Transaction Costs and the Gains from Trade in Water Markets

Nick Hagerty*

October 11, 2025

Abstract

This paper estimates the potential benefits of reducing transaction costs in California's surface water market. I develop an empirical framework to analyze welfare in water markets that uses transactions data, inferring the preferences of water districts from their prior behavior. I separate observed prices into demand and transaction costs and then simulate a market without transaction costs, assuming no capacity constraints. Regulatory review and market inexperience predict transaction costs. Eliminating transaction costs produces gains from trade worth 1-3% of statewide water expenditures per year. Even with zero transaction costs, physical conveyance costs prevent further gains, making some price dispersion inevitable.

---

[*] Montana State University, nicholas.hagerty@montana.edu. This paper previously circulated under the title "Liquid Constrained in California: Estimating the Potential Gains from Water Markets." I thank Clay Landry and WestWater Research, LLC for generously sharing the data that made this project possible, as well as Kavish Gandhi for excellent research assistance. Many people contributed helpful feedback that improved this paper; I particularly thank David Atkin, Ellen Bruno, Sydnee Caldwell, Josh Dean, Dave Donaldson, Esther Duflo, Michael Greenstone, Peter Hull, Donghee Jo, Chris Knittel, Jack Liebersohn, Matt Lowe, Rafe Meager, Ben Olken, Arianna Ornaghi, Doug Parker, Will Rafey, Gary Sawyers, Matt Zaragoza-Watkins, Ariel Zucker, two anonymous referees, and seminar participants at MIT, Harvard, EDF, UC Davis, UC Berkeley, PPIC, and the Tinbergen Institute. Funding for this research was provided by the Abdul Latif Jameel Water and Food Systems Lab (J-WAFS) at MIT.

1 Introduction

Water supplies are becoming scarcer and more variable in many parts of the world (UNDP 2006). Fueled by population pressure and climate change, water scarcity can increase poverty and conflict (Sekhri 2014; Burke et al. 2015) and is set to worsen in the coming decades (World Bank Group 2016). To help societies adapt, many observers advocate for greater use of water markets. Like other markets, water markets may yield benefits by allocating scarce resources to the most valuable uses and allowing participants to flexibly respond to changing conditions. But formal and large-scale water markets are still rare, even in wealthy countries with strong legal institutions (Brewer et al. 2008; Olmstead 2010). Why? A rich qualitative literature points to unique institutional features of water, complex political economy issues, and transaction costs (Leonard et al. 2019). But relatively little quantitative evidence is available to assess the magnitudes of these barriers and set priorities for reform.

This paper quantifies the importance of transaction costs in California's wholesale surface water market, the largest in the United States. I ask: How large are the potential gains from eliminating transaction costs? In other words, by how much is water misallocated at present, and by how much could it be improved through policy reform? To answer this question, I develop a new empirical framework to estimate demand in the presence of transaction costs that relies on transactions data. The idea is to infer the preferences of water districts – the local government agencies that supply water to farms and households and control most water rights – from their observed behavior in the market that already exists. After separating observed prices into demand and transaction costs, I use the demand model to simulate a market without transaction costs – but with the physical costs of conveyance – and estimate the gains from expanded trade.

Two reasons make California's water market a useful setting in which to study the role of transaction costs. First, transaction costs are likely binding constraints. Water transfers are already technologically feasible; a comprehensive network of canals, pipelines, rivers, and reservoirs connects nearly all water users in the state. Property rights are relatively strong: the basic legal basis for using and selling water is well-defined, and regulatory processes for transferring water exist

(Leonard et al. 2019). Second, anecdotal evidence suggests the potential gains from reallocating water could be large. California has a large and diverse economy, but most of it depends on water supplies that are imported over great distances and prone to droughts. Most of this water is quantity-rationed by historical rules, and retail prices can vary over more than two orders of magnitude: in 2023, according to their websites, commercial and industrial customers in the city of San Diego paid $2,855 per acre-foot,[1] while agricultural customers in the Imperial Valley, less than a two hours' drive away, paid just $20 per acre-foot.

My starting point is to interpret the high price dispersion and low transaction volume in California's water market as equilibrium outcomes of trade with transaction costs. Transaction costs insert a wedge between observed prices and districts' marginal value of water, and they prevent other trades from happening at all. These costs arise from a variety of sources: internal district decisions, search and contracting, regulatory approval processes, and transportation and delivery. Many transaction costs are unobserved or implicit, so rather than attempting to measure them directly, I infer them from market outcomes.

I model water districts that trade endowments of a homogeneous good with no production. Districts pay a fixed cost to enter the market and complete any transactions in a year, and they pay *ad valorem* variable costs to either sell or purchase water. Fixed costs vary across districts, while variable costs vary across buyer-seller pairs and are directionally asymmetric; both arise from observable and unobservable sources. Spatial arbitrage equalizes prices up to variable transaction costs. Using these assumptions, I derive an empirical model of water market prices that attributes price variation across counterparties to transaction costs and price variation over time to demand.

I estimate the model in four steps. First, I estimate variable transaction costs that arise from observable predictors, such as conveyance distance and regulatory reviews, by comparing prices across transactions made by the same district in the same year. Second, I estimate the extensive-margin decision of market entry as a function of observable predictors. Third, I estimate the inverse price elasticity of demand, which measures how districts' marginal value of water responds to

---

[1] An acre-foot, the standard unit of volume for water in the American West, is the amount of water that would cover one acre of land with twelve inches of water.

changes in quantity consumed, using variation in annual surface water endowments generated by allocation rules in the federal and state water projects. Fourth, I estimate district-specific demand intercepts and the unobserved component of variable transaction costs using random effects. From these estimates, I recover transaction costs and demand.

To estimate the model, I construct a comprehensive new dataset on California's water economy. For water transactions, I use a proprietary dataset that to my knowledge is the most complete in existence; crucially, it provides a nearly complete record of prices. I link the transactions data to the universe of wholesale yearly surface water deliveries in California, a new dataset assembled jointly for this paper and for Hagerty (2022). I also build a large crosswalk file to link users across datasets and years, a geospatial dataset on user locations and boundaries, and a basic hydrological model of California's water infrastructure.

In intermediate results, I find large variable transaction costs associated with specific predictors. For example, buyers receive a 45 percent discount for transactions that are subject to additional regulatory review because they originate from within a federal or state water project. Fixed costs are also important barriers to trade: Districts are less likely to trade in wetter years, when they have more recent market experience, and if they face fewer regulatory reviews across all potential transactions. I find an inverse elasticity of -1.4, implying a price elasticity of -0.69 for districts that trade. Accounting for both observed and unobserved sources, median variable transaction costs are 27 percent for sellers, 21 percent for buyers, and 54 percent of the transaction price in total. Marginal valuations remain highly dispersed, suggesting that gains are available from reducing transaction costs and increasing trade.

For my main result, I combine the demand model with the hydrological model to simulate an efficient market without transaction costs and calculate the resulting gains. The objective is to equalize marginal valuations across districts up to purely physical transportation costs, which are intrinsic to a water market and can never be eliminated. Two types of gains result: lower costs of transactions that already occur, and benefits from new transactions that would not otherwise occur. I estimate that efficient trading would reduce existing transaction costs by \$5 to \$85 million per year, and new trades would achieve gains of $211 to $559 million per year. Gains are disproportionately large in dry years, suggesting that markets can be especially valuable for risk management.

These gains are small relative to water supply expenditures or GDP in California. I discuss several explanations. First, water is costly to move, so much of the dispersion in prices is inevitable. Second, water demand is inelastic, so dispersion can be eliminated with relatively little reallocation. Third, water districts may be trading more conservatively than their own constituents or customers would prefer. Fourth, my analysis focuses on trade between water districts, but there may be greater potential gains from water reallocation among farmers and households within districts.

However, the gains are still considerable in absolute terms and may be worth policy reform efforts. Barriers to trade arise from both traditional transaction costs (i.e., search, negotiation, contracting) and regulatory reviews, showing up in both fixed and variable costs. These types of transaction costs may be reduced by increasing market information, improving coordination, and streamlining regulatory approval processes.

This analysis has several important limitations. First, the data are sparse because the market is thin – an inherent feature of seeking to learn from trading behavior in this context. Second, not all districts appear in the transactions data, so the simulations require some extrapolation. But the districts observed trading account for most of the available water and most of the gains from trade. Third, the simulations assume water infrastructure forms a fully connected network, which holds for the vast majority of water districts in California.[2] Fourth, I do not model most capacity constraints, as detailed infrastructure modeling is beyond the scope of this paper. But when I do impose two major constraints, total gains remain largely unchanged.

This paper makes several contributions. For one, it provides a method to estimate demand in water markets that uses transactions data from the existing market. Most prior analyses of the

---

[2] Some transfers might require a complicated series of conveyance, exchanges, and/or forbearance, but physically, it is possible to transfer water between nearly any two water users in the state, at least when canals are not at full capacity (Hanak and Stryjewski 2012). The only exceptions are small water users in remote areas that are not included in my simulation sample.

gains from water markets couple detailed engineering models of water infrastructure with models of water demand; examples include California (Howitt et al. 1999; Sunding et al. 2002; Jenkins et al. 2003; Medellín-Azuara et al. 2007), Australia (Peterson et al. 2005; Qureshi et al. 2009), and Chile (Rosegrant et al. 2000). The demand models in these studies are typically mathematical programming models of agricultural production that rely on a large number of calibrated parameters (Mérel and Howitt 2014). A recent variation is Rafey (2023), who uses modern production function estimation techniques to value the *ex post* gains from water trading in Australia. My approach is different because it models demand of the relevant actors: water districts, not individual farmers. It is more parsimonious than the programming models and arguably more realistic because it is estimated using observed water market transactions. Even where fine-scale agricultural data is available, working directly with the revealed preferences of water districts makes my approach more policy-relevant so long as districts continue to be the primary market participants.

My approach builds on a vast literature in international trade on measuring trade costs, particularly methods that infer costs from price gaps rather than trade flows (Eaton and Kortum 2002; Donaldson 2018; Atkin and Donaldson 2015). I share the goal of indirectly estimating frictions that are difficult or impossible to measure directly, but I adapt the setup because water is homogeneous, has no production, and transaction costs are bilaterally asymmetric. My approach also can be used more generally to analyze trade of a factor endowment in the presence of transaction costs, such as environmental permit markets (Stavins 1995; Liski 2001) or electricity markets.

I also empirically measure transaction costs in a specific market, which is uncommon outside of international trade. A handful of papers do so in financial markets (Kyle and Obizhaeva 2016), agricultural labor markets (Foster and Rosenzweig 2022), supplier contracting (Boehm 2022; MacKay 2022), and electricity markets (Jha and Wolak 2023). A few other papers evaluate the consequences of reducing transaction costs in financial services (Jack and Suri 2014; de Mel et al. 2022; Batista and Vicente 2023), the upgrade of durable goods (Hodgson 2023), and pollution permit markets (Gangadharan 2000; Cason and Gangadharan 2003). In water markets, Womble and Hanemann (2020) estimate explicit transaction costs from surveys, while Carey et al. (2002) and Regnacq et al. (2016) study the effects of transaction costs on trading quantities but not welfare. Ayres et al. (2018) study transaction costs in the governance of groundwater resources.

Finally, my results shed light on the factors constraining water markets: I show that transaction costs matter, quantify their consequences, and provide evidence on which types are most important. Empirical evidence on water markets has long been limited by both data and methods; I make progress on both. This paper also contributes to a broader literature on the costs of misallocation in markets such as housing (Glaeser and Luttmer 2003), capital (Hsieh and Klenow 2009), energy (Davis and Kilian 2011), labor (Bryan and Morten 2019; Adamopoulos et al. 2022), and land (Chen et al. 2023).

2 Context

2.1 California's water market

California's water market consists of voluntary sales or leases of the right to use surface water. Trading is allowed under the Water Code so long as the seller holds legal right to the water and the water would have been used otherwise (DWR and USBR 2019). Transactions are negotiated bilaterally, with no central clearinghouse, and may be temporary or permanent; my analysis focuses on within-year leases, the most common form. Most must obtain legal approval from one or more regulatory agencies.

Most trade involve local public agencies that supply water to farms and households and hold most water rights. These agencies are often referred to as water districts. Typical sellers are farm-serving districts with secure rights in water-rich areas. Typical buyers include urban districts, agricultural districts in water-scarce areas with higher-value crops, and environmental programs seeking to increase instream flows. Trading helps districts respond to short-term drought shocks and long-term shifts in water demand and urban growth (Hanak and Stryjewski 2012; DWR and SWRCB 2015).

Moving water is feasible thanks to one of the world's most sophisticated systems of water infrastructure. Canals, pipelines, and rivers together connect nearly all water districts in California.[3] Many canals have spare capacity, and reservoirs along routes provide storage. Canal owners must grant access for transfers unless negative effects would result (Water Code §1810). Legal and infrastructure barriers preclude trade with other states or countries. Groundwater is regulated separately and cannot be traded over long distances, so it falls outside the scope of my analysis.

2.2 Water districts

Most water rights are held by special districts, which include water districts, irrigation districts, flood control and water conservation districts, and municipal utility districts. A smaller share belongs to cities, county water agencies, nonprofit mutual water companies, and for-profit utilities. Individual farmers hold many water rights by count but few by volume, as most receive irrigation water through a district. For simplicity, I refer to all potential buyers and sellers as water districts, though my data include all these entities.

Districts hold and trade water rights on behalf of the farmers, households, and firms they serve; individual irrigators or households typically cannot trade outside of the district (Chong and Sunding 2006). Thus, I focus on trade among water districts rather than among users within a district.[4]

2.3 Water rights and allocation percentages

Many districts directly hold surface water rights, which come in two types: appropriative and riparian. Others hold long-term contracts to receive water from the federal and state water projects – California's State Water Project (SWP) and the federal Central Valley Project (CVP) and Lower Colorado River operations – with the project itself holding the underlying water right.

Surface water is scarce and rationed annually, creating year-to-year variation in districts' endowments based on winter precipitation and runoff. I use this variation to estimate the price elastic-

---

[3] On a shared river or canal, the seller simply takes less water upstream while the buyer takes more downstream. Other transactions are more complex, with water allowed to flow downstream later pumped into a canal and conveyed hundreds of miles.

[4] Local water trading among farmers within a district occurs in some irrigation districts, but data are scarce.

ity of demand. Most variation comes from the SWP and CVP. Deliveries from these projects equal a district's time-invariant maximum contract volume multiplied by that year's allocation percentage, which is determined separately for 13 categories of contracts grouped by history, geography, and sector. Some categories tend to have higher priority, but the ordering varies with regional water conditions (Stene 1995). Allocation percentages do not fully determine water deliveries, as districts can trade water, take less than their allocation, bank water for temporary storage, or apply for extra water under certain conditions.

Other sources of surface water are more stable. The Lower Colorado system saw no shortages before 2023. Appropriative rights follow a seniority rule: in dry years, senior rights-holders receive their full claim before junior holders receive any. The SWP and CVP are junior to all other rights, so cutbacks mainly affect project contractors (DWR and USBR, 2019).

3 Transaction costs

Trading surface water in California is generally legal and feasible but costly and difficult. Observers widely recognize the difficulties: Culp et al. (2014) write, "A myriad of physical, legal, and regulatory restrictions operate to inhibit the movement of water from one user to another and from one type of use to another." Sellers et al. (2016) note that "California has a water market, but its function is limited due to a patchwork of regulations that act as institutional barriers." Western Water Company (2000) complained that "The she[e]r complexity of the current water transfer process... is enough to stop most potential transfers from even being proposed."

Economists and legal scholars often describe these difficulties as transaction costs. Many authors interpret disparities in retail prices or marginal values as transaction costs that imply unexploited gains from trade (Young, 1986; Thompson, Barton H., Jr., 1993; Haddad, 2000; Brewer et al., 2007; Libecap, 2011). Bretsen and Hill (2009) argue that "transaction costs become a de facto tax on the transfer of water rights," creating a "chilling effect" that deters trades. A few studies provide empirical evidence that transaction costs impede trade (Archibald and Renwick, 1998; Carey et al., 2002; Regnacq et al., 2016), though they focus on quantities traded rather than monetizing the costs or foregone gains.

Transaction costs arise from many sources: districts must make internal decisions, find and bargain with trading partners, secure regulatory approvals, and physically transport the water. I follow much of the water markets literature in defining transaction costs broadly as all costs of completing a transfer, including both explicit payments and non-monetary frictions.[5] Regulatory hurdles in particular are not unique to water markets. Similar permitting and approval processes have been shown to impose large costs in sectors such as housing (Glaeser and Gyourko, 2018; Manville et al., 2023), transportation infrastructure (Brooks and Liscow, 2023; Bennon et al., 2024), and renewable energy (Jarvis, 2025).

3.1 Types of transaction costs

Internal decision costs. To trade water, a district's board must decide whether to participate, how much to trade, how to acquire water, and whether to proceed with specific transactions. These decisions are costly for several reasons: conflict over how to use the revenues (irrigators who give up water often cannot be directly compensated); uneven distribution of costs and benefits (transfers can harm other irrigators through reduced return flows or shifted tax burdens and harm the broader community by shrinking the local farm economy), and a preference for stability and consensus that leads many boards to avoid controversy (Bretsen and Hill, 2009; Edwards and Libecap, 2015).

Governance structure shapes these costs. Some boards are elected by popular vote, others by property-weighted vote.[6] Districts controlled by landowners tend to be more willing to sell water

---

[5] Leonard et al. (2019) define transaction costs as "the costs of defining, exchanging, and enforcing water rights." The traditional definition focuses on information, search, negotiation, and enforcement (Coase 1937; Demsetz 1968; Williamson 1979; Barzel 1982) and typically excludes regulatory and transportation costs. In contrast, the water markets literature often adopts a broader scope, including regulatory barriers (Colby, 1990; Archibald and Renwick, 1998; Bretsen and Hill, 2009) and conveyance costs (Thompson, Barton H., Jr., 1993; Rosegrant and Binswanger, 1994; McCann and Easter, 2004). This broader view parallels the concept of trade costs in international trade.

[6] Generally, irrigation districts use popular voting; water and water storage districts use landowner voting weighted by assessed value; reclamation districts by land area; and mutual water companies by shareholder holdings. But districts in California are authorized under more than 140 statutes so there are exceptions (Thompson, Barton H., Jr., 1993).

than those that respond to the broader community, as sales often benefit landowners while harming others. Decision costs are likely higher for districts that answer to more, and more heterogeneous, constituents than for those run by a small group of farmers growing similar crops (Hanak, 2003; Bretsen and Hill, 2009; Edwards and Libecap, 2015; Bruno et al., 2024).

Anecdotal evidence suggests intra-district costs are substantial. Hanak (2003) observes that "community resistance has soured a number of deals... and has likely prevented others from being proposed." Such resistance slows negotiations, adding cost and uncertainty. Libecap (2011) notes that in the 1990s and 2000s, the Palo Verde Irrigation District, controlled by landowners, transferred water to the Los Angeles area "with little controversy," while the neighboring Imperial Irrigation District, elected by popular vote, considered the same but was "mired in lengthy, complex negotiations." Sometimes harms are explicitly compensated: a large Imperial–San Diego transfer included a $20 million mitigation fund (Hanak, 2003).

Administrative costs. Next, districts must search for a trading partner, negotiate terms, and agree upon a contract. First, potential buyers and sellers must identify trading opportunities, find a willing trading partner, and learn about their preferences. Matches are usually made through word of mouth but sometimes with help from a professional broker. Second, the districts must negotiate price, quantity, timing, delivery point and route, payment terms, and responsibility for related costs (SWRCB, 1999). Third, they must draft, execute, and later enforce a contract.

These steps can be costly absent a centralized exchange, standardized contracts, or public data. They create explicit costs – broker fees, attorney fees, staff time – and may carry even larger implicit costs from delays and uncertainty, deterring all but the most motivated districts from attempting to trade (Gray et al., 2015). For example, a long-term transfer between Imperial and Metropolitan Water District required four years to reach an initial deal, followed by "five years of negotiations, three agreements, and a side letter" (Bretsen and Hill, 2009). By contrast, in Australia's Murray–Darling Basin, online platforms allow deals to clear and deliver within a day (Johnson, 2015).

Regulatory costs. After a contract is negotiated, districts must secure approvals from the State Water Resources Control Board (SWRCB), Department of Water Resources (DWR), and/or the U.S. Bureau of Reclamation (USBR), depending on source, destination, and delivery route. Transfers may also trigger review by state and federal wildlife agencies and require documentation of compliance with state and federal environmental and endangered species laws. Agencies must ensure that transfers do not (1) harm other legal water users, (2) unreasonably affect fish, wildlife, or instream uses, or (3) harm the origin county's economy (Water Code §§1725, 1810). Reviews focus on identifying "new" or "wet" water – the consumptive use that would have occurred absent the transfer (Hanak and Stryjewski, 2012).

