactual values of short- and long-run reallocation (that differ based on the age structure, through the optimal regenerative stopping).

In contrast, the model does not require that misallocation across water rights take any specific form. The reason is that, with microdata within water rights,

I do not have to use variation in crop choices across water rights to identify any structural parameters of the model. This flexibility is intended to allow for the inherent incompleteness of water markets. However, to identify counterfactual misallocation, I do (necessarily) take a stance on which water rights can, in some idealized counterfactual, be exchanged with one another. To do this, as I discuss in Section 5.4, I focus on four natural frictions or constraints on trade, (i) distortions within connected components of the network such as irrigation districts, rivers, and canals, due to the difficulty of transferring water rights; (ii) distortions across regions reflecting infrastructural or conveyance capacity constraints on trade; (iii) distortions over time reflecting infrastructural or water storage constraints on trade across years; and (iv) distortions in property rights to trade across reliability classes.

4 Data and empirical approach

I now describe the empirical analysis. A centerpiece of this project is very high-resolution satellite data that tracks irrigated crop choices at the field level in California for nearly forty different crops.²⁶ Yield and price data come from the California County Agricultural Commissioners, 1980–2022, collated by the USDA Census. I take evapotranspiration data from California’s Irrigation Management Information System, and crop coefficients from the DWR, to measure water efficiencies at the crop level at different locations, for observed as well as counterfactual crop choices. I then compare different places in the hydrological network, which allows me to partition California into various components of the river network that are relevant for the analysis.

4.1 Fields

California farms grow a wide range of crops under the current system of water property rights. I observe about forty distinct crops, $\{K_{\iota t}\}$, by field-year (ι, t) , in annual data from 2014–2023. Fields typically cover 10–20 acres; the data, summarized in Table A1, contains about 400,000 fields in each year. To

²⁶This data is costly to produce: one needs to train the model to recognize all of the land-use patterns, with extensive ground-truthing to verify the model’s classification. This data was funded through the Department of Water Resources and has been improved, corroborated, and verified, and is used to track water use for policymaking in California.

illustrate the granularity of this data, Figure A9 maps a subset of the field-level data near Fresno. The binary coloring emphasizes the abundance of perennial orchards (reliable rights) and annual crops (unreliable rights); in practice, the data distinguishes among about forty different crops.

To further explore this diversity, Table A5 reports long-run average shares of farmland statewide and for the two Central Valley regions, the Sacramento and San Joaquin Valleys, over this period, for all crops with at least a 0.5% land share. Given that 78% of the state’s agricultural land lies in the two Central Valley regions, statewide shares broadly approximate Central Valley totals, with a few small exceptions—for example, avocados, while iconically Californian, are mainly grown farther south; leafy vegetables like lettuce and broccoli are more commonly grown along the coast; and alfalfa is more frequently observed outside the Central Valley.

Table A5 also shows meaningful differences between the Sacramento and the San Joaquin River Valleys. Both regions harvest large quantities of almonds and walnuts, though San Joaquin irrigators grow more almond trees (21.6% of land relative to 11.6%), while Sacramento farmers grow more walnut trees (10.7% relative to 3.5%). San Joaquin growers also plant grapes (9.8% of irrigated land), pistachios (8.2%), sweet corn (9.7%) and orange trees (4.5%), at much higher rates than Sacramento farmers. Notably, Sacramento growers plant a great deal of rice (19.7% of irrigated land), a water-intensive but relatively low-value crop almost never grown in the San Joaquin.

4.2 Trees

As mentioned in the introduction, a key element of the data is novel field-level information on the year of planting for orchards observed in 2020, 2021, 2022, and 2023. Plant dates range from 1984–2022, as depicted in Figure 3. Table 3 documents the significance of replanting decisions across farms. Orchard investments exhibit extremely high persistence. Among the 2.97 million acres of fields growing perennial crops observed in my data in 2020, the unconditional annual probability that a tree is cut down between WY 2020 and 2023 is 3.1% statewide. Without correcting for differences in water consumption across tree types that correlate with optimal replanting times, this implies that, in the observed equilibrium, more than 96% of water used in a given year by perennial crops is “committed” to forms of capital with potentially costly adjustment.

Table 3 also shows important heterogeneity in replanting decisions across crop types and ages. Age is a key determinant of replanting; farmers cut down trees planted in the 1990s—between two and three decades old—at an annual rate of 9.0%, in contrast to 3.8% for trees in the second decade of their life, and 0.8% for the first decade. The broad pattern—given existing surface and groundwater rights—is that irrigators appear to have strong economic incentives not to cut down trees for at least two decades. Replacement rates also vary meaningfully across crops. Peaches, like other stone fruit trees, live the shortest productive lives, with annual hazard rates of 2.6% in the first decade of life, then climbing rapidly to 14–17%. Almond trees last longer, with low probabilities of annual replacement ($\approx 5\%$) in the first two decades, rising to 18% in the third decade. Walnut, orange, and olive tree replacement rates remain below 2–4% for the first three decades of their lives, but replacement rates rise in their fourth decade to 8.4%, 4.4%, and 3.1%, respectively. Pistachio trees appear to live forever, with annual replacement rates in the data rarely above 1%.²⁷

These differences across ages and crop types are important to explain the difference in unconditional replanting rates across the hydrological network; for example, 5.1% in the San Joaquin and 7.3% in the Sacramento River after three decades of production. Most of this aggregate difference appears to reflect the San Joaquin’s large share of orange and pistachio trees—which live much longer than the other tree types, and comprise about 13% of irrigated land in the San Joaquin in 2020 relative to less than 1% in the Sacramento—as well as the much larger share of peach and stone fruit trees in the Sacramento River Valley, which live the shortest fruitful lifespans.

4.3 Irrigation

The calculation of irrigation water demand, the functions $\zeta_{ict}(\tau)$, is conceptually very straightforward, but computationally somewhat involved. The surface of climate data allows us to calculate, in theory, how much water a plant needs in a given location i on day τ , what is known as “reference evapotranspiration,” $\eta_{it}(\tau)$. Since the 1980s, the state of California has published reference evapotranspiration for irrigators, who use the information in the same way that we might check the weather. Irrigators input this data into their irrigation scheduling algorithms to

²⁷Pistachio trees have been documented to maintain yields for up to three hundred years.

calculate how much water they should apply to their crops. These calculations use what are known as crop coefficients, $\vartheta_{ic}(\tau)$, which parametrize the demand of water for the crop over time given its location and growing season.²⁸

Figure A10 gives some examples of the daily gridded reference evapotranspiration data. I show maps at the end of October and in mid-July to emphasize how different the growing conditions can be at different points in the year. From the beginning of the California water year in October, it becomes less water-intensive to grow plants through December and January, and then as spring and summer arrive, one needs more water to grow the same crops. Inherent crop requirements also evolve through the growing cycle and differ through regional differences in optimal growing seasons and soil quality. Figure A11 plots monthly crop coefficients, $\vartheta_{ic}(\tau)$, for almonds in the San Joaquin Valley, from the government’s recommendations for optimal irrigation. Over time, as the almond tree grows more leaves and starts to bear almonds, it needs more water, so $\vartheta_{ic}(\tau)$ rises.

Irrigators combine reference evapotranspiration, $\eta_{it}(\tau)$, with crop coefficients, $\vartheta_{ic}(\tau)$, to determine per-acre crop water demand, defined as $\zeta_{ict}(\tau) = \eta_{it}(\tau)\vartheta_{ic}(\tau)$. In the example in Figure A11, the annual water that the almond tree requires is obtained by integrating under the (dotted) curve that multiplies reference evapotranspiration (black line, same variable depicted in Figure A10) by the crop coefficients. The same method applies to annual crops, with more seasonal variability because annual crops start from zero before growing up to be harvested.

4.4 Estimated water use

The model of irrigation scheduling then implies realized water demand, though the optimal irrigation weights $\int \zeta_{ict}(\tau)d\tau$ for each crop and the land allocated to grow these crops, via the identity (3) that follows from (1) and Assumptions A1 and A2. Figure A3, Panel A, shows the implied water use in California in 2020 across watersheds. As expected, most of the state’s water is used in the Central Valley, where large areas of land are used for agriculture and irrigators grow very water-intensive crops. Figure 5 zooms in to compare upstream and downstream regions. On the left is the water used North of the Delta, in the Sacramento

²⁸The methods developed here rely on irrigation scheduling algorithms applied to observed reference crop evapotranspiration from CIMIS, which allows for economically interpretable measures of water demand for irrigation and counterfactual water demand for crops now grown. Appendix B.4 discusses alternative approaches based on ensemble models of realized evapotranspiration.

River Basin, and the right total water used in the San Joaquin River Basin.

4.5 Crop prices, yields, input costs

In addition to observed crop choices and estimated water demand, annual marginal products require crop prices and yields. I use county-level realized prices P_{ict} and yields A_{ict} for all fields ι in a given county growing crop c and year t , substituting California-wide yields for counties where Census redacts yields (those without enough production). Table A6 considers a local HUC12 watershed in the Sacramento River Valley, showing the three largest crops that were produced by acreage. Yields and prices each vary quite a lot between years. The implied revenue-per-acre for different crops, also reported, underlie the estimates of the marginal products of water.

