What Holds Back Water Markets? Transaction Costs and the Gains from Trade

Nick Hagerty*

February 23, 2023

Abstract

This paper estimates the potential benefits of reducing transaction costs in California's wholesale surface water market. I develop an empirical framework to analyze welfare in water markets that uses transactions data, inferring preferences of water utilities from their behavior in the existing market. I separate observed prices into true valuations and transaction costs, estimate demand elasticities, and simulate a market without transaction costs. Gains from efficient trading across regions and sectors are less than 2% of statewide water expenditures. Reducing transaction costs may not achieve large gains without also reforming the policies and institutions that govern local water allocation.

*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, Rachael Meager, Ben Olken, Arianna Ornaghi, Doug Parker, Will Rafey, Gary Sawyers, Matt Zaragoza-Watkins, Ariel Zucker, 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). However, relatively little quantitative evidence is available to assess the magnitudes of these barriers and set priorities for reform.

In this paper, I quantify the importance of transaction costs in California's wholesale surface water market, the largest in the United States. I ask: What are the potential economic gains from eliminating transaction costs, holding other institutional features constant? To answer this question, I develop a new empirical framework to analyze welfare in water markets that relies on transactions data. The idea is to infer the preferences of market participants – which in this context are mostly water districts that supply water to retail customers – from their observed behavior in the market that already exists. I first separate observed transaction prices into transaction costs and true valuations. Then, I remove transaction costs, estimate demand using true valuations, and simulate a fully efficient market.

There are two reasons why California's water market is 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.

Two salient features of California's surface water market are high price dispersion and low transaction volume, relative to the number of water districts and other water users who might transact. To explain these facts as equilibrium outcomes, I model California's water market as an exchange economy with spatial arbitrage, in which water districts trade endowments of a homogeneous good.

[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.

Districts incur transaction costs that may be location pair-specific and directionally asymmetric. Transaction costs, defined broadly, may arise from many factors, some of which are observable (such as conveyance distance or a regulatory review) and others which are not (such as search or contracting difficulty). The model rationalizes observed trading outcomes with a set of demand curves and transaction costs, and it provides equilibrium conditions that I use to empirically recover these objects.

Then, I follow a four-step empirical procedure:

Step 1: **Estimate transaction costs arising from observable factors.** I directly estimate transaction costs associated with specific observable factors, such as a regulatory review or physical distance, using variation in prices within district and year. Many water districts sell to (or buy from) more than one other district in a given year, and often one transaction is subject to a particular cost determinant and another is not. In equilibrium, a seller is indifferent across buyers, so any difference in prices can be interpreted as marginal transaction costs.

Step 2: **Adjust prices to recover equilibrium marginal valuations of water.** I back out marginal valuations of water for all market participants using the condition that a buyer's marginal valuation must be at least as high as its highest price paid, and a seller's as low as its lowest price paid. I first adjust raw prices for *observed* determinants of transaction costs, as estimated in Step 1. Then, I adjust for *unobserved* determinants by taking the minimum (for sellers) or maximum (for buyers) across transactions within a year.

Step 3: **Estimate price elasticities of demand.** To extrapolate away from equilibrium, I estimate price elasticities of demand in the wholesale surface water market. I exploit California's historically-determined allocation rules, which translate a given rainfall shock into vastly different supply shocks for different water users. Because transaction costs lead each year's initial allocations to persist, I can measure how steeply marginal valuations respond to exogenous changes in quantity consumed.

Step 4: **Simulate an efficient market and calculate the gains from trade.** Marginal valuations and price elasticities together define demand curves for all districts. I use constrained optimization to find the efficient allocation without transaction costs and calculate the resulting gain in economic surplus. The objective is to equalize marginal valuations across regions and sectors, up to true physical transportation costs which can never be eliminated.

To conduct this empirical analysis, 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 mostly 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 hydrological model of California's water infrastructure to compute physical characteristics of water transfer routes.