Process. The seller files a petition following guidelines in SWRCB (1999) or DWR and USBR (2019). Petitions must justify legal authority, document compliance with relevant laws, and analyze effects on other users, the environment, and agency operations. Because California lacks standardized diversion metering, transfers must be based on verifiable forbearance actions. For example, for a transfer based on crop idling, sellers must submit 5–20 years of crop history, detailed maps, and a water budget with hydrological and agronomic estimates. Once filed, the petition enters public review. Other parties can protest, requiring the seller to respond. Conflicts may be resolved through negotiation; otherwise, the agency holds a public hearing. Agencies encourage pre-filing consultations with multiple agencies, environmental groups, and other water users to prevent conflicts (DWR and SWRCB, 2015). After approval, sellers face strict monitoring, including field inspections and even soil sensors (DWR and USBR, 2019).

Explicit costs. Costs of regulatory reviews include agency fees, staff time, attorney fees, and technical consultants. Agencies charge a petition fee plus full reimbursement for review and administration. For example, USBR charged $15,000 to review a 3,675 acre-foot transfer in 2015 (Scheer, 2016). Estimates of explicit regulatory costs in Western water markets range from 1% of price (Scheer, 2016) to 112% (Womble and Hanemann, 2020).

Implicit costs. Delays and uncertainty can be even greater burdens. Even short-term transfers take months (SWRCB, 1999); long-term transfers can take years (Bretsen and Hill, 2009). Sequential agency reviews can take so long that the transfer is no longer needed (Sellers et al., 2016). Third-party protests add risk and costs for negotiations and mitigation or compensation, as "literally anyone can... either block the transfer or make it too expensive to complete" (Western Water Company, 2000). Although outright rejections are rare (SWRCB approved 98% of petitions in the 1980s–90s), proposals are sometimes withdrawn. For example, a 2012 transfer between Modesto Irrigation District and San Francisco's utility was abandoned due to opposition from the City of Modesto and other local groups (Regnacq et al., 2016).

Delivery costs. Finally, the seller must physically move the water to the buyer. Even in an otherwise frictionless market, moving water is costly: some water is lost to evaporation and percolation, and energy is required to pump water through canals and over mountains. Districts must pay to use conveyance facilities, send more water than is delivered, and can face infrastructure constraints.

Wheeling charges. Many transfers require use of canals and pumps, whose owners – most often the DWR and USBR – charge fees known as wheeling charges. These fees typically exceed marginal costs to cover operations and sometimes capital costs, a practice some market participants call "barrier pricing" (Western Water Company, 2000). Charges can be so high as to effectively double the price of water, such as in 1998 when Metropolitan charged $262 per acre-foot to transport water between Imperial and San Diego (Chong and Sunding, 2006).

Conveyance and carriage losses. To account for evaporation and percolation during transport, agencies assess *conveyance losses* that reduce the quantity delivered. Transfers pumped across the Sacramento–San Joaquin Delta face additional *carriage losses* of 20 to 35%. This water is sent to sea to maintain outflows and prevent ocean salinity from harming ecosystems and other water rights (DWR and USBR, 2019). Both types of losses require the seller to forgo more water than the buyer receives, raising the effective price.

Delta pumping restrictions. Transfers that must cross the Delta also face effective losses and delivery risks due to environmental restrictions on DWR and USBR pumping. These rules protect water quality and species such as Delta smelt and Chinook salmon under the Endangered Species

Act (DWR and USBR, 2019). One cost is greater effective losses: transfers are limited to specific months, and if sellers cannot store water received outside this window, only part of the water they forgo using can actually be transferred. Hanak and Stryjewski (2012) estimate this can produce an effective surcharge of up to 40%. A separate cost is delivery risk: pumping may be curtailed unpredictably even during the transfer window, so water released by a seller may never reach the buyer. DWR and USBR (2019) warn that even fully approved transfers are not guaranteed to succeed, and Johnson (2015) reports that Metropolitan lost two-thirds of a $13 million water purchase in 2003.

3.2 Persistence of transaction costs

Why do transaction costs persist if there are gains to eliminating them? The literature offers two main explanations. First, some transaction costs are efficient. Many delivery costs are the unavoidable result of physics. Regulatory reviews clarify and protect property rights (Libecap, 2009) and reduce the harms of transfers by internalizing external costs (Colby, 1990; Culp et al., 2014). Still, overregulation is possible (Colby, 1990), and reviews introduce asymmetry: transfers face greater scrutiny than initial rights, and reviews seek to prevent all harm rather than weigh benefits against costs. Second, historical path dependence and distributional conflict impede reform. California's water institutions were built to move water into agriculture and are poorly suited to the reverse (Bretsen and Hill, 2009). Institutional change is costly, creates winners and losers, and generates uncertainty (Demsetz, 1967; Edwards and Libecap, 2015). Selling water often concentrates benefits among a subset of farmers while imposing costs on other farmers and community members (Leonard et al., 2019).

3.3 Representation in the analysis

Many transaction costs are unobservable. Prior methods, cited above, can estimate explicit costs like attorney and regulatory fees but not internal decision costs, administrative costs, or implicit regulatory and delivery costs. To account for all types of transaction costs, I use a model to infer their aggregate impact from price data. The model incorporates specific sources of transaction costs in two ways: by using observable predictors and by allowing for unobserved costs under theoretical assumptions. While this approach cannot measure each cost separately, it provides evidence on the sources and estimates their overall welfare impact.

Fixed vs. variable costs. Some costs are variable – scaling with transfer volume – such as per-acre-foot wheeling charges and fractional carriage losses. Others are more fixed – incurred for transfers of any size – such as search costs and attorney fees. Internal decisions likely include both components, as do regulatory costs: DWR and USBR (2019) note that review fees depend on transfer size, while survey evidence suggests that explicit legal and regulatory costs increase with volume but less than one-for-one (Colby, 1990; Scheer, 2016; Womble and Hanemann, 2020).

Incidence across buyers and sellers. Some transaction costs are nominally borne by sellers (e.g., regulatory review fees, monitoring and verification), others by buyers (most delivery costs and risks), and many by both (intra-district, administrative, implicit regulatory costs) (DWR and USBR, 2019). However, actual incidence depends on how contracts are structured. For example, sellers often invoice buyers for regulatory review fees (Scheer, 2016), and mitigation funds may be included in the transaction price (paid by sellers) or provided separately (paid by buyers) (Hanak, 2003). Details of contract provisions are scarce, so rather than assuming how transaction costs are shared between buyers and sellers, I estimate their incidence from the data.

Conveyance costs. Delivery costs that are purely physical – e.g., energy costs of pumping or water lost to evaporation and percolation – have no economic remedy and can never be eliminated. Demand estimation must adjust for all delivery costs, but the simulations include physical conveyance costs as calculated from a hydrological model.

4 Model

I model trade of surface water among water districts in the presence of transaction costs.

4.1 Assumptions

Each district j begins each year t with an endowment of water E_{jt}. There is no production, and water is a homogeneous good.[7] Districts make bilateral transactions indexed by i.

Assumption 1 (Demand)**.** *Each district has an inverse demand function $V_{jt}(Q_{jt})$, which gives marginal valuation V_{jt} as a function of quantity consumed Q_{jt}. Demand is isoelastic, varies across districts, and may shift over time:*

$$\ln V_{jt} = \eta_j \widetilde{\ln Q}_{jt} + \bar{v}_j + v_{jt} \tag{1}$$

where $\mathbb{E}[v_{jt}] = 0$, and log quantity consumed is centered on the district mean.[8]

Note this is demand for surface water consumption, not for market transactions. Demand describes not how trade itself responds to price, but rather how marginal valuations of water respond to surface water quantity. Water districts must have well-defined preferences, but I make no assumptions about how their preferences form.

The most substantive restriction is that a district's marginal valuation V_{jt} is constant across transactions i within the year t. This restriction follows from the plausible assumption that each year's transactions are planned simultaneously, and it is crucial for separating demand from transaction costs. The isoelastic functional form is important for tractability in estimation but not in extrapolation; in simulations I explore alternative functional forms.

Assumption 2 (Transaction costs)**.** *District j pays a fixed cost f_{jt} to enter the market and complete any transactions. Then, both sales and purchases incur variable transaction costs that are* ad valorem *(iceberg): District j must give up τ^s_{jk} units of water to sell 1 unit to another district k, and τ^b_{jk} units to buy 1 unit from district k, where $\tau^s_{jk}, \tau^b_{jk} \geq 1$ for all j, k.*

---

[7] Because a buyer receives all water through the same infrastructure, water quality does not vary across potential sellers. It is not literally the same molecules of water that are being sold but rather a series of changes in flows. Few contracts specify water quality (McCann and Easter 2004).

[8] $\widetilde{\ln Q}_{jt} \equiv \ln Q_{jt} - \overline{\ln Q}_j$, so that the intercept $\bar{v}_j$ represents mean log marginal valuation.

Fixed costs vary across districts and years but are constant across counterparties. This restriction is necessary for tractability, but fixed costs are plausibly more associated with market entry than with specific counterparties. Variable costs vary across buyer-seller pairs and are bilaterally asymmetric. The iceberg specification is uniquely tractable and common in international trade (Anderson and van Wincoop 2004), but unlike that literature, I distinguish between costs paid by the seller and the buyer. Variable costs are constant over time, a key restriction that allows me to separate transaction costs from demand.

Assumption 3 (Components of transaction costs)**.** *Variable transaction costs have observable and unobservable components that are multiplicatively separable:* $\ln \tau^s_{jk} = \rho^s \mathbf{B}_{jk} + \check{\tau}^s_{jk}$ *(for sales) and* $\ln \tau^b_{jk} = \rho^b \mathbf{B}_{jk} + \check{\tau}^b_{jk}$ *(for purchases), where* $\mathbf{B}_{jk}$ *is a vector of observable predictors, and unobservable sources load onto* $\check{\tau}^s_{jk}$ *and* $\check{\tau}^b_{jk}$*.*

Since variable transaction costs τ^s_{jk} and τ^b_{jk} are greater than 1, they can be rewritten with a mean-zero error term using a change of variables. For each side of the market $m \in \{s, b\}$:

$$\ln \tau^m_{jk} = \rho^m \mathbf{B}_{jk} + \bar{\tau}^m_j + \tilde{\tau}^m_{jk} \tag{2}$$

where $\bar{\tau}^m \equiv \mathbb{E}[\check{\tau}^m_{jk}]$ and $\tilde{\tau}^m_{jk} \equiv \check{\tau}^m_{jk} - \bar{\tau}^m$, such that $\mathbb{E}[\tilde{\tau}^m_{jk}] = 0$. The most substantive restriction here is that cost components do not interact – log total costs equals the sum of its parts.

Assumption 4 (No arbitrage)**.** *If a pair of districts trades, prices equalize marginal valuations up to variable transaction costs. For an observed transaction* i *in year* t *between any two districts, where the seller is* j *and the buyer* k*:*

$$\tau^s_{jk} V_{jt} = p_{ijkt} = \frac{1}{\tau^b_{kj}} V_{kt}. \tag{3}$$

Relative to marginal valuations, the negotiated price gives a premium to the seller, and a discount to the buyer, that is exactly large enough to compensate for the variable transaction costs that each incurs. In equilibrium, each pair of districts that trade equalize their marginal valuations up to variable transaction costs, and each district equalizes its marginal valuation across all districts it trades with. Fixed costs affect who participates in the market but are then sunk, so they do not affect equilibrium prices conditional on the set of participants.

No arbitrage is the crucial assumption underpinning most of the literature on transaction and trade costs.[9] The main reason it might not hold is if there is market power on either side of the market. I investigate this possibility in an appendix to an earlier working paper (Hagerty 2019) and find little evidence of market power by either buyers or sellers. A separate and contemporaneous analysis by Tomori et al. (2024) finds similar results in the same setting.

Not all pairs of districts trade with each other. Trade will not occur if variable transaction costs are greater than the gap in marginal valuations, or if fixed transaction costs are too large. But since the good is homogeneous, the absence of trade does not imply that transaction costs were too large. Each district might achieve Equation 3 by trading with a different partner instead, resulting in the sparse trading network seen in the data.

Not all districts trade at all. District j enters the market only if the potential surplus from trading makes up for the fixed cost of entry:

$$f_{jt} \leq S_{jt}(q_{jt}, Q_{jt}) \tag{4}$$

where q_{jt} is the potential trading quantity (i.e., net volume purchased in the market if entered), Q_{jt} is total quantity consumed, and potential surplus S_{jt} is the integral of inverse demand over the trading quantity less the sum of prices paid and transaction costs incurred.

---

[9] Appendix B.2 discusses possible motivations and derives Equation 3 from a setup involving water brokers.

4.2 Identification

To take the model to data, I take logs of Equation 3 and allow for a mean-zero residual term, $\mathbb{E}[\epsilon_{ijkt}] = 0$, which might represent errors in measurement or in districts' optimization:

$$\begin{aligned} (Sellers) \qquad \ln p_{ijkt} &= \ln V_{jt} + \ln \tau^{s}_{jk} + \epsilon_{ijkt} \\ (Buyers) \qquad \ln p_{ijkt} &= \ln V_{jt} - \ln \tau^{b}_{jk} + \epsilon_{ijkt}. \end{aligned} \tag{5}$$

For a given district, these equations attribute price variation across counterparties to differences in transaction costs, and price variation across years to changes in marginal valuations. Combining equations for buyers and sellers and substituting equations 1 and 2 yields a full model of water market prices:

$$\ln p_{ijkt} = \rho \mathbf{B}_{jk} + \eta_j \widetilde{\ln Q}_{jt} + \bar{v}_j + \bar{\tau}_j + \tilde{\tau}_{jk} + v_{jt} + \epsilon_{ijkt} \tag{6}$$

where $\rho \equiv (1-b)\rho^{s} - b\rho^{b}$, $\bar{\tau}_j \equiv (1-b)\bar{\tau}^{s}_{j} - b\bar{\tau}^{b}_{j}$, $\tilde{\tau}_{jk} \equiv (1-b)\tilde{\tau}^{s}_{jk} - b\tilde{\tau}^{b}_{jk}$, and b is an indicator that takes a value of 1 for buyers and 0 for sellers.

In Equation 6, the demand intercept $\bar{v}_j$ and the mean level of unobserved transaction costs $\bar{\tau}_j$ are not separately identified in transactions data without further restrictions. To identify them, I use a technique inspired by stochastic frontier analysis (Aigner et al. 1977; Kumbhakar et al. 2022).[10]

Assumption 5 (Unobserved transaction costs)**.** *For each district, trade with at least one counterparty carries no variable transaction costs:* $\min_k\{\tau^{s}_{jk}\} = 1$ *and* $\min_k\{\tau^{b}_{jk}\} = 1$ *in the observed data for each district* j.

This assumption shifts the distribution of unobserved transaction costs, setting their means to a level that guarantees that overall transaction costs are never a negative percentage of the transaction

---

[10] Stochastic frontier analysis was developed to estimate one-sided deviations from an unobserved frontier. Its usual application is firm inefficiency relative to a production frontier, but the basic setup also describes transaction costs relative to marginal valuations.

($\tau_{jk} \geq 1$), enforcing Assumption 2. For each side of the market $m \in \{s, b\}$, I can recover

$$\bar{\tau}_j^m = -\min_k \left\{ \rho^m \mathbf{B}_{jk} + \tilde{\tau}_{jk}^m \right\}. \tag{7}$$

The demand intercept $\bar{v}_j$ is then shifted (down for sellers, up for buyers) by the amount of these unobserved transaction costs. The mean shifter $\bar{\tau}_j^m$ creates distance between marginal valuations and equilibrium prices for unobserved costs the same way that the term $\rho^m \mathbf{B}_{jk}$ does for observed costs, and it captures transaction costs uncorrelated with the observable predictors.[11] Still, unobserved variable transaction costs are identified only relative to the least costly transaction, so any large variable costs incurred across all transactions will lead to an underestimate of the dispersion in marginal valuations and the potential gains from trade.

5 Estimation

The econometric model I take to data is a version of Equation 6 that combines a few terms. For transaction i made by district j with counterparty k in year t, log price p_{ijkt} is explained by observable predictors of transaction costs $\mathbf{B}_{jk}$, log quantity consumed $\widetilde{\ln Q}_{jt}$, a district-specific intercept $\delta_j \equiv \bar{v}_j + \bar{\tau}_j$, district-pair-specific transaction costs $\tilde{\tau}_{jk}$, and an error term $\varepsilon_{ijkt} \equiv v_{jt} + \epsilon_{ijkt}$:

$$\ln p_{ijkt} = \rho \mathbf{B}_{jk} + \eta_j \widetilde{\ln Q}_{jt} + \delta_j + \tilde{\tau}_{jk} + \varepsilon_{ijkt}. \tag{8}$$

To estimate Equation 8, I propose a four-step procedure that involves sequentially estimating parameters and moving their terms to the left-hand side. This multi-stage procedure allows me to isolate different sources of identifying variation best tailored to estimate each parameter. After estimation, I use Equation 7 to separate δ_j into $\bar{\tau}_j$ and $\bar{v}_j$, giving me all the terms I need to recover transaction costs (Equation 2) and demand (Equation 1).

---

[11] The international trade literature (e.g., Helpman et al. 2008) often assumes that unobserved costs are mean-zero distributed around the observable predictors, but here this assumption would likely miss important unobserved costs.

5.1 Step 1: Variable Transaction Costs from Observable Sources

I estimate observable variable transaction costs ρ incurred by both sellers and buyers by regressing observed prices on a vector of transaction cost predictors $\mathbf{B}_{jk}$:

$$\ln p_{ijkt} = \rho \mathbf{B}_{jk} + \alpha_{jt} + \epsilon^{1}_{ijkt}. \tag{9}$$

District-by-year fixed effects α_{jt} absorb the other terms from Equation 8, ensuring that coefficients ρ measure price gaps across transactions within district and year. For example, suppose a seller completes transactions with two different buyers in the same year, of which one requires the seller to undergo a particular regulatory review and the other does not. The seller is indifferent between buyers at the margin, so the price gap reveals the equilibrium price premium that compensates the seller for the regulatory review. If the review is also costly to the buyer, then this transaction will offer the buyer an equilibrium price discount relative to the buyer's other sellers.

Because the sample of transactions is not large, I avoid overfitting the model by using the least absolute shrinkage and selection operator (lasso) to perform variable selection (Belloni et al. 2014). Many predictors are *a priori* plausible, yet not all may be empirically important; others will be set to zero. I cluster standard errors by district to account for correlation over time, and also by transaction to adjust for the fact that each transaction appears multiple times in the dataset, once for each party. In some specifications, I use coarser fixed effects to explore robustness.

5.2 Step 2: Market Entry

I estimate the extensive-margin decision of whether to enter the water market in a given year using a probit model. I regress a binary indicator for whether district j completes any transactions in year t on a set of district-level explanatory variables $\mathbf{Z}_{jt}$:

$$\text{Enter}_{jt} = 1\left(\Pi \mathbf{Z}_{jt} + \epsilon^{2}_{jt} > 0\right). \tag{10}$$

I estimate Equation 10 separately in each year of data (to correct for selection in Step 3) and as a single pooled regression in the full dataset (to understand the overall determinants of market entry). Explanatory variables include observable predictors of fixed transaction costs f_{jt}, which consist of prior market experience and time-invariant district characteristics. They also include allocation percentages as exogenous determinants of potential surplus S_{jt}, using the full set of candidate instruments from Step 3. Negative coefficients reveal sources of either fixed transaction costs, since entry is decreasing in fixed cost f_{jt} (Equation 4), or of lower potential surplus.[12] I use the lasso to select from the full set of predictors, and I cluster standard errors by district.

5.3 Step 3: Demand Elasticity

I use instrumental variables to estimate the inverse price elasticity of demand η, the sensitivity of districts' marginal valuations to changes in quantity consumed (not traded). First, I adjust observed prices for the variable transaction costs estimated in Step 1. Then, I regress adjusted log prices on log quantity consumed:

$$\ln p_{ijkt} - \hat{\rho}\mathbf{B}_{jk} = \eta \widetilde{\ln Q_{jt}} + \alpha_j + \hat{\lambda}_{jt} + \theta t + \epsilon^3_{ijkt}. \tag{11}$$

District fixed effects α_j ensure that the elasticity is identified from within-district variation over time.[13]

Instruments. To overcome the joint determination of prices and quantities, I instrument for quantities using exogenous changes in annual endowments: allocation percentages in the federal and state water projects. Allocation percentages are determined by precipitation in the mountains during the previous winter, making their year-to-year variation exogenous to demand.[14] They are

---

[12] Directly quantifying fixed costs requires estimates of potential surplus, so I leave this task for Section 8.5. But the extensive margin allows me to model the *consequences* of fixed costs. I simulate a market with fixed costs by modeling trade among districts predicted to enter the market, and I simulate a market without fixed costs by setting the probability of market entry to 1 for all districts.

[13] In principle, the elasticity may be heterogeneous across districts; in practice, for precision and to avoid weak instrument problems, I estimate a single average elasticity for the market.

[14] Appendix B.2 gives a more extensive discussion of the identification assumptions.

relevant instruments because transaction costs create inertia in endowments. Appendix Figure A1 plots allocation percentages aggregated to several regional categories.

My instruments are log allocation percentages for the district itself, all other districts in the same hydrologic region, and all other districts in the state, and the full set of interactions between each contract type's log allocation percentage and hydrologic region indicator variables. These interactions allow allocations for each contract type (which are largely divided by region and sector) to have different effects on every other region of the state. They also yield a large number of candidate instruments.[15] To avoid overfitting and weak instruments, I estimate the model via post-lasso two-stage least-squares, following the IV lasso algorithm of Chernozhukov et al. (2015) as implemented in Stata by Ahrens et al. (2018).