To measure average planting costs, I use cost data from the USDA Agricultural Census, which reports county-level input costs in 2002, 2007, 2012, 2017, 2022. I construct average input cost shares χ_i for each county, including expenditure on seeds, plants, chemicals, fertilizer, fuels, utilities, supplies and repair, contract labor, hired labor, and machinery rental. I omit operating costs labeled taxes, interest, depreciation, livestock, feed, and custom and other agricultural services. I assume that these average planting costs are allocated across crops in proportion to that crop’s revenue, which can be microfounded with a model of flexible factor choice and constant output elasticities with respect to input expenditures (see Rafey, 2023, p. 451).

4.6 Marginal products of water

I now report the revenue per acre-foot of water corresponding to the marginal product of water for that specific field ι for that crop,

$$\alpha_{\iota ct}(x_{\iota t})P_{ict} \equiv \frac{A_{\iota ct}P_{ict}}{\int \zeta_{\iota ct}(\tau)d\tau}, \quad (16)$$

as well as the marginal product net of average county planting costs. Table A7 shows approximately the 95th percentile of revenue per unit water, on the order of about \$4,000 per acre-foot. This significantly exceeds the water prices reported earlier, which range no more than a thousand dollars per acre-foot in the most constrained case in 2024 in the southern Delta. Across crops, looking at the

bottom distribution of these shadow values of water in the lower half of Table A7, one sees very different crops than in the upper tail of the distribution. Wheat and pasture—annual crops that have much lower value per acre—use reasonable amounts of water. These crops use much less water per acre than pistachios, avocados, pears, et cetera, but even though they use less water, they still have much lower values per acre-foot, because their yields and prices are so low.

Average output is the marginal output of each field in the Leontief model specified above, but revenue per acre-foot is not a sufficient statistic for the marginal value when irrigators also adjust other inputs proportionally with greater irrigation volumes. I account for these average planting costs using the data from Section 4.5, multiplying revenue per acre-foot by $1 - \chi_i$, where χ_i is the long-run average revenue share of input costs in i 's county. Accounting for these costs has large consequences for the levels of the field-level water marginal products, as Figure 7, Panel B shows. The effect on dispersion is much more muted than the effect on levels, but in general, adjusting marginal products for average input costs works to tighten the empirical distribution, because average input cost shares in counties with higher-value crops rise greater than proportionally with the value of the crop.

4.7 Hydrology

I compile several federal and state data sources on the hydrological network—including rivers, canals, hydrological regions, rivers, canals, and groundwater wells (USGS, 2019; DWR, 2022, 2025)—which I then match to water rights in order to discipline the counterfactual analysis to interpretable and feasible real-locations.

I also calculate the state transition probabilities of random and persistent state variable, drought, use more than a century of historical California river inflows. For tractability, I use a binary indicator defined as whether the San Joaquin River Index is deemed Critical (C).²⁹ The data indicate that California transitions in

²⁹In practice, observed hydrological conditions include continuous and higher-dimensional processes; e.g., annual inflows, dam reservoir levels, et cetera. The challenge is that expanding the state space to include more general processes will introduce enormous computational complexity into the value function contraction mapping within the nested fixed point algorithm used to estimate the model and simulate counterfactuals. In addition, while I use more than a century of hydrological records to calculate drought state transitions, I only observe renewal decisions from 2020–2023, which limits my ability to estimate a richer model of how water owners respond to varying hydrological conditions.

and out of drought with a conditional probability that depends on the prior year’s drought state: while the probability of transitioning from non-drought to drought is about 0.12, but the probability of remaining in drought from one year to the next is close to 0.5.

Empirically, drought appears to affect several economic outcomes. In regressions of various observables on the drought indicator with crop-by-subwatershed fixed effects, drought lowers yields by about 5% and raises irrigation water requirements (evapotranspiration) by about 0.13 feet relative to a normal year (in wet years, water requirements are -0.24 feet less). At the same time, crop prices rise by about 16% in drought; in total, drought lowers the marginal product of water by an average of 4.7%.

Drought also affects irrigators’ replanting choices. The probability of cutting down an orchard rises by 0.008 in drought years, or an increase of $.008/.041 = 19.7\%$ from the average culling probability of 0.041, in a linear probability model that regresses an indicator for cutting down a given orchard on the annual drought indicator and crop-by-subwatershed fixed effects. This is consistent with an economic model of optimal renewal—drought is a stochastic, highly persistent state variable that affects crop yields, prices, and optimal irrigation levels—and motivates the inclusion of state transitions between drought and non-drought years.

4.8 Estimating unobserved planting costs

To tractably estimate the model from panel data on field-level crop choices, replanting decisions, optimal irrigation, and annual marginal products, I employ the additional structure of the model in Section 3.3. I estimate unobserved planting costs separately across the reliable and unreliable components of water rights, since the two types involve distinct decision problems.

The estimation procedure for reliable water rights aims to recover two sources of planting costs: first, within-field costs that differ over the lifecycle of a given orchard (time-to-build, the rate that yield declines with age, and the fixed cost of replanting); second, across-field costs that differ over fields within a water right (relative costs of planting different crops). I obtain these components in two steps. First, I run a standard full-solution Rust (1987) nested fixed point algorithm derived from (12) for each perennial crop $\times$ region pair to recover crop-region-specific yield decay rates θ_c , crop-region-drought-state-specific fixed renewal costs $\xi_0(s)$, and the value functions for committing water to a specific perennial crop

in a given region at each drought state. For time-to-build, newly-planted young perennials require between three to six years to mature and learn to bear fruit and this “time-to-build” follows known agronomic calendars, so I use reports from USDA to specify, for each crop, the years until first harvest, $\bar{\theta}_c$, with $\alpha_{ict}(x) = 0$ for $x < \bar{\theta}_c$. After maturing, orchard stock depreciates, with yield decaying over time with the age of the trees. I estimate a constant decay rate θ_c for each crop after the first decade of the orchard’s life, so that $\frac{\partial \alpha_{ict}(x)}{\partial x} = 0$ for $x \in (\bar{\theta}_c, 10)$ and $\frac{\partial \alpha_{ict}(x)}{\partial x} = -\theta_c e^{-(x-10)\theta_c}$ for $x > 10$. When the orchard is cut down, the owner incurs a fixed cost of replanting, ξ_c . This step is separable from the subsequent discrete crop choice problem at the time of planting (conditional on renewal) due to the assumptions discussed above.

Second, using value functions from the first step, I run GMM with a Berry (1994) contraction and lagged yield shocks as instruments for current crop prices to estimate relative planting costs across crops $c \in \mathcal{C}_i^k$ from the panel micro data on field-years (ι, t) within each water right i and reliability tier $k \in \{a, r\}$, scaling the flow payoffs for reliable water rights (perennials) with the Rust (1987) values obtained in the first step. Specifically, for each water right i , I construct discrete crop choice moments from fields $\iota \in i$ based on the score functions of the likelihoods implied by (11) and (13). For each candidate $\sigma_\epsilon^i = \sigma$, I run an inner-loop Berry (1994) contraction to obtain the average $\hat{\xi}_{ict}(\sigma)$ ’s from the discrete choice problem of optimal assignment of fields ι to crops c within a water right i that best match the water-right-level choice probabilities and—since only relative, not absolute, costs are identified—enforce that average planting costs within i are mean-zero (i.e., coincide with the level of average input costs in Section 4.5). Given these $\hat{\xi}_{ict}(\sigma)$ ’s, I then estimate the curvature $\hat{\sigma}_{\epsilon^k}^i$ for $k \in \{a, p\}$ by minimizing an objective function constructed by differentiating the likelihood functions given by (11) and (13) with respect to the σ to construct moments that should vanish for the true parameter value, $\mathbb{E}[Z'_{ict} h_\iota(c_{it}, P_{ict}; \sigma, \xi)] = 0$, using instruments Z_{ict} for the endogenous P_{ict} , so that I grid search for σ to minimize a quadratic form of the moments $\sum_\iota g_\iota(\sigma, \xi) \equiv \sum_\iota Z'_{ict} h_\iota(\sigma, \xi)$, where

$$h_\iota(\cdot; \sigma, \xi) = \frac{\partial \ell}{\partial \sigma} = \alpha_{ict}(x_{it}) P_{ict} (\mathbf{1}_{c_{it}=c} - \mathbb{P}(c_{it}(\iota) = c | \sigma, \xi))$$

is the (maximum likelihood) score function, which vanishes at the true (σ, ξ) . The estimator can be understood as a GMM version of a maximum likelihood

estimator for static multinomial logit discrete choice (over the value functions $\bar{V}_i(\iota, 0, 0)$, with choice sets $\mathcal{C}_i$ constructed as all crops observed to be grown in a given water right at any time during the panel), with two modifications. First, the inner-loop Berry (1994) contraction reduces the dimension of the parameter space for each water right so, for each i , I search for $\sigma_{\epsilon^k}^i \in \mathbb{R}_+$ instead of $(\sigma_{\epsilon}^i, \{\xi_{ict}\}_c) \in \mathbb{R}_+ \times \mathbb{R}^{|\mathcal{C}_i|}$. Second, I include price instruments $Z_{\iota ct} \equiv (1, P_{ic,t-1})$, which are necessary for identification even with the microdata (Berry and Haile, 2024). The estimate $\hat{\sigma}_{\epsilon}^i$ determines the magnitude of the unobserved logit shocks across units ι within the water right i . Together with ξ , I can construct the *de novo* value of a yet-uncommitted reliable water right i as

$$\bar{V}_i(0, 0, s) = \int \sigma_{\epsilon^p}^i \log \left(\sum_{c \in \mathcal{C}_i^p} \exp \left\{ \frac{1}{\sigma_{\epsilon^p}^i} V_i(\iota, c, 0) - \xi_{ict} \right\} \right) d\iota, \quad (17)$$

and its decomposition into units in a given year, $\{V_i(\iota, c_i(\iota), x_i(\iota))\}_{\iota \in i}$.