In intermediate results, I first document large price gaps resulting from specific, observable cost determinants. For example, I find that transactions that must cross the Sacramento–San Joaquin Delta (an environmentally sensitive juncture triggering additional regulatory reviews) are associated with a price premium for sellers of \$76 per acre-foot, and a price discount for buyers of \$31. I interpret these as marginal transaction costs, totaling \$107 per acre-foot – which is large as compared with the mean price in my data, \$221. Several other observable factors also each result in similarly large marginal transaction costs. Then, I find that water districts have fairly inelastic demand for wholesale water transfers. I estimate that the price elasticity of demand is -0.10 for the urban sector and -0.23 for the agricultural sector, magnitudes that are smaller than previous estimates that use individual household- or farm-level data.

In my main results, I find that observed trading in the existing market achieves welfare gains of \$10 to \$88 million per year, depending on water availability conditions. I estimate that an efficient market would result in additional gains of \$86 to \$278 million per year – values that are meaningful but still small when compared with overall water-related expenditures in California. These results suggest that reducing transaction costs alone may not unleash large economic gains without also reforming other institutional aspects of California's water economy.

I discuss three leading explanations for the results. First, permanent water transfers to date may have already achieved many of the available gains. Second, many of the potential benefits from water markets may be local rather than statewide. Third, water districts may be trading more conservatively than their own constituents or customers would prefer. Simulations using individual-level demand elasticities find substantially larger gains, suggesting that important parts of the explanation may be water districts' internal governance, local allocation policies, and regulatory structure.

This paper makes several contributions. First, I provide a method to analyze welfare in water markets that uses transactions data from the existing market. A large literature uses calibrated optimization models to estimate the prospective gains from water markets such as those in 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). These studies typically combine spatially-explicit agricultural production models with detailed hydrological constraints to simulate the profit-maximizing water allocation. Such models are convenient but rely heavily on imputed parameters and functional form assumptions (Mérel and Howitt 2014); my approach is more parsimonious and enables attention to causal inference. More recently, Rafey (2023) empirically estimates an agricultural production function from farm-level data and uses it to value the ex post gains from water trading in Australia. Unlike all of these, my approach avoids the challenge of modeling fundamental determinants of water demand, by inferring the preferences of market participants directly from observed transactions. Even where fine-scale agricultural data is available, working directly with the implied objective functions of water districts makes my approach more policy-relevant so long as districts continue to be the primary market participants.

More generally, this paper proposes a method to analyze the welfare impacts of transaction costs in thin asset markets. This may be particularly relevant to other settings in environmental economics, such as pollution permits or individual transferable quotas. There is a large literature in financial economics on thin markets and liquidity, but it is typically focused on strategic trading rather than transaction costs (Kyle 1989; Rostek and Weretka 2012, 2015). The literature on pollution permit markets covers the theoretical effects of market power (Hahn 1984; Malueg and Yates 2009; Liski and Montero 2011) and transaction costs (Stavins 1995; Liski 2001), with some empirical analysis of transaction costs (Gangadharan 2000; Cason and Gangadharan 2003), but there are few empirical studies of the welfare impacts of transaction costs. In water markets, Carey et al. (2002) and Regnacq et al. (2016) study the effects of transaction costs on trading quantities but not welfare; Womble and Hanemann (2020) estimate transaction costs by surveying market participants; and Ayres et al. (2017) study transaction costs in groundwater management decisions.

This paper also relates to a literature in international trade that estimates trade costs from price gaps (Donaldson 2012; Atkin and Donaldson 2015; Bergquist 2016). Furthermore, it contributes to a broad and growing literature on the costs of misallocation, in settings such as housing (Glaeser and Luttmer 2003), capital (Hsieh and Klenow 2009), energy (Davis and Kilian 2011), labor (Bryan and Morten 2015; Adamopoulos et al. 2017), and land (Restuccia and Santaeulalia-Llopis 2017). Finally, my intermediate results contribute to a literature in agricultural economics on the value of water in irrigated agriculture (Schlenker et al. 2007; Buck et al. 2014; Mukherjee and Schwabe 2014).