Sample selection correction. Because prices are only observed in years for which a district enters the market, the sample may be selected in ways that bias the elasticity estimate. For example, if a district only enters the market when its endowment is very low, I might only observe prices along a relatively elastic part of its demand curve, resulting in an overestimate of the average price elasticity. To correct for sample selection, I include inverse Mills ratios $\hat{\lambda}_{jt}$ calculated from year-specific estimates of Equation 10, following Wooldridge (2010), sections 19.6.2 and 19.9.2. The primary source of identification for the selection correction is the prior market experience variables, which I assume affect fixed costs but not demand or variable costs, so that they can be excluded from Equation 11.

Time effects. A time trend θt adjusts for changes in demand over time. I omit year fixed effects to avoid spillover bias. Since transaction prices in the same market in the same year are simultaneously determined, within-year comparisons give rise to mechanical spillovers that violate the stable unit treatment value assumption (SUTVA) and eliminate most of the useful variation. Appendix C.1 offers a proof of the potentially severe bias from year effects. Although I cannot control for

---

[15] There are 13 contract types and 10 hydrologic regions, plus 3 overall instruments, for a total of 133 potential instruments. All instruments are log-transformed to match the variation of the endogenous variable, log quantities.

unobserved shocks to demand, the most important time-varying factors relate to water availability, which is flexibly captured by the instruments.

Inference. To account for estimation error in both the outcome variable and the inverse Mills ratios, I calculate standard errors by bootstrapping Steps 1-3 together. I use the Bayesian bootstrap with 1000 iterations and $\exp(1)$ weights blocked by district to account for correlation over time.

5.4 Step 4: District Intercepts and Unobserved Transaction Costs

I estimate district intercepts δ_j and pair-specific unobserved variable transaction costs $\tilde{\tau}_{jk}$ using random effects. First, I adjust observed prices for the variable transaction costs estimated in Step 1 and the demand elasticity estimated in Step 3. Then, I regress adjusted log prices on district-level characteristics (the time-invariant predictors from Step 2) and indicators for both districts and district/counterparty pairs:

$$\ln p_{ijkt} - \hat{\rho}\mathbf{B}_{jk} - \hat{\eta}_j \widetilde{\ln Q_{jt}} = \Gamma\mathbf{X}_j + \mu_j + \tilde{\tau}_{jk} + \varepsilon_{ijkt}. \tag{12}$$

I estimate this three-level random-intercept model via maximum likelihood, with district-pair effects nested within districts. The model only directly estimates the variance of district effects μ_j and unobserved variable transaction costs $\tilde{\tau}_{jk}$, but it then can be used to generate best linear unbiased predictions of both. Random effects are preferred to fixed effects in this step because they are no longer nuisance parameters – I need to use their estimates in simulations – and fixed effect estimates would be excessively noisy. I recover district intercepts by summing fitted values of covariates and predictions of district effects: $\hat{\delta}_j = \hat{\Gamma}\mathbf{X}_j + \hat{\mu}_j$. I again calculate standard errors by bootstrapping Steps 1-4 together, following the process described in Step 3.

6 Data

6.1 Water transactions

No government agency maintains a centralized listing of water transactions in California. Instead, I use a proprietary dataset on water transactions between 1990 and 2015 compiled by WestWater Research, LLC. To my knowledge, this is the most complete dataset of water transactions in California, and it has been used in federal regulatory analyses. It includes many of the same transactions as other datasets previously used in research, but it has a more complete record of prices.[16]

I focus my analysis on spot-market transactions (i.e., within-year leases) as opposed to permanent transfers of rights, because their prices are much more easily interpretable. They can be linked to demand without strong assumptions over discount rates and risk premia. In addition, volumes in leases are fixed so per-unit prices are known with certainty; volumes in permanent transfers are stochastic so per-unit prices are uncertain (Hanak 2002).

The original dataset includes 6,264 transactions, but many are unrelated to the surface water market. I exclude transactions (a) of rights to pump groundwater within adjudicated basins, (b) of rights to store water in reservoirs or groundwater banks, (c) within programs that set prices administratively, and (d) longer than one year. Of the remaining transactions, price is observed for all but 28. I carry forward 705 spot-market transactions of surface water originating in appropriative or riparian rights, SWP or CVP contracts, or reservoir storage, in which the price is freely negotiated. Appendix D describes the data processing in full detail.

6.2 Water quantities and allocation percentages

For water quantity consumed, I assemble a complete accounting of wholesale surface water deliveries and diversions in California, by district, sector, and year, from 1993 through 2015.[17] De-

---

[16] One dataset was assembled by Gary Libecap at UC Santa Barbara (e.g., Brewer et al. 2008), but it lacks many prices and ends in 2008. Another is maintained by Ellen Hanak at the Public Policy Institute of California, but it is not publicly available, and it appears not to focus on prices (Hanak and Stryjewski 2012).

[17] Data is available at github.com/hagertynw/data-surface-water and updated through at least 2021.

liveries from the SWP, CVP, and Colorado River come from archives of the DWR and USBR. Diversions on the basis of appropriative and riparian rights come from diversion reports collected by the SWRCB. The data also includes maximum contract amounts.

Allocation percentages for CVP and SWP contractors also come from DWR and USBR archives; I set them to 100 percent for appropriative, riparian, and Lower Colorado rights. To link districts across datasets, I build a crosswalk file that accounts for variations and errors in names as well as mergers and name changes across time. This file has 28,764 entries (input names) pointing to 14,830 targets (output names). For water district locations, boundaries, and areas, I combine several publicly available shapefiles into a single geospatial dataset. Details of sources, cleaning, and processing of these datasets are described in Hagerty (2022).

6.3 Hydrologic network model

I construct a basic model of California's hydrological network to calculate (a) characteristics of transaction delivery pathways, and (b) purely physical conveyance costs for simulations. The model consists of a set of nodes and edges corresponding to all major water conveyance channels in California: rivers, canals, aqueducts, and pipelines. Channels come from the National Hydrography Dataset of the U.S. Geological Survey. Each node and edge is parameterized with physical transportation costs drawn from the literature: pumping costs (for the energy required to lift water), conveyance losses (to percolation and evaporation), and carriage losses required in the Sacramento–San Joaquin Delta. I run a graph-theory algorithm to obtain the least-cost delivery pathway for each unique pair of planning areas (55 geographic regions defined by the DWR). I assign districts to planning areas based on geography. As compared with more detailed engineering models such as CALVIN (e.g., Howitt et al. 1999), this model lacks comprehensive information on capacity constraints.

6.4 Observable predictors of transaction costs

I identify variables likely to reflect aspects of the transaction costs described in Section 3. Predictors of variable costs must vary across a district's counterparties, since these costs are identified by within-district variation. Predictors of fixed costs must be district-specific characteristics, since these costs are identified by between-district variation.

Variable costs. Delivery and some administrative costs are captured by physical characteristics of delivery routes. Using the hydrological model, I calculate five variables: elevation gain (pump lift); distance conveyed in rivers; distance conveyed in canals, aqueducts, and pipelines, distance of virtual movement (i.e., transfers against the direction of flow), and whether the transfer crosses the Sacramento–San Joaquin Delta. Elevation gain requires energy use, raising wheeling charges. Distance variables capture conveyance losses (rivers, canals), wheeling charges (canals), and administrative costs that rise with geographic separation (Regnacq et al., 2016). Crossing the Delta introduces carriage losses and delivery risks.

Regulatory costs are captured by institutional characteristics: whether water is imported into or exported from a federal or state water project, whether the transaction changes the place of use of a post-1914 appropriative water right, and whether the counterparty uses the majority of its water in agriculture. These reflect additional reviews by DWR, USBR, and the SWRCB and added complexity for agriculture-related transfers. They may also capture some administrative and internal decision costs.

Fixed costs. Administrative and some internal decision costs are captured by trading experience: whether the district completed trades in each of the past three years and the cumulative number of prior years with trades. Familiarity with the market is likely to make future transfers easier.

Regulatory costs are captured by institutional characteristics: whether the district holds CVP contracts, SWP contracts, Lower Colorado entitlements, or surface water rights, and whether it uses the majority of its water in agriculture. These capture differences in review requirements.

Internal decision costs are also captured by other district characteristics. To measure governance and voting structure, I construct indicators for California water districts, irrigation districts, mutual water companies, and other public agencies, with individual farmers and private companies as the omitted category. To measure stakeholder number and heterogeneity, I use log cropland area,[18] log maximum volume of water rights or contracts, and sectoral concentration, measured by a Herfindahl index (HHI) of shares of water across agriculture and municipal/industrial use.

6.5 Merged dataset and summary statistics

I construct a balanced panel of surface water quantity consumed for all 2,380 districts in California that consume at least 100 acre-feet per year.[19] To this dataset I merge the list of transactions observed for each district, such that each transaction is repeated for each buyer or seller involved. I keep transactions even when they do not successfully match the water quantity data, allowing me to use all available data at each step of the estimation: Steps 1-2 use data from 1990-2015, including districts without quantity data, while steps 3-4 use data from 1993-2015 for districts with quantity data.[20] The final dataset has 1,259 transaction-by-district observations, plus 59,165 records of water quantity and market nonparticipation for years in which districts do not trade and for districts that never trade. Districts covered by the transactions data represent most (63%) of the surface water in California.[21]

Table 1 shows summary statistics. Panel A shows that the mean transaction price is $236 per acre-foot (in 2010 dollars) and the distribution of volumes is highly skewed. Panel B shows that

---

[18] Calculated from district shapefiles and USDA's Cropland Data Layer, imputing 0 when a shapefile is unavailable.

[19] As measured by maximum water rights or average quantity consumed, whichever is greater. Water rights smaller than 100 AF/year constitute 67 percent of the count of observations but only 0.2 of total water consumption; they are generally held by individual farmers and rural households not served by any water or irrigation district.

[20] Unmatched districts represent 14% of the transaction-district observations and generally fall into two categories: associations that occasionally trade on behalf of multiple districts, and small districts that hold affiliations or agreements with larger districts instead of their own water rights or contracts. I drop a small number of districts not identified with enough specificity to determine the transaction cost variables. I also drop those that purchase water for instream flows and other environmental purposes, rather than for consumptive use in the agricultural and municipal sectors.

[21] The dataset covers few districts by count – only 7% of districts in the merged dataset ever trade – but most non-trading districts are very small. Because more gains from trade come from larger districts, weighting by surface water consumption is more informative. Appendix Figure A3(a) plots districts by water consumption and frequency of appearance in the transactions data.

most sales come from agriculture, but purchases are shared across sectors. Panel C shows that the most exports are from the Sacramento River hydrologic region and the most imports are from the South Coast and Tulare Lake regions. Figure 1(a) shows that the distribution of prices is highly dispersed even after adjusting for year.

7 Results

7.1 Step 1: Variable Transaction Costs from Observable Sources

Table 2, reports estimates of Equation 9. Positive coefficients reflect positive transaction costs on both sides of the market – price premiums for sellers and price discounts for buyers – because regressors for buyers are multiplied by -1, following the substitutions in Equation 6.

Columns (4)-(6) show the variable transaction costs associated with predictors selected by the lasso. Column (6) is my preferred specification, with seller-by-year and buyer-by-year fixed effects. Several predictors are indeed costly. For example, buyers pay approximately 45 percent less for transactions that involve exporting water from a federal or state water project, and sellers accept approximately 31 percent more for transactions that must cross the Sacramento–San Joaquin Delta.[22] Other sources of variable costs are distance conveyed in rivers and having an agricultural buyer (for sellers), and review by the State Water Board (for buyers).

Not all coefficients are statistically distinguishable from zero, but the fact they were selected by the lasso suggests they matter for the overall model. The importance of some of these predictors is also supported by the specifications with coarser fixed effects in columns (4)-(5), in which the Delta-crossing coefficient is larger and more precise. Across specifications, all lasso-selected coefficients are positive, supporting the assumptions of the model and the interpretation of these estimates as variable transaction costs.

Results without variable selection are given in columns (1)-(3). When including all predictors,

---

[22] The finding of transaction costs associated with the Delta is particularly reassuring given that it is consistent with the previous results of Regnacq et al. (2016), which studies trade frictions in the same setting using a model of trade flows rather than prices.

many coefficients are not statistically distinguishable from zero, and some are negative.[23] Still, the overall patterns are broadly consistent across specifications: Buyers require price discounts, and sellers require price premiums, to complete transactions that are subject to these cost sources.

7.2 Step 2: Market Entry

Table 3, column (1) reports estimates of Equation 10. The model is highly predictive of market entry: the pseudo-R^2 shows that just a handful of regressors explain 59 percent of the variation.

Results provide evidence that fixed costs are important barriers to trade. First, administrative costs matter: recent market experience makes districts more likely to trade, suggesting it reduces fixed costs (and provides independent variation for the sample selection correction). Second, regulatory costs matter: districts in the CVP and SWP are more likely to trade, likely because transactions within the projects are often streamlined and exempt from SWRCB review.

Evidence for internal decision costs is weaker. Districts with more stakeholders (proxied by water right volume) and those governed by broader constituencies (irrigation districts) are more likely to trade. Other measures of governance and stakeholder heterogeneity are insignificant or excluded by the lasso. These results appear contrary to predictions, but positive coefficients may reflect higher potential surplus in addition to lower fixed costs, so larger districts may simply have greater gains from trade.

Results also show that potential surplus affects market entry. The coefficient on statewide allocations shows that districts are more likely to trade in dry years, when the gains are large enough to outweigh the fixed costs.

7.3 Step 3: Demand Elasticity

The rest of Table 3 reports estimates of Equation 11. Columns (2) and (3) show first-stage and reduced-form regressions using only two instruments: the district's own allocation percentage and

---

[23] Negative coefficients may simply represent reduced costs relative to another predictor. For example, most transactions that involve importing into a project are also subject to State Water Board review, and the sum of those two coefficients is always positive.

the allocation percentage for all other districts. The instruments are strong, with an F-statistic above conventional thresholds. Coefficients have the expected signs: higher allocations increase quantity consumed and reduce adjusted prices (marginal valuations). Allocations to other districts matter more than own allocations, consistent with the theoretical model: once a district pays the fixed cost to enter the market, quantities and prices are set in statewide equilibrium.

Column (4) combines these relationships using IV lasso, which selects 3 instruments. The inverse price elasticity of demand is estimated at -1.4, implying that a 10% increase in quantity consumed lowers marginal valuations by 14%. The first-stage F-statistic is even higher than in column (2) because the instruments are now selected optimally. Table A2 shows the estimate is robust to more flexible time trends, the set of candidate instruments, dropping the sample selection correction, and using unadjusted prices.

The reciprocal of this estimate implies that districts' price elasticity of demand is -0.69, conditional on trading. This is the first estimate of the price elasticity of demand for water districts in California's surface water market. Previous studies of retail consumers generally find more inelastic demand.[24] This comparison is inexact but suggests that districts are more price-sensitive than their retail customers.

7.4 Step 4: District Intercepts and Unobserved Transaction Costs

Table A3 reports estimates of Equation 12. Many district-level covariates are statistically significant, and district and district-pair effects explain a large share of the residual variation. I use this estimated model to predict district intercepts δ_j and unobserved variable transaction costs $\tilde{\tau}_{jk}$. I then recover the mean of unobserved transaction costs $\bar{\tau}_j$ from $\hat{\tilde{\tau}}_{jk}$ and observed transaction costs $\hat{\rho}\mathbf{B}_{jk}$ (Equation 7), the mean district marginal valuations $\bar{v}_j$ from $\hat{\delta}_j$ and $\hat{\bar{\tau}}_j$ (Equation 8), and inverse demand from $\hat{\bar{v}}_j$ and the price elasticity $\hat{\eta}$ (Equation 1). I use this demand model to calculate

---

[24] For municipal water demand, my estimate is larger than Buck et al. (2016) in California (-0.14) and the US mean in Dalhuisen et al. (2003) (-0.41) but similar to Baerenklau et al. (2014) in Riverside County (-0.76). For agriculture, my estimate exceeds groundwater elasticities for California (Bruno and Sexton 2020, -0.19) and Kansas (Hendricks and Peterson 2012, -0.10), and the US mean in Scheierling et al. (2006) (-0.48). In these studies, even when quantity data is at the district level, the price under question is the retail price.

fitted values of marginal valuations for each district in each year.

Figure 1 plots the resulting distributions of estimated marginal valuations and variable transaction costs in dollar terms, with raw prices for comparison. There is still considerable dispersion in marginal valuations, suggesting that gains are available from reducing transaction costs and increasing trade. Median variable transaction costs are 27 percent on the seller side and 21 percent on the buyer side, for a total of 54 percent of the transaction price.[25] These costs are somewhat higher for sellers, suggesting substantial internal decision or regulatory costs. But buyers still face significant costs, consistent with shared administrative costs or high delivery costs.

8 Simulations

Using the demand model, I simulate counterfactual water market scenarios without transaction costs and calculate the resulting gains from trade. To make computation feasible, I focus on the 154 largest water districts, which together consume 85% of the water in California.[26] Most of these districts contribute to demand estimation: 53% appear in the transactions data, or 72% when weighted by water use.[27]

Still, the samples used for estimation and simulation are not the same, and district-specific demand intercepts $\bar{v}_j$ are not identified for districts never observed trading. To extrapolate the demand model, I project the estimated intercepts $\hat{\bar{v}}_j$ onto the district covariates from Step 4.[28] This step replaces the unrestricted heterogeneity in $\bar{v}_j$ with a parametric model based on observable district characteristics. I then generate fitted values and use them as the demand intercepts in simulations for all districts – including those observed trading, so that trading and non-trading

---

[25] Values do not sum due to log transformations. Mean values are different between sellers and buyers in a t-test.

[26] I select these districts using a cutoff in mean quantity consumed of 25,000 AF/year. Because surplus is approximately proportional to the square of trading quantity, these districts are likely responsible for more than 85 percent of the potential gains from trade.

[27] Districts that appear in transactions data do so in 30% of years on average. In an average year, 12% of all simulation districts appear in the transactions data. Appendix Figure A3(b) plots market participation rates by year. Appendix Table A4 shows that districts never observed trading are smaller (fewer water rights) and less likely to be in the CVP, but broadly similar in other ways.

[28] Covariates also include hydrologic region indicators. Regression results are shown in Appendix Table A5.

districts are treated consistently.

I run simulations under three water-supply scenarios: a wet year, a median year, and a dry year. These scenarios draw observed quantities and trading volumes from the years 2006, 2010, and 2014, respectively.[29] Figure 2 maps fitted marginal valuations in the median-year scenario for each district.

8.1 Surplus from the existing market

I first calculate the economic surplus achieved to date by observed transactions. Figure 3, Panel (a) maps the geographic patterns of these transactions. Sellers tend to be in the Sacramento Valley and northern San Joaquin Valley; buyers tend to be in urban coastal areas and the southern San Joaquin Valley. Table 4, Panel A, reports that the total volume of water traded ranges between 35,000 and 289,000 acre-feet per year depending on water conditions.

Surplus for each district S_j (suppressing time subscripts) is the difference in marginal valuations before and after trading, integrated over the quantity traded q_j (i.e., net purchases) and evaluated using the estimated demand model (Equation 1):

$$S(q_j, Q_j) = \int_{Q_j - q_j}^{Q_j} \Big[V_j(\varphi) - V_j(Q_j) \Big] d\varphi. \tag{13}$$

Table 4, Panel A reports that observed transactions result in economic surplus of \$8.3 million per year in a median year – a figure that is small relative to annual water-related expenditures in California.[30] In a dry year, total gains are larger but still relatively small, at \$55.2 million per year.

---

[29] Summing water quantities across the state, 2006 is the wettest year in my data, 2014 the least, and 2010 is the median year over the 1998-2015 period.

[30] Hanak et al. (2014) estimate that federal, state, and local agencies spent \$16.9 billion on water supply, and \$30.5 billion on all water-related spending (including pollution control, flood management, ecosystem management, and debt service) per year between 2008 and 2011.

8.2 Counterfactual simulations

To simulate an efficient market, I solve the social planner's problem in a constrained optimization problem. Because an ideal market could implement the efficient allocation, the increase in surplus represents the potential gains from trade. Because even an ideal market cannot avoid the physical costs of water conveyance, I include pair- and direction-specific conveyance costs from my hydrological network model as costs in the objective function.[31] I do not model constraints on flow capacity except in two specific scenarios.