The final estimator is computationally light relative to alternatives. The initial Rust (1987) algorithm takes about four hours. The subsequent Berry (1994) relative planting cost loop—a grid search with a gradient-descent stopping rule—can be easily parallelized across water rights and takes less than three days to run on three 32GB RAM machines to obtain parameter estimates and value functions for each of the approximately 6,000 water rights in the data.

5 Results

I now use the estimates of water use, annual marginal products, and planting costs to study water allocation in California. Section 5.1 constructs marginal products and discusses the primary drivers of average value; Section 5.2 reports the planting cost estimates; Section 5.3 studies dispersion within the empirical distribution of water values. Section 5.4 explains the approach taken to counterfactuals and some of its limitations, then reports the main findings. Section 5.5 discusses sensitivity of the main results to various assumptions.

5.1 Sources of water productivity

Table 6 reports deciles of the distribution of marginal water products and Figure 6 depicts the field-level marginal products of water (\$/acre-foot/yr) at the HUC12

level.³⁰ The lighter green colors are lower values, ranging from zero to about \$2,500/af. Three observable features of water rights strongly predict greater water productivity among fields. First, reliability. Table 6, Panel B shows large differences between productivities of reliable and unreliable water uses, i.e., water used to grow perennial and annual crops, respectively. The empirical distribution of marginal products of water applied to perennial crops nearly stochastically dominates the distribution for annual crops, with decile ratios between the 10% and 90%-iles ranging between 1.9–3.5.

Second, location relative to the Sacramento–San Joaquin Delta flow constraint. In Figure 6, lighter shades of green north of the Delta, particularly in the northern part of the Sacramento River Valley, relative to the much darker patterns below in the San Joaquin River Basin, illustrate the higher average values below the chokepoint in the San Joaquin, as well as substantial dispersion within each region. Table 6, Panel A shows this pattern extends to the deciles of the water value distribution. Much of the difference across watersheds comes after the sixth or seventh decile, places where, in the San Joaquin, irrigators grow high-value fruits and nuts rather than the annual crops in Sacramento.³¹

Third, historical priority. Table A15 regresses some measures of the empirical distribution of productivities—here, water productivity by water right in each watershed-year from 2014–2022—on fixed attributes of historical property rights: the share of “reliable” water rights used to grow perennials, location above or below the Delta, and the share of paper water rights at baseline with priority dates prior to 1914 (depicted in Figure A7), and an indicator for drought). These regressions are descriptive and should not be interpreted as causal, but reveal some interesting patterns. Watersheds with greater pre-1914 water rights have higher median and means, much higher bottom decile values, and less dispersion. These effects are smaller than for the share of perennial crops or location below the chokepoint, and do not survive controlling for region and perennial shares; the latter two fully explain the pre-1914 advantage. This indicates that places with more senior water rights invest more by planting perennial crops, provid-

³⁰A HUC12 subwatershed covers about 40 square miles; this is the finest partition of the USGS (2013) hydrological network.

³¹It is interesting to compare the ratios in these water values with the earlier differences discussed in water prices for trades above or below the Delta chokepoint, reported in Table 2. The Delta price ratio typically ranged from one-and-a-half to two or two-and-a-half, a price gradient that appears comparable to that implied by the microfoundations of the crop model, further indicating the importance of flow constraints (examined in Section 5.4.1).

ing evidence of the greater reliability—and higher value—of these water rights, and contravening a story of moral hazard where senior rightsholders waste water rather than risk expropriation of inherited water rights.

5.2 Planting cost estimates

The Rust (1987) estimates of yield curves and replacement (replanting) fixed costs appear in Panel A of Table 5. Overall, the structural estimates imply that the net-present-discounted-value of operating an orchard ranges from 66–90% of the value of an orchard producing the average yield in every year that lives forever; equivalently, replanting costs, time-to-build, and age-related yield declines comprise 10–34% of the long-run expected value of an eternally fruitful, immortal orchard. This value reflects two components that vary across crops: how crop c ’s yield evolves with age x (time-to-build and the rate of age-related yield declines) and its fixed cost of replanting.

The main heterogeneity in replacement cost estimates occurs across crops, as seen in Table A10 and the crop-specific value functions plotted in Figure A14. This reflects the different replanting patterns discussed Section 4.2 of Table 3. Almonds have total replanting and regeneration costs equal to about 32–35% of the long-run value of the orchard, whereas pistachio costs are closer to 20% of the long-run value, with negligible yield decays ($\hat{\theta}_c \approx 0.009$ and 0.006 in Sacramento and San Joaquin, respectively) relative to almonds ($\hat{\theta}_c \approx 0.047$ and 0.044). Like almonds, stone fruits such as apricots also have substantial yield declines ($\hat{\theta}_c \approx 0.014$ and 0.018), but lower replacement costs ($\hat{\xi}_{c0} \approx 0.80, 1.2$ relative to $4.07, 3.33$), so total costs over the lifecycle equal only about 18–20% of the orchard’s long-run value.³²

Drought lowers the fixed cost of replanting, consistent with the evidence above that orchard culling rises by about 20% in drought years. When yields fall and water requirements rise, the option value of maintaining an aging orchard declines, lowering the threshold for optimal replacement. The mean estimated replacement cost in drought ($\hat{\xi}_{c0}(1) = 1.11$) is about 21% lower than otherwise ($\hat{\xi}_{c0}(0) = 1.41$), with a starker gap in the bottom quartile (0.44 vs. 0.99). Despite this, the long-run value of a committed orchard along the optimal path, $V_i(\iota, c, x, s)$, is

³²Within the same crop, the estimates do not differ much across regions; for example, almond planting costs over the lifecycle are 32.2% in the San Joaquin and 34.5% in Sacramento, apricot planting costs over the lifecycle are 19.5% and 17.9% in the two regions, respectively.

nearly identical across drought states, differing by less than 1% for the median orchard, reflecting orchards’ multi-decade lifespans relative to the typical drought duration; the transition probability back to non-drought is about 0.5.

The Berry (1994) relative planting costs across crops appear in Panels B and C of Table 5, which show estimates from the discrete choice model in Section 3.3. Variation in relative planting costs across c within each water right i in year t is apparent. Mean relative planting costs, ξ_{ict} , exhibit a standard deviation over c for fixed i and t that ranges between 0.07–4.0 for perennial rights and 0.20–2.7 for annual rights between the first and third quartiles. Field-level heterogeneity, determined by the estimates of $\hat{\sigma}_\epsilon^i$, is also nontrivial and varies a great deal across rights; the estimated $\hat{\sigma}_{\epsilon^p}^i$ averages 4.2 across reliable rights i , ranging 0.6 and 6.0 between the first and third quartile, and averages 3.4 for unreliable rights, with an interquartile range of 0.6 to 5.2.

Elasticity calculations, of interest to give the estimates a more easily-interpretable form, follow standard logit formulas. The estimates $\hat{\sigma}_{\epsilon^p}^i$ and $\hat{\sigma}_{\epsilon^a}^i$, combined with the logit structure of (11) and (13), imply own-price elasticities of crop choice at the field level equal to $\frac{1}{\hat{\sigma}_\epsilon^i} [1 - \mathbb{P}(c_{it}(\iota) = c | \sigma, \xi)] \cdot \alpha_{ict}(x_{it}) P_{ict}$. Cross-price elasticities of choosing crop c with respect to the crop price $P_{ic't}$ of an alternative c' is proportional to the choice probability of c' and its marginal product per acre-foot, i.e., $\frac{1}{\hat{\sigma}_\epsilon^i} \mathbb{P}(c_{it}(\iota) = c' | \sigma, \xi) \cdot \alpha_{ic't}(x_{it}) P_{ic't}$. Both annual and perennial crops exhibit moderate average own-price elasticities, 0.161 and 0.245, with quite a lot of dispersion: interquartile ranges of 0.024–0.165 for annual and 0.041–0.261 for perennial crops. This reveals nontrivial planting cost heterogeneity across fields, which prevents water owners from reallocating all of their water to crops that experience positive yield or price shocks.

5.3 Sources of dispersion

As a prelude to analyzing allocative efficiency, Panel A of Table 7 shows how productivity dispersion among California water rights varies across and within various hydrological and institutional boundaries.

Specifically, I partition water rights into subsets indexed by k , so that $\mathcal{J} = \cup_k \mathcal{J}_k$ using observable characteristics of water rights that I obtain—location of the right within the hydrological network (hydrological region or subwatershed, river, and canal) and reliability of the right (reliable versus unreliable). I consider the following partitions: regions ($\mathcal{J}_{SV}$, $\mathcal{J}_{SJ}$, $\mathcal{J}_{other}$), counties (e.g., $\mathcal{J}_{Kern\ County}$,

$\mathcal{I}_{\text{Tulare County}}$, $\mathcal{I}_{\text{Fresno County}}$, etc), water districts (e.g., $\mathcal{I}_{\text{unassigned}}$, $\mathcal{I}_{\text{Westlands}}$, $\mathcal{I}_{\text{Semitropic}}$, etc), rivers (e.g., $\mathcal{I}_{\text{Kern River}}$, $\mathcal{I}_{\text{San Joaquin River}}$, $\mathcal{I}_{\text{Kings River}}$, $\mathcal{I}_{\text{Sacramento River}}$, etc), canals (e.g., $\mathcal{I}_{\text{California Aqueduct}}$, $\mathcal{I}_{\text{Tehama-Colusa Canal}}$, $\mathcal{I}_{\text{Mokelumne Aqueduct}}$, etc), and HUC12 subwatersheds (e.g., $\mathcal{I}_{180400010803}$, $\mathcal{I}_{180300031406}$, $\mathcal{I}_{180201580404}$, etc).