2 Background on Water in California

2.1 California is a closed and interconnected hydrological system

Water is scarce in most of California, the largest economy and most populous state in the United States. A majority of the state's population lives in Southern California, where there is little rainfall and no major rivers. Farms in the Central Valley, a major agricultural region, receive little rainfall during the summer growing season and instead rely on irrigation. Most precipitation in the state falls during the winter in mountain ranges in the north and east.

Moving water throughout the state is technologically feasible, thanks to an interconnected system of water infrastructure that is the world's most complex. Federal, state, and local authorities operate canals and pipelines that, together with the river system, form a fully connected hydrological network among the vast majority of water users in the state of California. Although there are capacity constraints, at the margin it is possible to transfer water between nearly any two consumers in the state.

Very little of California's precipitation flows to other states or countries. The only major water source it shares with other states is the Colorado River, which is governed by long-term interstate compacts and international treaties that I treat as fixed.

2.2 Water is initially allocated by fixed rules and environmental conditions

Property rights to water in California are distributed not according to private or social value but instead following a complex system of historical precedent. Some districts and consumers hold entitlements that are almost never curtailed, while others are rationed according to precipitation and runoff during the previous winter. In my empirical exercises, I exploit this rationing to estimate demand elasticities and the marginal value of water in agriculture.

Figure 1 summarizes the structure of water use in California. All lawful surface water use in the state derives from a legal framework of appropriative and riparian rights. Some independent consumers (such as rural households or isolated farmers) hold their own water rights; others obtain water from federal or state water projects. More commonly, retail consumers (including farms, households, and other consumers) obtain water from their local water district, which in turn either holds its own water rights or long-term contracts with the federal or state water projects. Water districts go by many legal classifications, such as irrigation district, water agency, or mutual water company, and may be public or private. They may also have multiple layers, in which a wholesale district sells to retail districts.

Most year-to-year variation in water endowments occurs within the federal Central Valley Project (CVP) and California's State Water Project (SWP). Contractors within these projects do not buy water at market-clearing rates; rather, they hold a multi-decade contract that specifies a fixed maximum volume. Actual yearly allocations vary from year to year, determined by regulatory agencies on the basis of weather in the mountains during the previous winter and other environmental conditions and regulations. These allocations are specified as percentages of maximum contract volumes, determined separately for each of 14 contract categories that are based on history, geography, and sector. Some categories tend to have priority over others, but the ordering is neither constant (due to regional differences in water availability) nor lexicographic (like appropriative rights) (Stene 1995). The amount of water actually consumed can differ from allocations for several reasons: a water user can choose to take less than its allocation, bank water for temporary storage, apply to receive extra water under certain circumstances, or transfer water to another user.

The other sources of water are more stable over time. Water rights (appropriative and riparian) follow a seniority rule determined by the date of first use; in droughts, senior rights-holders are entitled to their full claim before junior rights-holders are entitled to any. However, it is rare for this seniority system to substantially affect water diversions in major rivers, since the residual claimants are generally the high-volume federal and state water projects. Another major source of water is the Lower Colorado system managed by the federal government; prior to 2023, contractors had always received their full contract amounts. The final major source of water is local groundwater. Groundwater use is unmonitored throughout most of the state, and availability and pumping costs vary considerably across regions. In this paper, I treat local groundwater supplies as fixed.

2.3 Transaction costs inhibit secondary markets

Because water is not allocated according to a price mechanism, a robust secondary market might be expected. A surface water market does exist in California, but it is thin. Figure 2 plots the total volume of market transactions over time, as compared with total water supply, using data described in Section 4.

I focus on the statewide water market, which I define as transactions directly among districts and other independent consumers at freely negotiated prices. This definition excludes transactions between wholesale districts and retail districts, and between retail districts and retail consumers. Such transactions take place within fixed, long-term relationships in which neither prices nor quantities may be flexible. It also excludes transactions involving retail customers within a water district.[2]

Many factors may make transactions in this market costly. I offer the following typology of transaction costs, building on Regnacq et al. (2016), Scheer (2016), and others.