The social planner chooses the vector of bilateral transaction quantities q_{jk} (net volumes delivered from district j to district k, for all unique pairs $k > j$) that maximizes the aggregate value of water net of physical transaction costs, subject to the resource constraint. Aggregate value is the area under each district's demand curve between observed quantities Q_j and the simulated final quantity Q_j^f, summed over districts. Conveyance costs are per-unit and directionally asymmetric. The resource constraint requires that final quantities be nonnegative. Together, the full optimization problem is:

$$\max_{\{q_{jk}\}_{j,k>j}} \sum_j \int_{Q_j}^{Q_j^f} V_j(\varphi)d\varphi - \sum_j \sum_{k>j} \Big[1(q_{jk} > 0)c_{jk} - 1(q_{jk} < 0)c_{kj} \Big] q_{jk} \tag{14}$$

subject to

$$\text{(definition of final quantities)} \qquad Q_j^f = Q_j - \sum_{k>j} q_{jk} + \sum_{k<j} q_{kj} \ \forall j$$

$$\text{(resource constraint)} \qquad Q_j^f \geq 0 \ \forall j.$$

I solve this problem using the `patternsearch` solver in Matlab, using observed quantities as the initial conditions. Appendix C.2 proves that the solution to this planner's problem also satisfies the conditions of an efficient market. Efficient trading creates two types of benefits: reduced costs of transactions that already take place, and gains from new transactions that would not otherwise

---

[31] Appendix Figure A2 illustrates these costs. I use this model instead of empirically estimating conveyance costs because they are correlated with other sources of transaction costs and therefore difficult to isolate.

take place. I calculate both types of benefits and sum across districts.[32]

8.3 Gains from an efficient market

Table 4, Panel B reports simulation results. Quantities and gains are reported relative to observed trade. Scenario 2 represents my main scenario: an efficient market that accounts for physical costs of conveyance but eliminates other sources of transaction costs.

Surface water trade rises dramatically, by 3 to 4 million acre-feet per year. Figure 3 maps geographic patterns of trade. In a median year (Panel (c)), considerable amounts of water are sold by the Sacramento Valley, northern San Joaquin Valley, and agricultural regions of Southern California, and purchased by urban coastal areas and the southern San Joaquin Valley. These patterns intensify in a dry year (Panel (d)). The most purchases are by the Los Angeles area, which is consistent with previous findings of costly water shortages in urban Southern California (Buck et al. 2016).

Gains from these new transactions, as valued by the water districts, range from \$211 million in a median year to \$559 million in a dry year. In addition, districts avoid \$5 to \$85 million of transaction costs in pre-existing transactions. Larger gains in dry years suggest that water markets are most valuable in California when water is scarcest, and that short-term leases offer valuable flexibility for managing risk.

8.4 Extensions and sensitivity checks

With fixed transaction costs. Scenario 3 simulates a market that keeps the fixed costs of market entry but continues to eliminate variable costs other than those of physical conveyance. To implement this scenario, I simulate trade among districts predicted to enter the market by the model

---

[32] Gains from new transactions are calculated for each district in the same way as Equation 13, evaluated between Q_j and Q_j^f. Appendix C.3 proves that the sum of district-specific gains across districts is equal to the objective function in Equation 14. Gains from reduced costs of observed transactions are $R_j \equiv \left(V_j(Q_j) - V_j(Q_j^f)\right)q_j$. I restrict R_j to be non-negative, since negative values occasionally arise if the simulation predicts net trading in the opposite direction of observed trading.

estimated in Step 2.[33] Results show that both variable and fixed costs are important. Eliminating only variable costs (while keeping fixed costs) produces 35-46% of the trade volume, and 36-70% of the gains, as eliminating both variable and fixed costs.[34] In other words, eliminating fixed costs produces \$81 to \$166 million in gains, conditional on having eliminated variable costs. Results also suggest that much of the potential gains from trade can be attained by relatively few market participants and relatively little new transaction volume.

Without conveyance costs. I also investigate the consequences of purely physical conveyance costs. How much more trade would occur if water could be traded costlessly? Scenario 4 simulates an efficient market with all transaction costs c_{kl} set to zero, implausibly. More trading occurs, especially in non-dry years. Relative to Scenario 2, gains are about double in a dry year, at \$1.1 billion, and triple in median and wet years. These results suggest that much of the observed dispersion in water valuations in different parts of California is unavoidable and due to the fact that water is costly to move.

Environmental and infrastructure constraints. Because my main scenario does not account for hydrological capacity constraints, it is natural to ask how important they might be. I focus on the most important juncture in California's hydrology, the Sacramento–San Joaquin Delta. Increasing flows through the Delta may be infeasible due to constraints to protect water quality and endangered species. How much would maintaining these constraints affect the gains from trade? There is reason to expect the answer to be large, as Step 1 showed the Delta is associated with large variable transaction costs, and the new trade volume moving across the Delta in Scenario 2 exceeds the capacity of the state's proposed Delta Conveyance Project.[35]

To find out, Scenario 5 imposes a hard constraint at the Delta. Specifically, I simulate separate markets north and south of the Delta and then combine results. Trading volumes and economic

---

[33] Only 12 districts enter in the wet year, 28 in the median year, and 43 in the dry year.

[34] This does not imply that fixed costs are more important than variable costs, since a scenario that eliminated fixed but not variable costs might also achieve more than half the available gains.

[35] Estimates suggest this project would enable 0.1 to 1.0 million acre-feet per year to be transferred across the Delta (DWR, 2025), as compared with 0.9 to 1.8 million in Scenario 2.

gains are slightly lower than in Scenario 2, but not by much. The lost trade across the Delta is replaced by an near-equal volume of trade within regions (i.e., north and south of the Delta). The reason the gains are relatively unaffected is that marginal valuations have greater variation within region than between them. Most gains from trade come from eliminating this variation within each region rather than joining the two regions.[36] This result suggests that most of the gains from a more efficient surface water market could be achieved without relaxing environmental regulations or building major new infrastructure.

I also consider capacity constraints at the Colorado River Aqueduct, the only link between the southeastern corner of California and the rest of the state. This aqueduct operates near capacity in some years, so large quantities of new sales from Colorado River districts to urban Southern California may not be feasible. To exclude this possibility, Scenario 6 simulates three separate markets, splitting the south-of-Delta market into two at the Colorado River Aqueduct. Again trading volumes and gains are very similar.

Among only observed traders. How much of the gains from trade rely on extrapolating the demand model out of sample? Scenario 7 simulates trade among only the 82 districts observed in the transactions data. Most gains are still achieved: 63 to 83% of those in Scenario 2. This is because most gains from trade come from large districts, which are more often observed to trade.

Using estimated demand intercepts instead of fitted values. How sensitive are results to applying the extrapolation model universally, rather than only to districts lacking estimates? Appendix Table A6 repeats all simulations using demand intercepts taken directly from the estimated demand model for prior traders. Districts not observed in the transactions data still use fitted values from the extrapolation model. Gains are somewhat larger in most scenarios – for example, 23 to 118% greater in the base scenario – but still within the same order of magnitude. Relative patterns are similar across scenarios.

---

[36] Appendix Figure A4 provides a visualization.

Functional form. How sensitive are results to the functional form of demand? Scenario 8 optimizes trade and evaluates gains using linear demand instead of isoelastic. Gains are lower – 64 to 76% of Scenario 2 – but otherwise broadly similar. This suggests that the qualitative implications of my results are not driven by the choice of functional form.

8.5 Quantifying fixed costs

As a final exercise, I quantify fixed costs using bounds implied by Equation 4. For tractability, I now assume fixed costs are constant in time. For each district observed in the transactions data, the fixed cost must be no greater than surplus attained ($f_j \leq S_{jt}$ for all t), so the minimum nonzero surplus observed in any year estimates an upper bound. For each district never observed trading, the fixed cost must be greater than surplus forgone ($f_j > S_{jt}$ for all t), so the maximum surplus attained in any year of the efficient-trading simulation (Scenario 2) estimates a lower bound. I then assume these bounds hold with equality.[37]

Appendix Figure A5 plots the distribution of estimated fixed costs across districts.[38] As expected, fixed costs are greater for districts that never trade, but the two methods produce overlapping densities, so it is reasonable to believe they estimate different parts of the same underlying distribution. Aggregate fixed costs are $221 million per year, suggesting that between 40 and 100 percent of simulated gains are attributable to fixed costs, depending on the year. This result is broadly consistent with the indirect method of assessing fixed costs in Scenario 3, confirming that both fixed and variable costs are important in preventing gains from trade.

---

[37] This assumption is plausible for observed traders, few of whom trade every year, so if fixed costs were lower, they would trade more often. It is less plausible for never-traders, but a countervailing factor is that the efficient-trading simulation overstates forgone surplus, so perhaps on net it is a reasonable approximation. (The ideal measure of forgone surplus would hold constant the set of observed traders and simulate gains from adding one new district at a time, but running that many simulations would be computationally prohibitive.)

[38] The median is $55,600 and the mean is $973,000.

9 Discussion and Conclusion

My main results estimate that the gains available from reducing transaction costs in California's water market, including both direct cost savings and the gains from trade, are \$225 to \$644 million per year, depending on water supply conditions. These estimates can be interpreted as the potential benefits of legal and policy reforms that improve the allocative efficiency of water use across water districts.

Are the potential gains large or small? I suggest they are moderate. On one hand, annual trading volumes would increase by two to three orders of magnitude, and the dollar value of benefits is considerable, so if the costs of achieving them are low, they are worth pursuing. On the other hand, only 10 to 15 percent of California's water consumption would change hands, and the gains from trade are no more than 2 percent of California's agricultural GDP or 4 percent of the state's water supply expenditures (Hanak et al. 2014).

Why are the potential gains smaller than anecdotal evidence suggests? Several reasons may contribute. First, water is costly to move: Simulations without conveyance costs generate 2 to 3 times the gains, suggesting that much of the differences in water valuations across California are the inevitable result of physics. Second, some transaction costs may remain unmeasured, especially internal decision costs. Water districts may value water differently than individual farmers or households, leading them to trade less than their constituents would (Libecap 2009; Ayres and Bigelow 2022). Third, gains may be mostly local, available within districts rather than between them. They could be achieved through district-level reforms to allocation rules, pricing, or strengthening within-district water markets.[39] Fourth, analyzing the lease market may understate the potential gains. Urban districts and perennial crop growers place a high value on long-term supply certainty, so they may value one-year leases in the data below either (a) the annualized

---

[39] Local transactions among irrigators are not uncommon; brokers and investors actively facilitate these kinds of transfers. However, markets for surface water within agricultural water districts likely have significant room to expand. For example, districts that supply surface water could set up the kinds of formal, centralized exchanges that operate for groundwater pumping rights in several areas, such as the Fox Canyon Groundwater Management Agency (Heard et al. 2021) and the Mojave Water Agency (Ayres et al. 2021).

value of a permanent water right or (b) leases in a more robust, predictable market.[40] Future work could quantify these factors, but my approach remains relevant to realistic policy settings in which districts retain control of water rights.

Which types of transaction costs matter most? Evidence suggests that regulatory costs, administrative costs, and delivery costs are all important barriers to trade. Regulatory reviews are associated with large transaction costs, both variable and fixed. Administrative costs are supported by two findings. First, market inexperience strongly predicts fixed costs. Second, fixed costs remain high even for districts subject to fewer regulatory reviews – many districts never trade, and even active districts do not trade every year. Delivery costs appear in the large variable costs of crossing the San Francisco–San Joaquin Delta. However, delivery costs for short-term transfers can be largely attributed to conveyance costs without invoking environmental restrictions. In simulations, conveyance costs are substantial – removing them generates more gains than eliminating all other transaction and delivery costs combined. By contrast, adding pumping restrictions at the Delta – after including carriage losses and other conveyance costs – reduces gains only slightly.[41]

Finally, how can transaction costs be reduced? Not all can or should be eliminated, but many can be lowered. Administrative costs can be remedied with better information: The state could create a centralized online exchange, as in Australia, or at least collect and publish standardized data about water transfers, including prices (Gray et al., 2015). Internal decision costs could be reduced through legal reforms that allow individual farmers to transfer unused water outside their district (Culp et al., 2014). Delivery costs could be lowered by requiring predictable access to conveyance at marginal cost (Association of California Water Agencies, 2016). Reducing Delta pumping restrictions would require either loosening environmental protections or building new conveyance – both bringing costs that must be weighed against the potential gains.

---

[40] Evidence for this point, as suggested by a referee, is seen in the tendency of Southern California water districts in particular to invest in expensive desalination plants, stormwater and wastewater recycling plants, and water efficiency regulations on new construction. The levelized costs of water from these methods is typically in the range of $1,000 to 3,000 per acre-foot (Cooley et al. 2019), well above the mean marginal valuations I estimate for urban water districts.

[41] Risk from environmental pumping restrictions may still be costly in longer-term transactions. Transfers that cross the Delta have a higher lease-to-sale price ratio (0.48) than transfers within the south-of-Delta region (0.15). This comparison – based on Grainger and Costello (2014) and suggested by an anonymous referee – suggests that long-term risk lowers the value of a permanent right more than a lease.

Regulatory costs can be reduced by streamlining review processes. Proposed reforms tend to fall into three categories. First, the number of reviews could be reduced by combining processes across agencies, consolidating places of use within watersheds and projects, and pre-approving small or temporary transfers (Sellers et al., 2016). Second, review processes could be simplified and shortened by establishing an authoritative water accounting system and standardized consumptive use rates (Gray et al., 2015). Finally, risks and costs of objections could be reduced by replacing hearings with staff-level decisions, more precisely defining thresholds for injury, and standardizing methods for estimating third-party effects (Western Water Company, 2000; Hanak and Stryjewski, 2012).

References

Adamopoulos, Tasso, Loren Brandt, Jessica Leight, and Diego Restuccia, "Misallocation, Selection, and Productivity: A Quantitative Analysis With Panel Data From China," *Econometrica*, 2022, *90* (3), 1261–1282.

Ahrens, Achim, Christian B. Hansen, and Mark E Schaffer, "PDSLASSO: Stata module for post-selection and post-regularization OLS or IV estimation and inference," Statistical Software Components, Boston College Department of Economics February 2018.

Aigner, Dennis, C. A. Knox Lovell, and Peter Schmidt, "Formulation and estimation of stochastic frontier production function models," *Journal of Econometrics*, July 1977, *6* (1), 21–37.

Anderson, James E. and Eric van Wincoop, "Trade Costs," *Journal of Economic Literature*, September 2004, *42* (3), 691–751.

Archibald, Sandra O. and Mary E. Renwick, "Expected Transaction Costs and Incentives for Water Market Development," in K. William Easter, Mark W. Rosegrant, and Ariel Dinar, eds., *Markets for Water: Potential and Performance*, Boston, MA: Springer US, 1998, pp. 95–117.

Association of California Water Agencies, "Recommendations for Improving Water Transfers and Access to Water Markets in California," Technical Report, Association of California Water Agencies April 2016.

Atkin, David and Dave Donaldson, "Who's Getting Globalized? The Size and Implications of Intra-national Trade Costs," Working Paper 21439, National Bureau of Economic Research July 2015.

Ayres, Andrew and Daniel Bigelow, "Engaging Irrigation Districts in Water Markets," in "The Future of Water Markets: Obstacles and Opportunities" PERC Policy Reports September 2022, pp. 53–60.

Ayres, Andrew B., Eric C. Edwards, and Gary D. Libecap, "How transaction costs obstruct collective action: The case of California's groundwater," *Journal of Environmental Economics and Management*, 2018, *91*, 46–65.

_, Kyle C. Meng, and Andrew J. Plantinga, "Do Environmental Markets Improve on Open Access? Evidence from California Groundwater Rights," *Journal of Political Economy*, 2021, *129* (10), 2817–2860.

Baerenklau, Kenneth A, Kurt A Schwabe, and Ariel Dinar, "The Residential Water Demand Effect of Increasing Block Rate Water Budgets," *Land Economics*, 2014.

Barzel, Yoram, "Measurement Cost and the Organization of Markets," *The Journal of Law & Economics*, 1982, *25* (1), 27–48.

Batista, Càtia and Pedro C. Vicente, "Is Mobile Money Changing Rural Africa? Evidence from a Field Experiment," *The Review of Economics and Statistics*, 05 2023, pp. 1–29.

Belloni, Alexandre, Victor Chernozhukov, and Christian Hansen, "Inference on Treatment Effects after Selection among High-Dimensional Controls," *The Review of Economic Studies*, April 2014, *81* (2), 608–650.

Bennon, Michael, Daniel De La Hormaza, and R. Richard Geddes, "A hazard analysis of federal permitting under the national environmental policy act of 1970," *Journal of Regulatory Economics*, June 2024, *65* (1-3), 154–183.

Boehm, Johannes, "The Impact of Contract Enforcement Costs on Value Chains and Aggregate Productivity," *The Review of Economics and Statistics*, 01 2022, *104* (1), 34–50.

Bretsen, Stephen N. and Peter J. Hill, "Water Markets as a Tragedy of the Anticommons," *William & Mary Environmental Law and Policy Review*, 2009, *33* (3), 723–783.

Brewer, Jedidiah, Robert Glennon, Alan Ker, and Gary Libecap, "Transferring Water in the American West: 1987–2005," *University of Michigan Journal of Law Reform*, 2007, *40* (4), 1021–1053.

_, _, _, and _, "2006 Presidential Address: Water Markets in the West: Prices, Trading, and Contractual forms," *Economic Inquiry*, 2008, *46* (2), 91–112.

Brooks, Leah and Zachary Liscow, "Infrastructure Costs," *American Economic Journal: Applied Economics*, April 2023, *15* (2), 1–30.

Bruno, Ellen M. and Richard J. Sexton, "The Gains from Agricultural Groundwater Trade and the Potential for Market Power: Theory and Application," *American Journal of Agricultural Economics*, 2020, *102* (3), 884–910.

_, Nick Hagerty, and Arthur R. Wardle, "The Political Economy of Groundwater Management: Descriptive Evidence from California," in Gary D. Libecap and Ariel Dinar, eds., *American Agriculture, Water Resources, and Climate Change*, National Bureau of Economic Research: University of Chicago Press, 2024, pp. 343–365.

Bryan, Gharad and Melanie Morten, “The Aggregate Productivity Effects of Internal Migration: Evidence from Indonesia,” *Journal of Political Economy*, 2019, *127* (5), 2229–2268.

Buck, Steven, Maximilian Auffhammer, Stephen Hamilton, and David Sunding, “Measuring Welfare Losses from Urban Water Supply Disruptions,” *Journal of the Association of Environmental and Resource Economists*, 2016, *3* (3), 743–778.

Burke, Marshall, Solomon M. Hsiang, and Edward Miguel, “Climate and Conflict,” *Annual Reviews of Economics*, 2015, *7*, 577–617.

Carey, Janis, David L. Sunding, and David Zilberman, “Transaction costs and trading behavior in an immature water market,” *Environment and Development Economics*, 2002, *7* (04), 733–750.

Cason, Timothy N. and Lata Gangadharan, “Transaction Costs in Tradable Permit Market: An Experimental Study of Pollution Market Design,” *Journal of Regulatory Economics*, 2003, *23* (2), 145–165.

Chen, Chaoran, Diego Restuccia, and Raül Santaeulàlia-Llopis, “Land Misallocation and Productivity,” *American Economic Journal: Macroeconomics*, April 2023, *15* (2), 441–65.

Chernozhukov, Victor, Christian Hansen, and Martin Spindler, “Post-Selection and Post-Regularization Inference in Linear Models with Many Controls and Instruments,” *American Economic Review: Papers & Proceedings*, 2015, *105* (5), 486–490.

Chong, Howard and David Sunding, “Water Markets and Trading,” *Annual Review of Environment and Resources*, 2006, *31*, 239–264.

Coase, R. H., “The Nature of the Firm,” *Economica*, November 1937, *4* (16), 386–405.

Colby, Bonnie G., “Transactions Costs and Efficiency in Western Water Allocation,” *American Journal of Agricultural Economics*, 1990, *72* (5), 1184–1192.

Cooley, Heather, Rapichan Phurisamban, and Peter Gleick, “The cost of alternative urban water supply and efficiency options in California,” *Environmental Research Communications*, May 2019, *1* (4), 042001.

Culp, Peter W., Robert Glennon, and Gary Libecap, “Shopping for Water: How the Market Can Mitigate Water Shortages in the American West,” Discussion Paper, The Hamilton Project & Stanford Woods Institute for the Environment October 2014.

Dalhuisen, Jasper M., Raymond J. G. M. Florax, Henri L. F. de Groot, and Peter Nijkamp, “Price and Income Elasticities of Residential Water Demand: A Meta-Analysis,” *Land Economics*, 2003, *79* (2), 292–308.

Davis, Lucas W. and Lutz Kilian, “The Allocative Cost of Price Ceilings in the U.S. Residential Market for Natural Gas,” *Journal of Political Economy*, 2011, *119* (2), 212–241.

de Mel, Suresh, Craig McIntosh, Ketki Sheth, and Christopher Woodruff, "Can Mobile-Linked Bank Accounts Bolster Savings? Evidence from a Randomized Controlled Trial in Sri Lanka," *The Review of Economics and Statistics*, 03 2022, *104* (2), 306–320.

Demsetz, Harold, "Toward a Theory of Property Rights," *The American Economic Review*, 1967, *57* (2), 347–359.

_ , "The Cost of Transacting," *The Quarterly Journal of Economics*, 1968, *82* (1), 33–53.

Donaldson, Dave, "Railroads of the Raj: Estimating the Impact of Transportation Infrastructure," *American Economic Review*, April 2018, *108* (4-5), 899–934.

DWR, "The Delta Conveyance Project in Practice," Technical Report, California Department of Water Resources 2025.

_ **and SWRCB**, "Background and Recent History of Water Transfers in California," Technical Report, California Department of Water Resources & State Water Resources Control Board July 2015.

_ **and USBR**, "Draft Technical Information for Preparing Water Transfer Proposals (Water Transfer White Paper)," Technical Report, California Department of Water Resources & U.S. Bureau of Reclamation, California–Great Basin Region December 2019.