Table 7 shows the average volume of water rights in each of these partitions and reports measures of dispersion (interquartile, interdecile ranges). As expected, the dispersion in water values among water rights falls as we look at narrower collections of water rights. The interquartile range across all water rights in the Central Valley is about \$380/af, and nearly identical across rights within each year; the average range falls to \$220/af within regions. Restricting to within irrigation districts nearly halves the interquartile range, to \$115/af on average among districts, with more muted effects within-river (\$180/af) or within-canal (\$130/af). Within a local subwatershed, the average interquartile range among water rights of the same reliability class is only \$100/af.

5.4 Counterfactuals

Finally, with the empirical distribution of water rights and their values, I explore four questions related to water market design: one, I consider traditional allocative efficiency by revisiting the standard question of trade across water rights in the hydrological network; two, I explore the value of water storage by assessing trade across years; three, I explore the importance of stochastic attributes of water rights by assessing trade across states of the world; four, I explore the importance of the investment channel by contrasting the (extremely limited) short-run water reallocation opportunities with the (much larger) long-run water reallocation opportunities.

The main challenge for counterfactuals is that identifying misallocation requires more assumptions than recovering the values of water across users: the dispersion in true water values at the observed allocation can reflect a combination of flow constraints, trade frictions, and allocative inefficiency. The work above to construct the empirical distribution of the true value of water rights did not restrict trade constraints or frictions, but held fixed—i.e., did not endogenize—water ownership, $\{\widetilde{W}_{it}\}$. For counterfactuals to identify misallocation—that is, to disentangle dispersion that reflects feasibility constraints or trade costs from dispersion that evinces true misallocation—one must take a stance on what (counterfactual) reallocations can occur and at what cost.

To make progress, I draw on additional observables—related to the hydrological network and institutional and jurisdictional boundaries—to follow the comparative approach of *Asker et al.* (2019), which can decompose allocative efficiency into various institutional and hydrological constraints and frictions. Their method is well-suited to study misallocation in settings with many overlapping sources of inefficiency—e.g., unobserved distortions due to market structure, political economy, et cetera—and for which the researcher observes certain characteristics relevant to plausible sources of misallocation (e.g., in *Asker et al.*, 2019, the identity of the cartel member countries and the country of each oil well) alongside the microdata on each unit of production. The approach involves two steps. First, calculate counterfactual values of constrained-efficient allocations under different constraints: i.e., for each k , obtain various values of reallocating different quantities of water among water rights $i \in \mathcal{J}_k$. Second, compare output under the constrained-efficient allocations to decompose distortions into within- and across-components, by subtracting the value of a given constrained-efficient allocation (e.g., within-year, within-region) from the value of a slightly less-constrained allocation (e.g., within-year, across-region), to attribute misallocation to a given distortion (e.g., across-region).

5.4.1 Reallocation across hydrological components

Evaluating traditional reallocation across space—trade in water rights within the hydrological network—is not difficult because I observe the location of each water right within the hydrological network, as well as other institutional attributes that largely correspond to geography (e.g., irrigation district).

A natural place to look for opportunities to reallocate based on the discussion in Section 2.4 is across regions with known flow constraints, such as the upper and lower components of the Central Valley. The interpretation of these calculations is less about misallocation and trading frictions per se that help to explain the distortions above, but instead to speak to the prospective value of relaxing the existing water transport constraints through greater infrastructure investments.

Panel A of Table 7 takes three approaches to obtain constrained-efficient allocations: `gft_q50up` (`gft_q75up`) reallocates water assuming that the land constraint of each existing water right is twice (thrice) that of observed; `gft_q90` assumes all water rights in a given subset produce at the 90%-ile productivity.

Panel B(i) of Table 7 then shows that water reallocation across the Sacramento-

San Joaquin Delta—the value of reallocating water rights within California relative to the value of reallocating water within the two regions separately (i.e., A–C)—equal to 9.9% of total surplus, or between 8.9–10.9% with a 95% bootstrapped confidence interval, when land constraints are twice the observed allocation. Similarly, the estimates indicate that reallocating water across regions, rivers, or canals increases surplus by 10–17% in the baseline land constraint case, with the largest gains across canals (17.0%) and the least across counties (10.3%).

Another natural place to look for opportunities to reallocate water across space among water rights that are likely to be fully hydrologically connected, such as water rights in a given year and reliability class within the same irrigation district or along the same river or canal. The predicted gains from within-district water reallocation—the value of reallocating water rights within districts (in the same year and reliability class) relative to the value of reallocating water within water rights in the same year and reliability class, when land constraints are twice the observed allocation—equal 17.7%, or between 15.0–19.9% with a 95% bootstrapped confidence interval (Table 7B(iii)).

These within-district distortions imply that the combination of observed hydrological constraints, water reliability tiers, and the estimated within-water-right planting costs cannot fully explain water misallocation, consistent with the work of Hagerty (2023) and others. The distortions could either indicate frictions or prospective gains from water trade. The water used within these jurisdictions should not typically be subject to hydrological constraints, but other frictions (such as those mentioned in Section 2.4) may be present.

5.4.2 Reallocation over time

Water reallocation over time can be accomplished through large-scale water storage; California has dam capacity roughly equal to one year of the entire state’s water consumption, and has experimented extensively with groundwater recharge (i.e., water storage in existing aquifers). Interestingly, the intertemporal allocation of water seems efficient: reallocating additional water across years does not seem to create much additional value (Table 7B(ii)). The reason is the enormous persistence in water use from year to year documented in Figure 5. This finding—limited inter-annual gains from trade, despite perfect foresight about future water abundance and scarcity—contrasts sharply with the interregional, within-river, and within-canal prospective gains from trade across water rights.

It indicates limited value for the construction of new water storage projects.

5.4.3 Reallocation across states of the world

The natural way to reallocate water across states of the world is to bundle several unreliable water rights into reliable water rights. Here, a natural counterfactual is how valuable it would be to convert less-reliable to reliable rights (by bundling unreliable rights into smaller, reliable rights). The comparative analysis indicates that bundling unreliable water rights into more reliable water rights within the same local watershed can, in the long run, increase total surplus by 5–12% under the `q50up` land constraint case. This result, reported in Table 7B(iv), is obtained by comparing the value of the efficient allocation within a local watershed within a `huc12`-year and within a water-right-year, with (J, L) and without (K, M) the additional constraint on reliability class (J–K and L–M).

5.4.4 Reallocation in the short- and long-run

The analysis above takes the observed distribution of water rights and orchard capital as a measure of the steady-state value of water use in each location in order to value long-run water reallocation. In addition to the costs of optimal regeneration within a water right, the Rust (1987) structural parameters also allow us to calculate how costly it would be to deviate from the optimal path, e.g., by cutting down all of the orchards in California in a single year or by valuing water transfers to new reliable uses at their initial *de novo* value rather than the long-run value to which they converge. To study how rigidities in water-reliant capital and investment create meaningful differences across short- and long-run water reallocation schemes (which are particularly important to explore given the additional legal barriers in California to permanent or multi-year water trades; cf. fn. 6), I consider “accelerated” counterfactuals that explore cases where water rights are reallocated more quickly.

These abrupt transitions turn out to involve much larger costs than those incurred on the optimal path. Table 8, Panel B, contrasts baseline long-run `gft_q50up` and `gft_q75up` calculations with those that adjust the newly-allocated water to take into account cutting down and replanting an orchard. These costs—about 13% of total surplus in `gft_q50up` and 22% in `gft_q75up`—considerably diminish the prospective gains from reallocation within subsets of water rights. In the comparative analysis—to decompose misallocation into its components in

Panel B—the gains from regional reallocation (A–C) fall by one-fifth and one-third across the two land constraint scenarios (16% and 33%). The rest of the comparisons do not change much, for similar reasons that the *Asker et al. (2019)* method appears robust to planting costs as discussed below in footnote 33: land constraints mean that least-cost allocations each involve the same share of water reallocated, leading the aggregate costs of forced replanting to be similar across most of the least-cost calculations. Interestingly, as Table A16 shows, while replanting costs are lower in drought, this qualitative finding does not seem quantitatively important for the aggregate loss from forced replanting, which is nearly identical whether or not it occurs during drought.

5.5 Caveats, sensitivity, and extensions

A final result is the importance of unobserved planting costs (within water rights) to the bottom-line estimates. For example, row M1 in Table 8, Panel A, shows that the average interquartile and interdecile ranges of values within a water right are \$155/af and \$280/af, respectively, or nearly half of the overall dispersion in water values across water rights (row A in Table 7) or one-fourth to one-fifth of overall dispersion in water values across fields (row A1 in Table A11). If relative planting costs were truly zero, then the estimates would imply a value of “reallocation” within a water right of 29.6% and 53.3% of the overall value (Table A11, row M), or evidence of considerable inefficiency across uses within the same water right.

Misattributing this underlying heterogeneity to productivity instead of planting costs substantially inflates measures of productivity dispersion and predicted gains from trade. Without within-water-right planting costs, interquartile and interdecile dispersion rises to \$600/af and \$1450/af (Table A11, row A), respectively, from \$380/af and \$650/af in the preferred specification that includes within-water-right planting costs. Likewise, the implied distances between the current and “efficient” allocation (Table 8, Panel A, columns 5–7) under a specification that does not adjust for within-water-right planting costs are about 1.5–2 times the gains implied across water rights in the benchmark.

Interestingly, Table A11 also implies that the predicted *Asker et al. (2019)*-style gains (constructed by comparing constrained-efficient allocations across constraints—used to assess the value of relaxing some, but not all, constraints—to value more plausible reallocations) are not much affected by the inclusion

of within-water-right planting costs.³³ This would be the case if the unobserved within-water-right planting costs are not so correlated with the various constraints analyzed. This shows the robustness of comparative analysis in the style of *Asker et al.* (2019), relative to comparisons only of the least-cost or efficient allocation with the observed allocation.

On groundwater. Reliable water rights with greater groundwater shares involve greater irrigation costs because groundwater extraction costs more than diverting surface water. Average groundwater costs are captured by the average variable costs in Section 4.5, which includes utility expenses such as electricity. In addition, meaningful differences exist in groundwater extraction costs across space. To measure groundwater costs, I use well reports to construct average watershed aquifer depths and then follow standard engineering calculations to convert depth into water pumping costs using average electricity rates and pump efficiencies from data on Central Valley irrigators (CEC, 2023). Table A9 reports the empirical distributions of well depths and implied pumping costs in the San Joaquin and the Sacramento Valleys, respectively. The typical San Joaquin well is about twice as deep as the typical well in Sacramento. This translates into differences in variable pumping costs, reported as a ratio in the rightmost column, that are meaningful in terms of groundwater extraction—about a factor of two—although in levels they do not come close to the magnitude of the differences in marginal products from the shadow values. Accounting for these groundwater costs in the valuation of a region’s water rights leads to a smaller gap in water values across regions, as irrigators in the San Joaquin and Tulare River Basins incur higher pumping costs to maintain the reliability of their water rights.³⁴

As an aside, another interesting aspect of the distribution of groundwater costs in Table A9 is that it suggests another kind of water trading, among conjunctive users (i.e., users of both ground and surface water). A conjunctive user in the Sacramento Valley can sell a unit of surface water to a downstream conjunctive user in the San Joaquin, then extract an additional unit of groundwater. If the

³³ E.g., without planting costs the relative misallocation within districts, rivers, and canals ranges from 7–13% (compared with 15–21%) in `gft_q50up` and 13–25% (compared with 28–40%) in `gft_q75up` (see Panel B of Table A11).

³⁴ An alternative interpretation of this difference—in a model in which water is not a binding constraint, but instead farms continuously extract groundwater up to their marginal value for the water—would be that the farmers in the San Joaquin are willing to pull groundwater for a cost of almost twice that of the Sacramento River Basin. In that case, a farmer’s value for water is inferred directly from groundwater pumping costs, and the data directly indicates misallocation of surface water across the two regions.

downstream user in San Joaquin obtaining this surface water then extracts one fewer unit of groundwater, both users will maintain the same levels of irrigation, but the gains from trade equal the difference in groundwater extraction costs, highlighted in Table A9, across the two locations.

Finally, on robustness to irrigation scheduling algorithm assumptions. The results on water misallocation that I find above are made based on comparisons across and between water rights; as these assumptions are sustained across all users and all fields, if they are violated, the bottom-line results on water misallocation will be affected to the extent that the assumption fails systematically for certain subsets of water rights in ways that correlate with the water rights' characteristics used to delineate trading opportunities.³⁵

6 Conclusion

This paper outlined some economics of water property rights and a framework for evaluating water misallocation. The general approach captures three important sources of water value not present in some prior empirical models. First, equilibrium investment in complementary forms of capital, such as trees in orchards, will raise the productivity of water use; these investment costs need to be subtracted from the marginal product of water to determine the water owner's reservation value, and they also imply rigidities that limit short-run water reallocation and create value from commitment to longer-run water trades. Second, water rights often have random components, and the probability of rationing affects the value of the water right and comparisons across water owners. When annual spot markets are incomplete and investment gives rise to adjustment costs, the probability of rationing will affect the value of the right nonlinearly, with more reliable rights creating more value than a mean-preserving spread of those rights across multiple less-reliable tiers. Third, inherent differences in water productivity across locations, correlated with the observed use of water rights, make learning about the counterfactual values of water not straightforward.

The empirical work also delivers several concrete findings. First, empirically,

³⁵For example, if the irrigation scheduling algorithm underpredicts actual irrigation by a constant factor across all fields and crops, the proportional gains from trade will not change because the results are scale-invariant. Alternatively, if walnut orchards deviate more often from irrigation scheduling algorithms than almond orchards, this will be a problem for implied gains from reallocating water between walnuts and almonds.

there seem to be significantly greater marginal products of water used in places below the network bottleneck—where water can be scarce due to flow constraints—and places endowed with senior water rights that are more reliable during drought. These complementary investments are what we might expect in most models of factor-augmenting technical change. Second, this gradient in marginal values has some implications for welfare that depend on the extent to which the gradient reflects planting costs consistent with more senior appropriative rights, or more welfare-relevant constraints on trade. The main specifications indicate (a) large interregional gains from trade; (b) sizeable remaining frictions within regions, irrigation districts, rivers, and canals; (c) value in repackaging or “re-tranching” riskier, transient rights into higher-value, reliable water rights; and (d) very little, if any, gains from trade across years. About one-fifth to one-third of the prospective interregional gains from trade cannot be realized without a gradual transition (on the order of one to two decades) involving optimal replanting, because long-lived capital investments are intertwined with the existing distribution of water rights and steady-state equilibrium water values can be much larger than an abrupt reallocation. Accounting for competing explanations for the dispersion in marginal products, including—but not limited to—estimation of planting costs within water rights and accounting for reliability and hydrological constraints—are crucial to the bottom-line calculations.

References

- ACEMOGLU, D. (2002). Technical change, inequality, and the labor market. *Journal of Economic Literature*, **40** (1), 7–72.
- ADAMOPOULOS, T. and RESTUCCIA, D. (2014). The size distribution of farms and international productivity differences. *American Economic Review*, **104** (6), 1667–1697.
- ARONOFF, D. and RAHEY, W. (2023). Conservation priorities and environmental offsets: Markets for Florida wetlands. *NBER Working Paper*, **31495**.
- ASKER, J., COLLARD-WEXLER, A. and DE LOECKER, J. (2014). Dynamic inputs and resource (mis)allocation. *Journal of Political Economy*, **122** (5), 1013–1063.
- , — and DE LOECKER, J. (2019). (Mis)allocation, market power, and global oil extraction. *American Economic Review*, **109** (4), 1568–1615.
- ATKESON, A. and KEHOE, P. (1999). Models of energy use: Putty-putty versus putty-clay. *American Economic Review*, **89** (4), 1028–1043.

- AYRES, A., MENG, K. and PLANTINGA, A. (2021). Do environmental markets improve on open access? Evidence from California groundwater rights. *Journal of Political Economy*, **129** (10).
- BALBONI, C., BURGESS, R. and OLKEN, B. (forthcoming). The origins and control of forest fires in the tropics. *Review of Economic Studies*.
- BAQAEE, D. and FARHI, E. (2020). Productivity and misallocation in general equilibrium. *Quarterly Journal of Economics*, **135** (1), 105–163.
- and SANGANI, K. (2025). Quotas in general equilibrium. *Working Paper*.
- BERRY, S. (1994). Estimating discrete-choice models of product differentiation. *RAND Journal of Economics*, **25** (2), 242–262.
- and HAILE, P. (2024). Nonparametric identification of differentiated products demand using micro data. *Econometrica*, **92** (4), 1135–1162.
- , LEVINSOHN, J. and PAKES, A. (1995). Automobile prices in market equilibrium. *Econometrica*, **63** (4), 841–890.
- BORENSTEIN, S. *et al.* (2002). Measuring market inefficiencies in California’s restructured wholesale electricity market. *American Economic Review*, **92** (5), 1376–1405.
- BOSER, A., CAYLOR, K., LARSEN, A., PASCOLINI-CAMPBELL, M., REAGER, J. and CARLETON, T. (2022). Field scale crop water consumption estimates reveal potential water savings in California agriculture. *Nature Communications*, **15** (2366).
- BURGESS, R., HANSEN, M., OLKEN, B., POTAPOV, P. and SIEBER, S. (2012). The political economy of deforestation in the tropics. *Quarterly Journal of Economics*, **127** (4), 1707–1754.
- BURLIG, F., PREONAS, L. and WOERMAN, M. (2024). Groundwater and crop choice in the short and long run. *Energy Institute at Haas Working Paper*.
- BURNESS, H. S. and QUIRK, J. (1979). Appropriative water rights and the efficient allocation of resources. *American Economic Review*, **69** (1), 25–37.
- CARLETON, T., CREWS, L. and NATH, I. (2025). Agriculture, trade, and the spatial efficiency of global water use. *Working Paper*.
- CEC (2023). Appendix A for the CEC Final Project Report titled Estimates of Groundwater Pumping Electricity Use and Costs in California.
- COSTINOT, A., DONALDSON, D. and SMITH, C. (2016). Evolving comparative advantage and the impact of climate change in agricultural markets: Evidence from 1.7 million fields around the world. *Journal of Political Economy*, **124** (1), 205–248.
- D’ODORICO, P., CHIARELLI, D., ROSA, L., BINI, A., ZILBERMAN, D. and RULLI, M. C. (2020). The global value of water in agriculture. *Proceedings of the National Academy of Sciences*, **117** (36), 21985–21993.
- DWR (2022). i12 Canals and Aqueducts local.
- DWR (2023). *State Water Project drives California’s economy*. Tech. rep., California Department of Water Resources.
- DWR (2025). Well Completion Reports.
- EDELSTEIN, B. (2025). Water scarcity management and housing markets: Evidence from water impact fees in Colorado. *Job Market Paper, UPenn.*
- FERGUSON, B. (2024). Trade frictions in surface water markets. *Job Market Paper, Stanford GSB*.