Eaton, Jonathan and Samuel Kortum, "Technology, Geography, and Trade," *Econometrica*, 2002, *70* (5), 1741–1779.

Edwards, Eric C. and Gary D. Libecap, "Water institutions and the law of one price," in "Handbook on the Economics of Natural Resources," Edward Elgar Publishing, March 2015, pp. 442–473.

Foley, Duncan K., "Economic equilibrium with costly marketing," *Journal of Economic Theory*, 1970, *2* (3).

Foster, Andrew D. and Mark R. Rosenzweig, "Are There Too Many Farms in the World? Labor Market Transaction Costs, Machine Capacities, and Optimal Farm Size," *Journal of Political Economy*, 2022, *130* (3), 636–680.

Gangadharan, Lata, "Transaction Costs in Pollution Markets: An Empirical Study," *Land Economics*, 2000, *76* (4), 601–614.

Glaeser, Edward and Joseph Gyourko, "The Economic Implications of Housing Supply," *Journal of Economic Perspectives*, February 2018, *32* (1), 3–30.

Glaeser, Edward L. and Erzo F. P. Luttmer, "The Misallocation of Housing Under Rent Control," *American Economic Review*, 2003, *93* (4), 1027–1046.

Grainger, Corbett A. and Christopher J. Costello, "Capitalizing property rights insecurity in natural resource assets," *Journal of Environmental Economics and Management*, 2014, *67* (2), 224–240.

Gray, Brian, Ellen Hanak, Richard Frank, Richard Howitt, Jay Lund, Leon Szeptycki, and Barton "Buzz" Thompson, “Allocating California’s Water: Directions for Reform,” Technical Report, Public Policy Institute of California Water Policy Center November 2015.

Haddad, Brent M., *Rivers of Gold: Designing Markets To Allocate Water In California*, Island Press, 2000.

Hagerty, Nick, “Liquid Constrained in California: Estimating the Potential Gains from Water Markets,” October 2019. Working paper.

_ , “Adaptation to Surface Water Scarcity in Irrigated Agriculture,” 2022. Working paper.

Hanak, Ellen, “California’s Water Market, by the Numbers,” PPIC Occasional Paper OP 10-02, Public Policy Institute of California October 2002. Occasional Paper, October 2002.

_ , *Who Should Be Allowed to Sell Water in California? Third-Party Issues and the Water Market* 2003.

_ **and Elizabeth Stryjewski**, “California’s Water Market, By the Numbers: Update 2012,” PPIC Report, Public Policy Institute of California November 2012.

_ **, Brian Gray, Jay Lund, David Mitchell, Caitrin Chappelle, Andrew Fahlund, Katrina Jessoe, Josué Medellín-Azuara, Dean Misczynski, James Nachbaur, and Robyn Suddeth**, “Paying for Water in California,” PPIC Report, Public Policy Institute of California March 2014.

Heard, Sarah, Matthew Fienup, and EJ Remson, “The first SGMA groundwater market is trading: The importance of good design and the risks of getting it wrong,” *California Agriculture*, 2021, *75* (2).

Helpman, Elhanan, Marc Melitz, and Yona Rubinstein, “Estimating Trade Flows: Trading Partners and Trading Volumes,” *Quarterly Journal of Economics*, May 2008, *123* (2), 441–487.

Hendricks, Nathan P. and Jeffrey M. Peterson, “Fixed effects estimation of the intensive and extensive margins of irrigation water demand,” *Journal of Agricultural and Resource Economics*, 2012, *37* (12), 1–19.

Hodgson, Charles, “Trade-ins and Transaction Costs in the Market for Used Business Jets,” *American Economic Journal: Microeconomics*, November 2023, *15* (4), 350–91.

Howitt, Richard E., Jay R. Lund, K.W. Kirby, Marion W. Jenkins, and Andrew J. Draper, “Integrated Economic-Engineering Analysis of California’s Future Water Supply: a report for the State of California Resources Agency,” 1999.

Hsieh, Chang-Tai and Peter J. Klenow, “Misallocation and Manufacturing TFP in China and India,” *Quarterly Journal of Economics*, 2009, *124* (4).

Jack, William and Tavneet Suri, “Risk Sharing and Transactions Costs: Evidence from Kenya’s Mobile Money Revolution,” *American Economic Review*, January 2014, *104* (1), 183–223.

Jarvis, Stephen, "The Economic Costs of NIMBYism: Evidence from Renewable Energy Projects," *Journal of the Association of Environmental and Resource Economists*, 2025, *12* (4), 983–1022.

Jenkins, Marion W., Jay R. Lund, and Richard E. Howitt, "Using Economic Loss Functions to Value Urban Water Scarcity in California," *American Water Works Association*, 2003, *95* (2), 395–429.

Jha, Akshaya and Frank A. Wolak, "Can Forward Commodity Markets Improve Spot Market Performance? Evidence from Wholesale Electricity," *American Economic Journal: Economic Policy*, May 2023, *15* (2), 292–330.

Johnson, Nathanael, "California Has a Real Water Market — But It's Not Exactly Liquid," *Grist*, May 2015.

Kumbhakar, Subal C., Christopher F. Parmeter, and Valentin Zelenyuk, "Stochastic Frontier Analysis: Foundations and Advances II," in Subhash C. Ray, Robert Chambers, and Subal Kumbhakar, eds., *Handbook of Production Economics*, Singapore: Springer, 2022, pp. 1–38.

Kyle, Albert S. and Anna A. Obizhaeva, "Market Microstructure Invariance: Empirical Hypotheses," *Econometrica*, 2016, *84* (4), 1345–1404.

Leonard, Bryan, Christopher Costello, and Gary D. Libecap, "Expanding Water Markets in the Western United States: Barriers and Lessons from Other Natural Resource Markets," *Review of Environmental Economics and Policy*, 2019, *13* (1), 43–61.

Libecap, Gary D., "Chinatown Revisited: Owens Valley and Los Angeles—Bargaining Costs and Fairness Perceptions of the First Major Water Rights Exchange," *The Journal of Law, Economics, and Organization*, October 2009, *25* (2), 311–338.

Libecap, Gary D, "Institutional Path Dependence in Climate Adaptation: Coman's "Some Unsettled Problems of Irrigation"," *American Economic Review*, 2011, *101* (1), 64–80.

Liski, Matti, "Thin versus Thick CO2 Market," *Journal of Environmental Economics and Management*, 2001, *41*, 295–311.

MacKay, Alexander, "Contract Duration and the Costs of Market Transactions," *American Economic Journal: Microeconomics*, August 2022, *14* (3), 164–212.

Manville, Michael, Paavo Monkkonen, Nolan Gray, and Shane Phillips, "Does Discretion Delay Development?," *Journal of the American Planning Association*, 2023, *89* (3), 336–347.

McCann, Laura and K. William Easter, "A framework for estimating the transaction costs of alternative mechanisms for water exchange and allocation," *Water Resources Research*, September 2004, *40* (9), 2003WR002830.

Medellín-Azuara, Josué, Julien J. Harou, Marcelo A. Olivares, Kaveh Madani, Jay R. Lund, Richard E. Howitt, Stacy K. Tanaka, Marion W. Jenkins, and Tingju Zhu, "Adaptability and adaptations of California's water supply system to dry climate warming," *Climatic Change*, 2007, *87*.

Mérel, Pierre and Richard Howitt, "Theory and Application of Positive Mathematical Programming in Agriculture and the Environment," *Annual Review of Resource Economics*, 2014, *6* (1), 451–470.

Olmstead, Sheila M., "The Economics of Managing Scarce Water Resources," *Review of Environmental Economics and Policy*, 2010, *4* (2), 179–198.

Peterson, Deborah, Gavan Dwyer, David Appels, and Jane Fry, "Water trade in the southern Murray-Darling basin," *Economic Record*, 2005, *81* (255), S115–S127.

Qureshi, M. Ejaz, Tian Shi, Sumaira E. Qureshi, and Wendy Proctor, "Removing barriers to facilitate efficient water markets in the Murray-Darling Basin of Australia," *Agricultural Water Management*, 2009, *96* (11), 1641–1651.

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

Regnacq, Charles, Ariel Dinar, and Ellen Hanak, "The Gravity of Water: Water Trade Frictions in California," *American Journal of Agricultural Economics*, 2016, *98* (5), 1273–1294.

Rosegrant, M. W., C. Ringler, D. C. McKinney, X. Cai, A. Keller, and G. Donoso, "Integrated economic-hydrologic water modeling at the basin scale: The Maipo river basin," *Agricultural Economics*, 2000, *24* (1), 33–46.

Rosegrant, Mark W. and Hans P. Binswanger, "Markets in tradable water rights: Potential for efficiency gains in developing country water resource allocation," *World Development*, 1994, *22* (11), 1613–1625.

Scheer, Jennifer B., "Participation, Price, and Competition in the Sacramento Valley Water Market." PhD dissertation, University of California, Davis 2016.

Scheierling, Susanne M., John B. Loomis, and Robert A. Young, "Irrigation water demand: A meta-analysis of price elasticities," *Water Resources Research*, 2006, *42* (1), 1–9.

Sekhri, Sheetal, "Wells, Water, and Welfare: The Impact of Access to Groundwater on Rural Poverty and Conflict," *American Economic Journal: Applied Economics*, 2014, *6* (3), 76–102.

Sellers, Scott, Matthew Zaragoza-Watkins, Christina Babbitt, Ana Lucía García Briones, Ann Hayden, and David Festa, "Better Access. Healthier Environment. Prosperous Communities. Recommended Reforms for the California Water Market," Technical Report, Environmental Defense Fund 2016.

Stavins, Robert N., "Transaction Costs and Tradeable Permits," *Journal of Environmental Economics and Management*, 1995, *29*, 133–148.

Stene, Eric A., "Central Valley Project Overview," Technical Report, U.S. Bureau of Reclamation 1995.

Sunding, David L., David Zilberman, Richard Howitt, Ariel Dinar, and Neal MacDougall, "Measuring the costs of reallocating water from agriculture: A multi-model approach," *Natural Resources Modeling*, 2002, *15* (2), 201–225.

SWRCB, "Guide to Water Transfers in California," Technical Report, State Water Resources Control Board, Division of Water Rights 1999.

Thompson, Barton H., Jr., "Institutional Perspectives on Water Policy and Markets," *California Law Review*, 1993, *81*, 671.

Tomori, Françeska, Erik Ansink, Harold Houba, Nick Hagerty, and Charles Bos, "Market power in California's water market," *American Journal of Agricultural Economics*, 2024, *106* (3), 1274–1299.

UNDP, *Human Development Report 2006: Beyond Scarcity: Power, Poverty and the Global Water Crisis*, New York: Palgrave MacMillan, January 2006.

Western Water Company, "Statement for the State Water Resources Control Board," Statement, Public Workshop on Water Transfers May 2000.

Williamson, Oliver E., "Transaction-Cost Economics: The Governance of Contractual Relations," *The Journal of Law and Economics*, October 1979, *22* (2), 233–261.

Womble, Philip and W. Michael Hanemann, "Water Markets, Water Courts, and Transaction Costs in Colorado," *Water Resources Research*, 2020, *56* (4), 1–28.

Wooldridge, Jeffrey M., *Econometric Analysis of Cross Section and Panel Data*, 2nd ed., Cambridge, MA: MIT Press, 2010.

World Bank Group, "High and Dry: Climate Change, Water, and the Economy," Technical Report, World Bank Group May 2016.

Young, Robert A, "Why are there so few transactions among water users?," *American Journal of Agricultural Economics*, 1986, *68* (5), 1143–1151.

Figures

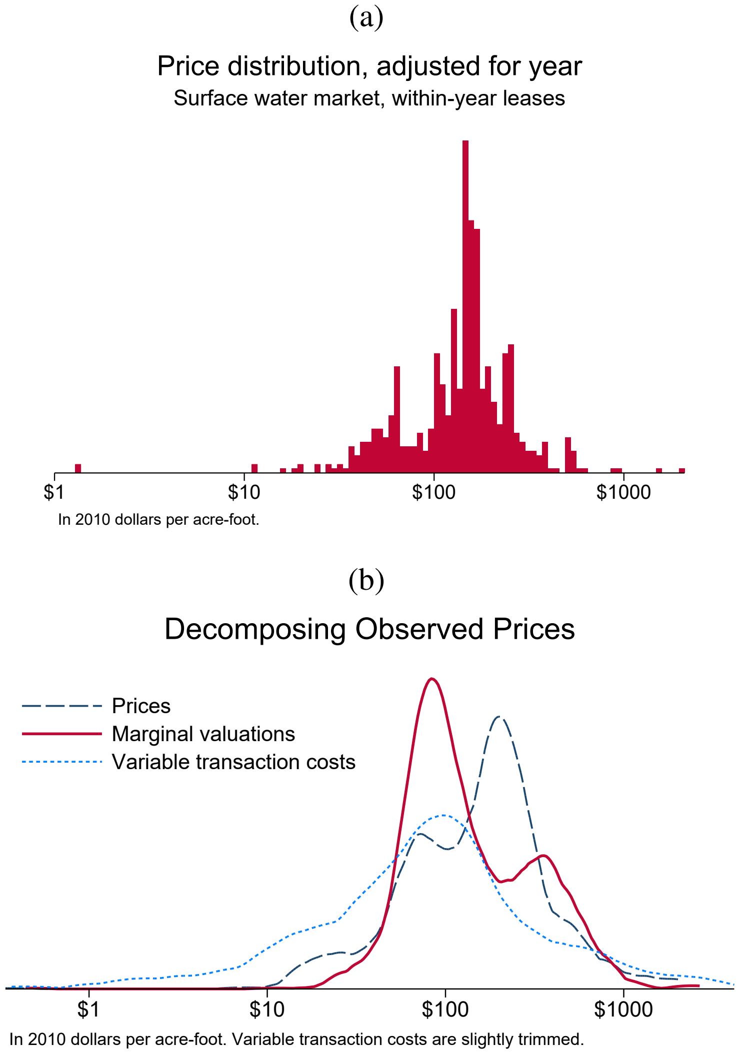

Figure 1: Panel (a): Distribution of surface water transaction prices in the data used for analysis, controlling for year, logarithmic scale. Observations are transactions of within-year leases of surface water observed in California, 1992-2015. The overall sample mean is added to the residual from a regression of log price on year fixed effects. Panel (b): Kernel density estimates of observed prices and estimated marginal valuations for the same transactions, logarithmic scale. Observations are surface water transactions made by each water district.

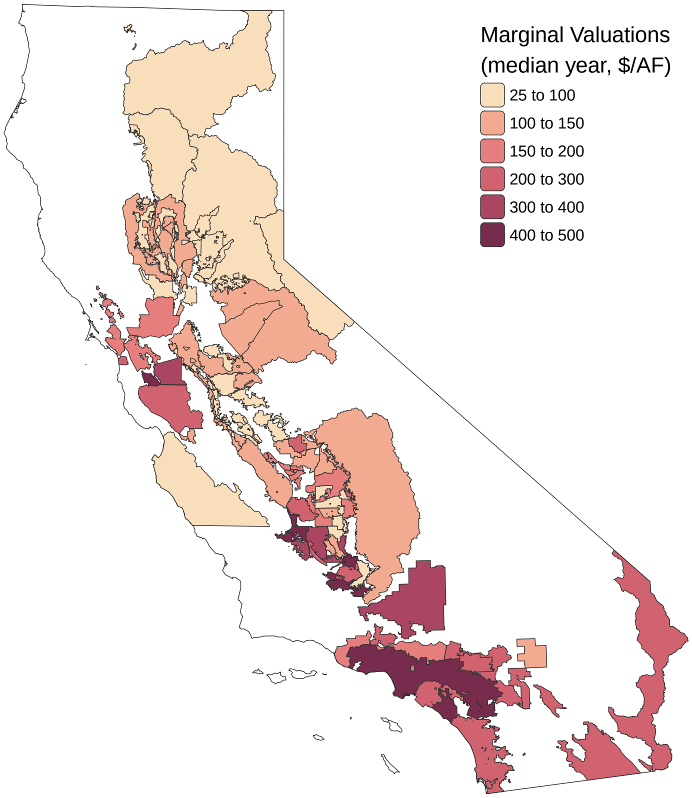

Figure 2: Mean estimated marginal valuations (per acre-foot) by water district for a median-year scenario. Estimates are obtained by evaluating the demand model at quantities observed in 2010. Boundaries are available for districts accounting for 89% of surface water in the simulation sample. Other districts are aggregated to planning areas, geographical regions defined by the California Department of Water Resources, taking means weighted by surface water quantity consumed. Means are taken over districts, weighting by surface water quantity consumed.

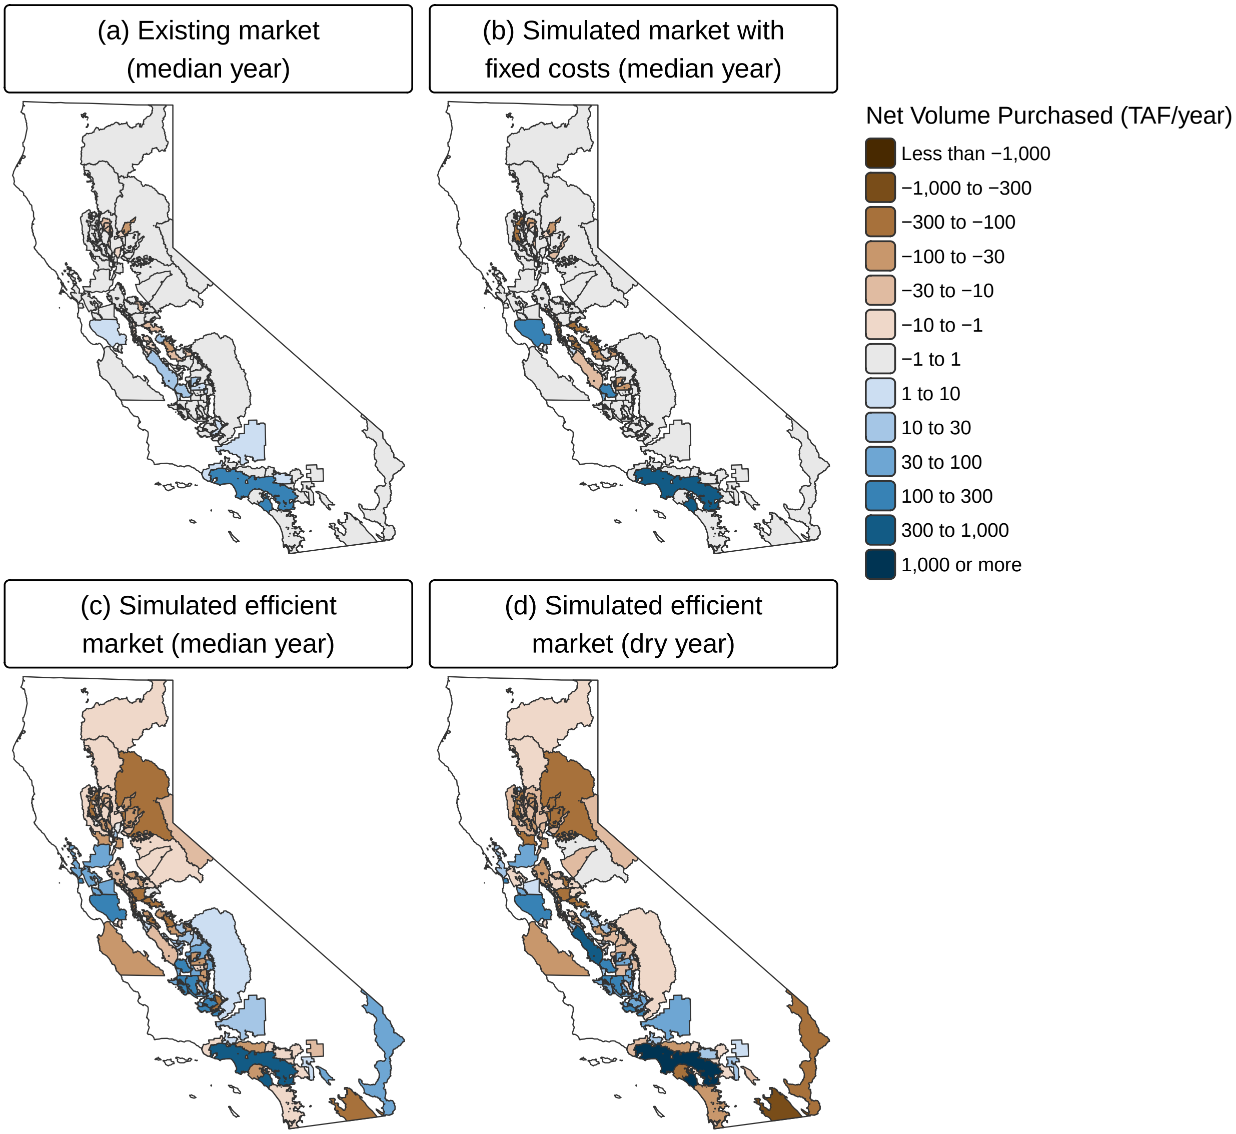

Figure 3: Net annual surface water quantities traded (TAF = thousand acre-feet) in observed transactions (Scenarios 1-2) and additional quantities traded in counterfactual simulations (Scenarios 3-4), by planning area. The efficient market scenarios use the estimated demand model to maximize aggregate surplus, subject to physical conveyance costs calibrated from a hydrological network model. The fixed cost scenario does the same, with trading restricted to districts predicted to enter the market.

Tables

Table 1: Transactions data

Panel A: Summary statistics

Variable	Mean	SD	Obs.
Volume (acre-feet)	9,571.7	28,224.7	1248
Price (2010$/acre-foot)	235.5	297.6	1259
Distance conveyed (km)	206.1	269.7	1247
Distance conveyed in rivers (km)	80.0	103.1	1247
Distance conveyed in canals (km)	126.0	203.8	1247
Distance of virtual conveyance (km)	86.4	118.5	1247
Elevation lift (ft)	421.5	1,029.5	1247
Crosses the Sacramento-San Joaquin Delta (=1)	0.27	0.444	1247
Reviewed by the State Water Boards (=1)	0.411	0.492	1259
Within project (=1)	0.422	0.494	1259

Panel B: Transactions by sector

	Sales		Purchases		Net Purchases
	Count	Volume (TAF)	Count	Volume (TAF)	Volume (TAF)
Agricultural	639	7,130.4	330	1,785.9	-5,344.5
Urban (Municipal & Industrial)	100	819.3	190	2,209.9	1,390.6

Panel C: Transactions by hydrologic region

	Sales Count	Sales Volume (TAF)	Purchases Count	Purchases Volume (TAF)	Net Purchases Volume (TAF)
Central Coast	27	18.5	33	18.7	0.2
Colorado River	20	824.4	6	69.5	-755.0
North Coast	2	3.7	2	0.7	-2.9
Sacramento River	275	4,219.5	37	89.5	-4,130.1
San Francisco Bay	16	188.6	46	530.1	341.5
San Joaquin River	298	2,088.2	104	576.6	-1,511.6
South Coast	10	154.3	68	1,573.3	1,419.0
South Lahontan	2	14.4	30	11.1	-3.2
Tulare Lake	89	438.2	194	1,126.3	688.1

Dataset of spot-market transactions (i.e., within-year leases) of surface water in California made by each district (so transactions are repeated for each party). Panel A reports summary statistics of transaction characteristics. Panels B and C reports the count and total volume of transactions sold and purchased in each sector or hydrologic region. Net purchases within category do not sum to zero because inclusion criteria are applied separately to each side of a transaction (e.g., many agricultural sales are purchased by environmental or government entities, which do not appear in the table as buyers). TAF = thousand acre-feet.

Table 2: Variable Transaction Costs from Observed Determinants (Step 1)

	Dependent variable: Log Price					
	All determinants			Lasso-selected determinants		
	(1)	(2)	(3)	(4)	(5)	(6)
Seller × River dist. (km, 1000s)	–0.81	–0.84	0.22			0.45
	(0.76)	(0.63)	(0.60)			(0.64)
Seller × Canal dist. (km, 1000s)	–0.31	–0.30	0.042			
	(0.68)	(0.73)	(0.80)			
Seller × Virtual dist. (km, 1000s)	–0.73	–0.68	–0.50			
	(0.47)	(0.50)	(0.44)			
Seller × Pumping lift (ft, 1000s)	0.025	0.075	0.027			
	(0.10)	(0.10)	(0.12)			
Seller × Delta crossing (=1)	1.0***	0.94***	0.42	0.66***	0.68***	0.31
	(0.26)	(0.27)	(0.30)	(0.14)	(0.20)	(0.20)
Seller × State Boards review (=1)	0.53**	0.51**	–0.055	0.59**	0.49*	
	(0.22)	(0.23)	(0.27)	(0.26)	(0.26)	
Seller × Import into project (=1)	–0.0079	0.095	0.058			
	(0.064)	(0.091)	(0.059)			
Seller × Export from project (=1)	0.39***	0.28*	0.35	0.31***	0.26*	
	(0.13)	(0.17)	(0.28)	(0.096)	(0.14)	
Seller × Ag counterparty (=1)	0.0034	0.061	0.25***			0.20**
	(0.096)	(0.11)	(0.079)			(0.080)
−Buyer × River dist. (km, 1000s)	–0.91	–0.71	–0.65			
	(0.59)	(0.53)	(0.50)			
−Buyer × Canal dist. (km, 1000s)	–0.48	–0.15	0.25			
	(0.66)	(0.72)	(0.66)			
−Buyer × Virtual dist. (km, 1000s)	0.095	0.017	–0.23			
	(0.36)	(0.49)	(0.32)			
−Buyer × Pumping lift (ft, 1000s)	–0.20*	–0.26*	–0.36***			
	(0.11)	(0.14)	(0.13)			
−Buyer × Delta crossing (=1)	0.28	0.22	0.17			
	(0.21)	(0.16)	(0.17)			
−Buyer × State Boards review (=1)	0.31**	0.31**	0.30**			0.17**
	(0.13)	(0.14)	(0.13)			(0.076)
−Buyer × Import into project (=1)	–0.31**	–0.26**	–0.13			
	(0.14)	(0.13)	(0.10)			
−Buyer × Export from project (=1)	0.56***	0.55***	0.54***		0.48	0.45**
	(0.20)	(0.21)	(0.19)		(0.31)	(0.22)
−Buyer × Ag counterparty (=1)	–0.20	–0.21*	–0.16			
	(0.12)	(0.12)	(0.13)			
Side × District FE	✓	✓		✓	✓	
Side × Hydro. region × Year FE	✓			✓		
Side × Planning area × Year FE		✓			✓	
Side × District × Year FE			✓			✓
Observations	1,247	1,247	1,247	1,247	1,247	1,247
Observations excluding singletons	1,032	966	689	1,032	966	689
Clusters	113	105	76	113	105	76

Regressions of transaction price on observable cost determinants. Observations are transactions made by each water district, stacked such that each transaction is repeated for each district involved. Cost determinants are interacted with the district's side of the market (buyer or seller). On both sides of the market, positive coefficients indicate positive transaction costs. Seller-side coefficients compare prices across buyers, so a positive coefficient reflects a price premium. Buyer-side coefficients compare prices across sellers, but regressors are multiplied by -1, so a positive coefficient reflects a price discount. California is divided into 10 hydrological regions and 56 planning areas. Standard errors in parentheses are clustered by transaction and district. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 3: Market Entry and Inverse Price Elasticity of Demand (Steps 2-3)

	(1)	(2)	(3)	(4)
	Probit	First stage	Reduced form	IV Lasso
	Any Trades	Log Quantity	Adjusted Log Price	Adjusted Log Price
Log quantity consumed				–1.4**
				(0.46)
Log allocation %, own	0.021	0.16***	–0.11	
	(0.040)	(0.045)	(0.075)	
Log allocation %, rest of state	–1.9***	0.96***	–3.1***	
	(0.44)	(0.29)	(0.91)	
Inverse Mills ratio		–0.0030	–0.045	–0.092
		(0.027)	(0.061)	(0.069)
Any trades, year $t-1$	0.73***			
	(0.097)			
Any trades, year $t-2$	0.92***			
	(0.12)			
Any trades, year $t-3$	0.25***			
	(0.092)			
Any trades, sum of prior years	0.13***			
	(0.017)			
Log maximum water rights	0.18***			
	(0.030)			
Log cropland area	0.016			
	(0.027)			
Central Valley Project (=1)	0.50***			
	(0.077)			
State Water Project (=1)	0.89***			
	(0.15)			
Irrigation District (=1)	0.29***			
	(0.098)			
California Water District (=1)	0.30***			
	(0.10)			
Allocation variables (lasso-selected)	✓			
District fixed effects		✓	✓	✓
Linear time trend		✓	✓	✓
Observations	54,498	972	972	972
Clusters	2,355	155	155	155
Pseudo-R^2	0.59			
Lasso-selected instruments				3
First-stage F-statistic		13		46
Sup-score weak-ID test				reject

Observations are transactions made by each water district. Column (1) reports pooled probit estimates of the extensive-margin decision of market entry (i.e., the probability of any surface water transactions); the data also includes years when districts do not trade and districts that never trade. Columns (2) and (3) report first-stage and reduced-form effects of two allocation instruments on quantity consumed (of surface water) and marginal valuations (i.e, prices adjusted for variable transaction costs estimated in Step 1). Column (4) reports instrumental variable lasso estimates of the elasticity of marginal valuations with respect to quantity consumed, using the full set of candidate instruments. Columns (2)-(4) correct for sample selection using inverse Mills ratios from year-specific estimates of the model in column (1). Standard errors in parentheses are clustered by district in column (1) and calculated by Bayesian bootstrap with weights blocked by district in columns (2)-(4). * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 4: Annual Economic Gains from Surface Water Markets

Panel A: Observed Transactions in the Existing Market

Conditions	Volume traded (TAF)	Average marginal valuation ($/AF)		Gains (millions)
Scenario 1. Observed market (spot market transactions)				
Dry	123	$ 285		$ 55.2
Median	289	$ 161		$ 8.3
Wet	35	$ 129		$ 1.5

Panel B: Simulated Markets Without Transaction Costs

Conditions	Additional volume traded (TAF)	Average equilibrium price ($/AF)	Avoided costs of observed transactions (millions)	Gains from new transactions (millions)
Scenario 2: Efficient market (with only physical costs of conveyance)				
Dry	3,929	$ 176	$ 84.7	$ 559.0
Median	3,338	$ 137	$ 45.1	$ 211.3
Wet	3,502	$ 126	$ 4.6	$ 220.3
Scenario 3: With fixed transaction costs				
Dry	1,823	$ 238	$ 71.1	$ 393.1
Median	1,157	$ 161	$ 38.8	$ 76.8
Wet	1,219	$ 129	$ 3.0	$ 139.6
Scenario 4: No conveyance costs				
Dry	4,698	$ 205	$ 93.6	$ 1,099.5
Median	4,982	$ 165	$ 82.0	$ 635.2
Wet	5,283	$ 151	$ 11.0	$ 721.0
Scenario 5: With a key environmental constraint (Sacramento River outflow held fixed)				
Dry	3,402	$ 188	$ 66.3	$ 505.4
Median	3,016	$ 139	$ 43.3	$ 199.1
Wet	3,251	$ 127	$ 4.2	$ 213.7
Scenario 6: With more capacity constraints (Sac. River and Colorado River Aqueduct held fixed)				
Dry	3,325	$ 189	$ 66.3	$ 503.0
Median	3,015	$ 139	$ 43.3	$ 198.9
Wet	3,233	$ 128	$ 4.2	$ 213.3
Scenario 7: Among only prior traders (districts observed in the transactions data)				
Dry	3,138	$ 211	$ 69.4	$ 462.9
Median	2,340	$ 154	$ 43.8	$ 132.7
Wet	2,701	$ 126	$ 3.3	$ 177.6
Scenario 8: Linear demand (instead of isoelastic)				
Dry	2,658	$ 133	$ 80.4	$ 360.5
Median	2,825	$ 114	$ 51.0	$ 159.4
Wet	2,980	$ 105	$ 2.5	$ 167.0

Per-year welfare analysis of the existing market (Panel A) and counterfactual simulations (Panel B). Each scenario is run under dry, median, and wet conditions, which draw quantities and trading volumes from the years 2014, 2010, and 2006, respectively. All dollar figures are in 2010 USD; gains are per year. TAF = thousand acre-feet.

A Online Appendix: Tables and Figures

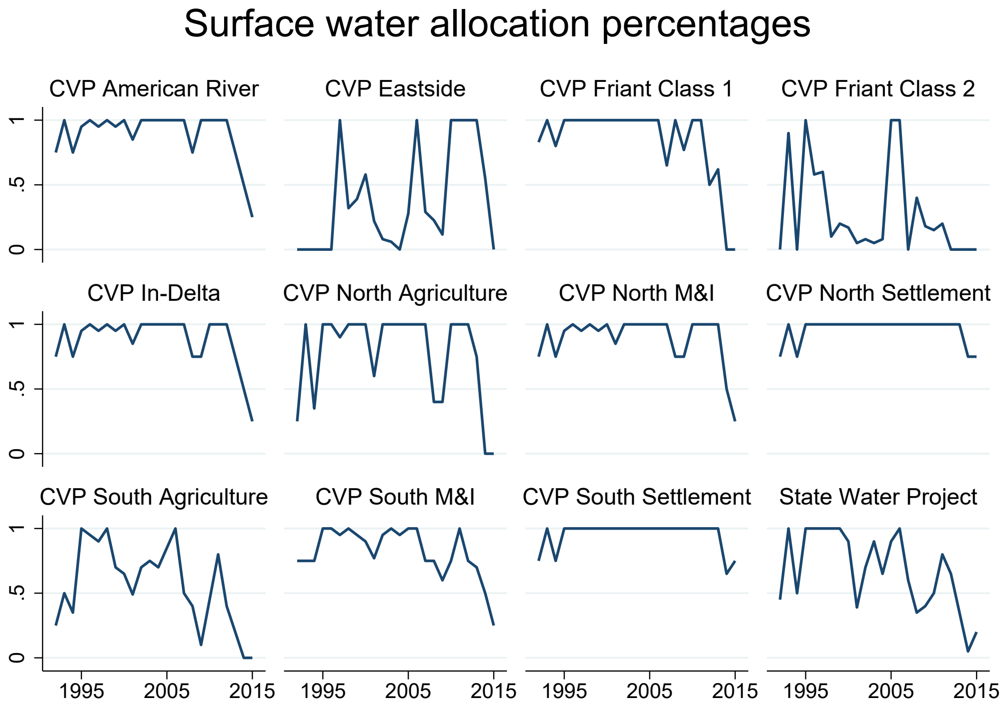

Figure A1: Variation over time in allocation percentages for each category of contracts with the federal and state water projects (CVP = Central Valley Project, M&I = Municipal & Industrial). These allocation percentages are used as instruments to estimate demand elasticities. This figure combines agricultural and municipal contracts in the State Water Project because they are equal during the years used in analysis.

(a) Modeled conveyance costs from selected origins

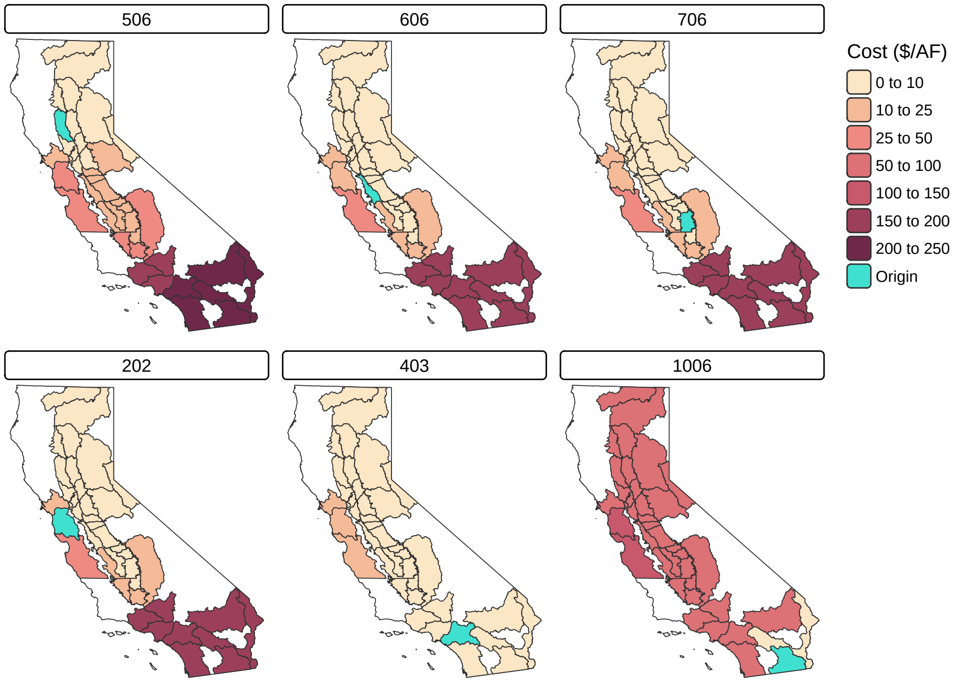

(b) Modeled conveyance costs to selected destinations

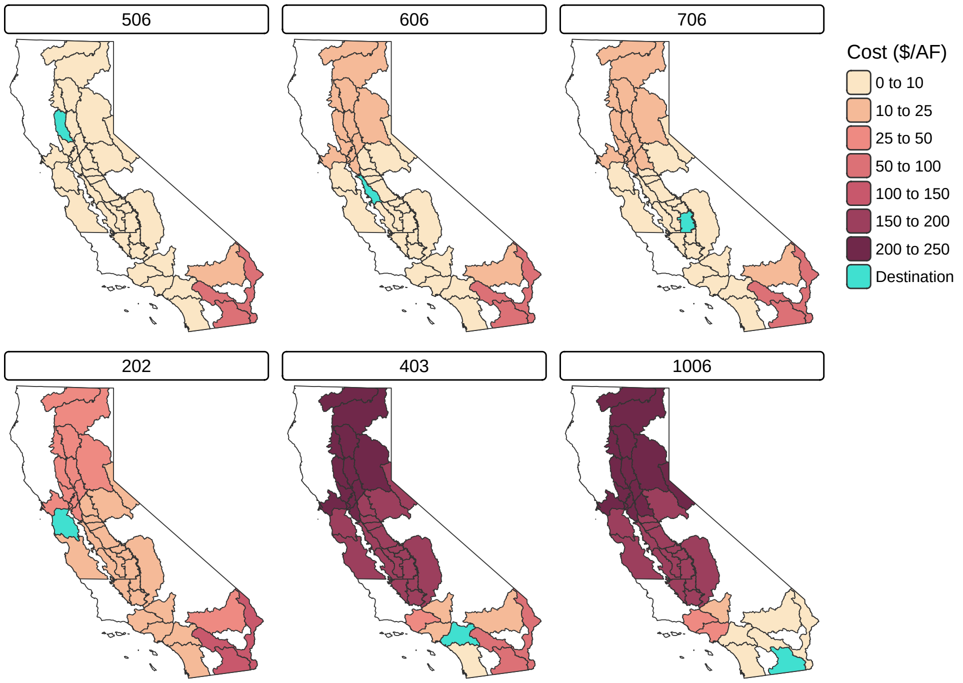

Figure A2: Conveyance costs in the hydrologic network model used in simulations. Panel (a) shows costs of transporting water from selected origins (in turquoise) to all other planning areas in the data; Panel (b) shows costs of transporting water to selected destinations from all other planning areas. The six origins/destinations selected for these maps are representative planning areas within each of the six hydrological regions with the most water consumption.