- GARTRELL, G., MOUNT, J., HANAK, E. and GRAY, B. (2017). *A new approach to accounting for environmental water: Insights from the Sacramento-San Joaquin Delta*. Tech. rep., Public Policy Institute of California.
- GATES, P. (1991). *Land and law in California: Essays on land policies*. Ames, Iowa: Iowa State University Press.
- GOWRISANKARAN, G. and RYSMAN, M. (2012). Dynamics of consumer demand for new durable goods. *Journal of Political Economy*, **120** (6), 1173–1219.
- HAGERTY, N. (2022). Adaptation to surface water scarcity in irrigated agriculture. *Working Paper, Montana State University*.
- (2023). What holds back water markets? Transaction costs and the gains from trade. *Working Paper, Montana State University*.
- HANAK, E., SENCAN, G. and AYRES, A. (2021). *California’s water market*. Tech. rep., Public Policy Institute of California.
- HSIAO, A. (2026). Coordination and commitment in international climate action: Evidence from palm oil. *Econometrica*, **94** (1), 1–33.
- HSIEH, C.-T. and KLENOW, P. (2009). Misallocation and manufacturing TFP in China and India. *The Quarterly Journal of Economics*, **74** (4), 1403–1448.
- LEONARD, B., COSTELLO, C. and KROETZ, K. (2025). The economics of contracting for conservation. *Working Paper*.
- LIBECAP, G. (2011). Institutional path dependence in climate adaptation: Coman’s ‘Some unsettled problems of irrigation’. *American Economic Review*, **101** (1), 64–80.
- MEDELLIN-AZUARA, J. *et al.* (2022). *Economic impacts of the 2020–22 drought on California agriculture*. Tech. rep., U.C. Merced Water Systems Management Lab Report for the California Department of Food and Agriculture.
- MIDRIGAN, V. and XU, D. Y. (2014). Finance and misallocation: Evidence from plant-level data. *American Economic Review*, **104** (2), 422–458.
- RAFEY, W. (2023). Droughts, deluges, and (river) diversions: Valuing market-based water reallocation. *American Economic Review*, **113** (2).
- REGNACQ, C., DINAR, A. and HANAK, E. (2016). The gravity of water: Water trade frictions in California. *American Journal of Agricultural Economics*, **98** (5), 1273–1294.
- RUST, J. (1987). Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher. *Econometrica*, **55** (5), 999–1033.
- SCOTT, P. (2013). Dynamic discrete choice estimation of agricultural land use. *Working Paper*.
- SOUZA-RODRIGUES, E. (2019). Deforestation in the Amazon: A unified framework for estimation and policy analysis. *Review of Economic Studies*, **86** (6), 2713–2744.
- SUM, S. (2024). Trading water and work: Labor and production shifts under California’s groundwater markets. *Job Market Paper, UCSD*.
- SYVERSON, C. (2004). Market structure and productivity: A concrete example. *Journal of Political Economy*, **112** (6), 1181–1222.
- TIMMINS, C. (2002). Measuring the dynamic efficiency costs of regulators’ preferences: Municipal water utilities in the arid West. *Econometrica*, **70** (2), 603–629.
- TOLKIEN, J. (1965). *The Lord of the Rings*, vol. 2. Ballantine Books.

- USGS (2013). *Federal standards and procedures for the National Watershed Boundary Dataset (WBD), Fourth edition*. Tech. rep., USGS and USDA, Natural Resource Conservation Service.
- USGS (2019). *National Hydrography Dataset (NHD)*. Tech. rep., U.S. Geological Survey.
- WONG, A. J., JIN, Y., MEDELLÍN-AZUARA, J., PAW, K. T., KENT, E. R., CLAY, J. M., GAO, F., FISHER, J. B., RIVERA, G., LEE, C. M., HEMES, K. S., EICHELMANN, E., BALDOCCHI, D. D. and HOOK, S. J. (2021). Multiscale assessment of agricultural consumptive water use in California’s Central Valley. *Water Resources Research*, **57** (9).
- WONG, C., LOBELL, D. and MAUTER, M. (2025). Field-scale mapping of California crop water productivity to inform water management in critically overdrafted groundwater basins. *Environmental Research Letters*, **20** (074034).
- ZEFF, H., HAMILTON, A. L., MALEK, K., HERMAN, J. D., COHEN, J. S., MEDELLIN-AZUARA, J., REED, P. M. and CHARACKLIS, G. W. (2021). California’s food-energy-water system: An open source simulation model of adaptive surface and groundwater management in the Central Valley. *Environmental Modelling and Software*, **141** (105052).

TABLE 1. NATURAL VARIABILITY AND WATER RATIONING

	Wet Year	Dry Year	Dry, pre-1914	Dry, post-1914
0	0	0	0	0
0.1	0	0	0	0
0.2	0.0002	0	0.001	0
0.3	0.031	0.002	0.032	0
0.4	0.105	0.040	0.093	0.016
0.5	0.234	0.120	0.191	0.086
0.6	0.426	0.251	0.356	0.203
0.7	0.676	0.475	0.593	0.400
0.8	0.935	0.774	0.890	0.714
0.9	1	1	1	1
1	1	1	1	1

Central Valley irrigators' reported annual water diversions divided by reported face value — deciles of the **water right** $\times$ **year** distribution, 2010–2023.

Columns (1) and (2) split the sample by water-year designation. Columns (3) and (4) split the sample in column (2) by water right priority date, into rightsholders whose claims date before (“pre-1914”) and after (“post-1914”) the 1914 Water Commission Act. Wet and (critically) dry water-year designations $\equiv$ the California Eight River Index (Sacramento River + San Joaquin River Runoff).

Source. Author's calculation using SWRCB WRIMS data. 27,010 wet year diversion statements; 26,750 critically dry year diversion statements.

TABLE 2. SURFACE WATER PRICE GRADIENTS ACROSS THE DELTA

	San Joaquin (SOD)	Sacramento (NOD)	diff	ratio
0	5.2	1.1	4.1	4.7
0.1	30.8	43.6	−12.9	0.7
0.2	48.8	51.3	−2.5	1.0
0.3	71.2	60.6	10.6	1.2
0.4	102.5	65	37.5	1.6
0.5	146.7	65	81.7	2.3
0.6	166	83.1	82.9	2.0
0.7	241.6	107	134.6	2.3
0.8	278.4	133.5	145.0	2.1
0.9	983	250	733	3.9
1	1, 550	342.3	1, 207.7	4.5

Distribution of annual surface water allocation prices (2024 \$/af), trade-level, over assembled trades 1987–2009 (Libecap 2011) and 2022–2024 (WestWater, LLC).

TABLE 3. HAZARD RATE BY DECADE PLANTED, 2020–2023

tree_type	acres_2020	haz_CA	haz_SV	haz_SJ
all / all	2,946,282	0.031	0.029	0.031
1980s /	340,265	0.056	0.073	0.051
1990s /	353,317	0.090	0.092	0.098
2000s /	879,380	0.038	0.028	0.041
2010s /	1,373,320	0.008	0.007	0.009
/ citrus	299,035	0.022	0.035	0.020
/ olives	48,271	0.013	0.009	0.026
/ almonds	1,390,560	0.042	0.037	0.044
/ walnuts	426,517	0.026	0.020	0.032
/ pistachios	460,901	0.001	0.002	0.001
/ plums, prunes and apricots	70,817	0.065	0.058	0.078
1980s / citrus	79,441	0.044	0.051	0.042
1990s / citrus	58,232	0.021	0.042	0.019
2000s / citrus	72,491	0.024	0.012	0.020
2010s / citrus	88,871	0.004	0.045	0.007
1980s / olives	14,056	0.031	0.023	0.047
1990s / olives	5,770	0.022	0.008	0.046
2000s / olives	16,121	0.018	0.016	0.036
2010s / olives	12,324	0.003	0.011	0.002
1980s / almonds	52,930	0.110	0.103	0.102
1990s / almonds	137,416	0.189	0.182	0.197
2000s / almonds	464,391	0.052	0.037	0.055
2010s / almonds	735,823	0.010	0.004	0.011
1980s / walnuts	72,681	0.084	0.085	0.083
1990s / walnuts	66,514	0.041	0.031	0.051
2000s / walnuts	104,075	0.017	0.010	0.026
2010s / walnuts	183,247	0.011	0.008	0.014
1980s / pistachios	50,968	0.006	0.013	0.006
1990s / pistachios	40,436	0.004	0.007	0.003
2000s / pistachios	127,951	0.001	0.001	0.001
2010s / pistachios	241,546	0.001	0.002	0.001
1980s / plums, prunes and apricots	12,701	0.076	0.070	0.108
1990s / plums, prunes and apricots	16,779	0.143	0.142	0.143
2000s / plums, prunes and apricots	15,987	0.074	0.049	0.103
2010s / plums, prunes and apricots	25,349	0.014	0.009	0.026

Average share of perennial trees cut down (annual “hazard rate”) between WYs 2020–21, WYs 2021–22, and WYs 2022–23 by decade of original planting date and crop type. Reported for the top six perennial crops, by acreage among Central Valley irrigators in 2020. Calculated for all of California, for the Sacramento Valley (SV) and San Joaquin and Tulare River Basins (SJ). Figures 3 and A13 plot related information.

TABLE 4. NEW PERENNIAL PLANTING DECISIONS, 2004–2022

Dependent Variable: Model:	(1)	(2)	chosen (3)	(4)	(5)
<i>Variables</i>					
$\ln(\alpha_{ict}P_{ict})$	−0.0696 (0.0164)	0.0309 (0.0076)	0.0287 (0.0078)	0.0733 (0.0140)	0.0595 (0.0155)
Instruments				✓	✓
<i>Fixed-effects</i>					
region-year	✓	✓		✓	
region-crop		✓		✓	
region-crop-year			✓		✓
<i>Fit statistics</i>					
Observations	2,337,203	2,337,203	2,337,203	2,337,176	2,337,176
R ²	0.01587	0.18235	0.19766	0.18036	0.19704
Within R ²	0.01425	0.00129	0.00067	−0.00115	−0.00011

Clustered (crop-year) standard-errors in parentheses

Illustrative linear regression at the field-crop-year-planted level for newly-planted perennial orchards from 2004–2022. Outcome is 1 if the crop planted equals c and 0 otherwise.

Columns (1)–(3) regress choice on log annual marginal water products and covariates. Columns (4)–(5) instrument for log annual marginal water products using lagged yields and prices, i.e., $\ln(A_{ic,t-1}P_{ic,t-1})/\zeta_{ict}$ for $\ln(A_{ict}P_{ict}/\zeta_{ict})$.

TABLE 5. MLE/GMM ESTIMATES OF PLANTING COSTS

A. Optimal Regeneration (Rust)

Statistic	q25	Median	Mean	q75
N	361.2	1,590.5	9,306.833	5,759.8
theta	0.003	0.006	0.010	0.012
xi0	0.988	1.650	1.409	2.000
xi0_drought	0.438	1.100	1.106	1.950
V_1.c	16.007	16.744	16.640	17.484
V_1.c_drought	16.027	16.887	16.683	17.520

B. New Planting (Perennial)

Statistic	q25	Median	Mean	q75
vol	1,464.007	5,975.932	40,639.000	28,629.440
N	15	55	322.391	252
N_c	2	3	3.661	5
sig	0.600	3.600	4.200	6.000
msqe	0.00001	0.0001	0.008	0.001
xi_mean	-1.251	-0.405	-0.660	-0.105
xi_sd	0.070	1.314	2.875	4.000
E_pred	189.580	362.074	1,209.136	515.898

C. New Planting (Annual)

Statistic	q25	Median	Mean	q75
vol	220.429	1,399.805	10,266.100	4,878.628
N	11	30	135.739	95
N_c	2	4	4.817	6
sig	0.600	2.800	3.407	5.200
msqe	0.00000	0.00000	0.001	0.0001
xi_mean	-1.180	-0.452	-0.656	-0.118
xi_sd	0.197	1.123	1.902	2.691
E_pred	248.377	415.233	987.284	706.351

Panel A. Summary of Rust (1987) maximum likelihood estimates of planting cost parameters over all (51) region×tree types. N = number of field-level observations from 2020–2023 within each region×tree type subsample; year-planted data starts in 2020.

Panel B. Summary of GMM estimates of relative crop planting cost parameters over all (1,761) reliable Central Valley water rights. N = number of field-level observations from 2004–2022 within each water right reconstructed from planting dates from field-level data in 2020–2023; evapotranspiration data starts in 2004.

Panel C. Summary of GMM estimates of planting cost parameters over all (4,270) unreliable Central Valley water rights. N = number of field-level observations from 2014–2022 within each water right; field-level data starts in 2014 and yield data stops in 2022.

See Appendix C for additional details and Table A10 for additional results.

TABLE 6. AVERAGE WATER PRODUCTIVITIES, 2020

A. Distribution of water marginal products per acre-foot by region

	San Joaquin (SOD)	Sacramento (NOD)	diff	ratio
0	25.60	23.48	2.12	1.09
0.1	140.09	87.25	52.84	1.61
0.2	189.56	119.54	70.02	1.59
0.3	198.09	149.04	49.06	1.33
0.4	243.08	156.08	87.00	1.56
0.5	276.14	179.17	96.97	1.54
0.6	299.55	184.29	115.26	1.63
0.7	344.43	213.12	131.31	1.62
0.8	505.02	220.52	284.50	2.29
0.9	1,188.41	254.54	933.88	4.67
1	9,207.11	897.94	8,309.18	10.25

B. Distribution of water marginal products per acre-foot by reliability

	Perennial crops	Annual crops	diff	ratio
0	19.19	25.60	-6.41	0.75
0.1	492.25	140.09	352.17	3.51
0.2	528.54	189.56	338.98	2.79
0.3	560.93	198.09	362.83	2.83
0.4	598.99	243.08	355.91	2.46
0.5	641.45	276.14	365.31	2.32
0.6	776.81	299.55	477.27	2.59
0.7	1,045.28	344.43	700.86	3.03
0.8	1,540.72	505.02	1,035.70	3.05
0.9	2,257.93	1,188.41	1,069.51	1.90
1	3,103.81	9,207.11	-6,103.31	0.34

Annual marginal products of water (\$/af), Central Valley irrigators, WY 2020.

TABLE 7. MEASURING MISALLOCATION IN CALIFORNIA

A. Dispersion

	N	avg_vol	IQR	IDR	gft_q50up	gft_q75up	gft_q90
all [A]	1	115,382,621	381.175	653.070	0.443	0.766	0.668
year [B]	7	16,483,231	373.452	661.792	0.440	0.754	0.683
region [C]	2	57,691,310	218.843	460.163	0.344	0.638	0.544
region-year [D]	14	8,241,615	217.722	447.013	0.339	0.625	0.531
ry-class [E]	28	4,120,807	212.565	432.987	0.319	0.600	0.523
ryclass-county [F]	471	244,973	218.114	458.183	0.216	0.408	0.361
ryclass-district [G]	1,073	107,532	116.091	248.591	0.177	0.358	0.317
ryclass-river [H]	1,063	108,544	179.418	346.051	0.209	0.399	0.372
ryclass-canal [I]	1,525	75,660	134.267	272.988	0.148	0.282	0.277
ry-huc12 [J]	4,235	27,245	131.875	244.170	0.144	0.258	0.258
ryclass-huc12 [K]	6,501	17,748	100.530	186.421	0.095	0.161	0.172
ry-wr [L]	26,991	4,274	63.267	111.890	0.127	0.154	0.223
ryclass-wr [M]	38,029	3,034	0	0	0	0	0
year-pre1914 [N]	14	8,241,615	357.067	638.384	0.432	0.743	0.664
ry-pre1914 [O]	28	4,120,807	253.562	525.597	0.336	0.618	0.546
ryclass-pre1914 [P]	56	2,060,403	257.932	505.251	0.321	0.591	0.533

Panel A. Estimated distribution and long-run values of water rights owned by Central Valley irrigators, WYs 2014–2022. Partitions indexed by year (y), region (r: Sacramento or San Joaquin), reliability class (c: reliable or unreliable), and hydrological or institutional boundaries: county, irrigation district, river, canal, HUC12 subwatershed (huc12), and individual water right (wr). Row [M] (ryclass-wr) assigns each water right to its own group, so within-group dispersion is zero by construction (Assumption A3). Rows [N]–[P] replace the regional partition with a pre-1914/post-1914 priority date partition.

N reports the number of partitions, **avg_vol** the mean volume of water rights in each partition (acre-feet), and **IQR** and **IDR** average measures of productivity dispersion (volume-weighted interquartile and interdecile ranges, in \$/af/yr) within the indicated partitions. **gft_q50up** and **gft_q75up** report total gains from reallocating water within each group as a share of the total pre-trade value; **gft_q50up** (**gft_q75up**) assumes that land constraints bind at twice (thrice) each right’s observed use. **gft_q90** reports the gain if all water rights within each group had the 90th-%ile productivity.