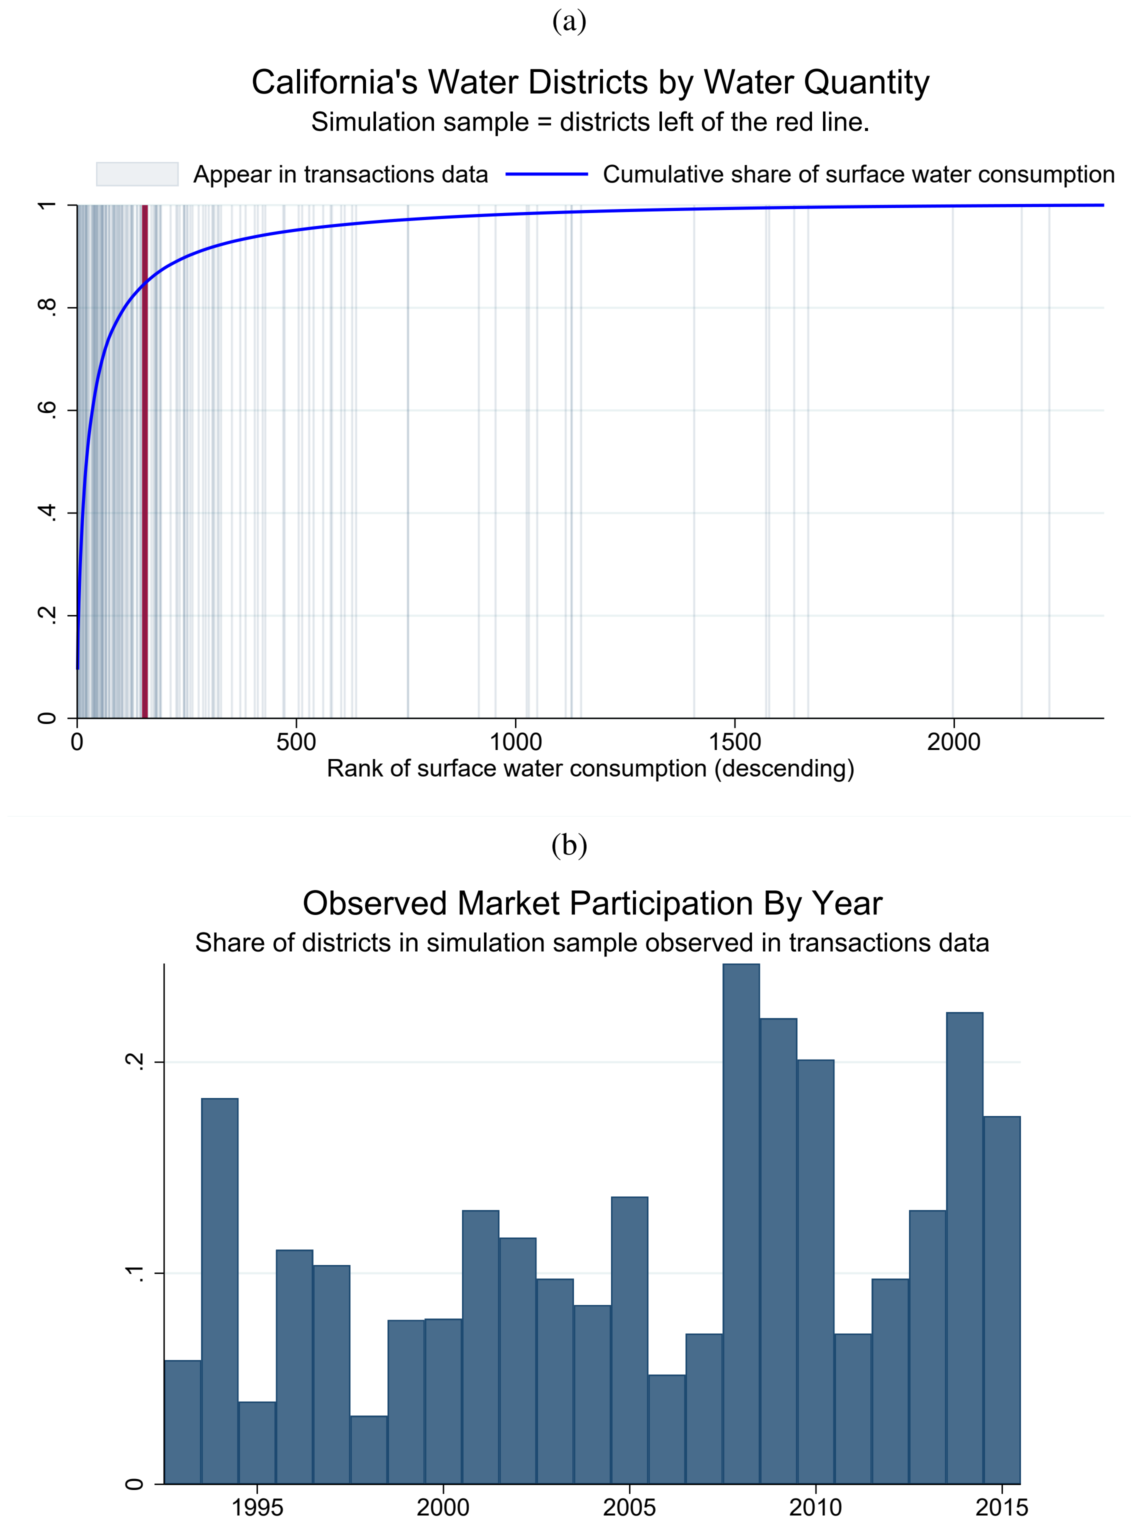

Figure A3: Districts in the simulation sample, by surface water quantity consumed and appearance in the transactions data. Panel (a) plots the cumulative density of surface water use in the state for the full analysis sample (i.e., all water districts in California that consume at least 100 acre-feet per year). The background is shaded for districts that also appear in the transactions data. The simulation sample lies to the left of the vertical red line (i.e., the 154 districts with more than 25,000 acre-feet/year). Panel (b) plots the share of districts in the simulation sample that appear in the transactions data in each year.

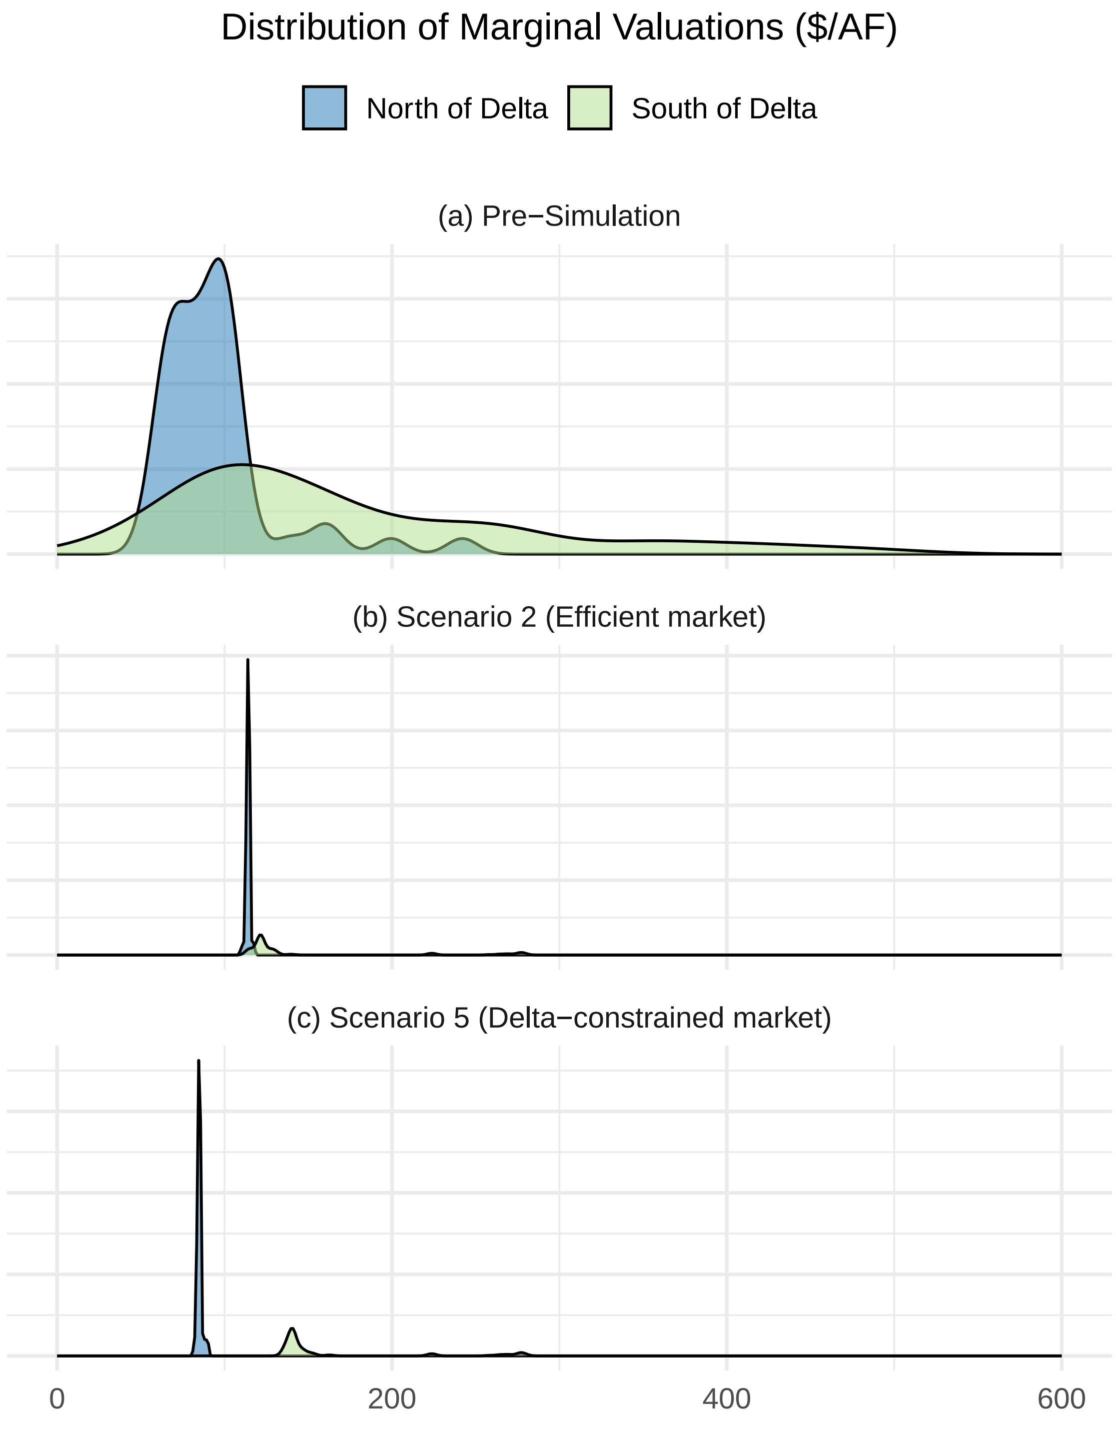

Figure A4: Kernel density estimates of marginal valuations for median-year scenarios. Panel (a) plots estimated marginal valuations before simulations. Panels (b) and (c) plot final marginal valuations after simulations in Scenarios 2 and 5 – with and without the environmental constraint at the Delta. Observations are water districts.

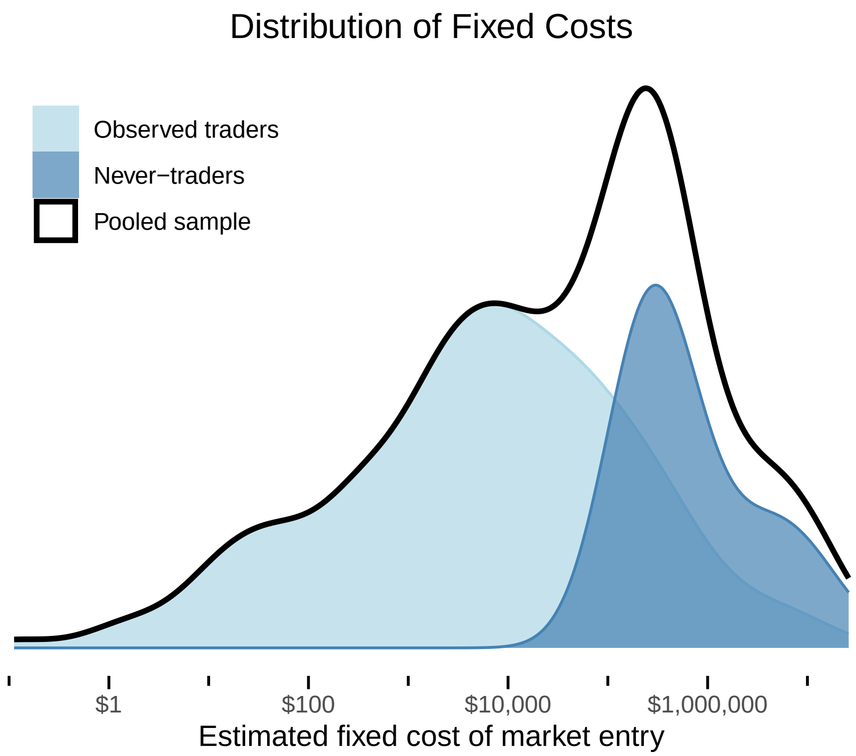

Figure A5: Kernel density estimates of estimated fixed transaction costs, logarithmic scale. Observations are water districts. "Observed traders" are districts ever observed in transactions data; fixed costs are estimated as the minimum surplus observed in any year of the data. "Never-traders" are districts included in the simulation sample but never observed in transactions data; fixed costs are estimated as the maximum surplus simulated in the efficient-market scenario. "Pooled sample" includes all districts from both samples.

Table A1: Variable Transaction Costs from Observed Determinants: Robustness Checks

	Dependent variable: Log Price					
	All determinants			Lasso-selected determinants		
	(1)	(2)	(3)	(4)	(5)	(6)
Seller × River dist. (km, 1000s)	–0.62	–0.95	–0.16			–0.25
	(0.60)	(0.62)	(0.72)			(0.51)
Seller × Canal dist. (km, 1000s)	0.094	–0.79	–1.2*			
	(0.52)	(0.73)	(0.65)			
Seller × Virtual dist. (km, 1000s)	–0.44	–0.18	–0.032			
	(0.51)	(0.45)	(0.47)			
Seller × Pumping lift (ft, 1000s)	–0.019	0.15	0.22*			
	(0.085)	(0.12)	(0.12)			
Seller × Delta crossing (=1)	0.77***	0.99***	0.55**	0.60***	0.73***	0.53**
	(0.16)	(0.19)	(0.24)	(0.13)	(0.15)	(0.22)
Seller × State Boards review (=1)	0.28	0.31	–0.019		0.29	
	(0.27)	(0.27)	(0.27)		(0.25)	
Seller × Import into project (=1)	0.11	0.11	0.0065			
	(0.15)	(0.15)	(0.075)			
Seller × Export from project (=1)	0.33*	0.36**	0.47*	0.47***	0.26	0.45***
	(0.17)	(0.17)	(0.26)	(0.13)	(0.18)	(0.13)
Seller × Ag counterparty (=1)	0.063	0.14	0.29***			0.22***
	(0.085)	(0.11)	(0.089)			(0.079)
−Buyer × River dist. (km, 1000s)	–0.48	–0.25	–0.43			
	(0.60)	(0.53)	(0.54)			
−Buyer × Canal dist. (km, 1000s)	–1.3	–1.5	–1.3	–1.2***	–1.4***	–1.0***
	(1)	(1.1)	(1.1)	(0.44)	(0.43)	(0.35)
−Buyer × Virtual dist. (km, 1000s)	–0.44	–0.51	–0.69**			
	(0.42)	(0.46)	(0.32)			
−Buyer × Pumping lift (ft, 1000s)	–0.032	–0.070	–0.15			
	(0.14)	(0.16)	(0.16)			
−Buyer × Delta crossing (=1)	0.19	0.22	0.26*			
	(0.15)	(0.15)	(0.15)			
−Buyer × State Boards review (=1)	0.11	0.086	0.0086			
	(0.084)	(0.086)	(0.066)			
−Buyer × Import into project (=1)	–0.23	–0.17	0.0059			
	(0.15)	(0.15)	(0.13)			
−Buyer × Export from project (=1)	0.22	0.19	0.011			
	(0.15)	(0.17)	(0.084)			
−Buyer × Ag counterparty (=1)	–0.12	–0.096	0.024			
	(0.12)	(0.13)	(0.14)			
Side × District FE	✓	✓		✓	✓	
Side × Hydro. region × Year × Month FE	✓			✓		
Side × Planning area × Year × Month FE		✓			✓	
Side × District × Year × Month FE			✓			✓
Observations	1,233	1,233	1,233	1,233	1,233	1,233
Observations excluding singletons	951	863	605	951	863	605
Clusters	104	97	67	104	97	67

Regressions of transaction price on observable cost determinants, as in Table 2. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A2: Inverse Price Elasticity of Demand: Robustness Checks

	Adjusted log price				Unadjusted	Adjusted
	(1)	(2)	(3)	(4)	(5)	(6)
Log quantity consumed	–1.9***	–1.7***	–1.4**	–1.5**	–1.5**	–0.91*
	(0.48)	(0.46)	(0.62)	(0.46)	(0.48)	(0.52)
Inverse Mills ratio	–0.22*	–0.085	–0.10		–0.069	–0.17
	(0.072)	(0.075)	(0.069)		(0.065)	(0.066)
District fixed effects	✓	✓	✓	✓	✓	✓
Linear time trend		✓	✓	✓	✓	✓
Planning area X time trend	✓					
Year fixed effects						✓
Candidate instruments	All	Overall	Regional	All	All	All
Lasso-selected instruments	3	2	4	3	3	3
Observations	972	972	976	981	978	972
Clusters	155	155	156	155	156	155
First-stage F-statistic	53	49	39	47	50	12
Sup-score weak-ID test	reject	reject	reject	reject	reject	fail to reject

Instrumental variable lasso estimates of the inverse price elasticity of surface water demand. Observations are transactions made by each water district. Column (1) interacts the linear time trend with indicators for planning areas (56 geographic regions). Column (2) uses only the 3 overall water allocation instruments; column (3) uses only the remaining 130 region-interacted water allocation instruments. Column (4) omits the sample selection correction. Column (5) uses unadjusted prices as the dependent variable. Column (6) uses year fixed effects, which create spillover bias due to the joint determination of prices and quantities throughout the market. Standard errors in parentheses are calculated by Bayesian bootstrap with weights blocked by district. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A3: District Intercepts of Marginal Valuations (Step 4)

	Adjusted Log Price
Serves primarily municipal customers (=1)	0.17
	(0.12)
Central Valley Project contractor (=1)	–0.096
	(0.11)
State Water Project contractor (=1)	0.53***
	(0.15)
Lower Colorado rights holder (=1)	0.73**
	(0.17)
Surface water rights holder (=1)	0.51***
	(0.096)
Log maximum water rights (centered)	–0.14***
	(0.039)
Latitude (degrees, centered)	–0.12*
	(0.061)
Constant	4.8***
	(0.14)
Var(District indicators)	0.55
	(0.24)
Var(District/counterparty pair)	0.35
	(0.067)
Var(Residual)	0.27
	(0.068)
Observations	965
Clusters	154

Random effects estimates of district and district-pair intercepts of marginal valuations (i.e., prices adjusted both for variable transaction costs estimated in Step 1 and the demand elasticity estimated in Step 3). Observations are transactions made by each water district. Standard errors in parentheses are calculated by Bayesian bootstrap with weights blocked by district. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A4: Summary statistics for water districts used in simulations

Panel A: Districts observed in transactions data

Variable	Mean	SD	Obs.
Serves primarily municipal customers (=1)	0.12	0.33	82
Central Valley Project contractor (=1)	0.59	0.5	82
State Water Project contractor (=1)	0.18	0.39	82
Lower Colorado rights holder (=1)	0.061	0.24	82
Surface water rights holder (=1)	0.65	0.48	82
Maximum water rights (TAF/year)	283.2	515.2	82
Latitude (degrees)	37.1	1.8	82
Marginal valuation of water, estimated mean (2010$/AF)	144.7	95.2	82

Panel B: Districts not observed in transactions data

Variable	Mean	SD	Obs.
Serves primarily municipal customers (=1)	0.17	0.38	72
Central Valley Project contractor (=1)	0.11	0.32	72
State Water Project contractor (=1)	0.12	0.33	72
Lower Colorado rights holder (=1)	0.014	0.12	72
Surface water rights holder (=1)	0.86	0.35	72
Maximum water rights (TAF/year)	107.1	137.1	72
Latitude (degrees)	37.2	2.0	72
Marginal valuation of water, estimated mean (2010$/AF)	155.9	82.9	72

Summary statistics for the set of water districts selected for simulations (i.e., all water districts that consume more than 25,000 AF/year), tabulated separately by whether districts are ever observed in the transactions dataset. Districts in Panel A are observed in the transactions data and therefore contribute to model estimation. Districts in Panel B are not, so their demand parameters are extrapolated from the estimated model on the basis of district covariates. TAF = thousand acre-feet.

Table A5: Extrapolation Model of District-Specific Demand Intercepts

	Estimated Mean Marginal Valuation
Serves primarily municipal customers (=1)	0.11
	(0.44)
Central Valley Project contractor (=1)	–0.26**
	(0.12)
State Water Project contractor (=1)	0.51**
	(0.26)
Lower Colorado rights holder (=1)	0.36
	(0.42)
Surface water rights holder (=1)	–0.20
	(0.17)
Log maximum water rights (centered)	–0.079
	(0.049)
Latitude (degrees, centered)	–0.083
	(0.12)
Hydrologic region: San Francisco Bay (=1)	0.85
	(0.68)
Hydrologic region: Central Coast (=1)	0.041
	(0.78)
Hydrologic region: South Coast (=1)	0.51
	(0.94)
Hydrologic region: Sacramento River (=1)	0.17
	(0.77)
Hydrologic region: San Joaquin River (=1)	0.13
	(0.81)
Hydrologic region: Tulare Lake (=1)	0.38
	(0.86)
Hydrologic region: South Lahontan (=1)	0.40
	(0.82)
Hydrologic region: Colorado River (=1)	0.011
	(0.86)
Constant	5.0***
	(0.75)
Observations	1,061
Clusters	154
R-squared	0.45

Regression of district-specific demand intercepts, as estimated in Step 4, on district-level variables. Observations are transactions made by each water district. Robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A6: Annual Economic Gains from Surface Water Markets: Robustness Checks

Using Estimated Demand Intercepts for Observed Traders				
Conditions	Additional volume traded (TAF)	Average equilibrium price ($/AF)	Avoided costs of observed transactions (millions)	Gains from new transactions (millions)
Scenario 2: Efficient market (with only physical costs of conveyance)				
Dry	4,497	$ 201	$ 105.2	$ 685.2
Median	4,323	$ 154	$ 37.9	$ 460.1
Wet	4,271	$ 140	$ 5.7	$ 382.2
Scenario 3: With fixed transaction costs				
Dry	1,798	$ 250	$ 70.8	$ 437.7
Median	1,331	$ 175	$ 33.4	$ 93.2
Wet	908	$ 146	$ 2.5	$ 92.4
Scenario 4: No conveyance costs				
Dry	5,341	$ 231	$ 116.1	$ 1,333.3
Median	5,898	$ 189	$ 77.2	$ 988.0
Wet	6,197	$ 171	$ 13.4	$ 1,005.7
Scenario 5: With a key environmental constraint (Sacramento River outflow held fixed)				
Dry	3,953	$ 204	$ 85.9	$ 632.9
Median	3,986	$ 154	$ 35.9	$ 448.7
Wet	4,089	$ 140	$ 5.3	$ 373.0
Scenario 6: With more capacity constraints (Sac. River and Colorado River Aqueduct held fixed)				
Dry	3,953	$ 204	$ 85.9	$ 632.9
Median	3,991	$ 153	$ 39.2	$ 445.7
Wet	4,066	$ 139	$ 5.3	$ 371.7
Scenario 7: Among only prior traders (districts observed in the transactions data)				
Dry	3,248	$ 230	$ 94.3	$ 519.4
Median	3,393	$ 167	$ 41.3	$ 374.2
Wet	3,463	$ 137	$ 5.5	$ 382.8
Scenario 8: Linear demand (instead of isoelastic)				
Dry	3,520	$ 135	$ 103.0	$ 435.2
Median	4,026	$ 110	$ 38.7	$ 325.1
Wet	3,996	$ 103	$ 3.4	$ 299.0

Per-year welfare analysis of counterfactual simulations in which demand intercepts are taken directly from the estimated demand model for districts observed in the transactions data. Districts not observed trading still use fitted values from an extrapolation model. Each scenario is run under dry, median, and wet conditions, which draw quantities and trading volumes from the years 2014, 2010, and 2006, respectively. All dollar figures are in 2010 USD; gains are per year. TAF = thousand acre-feet.