Table is continued on the next page $\hookrightarrow$

TABLE 7 (cont'd). MEASURING MISALLOCATION IN CALIFORNIA

B. Distortions

(i) infrastructure	$\Delta\text{gft_q50up}$		$\Delta\text{gft_q75up}$		$\Delta\text{gft_q90}$	
across_regions [A-C]	0.099	[0.089, 0.109]	0.128	[0.117, 0.140]	0.124	[0.098, 0.162]
across_county [E-F]	0.103	[0.068, 0.125]	0.191	[0.188, 0.230]	0.162	[0.132, 0.215]
across_district [E-G]	0.141	[0.112, 0.167]	0.241	[0.239, 0.281]	0.205	[0.191, 0.273]
across_river [E-H]	0.109	[0.085, 0.142]	0.200	[0.211, 0.255]	0.151	[0.130, 0.217]
across_canal [E-I]	0.170	[0.159, 0.209]	0.318	[0.325, 0.369]	0.246	[0.224, 0.313]
(ii) storage						
across_years [A-B]	0.003	[-0.002, 0.012]	0.012	[0.008, 0.029]	-0.015	[-0.034, 0.023]
(iii) local frictions						
within_county [F-M]	0.216	[0.192, 0.240]	0.408	[0.361, 0.401]	0.361	[0.324, 0.372]
within_district [G-M]	0.177	[0.150, 0.199]	0.358	[0.314, 0.352]	0.317	[0.268, 0.311]
within_river [H-M]	0.209	[0.175, 0.226]	0.399	[0.339, 0.379]	0.372	[0.322, 0.374]
within_canal [I-M]	0.148	[0.110, 0.151]	0.282	[0.232, 0.260]	0.277	[0.229, 0.277]
across_pre1914_ry [D-O]	0.003	[-0.012, 0.020]	0.007	[0.000, 0.022]	-0.015	[-0.033, 0.019]
across_pre1914_ryc [E-P]	-0.002	[-0.015, 0.023]	0.009	[0.006, 0.028]	-0.010	[-0.032, 0.025]
(iv) bundling						
bundle_within_hy [J-K]	0.049	[0.034, 0.060]	0.097	[0.070, 0.089]	0.086	[0.062, 0.102]
bundle_within_wry [L-M]	0.127	[0.046, 0.085]	0.154	[0.073, 0.097]	0.223	[0.109, 0.149]

Continuation of Table 7.

Panel B. Distortions attributed to specific constraints, obtained by differencing Panel A rows as discussed in Section 5.4. For example, **across_regions** [A-C] measures the value of cross-regional reallocation by subtracting line C from line A; **within_district** [G-M] measures within-district misallocation net of within-water-right planting costs; **bundle_within_hy** [J-K] measures the value of converting unreliable water rights into reliable rights within local subwatersheds.

Brackets report 95% confidence intervals constructed from 1000 bootstrap iterations.

TABLE 8. MEASURING MISALLOCATION – ROLE OF PLANTING COSTS

A. Within-Water-Right

	N	IQR	IDR	gft_q50up	gft_q75up	gft_q90
A	1	381.175	653.070	0.443	0.766	0.668
A1	1	589.009	1,439.722	0.688	1.423	1.242
B	7	373.452	661.792	0.440	0.754	0.683
B1	7	544.397	1,453.526	0.684	1.405	1.269
C	2	218.843	460.163	0.344	0.638	0.544
C1	2	456.056	1,061.785	0.603	1.245	1.125
D	14	217.722	447.013	0.339	0.625	0.531
D1	14	442.489	1,052.141	0.583	1.229	1.116
E	28	212.565	432.987	0.319	0.600	0.523
E1	28	308.514	783.514	0.476	1.071	0.987
F	471	218.114	458.183	0.216	0.408	0.361
F1	482	204.643	455.043	0.413	0.758	0.758
G	1,073	116.091	248.591	0.177	0.358	0.317
G1	1,331	337.582	617.506	0.382	0.772	0.658
H	1,063	179.418	346.051	0.209	0.399	0.372
H1	1,161	266.767	531.969	0.423	0.783	0.708
I	1,525	134.267	272.988	0.148	0.282	0.277
I1	2,116	305.379	627.690	0.364	0.665	0.611
J	4,235	131.875	244.170	0.144	0.258	0.258
J1	5,729	360.822	649.530	0.381	0.672	0.672
K	6,501	100.530	186.421	0.095	0.161	0.172
K1	9,476	252.766	472.372	0.288	0.493	0.545
L	26,991	63.267	111.890	0.127	0.154	0.223
L1	27,254	205.819	383.662	0.393	0.701	0.668
M	38,029	0	0	0	0	0
M1	38,376	155.213	280.198	0.296	0.533	0.533
N	14	357.067	638.384	0.432	0.743	0.664
N1	14	572.448	1,468.583	0.677	1.397	1.250
O	28	253.562	525.597	0.336	0.618	0.546
O1	28	452.800	1,073.084	0.582	1.227	1.108
P	56	257.932	505.251	0.321	0.591	0.533
P1	56	317.088	805.266	0.476	1.064	0.980

A. Baseline result (attributing dispersion within a water right to heterogeneous planting costs; Table 7, Panel A) compared to no heterogeneity in planting costs within a water right. Rows with “1” suffix (A1, B1, C1, ...) are the no-planting-cost specification.

TABLE 8 (cont'd). MEASURING MISALLOCATION – ROLE OF PLANTING COSTS

B. Within-Orchard				
	N	avg_vol	gft_q50up	gft_q75up
A	1	115,382,621	0.443	0.766
A1	1	115,382,621	0.315	0.543
B	7	16,483,231	0.440	0.754
B1	7	16,483,231	0.312	0.532
C	2	57,691,310	0.344	0.638
C1	2	57,691,310	0.233	0.456
D	14	8,241,615	0.339	0.625
D1	14	8,241,615	0.226	0.447
E	28	4,120,807	0.319	0.600
E1	28	4,120,807	0.214	0.422
F	471	244,973	0.216	0.408
F1	471	244,973	0.115	0.242
G	1,073	107,532	0.177	0.358
G1	1,073	107,532	0.059	0.186
H	1,063	108,544	0.209	0.399
H1	1,063	108,544	0.098	0.228
I	1,525	75,660	0.148	0.282
I1	1,525	75,660	0.023	0.113
J	4,235	27,245	0.144	0.258
J1	4,235	27,245	0.003	0.085
K	6,501	17,748	0.095	0.161
K1	6,501	17,748	−0.048	−0.008
L	26,991	4,274	0.127	0.154
L1	26,991	4,274	−0.032	−0.015
M	38,029	3,034	0	0
M1	38,029	3,034	−0.161	−0.161
N	14	8,241,615	0.432	0.743
N1	14	8,241,615	0.306	0.520
O	28	4,120,807	0.336	0.618
O1	28	4,120,807	0.222	0.437
P	56	2,060,403	0.321	0.591
P1	56	2,060,403	0.214	0.412

B. Baseline result (long-run water values; Table 7, Panel A) compared with forced replanting (reliable water reallocated within the year). Rows with “1” suffix (A1, B1, C1, ...) are the forced replanting specification. Table A16 shows drought / no drought.





FIGURE 1. TOTAL USABLE INFLOWS + APPROPRIATIVE RIGHTS, 1980–2021

Source. Author's calculations from data in Gartrell, et al. (2022).

NB. Vertical axes differ in the two panels.

(a) Irrigation Water Rights, by Priority Date



(b) Irrigation Water Rights, by Priority and Point of Diversion





FIGURE 2. THE LEGACY OF APPROPRIATIVE WATER RIGHTS

Source. Author's calculations using the California State Water Resources Control Board Water Rights Information Management System. Central Valley irrigation water rights only. 2,571 water rights claimed prior to 1914; 3,069 water rights claimed after 1914. See Figure A7 for richer spatial distribution of water rights.

A. Age Distribution, 2022



B. Hazard Rates by Age, 2020–2021




FIGURE 3. ORCHARD DEMOGRAPHICS

A. Age distribution distribution of all Central Valley orchards, 2022. Average orchard in 2022 was planted 14.7 years ago. The data does not distinguish between trees planted in 1984 or in earlier years; read “1984” as “ ≤ 1984 ” (but also consider reading *1984...*).

B. Hazard rates by age; all Central Valley orchards. See Figure A13 for crop-specific hazard rates for citrus, almonds, peaches, and pistachios.

Source. Author’s calculations from data described in Section 4.2.





FIGURE 4. IMPLIED AGRICULTURAL WATER RIGHTS, BY LOCATION

Implied agricultural water rights by subwatershed (HUC12), $\ln(\text{acre-feet})$, 2020. See Figure A3 for map of all California.

A. North of the Delta



B. South of the Delta






FIGURE 5. IMPLIED AGRICULTURAL WATER RIGHTS, BY RELIABILITY CLASS

Implied water rights reliability, 2014, 2016, and 2018–2023, inferred from agricultural uses. Reliable water rights $\equiv$ water used for perennial crops. Variable $\equiv$ annual crops.



(a) Sacramento River Basin
(North of the Delta)




(b) San Joaquin River Basin
(South of the Delta)

FIGURE 6. ANNUAL MARGINAL PRODUCTS OF WATER, \$/AF, 2020

Average estimated annual marginal products of water by subwatershed (HUC12), 2020.

A. Water Values



B. Marginal Products and Average Planting Costs





FIGURE 7. VALUE OF WATER RIGHTS, \$/AF/YR, 2014–2022


A. Estimated values of water rights used to irrigate Central Valley fields (net of average and estimated planting costs), 2014–2022. Rightmost bin includes all values $\geq \$2000/\text{af}$.

B. Estimated marginal products with (blue) and without (purple) average planting costs, 2014–2022. Rightmost bin includes all values $\geq \$2500/\text{af}$.

A. Reliability



B. Region



C. Historical Priority




FIGURE 8. WATER VALUES BY RELIABILITY, REGION, AND HISTORICAL PRIORITY

Estimated values of water rights, \$/af/yr, 2014–2022.

Reliable water rights $\equiv$ water used for perennial crops. Variable $\equiv$ annual crops.