B Online Appendix: Theory and Estimation

B.1 Motivating no-arbitrage

The no-arbitrage assumption can be motivated in at least two ways. The first appeals to competition across districts: If a district k buys water at one price from seller j, but the price is higher than what another seller j' can offer, it will be profitable for j' to undercut the price (Anderson and van Wincoop 2004). The second motivation, which dates back to Foley (1970), relies on intermediaries who buy goods from some consumers and sell them to others, use resources, and operate in perfect competition. In a water market the intermediaries might be thought of as brokers. I now derive Equation 3 using this second motivation.

Each district is comprised of a continuum of consumers whose preferences are aggregated to the district's demand function. Spatial arbitrage across districts is conducted by two layers of intermediaries, which might be thought of as brokers representing each district. In each district, selling intermediaries ("sellers") can buy units of water from consumers (at their marginal valuations V_j) and sell to buying intermediaries ("buyers") in another district. Buyers, in turn, buy water from sellers and can sell to consumers in their own district (at their marginal valuations V_k). Sellers and buyers meet at exchange points unique to each pair of districts, where prices p_{ijk} are determined. (In this section I suppress time subscripts.)

Assumption 6 (Perfect competition)**.** *Each district has enough intermediaries that they behave as price takers. That is, the quantity sold* q^s_{ijk} *or purchased* q^b_{ijk} *by any one intermediary does not affect the equilibrium price for any district pair:* $dp_{ijk}/dq^m_{ijk} = 0$ *for all* j *and* k*.*

Sellers and buyers choose non-negative quantities for each buyer-seller pair jk to maximize net profits:

$$\begin{aligned} &(Sellers) \quad \max_{q^s_{ijk}} p_{ijk}q^s_{ijk} - V_j(Q_j)q^s_{ijk}\tau^s_{jk} \quad \text{s.t. } q^s_{ijk} \geq 0 \\ &(Buyers) \quad \max_{q^b_{ikj}} V_k(Q_k)q^b_{ikj} - p_{ikj}q^b_{ikj}\tau^b_{kj} \quad \text{s.t. } q^b_{ikj} \geq 0 \end{aligned} \tag{15}$$

Each problem has two candidate solutions. First, sellers and buyers may not trade at all. In order for trading quantities to be positive, there must be non-negative marginal surplus between the seller and buyer: $V_j(Q_j)\tau^s_{jk} \geq V_k(Q_k)/\tau^b_{kj}$. Second, if buyers and sellers do trade, first-order conditions yield the no-arbitrage condition in Equation 3.

B.2 Identification assumptions for Step 3

Besides the sample selection, this step has two key identification assumptions. One is conditional independence: changes in allocation percentages are not correlated with any other time-varying factors that independently affect prices or quantities. The other is the exclusion restriction: changes in allocation percentages affect prices only through movements along demand curves, not through shifts in demand curves. Conditional independence ensures that the first stage and reduced form relationships are free from omitted variables bias, and the exclusion restriction ensures that the IV estimate can be interpreted as a causal relationship.

Conditional independence is a plausible assumption. Unit fixed effects absorb the influence of typical water availability, so the elasticity is estimated using only year-to-year variation in allocation percentages. In a different setting, local weather patterns might be an omitted variable, but in California, local rainfall meets a vanishingly small proportion of water demand. If a unit's own allocation percentage were the only instrument, another omitted variable might be water supplies in other parts of the state, since they are correlated and can all affect equilibrium outcomes, but I avoid this problem by using the full set of allocation percentages as instruments for prices faced by each unit.

The exclusion restriction is also plausible in this setting. Allocation percentages are pure supply shocks. Increasing one unit's surface water allocation will increase quantities and decrease equilibrium prices, moving along demand curves without changing underlying preferences. Increasing other units' allocations will lower their marginal valuations, raising quantities traded, decreasing equilibrium prices and again increasing the first unit's quantity. Substitution to groundwater or storage does not violate the exclusion restriction, which simply requires that any allocations-driven changes in quantities also be reflected in prices.[42]

---

[42] A decrease in surface water endowments may lead a unit to extract more groundwater, reducing surface water quantity by less than would occur otherwise. In a simple model, agents extract groundwater until the marginal cost equals the marginal valuation of water. Suppose the cost of groundwater extraction does not depend on surface water endowments, and groundwater is perfect substitute for surface water in the short run. Under these plausible conditions, year-to-year changes in groundwater quantities are fully determined by changes in the marginal valuations of water, and they need not enter demand as a separate term. Estimated elasticities measure the response of market prices to surface water quantities, net of any groundwater extraction response.

C Online Appendix: Proofs

C.1 Year effects bias elasticity estimates when comparing within a common market

Year effects introduce bias by comparing across units within the same water market. Trade creates mechanical spillovers: changes to any one unit's prices and quantities alter the equilibrium and affect the prices and quantities of others, violating the stable unit treatment value assumption (SUTVA). The intuition is that within a single interconnected market, year effects difference out statewide average prices and quantities, but these averages are themselves endogenous.

To keep the proof as simple as possible, I focus on the first-stage relationship, showing that the effect of the instrument on quantities is biased. Consider a regression of quantities on prices estimated by two-stage least squares with a single instrument. Because the two-stage least squares (2SLS) estimate is equal to the ratio of the first-stage and reduced form coefficients, if the reduced form is biased, the 2SLS estimate is also biased.

Consider a simple data generating process involving two agents. In each year, each agent receives a fixed quantity α_i plus an observed time-varying endowment z_{it}, of which they keep a fraction β and trade away the remainder to the other agent. Each agent also consumes an idiosyncratic shock ε_{it} that is uncorrelated with endowments. Total quantities are:

$$
\begin{aligned}
q_{1t} &= \alpha_1 + \beta z_{1t} + (1-\beta) z_{2t} + \varepsilon_{1t} \qquad \mathbb{E}[\varepsilon_{1t} | z_{1t}, z_{2t}] = 0 \\
q_{2t} &= \alpha_2 + \beta z_{2t} + (1-\beta) z_{1t} + \varepsilon_{2t} \qquad \mathbb{E}[\varepsilon_{2t} | z_{1t}, z_{2t}] = 0
\end{aligned}
$$

This model captures a market with inertia; $\beta = 1$ would correspond to autarky while $\beta = 0.5$ would suggest no inertia, since endowments given to either agent would be allocated evenly. For simplicity, the model is linear and the coefficient β is constant across the two agents.

First, under this data generating process, a simple fixed effects regression that includes both endowments (each agent's own endowment and the other agent's endowment) would recover the correct, unbiased parameter β, because the econometric model would be identical to the data generating process.

Second, in general, a regression measuring the effect of agents' own endowment must also control for the other agent's endowment. An estimate of β from a regression containing only each agent's own endowment would suffer from omitted variables bias unless the other agent's endowment z_{-it} is uncorrelated with own endowment z_{it}.

Third, using year effects will produce a biased estimate of β. Consider the regression

$$q_{it} = \alpha_i + \beta z_{it} + \theta_t + v_{it}.$$

Year effects are incidental parameters, so they can be eliminated by differencing the two agents:

$$\begin{aligned}(q_{1t}-q_{2t}) &= (\alpha_1-\alpha_2)+\beta(z_{1t}-z_{2t})+(\theta_t-\theta_t)+(v_{1t}-v_{2t}).\\ \Delta q_t &= \Delta\alpha+\beta\Delta z_t+\Delta v_t.\end{aligned}$$

This is now a simple one-variable ordinary least squares model, so the coefficient estimate $\hat{\beta}$ can be expressed as a ratio of covariances:

$$\begin{aligned}\hat{\beta} &= \frac{\operatorname{cov}(\Delta z_t,\Delta q_t)}{\operatorname{var}(\Delta z_t)}=\frac{\operatorname{cov}(\Delta z_t,q_{1t}-q_{2t})}{\operatorname{var}(\Delta z_t)}\\ &= \frac{\operatorname{cov}\big(\Delta z_t,(\alpha_1+\beta z_{1t}+(1-\beta)z_{2t}+\varepsilon_{1t})-(\alpha_2+\beta z_{2t}+(1-\beta)z_{1t}+\varepsilon_{2t})\big)}{\operatorname{var}(\Delta z_t)}\\ &= \frac{\operatorname{cov}\big(\Delta z_t,(2\beta-1)z_{1t}-(2\beta+1)z_{2t}\big)}{\operatorname{var}(\Delta z_t)}=(2\beta-1)\frac{\operatorname{cov}(\Delta z_t,\Delta z_t)}{\operatorname{var}(\Delta z_t)}=2\beta-1\\ &= \beta-(1-\beta)\end{aligned}$$

which is not equal to β. Thus, year effects introduce a mechanical relationship such that the estimate $\hat{\beta}$ captures not only the correct effect of the endowment on the agent's own quantities (β), but also the effect of the endowment on the other agent's quantities ($1-\beta$). In autarky ($\beta=1$) there would be no market spillovers and $\hat{\beta}$ would be unbiased. In the no-inertia case of $\beta=0.5$, the estimated effect would be zero – falsely suggesting that raising endowments does not increase quantities.

C.2 Solution to the planner's problem has the same necessary conditions as the market equilibrium

First, expand the first term of the maximand and rearrange it:

$$\begin{aligned}\sum_j\int_{Q_j}^{Q_j^f}V_j(\varphi)d\varphi &= \sum_j\exp\Big(-\eta\psi_j\Big)\Big(\eta+1\Big)^{-1}\big((Q_j^f)^{\eta+1}-(Q_j)^{\eta+1}\big)\\ &= \sum_j\exp\Big(-\eta\psi_j\Big)\Big(\eta+1\Big)^{-1}(Q_j)^{\eta+1}\left[\Big(\frac{Q_j^f}{Q_j}\Big)^{\eta+1}-1\right]\\ &= \sum_j\exp\Big(-\eta\psi_j\Big)\Big(\eta+1\Big)^{-1}(Q_j)^{\eta+1}\left[\Big(1+\frac{-\sum_{k>j}q_{jk}+\sum_{k<j}q_{kj}}{Q_j}\Big)^{\eta+1}-1\right].\end{aligned}$$

Then, take a first-order condition with respect to q_{od} by setting the derivative of the entire maximand equal to zero (assume district o sells to district d, without loss of generality):

$$\begin{aligned}
0 &= \frac{d}{dq_{od}}\left\{\exp\left(-\eta\psi_o\right)\left(\eta+1\right)^{-1}(Q_o)^{\eta+1}\left[\left(1+\frac{-\sum_{k>o}q_{ok}+\sum_{k<o}q_{ko}}{Q_o}\right)^{\eta+1}-1\right]\right\} \\
&\quad +\frac{d}{dq_{od}}\left\{\exp\left(-\eta\psi_d\right)\left(\eta+1\right)^{-1}(Q_d)^{\eta+1}\left[\left(1+\frac{-\sum_{k>d}q_{dk}+\sum_{k<d}q_{kd}}{Q_d}\right)^{\eta+1}-1\right]\right\} \\
&\quad -\frac{d}{dq_{od}}\left\{\sum_{k>o}c_{ok}q_{ok}\right\} \\
&= \exp\left(-\eta\psi_o\right)\left(\eta+1\right)^{-1}(Q_o)^{\eta+1}\frac{d}{dq_{od}}\left(1+\frac{-\sum_{k>o,k\neq d}q_{ok}+\sum_{k<o}q_{ko}-q_{od}}{Q_o}\right)^{\eta+1} \\
&\quad +\exp\left(-\eta\psi_d\right)\left(\eta+1\right)^{-1}(Q_d)^{\eta+1}\frac{d}{dq_{od}}\left(1+\frac{-\sum_{k>d}q_{dk}+\sum_{k<d,k\neq o}q_{kd}+q_{od}}{Q_d}\right)^{\eta+1} \\
&\quad -c_{od}\frac{d}{dq_{od}}q_{od} \\
&= \exp\left(-\eta\psi_o\right)\left(\eta+1\right)^{-1}(Q_o)^{\eta+1}\left(\eta+1\right)\left(\frac{Q_o^f}{Q_o}\right)^{\eta}\frac{-1}{Q_o} \\
&\quad +\exp\left(-\eta\psi_d\right)\left(\eta+1\right)^{-1}(Q_d)^{\eta+1}\left(\eta+1\right)\left(\frac{Q_d^f}{Q_d}\right)^{\eta}\frac{1}{Q_d}-c_{od} \\
&= -\exp\left(-\eta\psi_o\right)\left(Q_o^f\right)^{\eta}+\exp\left(-\eta\psi_d\right)\left(Q_d^f\right)^{\eta}-c_{od}.
\end{aligned}$$

Next, rearrange the demand model $Q_j^f = \left(V_j^f\right)^{\eta_j} e^{\psi_j}$ and substitute it for the parameters ψ_o and ψ_d:

$$\begin{aligned}
0 &= -V_o^f\left(Q_o^f\right)^{-\eta}\left(Q_o^f\right)^{\eta}+V_d^f\left(Q_d^f\right)^{-\eta}\left(Q_d^f\right)^{\eta}-c_{od} \\
&= -V_o^f+V_d^f-c_{od}.
\end{aligned}$$

Rearranging, and splitting transaction costs into seller and buyer components ($c_{od} = \tau_{od}^s + \tau_{od}^b$), the first-order conditions are:

$$V_d^f - V_o^f = \tau_{od}^s + \tau_{od}^b \quad \forall o, d \text{ s.t.} \quad q_{od} > 0.$$

This is the no-arbitrage condition for the case of additive transaction costs.

C.3 Sum of district-specific gains equals the maximand

District-specific gains are defined as:

$$H_k \equiv \int_{Q_k^0}^{Q_k^f} \left[V_k(\varphi) - V_k(Q_k^f) \right] d\varphi. \tag{16}$$

I need to prove that the sum of district-specific gains in Equation 16 equals the maximand in Equation 14. The first term is identical in each expression, so it suffices to prove that the second terms are equal. Beginning with the second term of Equation 16 summed over all units j, I rearrange, switch indices twice, and expand:

$$\begin{aligned}
-\sum_j \int_{Q_j}^{Q_j^f} V(Q_j^f) d\varphi &= -\sum_j (Q_j^f - Q_j) V(Q_j^f) \\
&= -\sum_j \Big(-\sum_{k>j} q_{jk} + \sum_{k<j} q_{kj} \Big) V_j^f \\
&= \sum_j \sum_{k>j} q_{jk} V_j^f - \sum_j \sum_{k<j} q_{kj} V_j^f \\
&= \sum_j \sum_{k>j} q_{jk} V_j^f - \sum_k \sum_{j<k} q_{jk} V_k^f \\
&= \sum_j \sum_{k>j} q_{jk} V_j^f - \sum_j \sum_{k>j} q_{jk} V_k^f \\
&= \sum_j \sum_{k>j} q_{jk} (V_j^f - V_k^f) \\
&= \sum_j \sum_{k>j} q_{jk} \left[1(q_{jk} > 0) + 1(q_{jk} < 0) \right] (V_j^f - V_k^f) \\
& \quad \sum_j \sum_{k>j} q_{jk} \left[1(q_{jk} > 0)(V_j^f - V_k^f) + 1(q_{jk} < 0)(V_j^f - V_k^f) \right]
\end{aligned}$$

Inserting the first-order conditions from the previous proof (i.e., $V_d - V_o = c_{od}$ for all o and d such that $q_{od} > 0$):

$$\begin{aligned}
-\sum_j \int_{Q_j}^{Q_j^f} V(Q_j^f) d\varphi &= \sum_j \sum_{k>j} q_{jk} \left[1(q_{jk} > 0)(-c_{jk}) + 1(q_{jk} < 0) c_{kj} \right] \\
&= -\sum_j \sum_{k>j} q_{jk} \left[1(q_{jk} > 0) c_{jk} - 1(q_{jk} < 0) c_{kj} \right]
\end{aligned}$$

which is the second term of Equation 14.

D Online Appendix: Processing of Water Transactions Data

WestWater Research, LLC provided a dataset containing 6,264 water transactions in California between 1990 and 2015. Only a fraction of these transactions pertain to the surface water market; most are unrelated. Variables include date, volume, price, and duration of the transaction; type of water right the transaction is based on; and name and latitude and longitude of the origin and destination parties.

Cleaning. I calculate price per acre-foot and deflate to 2010 dollars using the CPI. I reshape the data so there is one observation per party per transaction, creating 12,975 observations. For transactions with multiple buyers or multiple sellers, if more specific information is not available, I divide transaction volume across parties in proportion to their maximum water rights.

Sample restrictions. To focus on transactions of actual surface water, I drop transactions of rights to pump groundwater within adjudicated basins, rights to store water in reservoirs or groundwater banks, and desalinated water. To focus on transactions in which the price is freely negotiated, I drop transactions within programs in which prices are set administratively. I drop one transaction that never cleared. This leaves 2,584 transaction-by-party observations from 1,104 unique transactions.

Location. I classify all observations into one of 10 hydrologic regions (defined by California's Department of Water Resources), and I generate latitude and longitude coordinates whenever possible. I first attempt to use centroids from my user location file, matching parties to users via my crosswalk file. This matches 1,827 transaction-by-party observations. Second, I manually geolocate 65 users that do not appear in the user location file but which are common in either transactions or water quantity data. For these users, I generate coordinates based on addresses, towns, or maps found via user websites and other publicly available documents; they match 222 additional observations. For remaining unmatched observations, I use location information from the original WestWater dataset; this assigns hydrologic region for all remaining observations and coordinates for 138 additional observations. This process leaves 398 observations for which hydrologic region is known but specific location coordinates are unavailable. Finally, I spatially join coordinates to 8-digit watershed (hydrological unit code, as defined by the U.S. Geological Survey), sub-sub-region (detailed analysis unit, as defined by DWR), and county. These shapefiles are available from DWR's California Water Plan: http://www.water.ca.gov/waterplan/gis/index.cfm.

Sector. I classify all parties into one of three sectors: agriculture, urban/municipal, or environmental. I use the first successful method in the following order of priority:

1. Water rights category (67% of observations). If the party appears in the surface water quantity dataset, I assign to agriculture or municipal depending on which sector holds a majority of the maximum water rights, including project contracts.

2. District associations and state and federal government agencies (20%). I classify as agricultural several known associations of agricultural water districts: San Joaquin River Group Authority, San Joaquin Water Conservation District, San Luis & Delta-Mendota Water Authority, Tehama-Colusa Canal Authority, Westside Districts, West Coast Basin Water Right Holders. I classify as environmental several state and federal agencies: California Dept. of Fish & Wildlife, California Dept. of Water Resources, Environmental Water Account, U.S. Bureau of Land Management, U.S. Bureau of Reclamation.

3. Keywords (9%). I classify users based on the following keywords in their name. Agriculture: almonds, citrus, contractors, dairies, dairy, famers, farm, farmers, farming, grower, irrigating, irrigation, nurseries, nursery, orchard, ranch, trust. Municipal: cement, cemetery, chemical, Chevron, church, city of, college, communities, community services, companies, company, corporations, developer, development, electric, energy, estate, gardens, golf, gravel, homeowners, homes, housing, inc., Indians, industries, investment, landscaper, leasing, LLC, LP, military, municipal, mutual water, non-ag, oil, owners, park, paving, power authority, properties, property, railway, real estate, realty, recycled, refining, replenishment district, retail, sanitation, school, Texaco, town of, tribe, university, ventures. Environmental: conservancy, duck hunting, ducks, fish & wildlife, forest service, forestry & fire prevention, water bank.

4. Original use categories (4%). I apply WestWater's original water use categories, based on agriculture, irrigation, and environmental, counting all other categories as urban.

5. Remainder (1%). I assume remaining users called water districts are agricultural, and that all observations remaining after that are urban.

Characteristics of delivery path. I use the hydrologic model to calculate characteristics of the delivery path from buyer to seller. When there is more than one counterparty on the opposite side of the transaction, I define the delivery path by choosing the counterparty nearest to the geographic center of all the counterparties. For instream (environmental) transfers, I use the same location for the party and counterparty.