Administrative transaction costs To trade water, a buyer or seller must first search for a potential trading partner. Without a central exchange, this happens mostly by word of mouth in social networks; sometimes a professional broker helps with matchmaking. There is no single standard contract for water transactions. The buyer and seller must negotiate over the quantity, duration, price, payment terms, delivery date, point of delivery, and delivery pathway. Transaction durations fall into two basic types: (1) permanent sales of water rights or contract entitlements, and (2) intra-year leases, in which the seller transfers a certain quantity of water while retaining the underlying entitlement. Together, search and contracting processes may create considerable administrative transaction costs, both explicit (i.e., attorney fees) and implicit (e.g., hassle costs), for both buyers and sellers.

Physical transaction costs Water is heavy; moving it from one place to another is costly. Water is lost in conveyance to evaporation and percolation. Pumping water uphill into canals requires energy to run turbines. Not all transactions incur these costs: upstream transfers may not incur any conveyance losses, while downstream transfers on a river may not incur any pumping costs. However, these pure physical costs may differ from the costs directly incurred by buyers or sellers. Transactors pay "wheeling" charges to the owners of the intermediate conveyance facilities along the delivery pathway, including canals, pumping stations, and reservoirs. Wheeling charges are not generally equal to the true physical marginal cost of conveyance and pumping; some stakeholders believe they are often substantially higher (Western Water Company 2000).

Regulatory transaction costs Proposed transactions can be subject to regulatory review by three main agencies: California's State Water Resources Control Board (SWRCB), California's

[2]Intra-district transactions between consumers are rare in urban water districts; they do occur in some irrigation districts, but data is scarce: even when districts keep records of these transactions, they rarely record prices. Retail consumers are not generally allowed or able to negotiate their own transactions with external districts; instead they must rely on their own district to represent their interests on the statewide water market.

Department of Water Resources (DWR), and the U.S. Bureau of Reclamation (USBR). Some transactions are reviewed by counties. Depending on the proposed source and destination, transactions may be reviewed by more than one of these agencies, or none of them.

In these reviews, agencies (1) require sellers to compile records proving they have legal entitlement to the water and physical ability to transfer it; (2) carefully estimate consumptive use (the amount not returned to the water system); (3) set up monitoring systems to verify sellers do not continue using water once transferred; (4) conduct extensive environmental impact analyses to meet requirements of the National Environmental Policy Act (NEPA) and the California Environmental Quality Act (CEQA); (5) estimate impacts on the local economy; and (6) schedule delivery for a time with available capacity (California State Water Resources Control Board 1999; California Department of Water Resources and U.S. Bureau of Reclamation 2015).

Transactions that move water across the Sacramento–San Joaquin Delta, the biggest bottleneck in the system, must meet additional environmental regulations concerning outflow volumes, salinity levels, and endangered fish. These transactions can be risky; because the DWR and USBR are sometimes not allowed to pump water into their canals, there is no guarantee that the seller's water will actually reach the buyer. In addition, these transactions are assessed "carriage losses" to satisfy environmental goals and regulatory constraints.

In sum, regulatory reviews may create substantial policy-induced transaction costs, which again may be explicit or implicit. Explicit costs include agency review fees and the engineering fees to prepare documents. Implicit costs include the hassle or general disutility from the process, the time costs of the review plus its public notice and public comment periods, and the risks of disapproval or delivery failure. Both buyers and sellers may bear these costs.

Political-economy transaction costs Implicit transaction costs may also arise from political economy effects. First, water districts imperfectly represent the interests of their retail customers, driving a wedge between statewide water prices and an individual farmer's or household's willingness to pay or to accept. This is especially relevant in districts controlled by popular vote rather than property value or land area. Second, farmers may be reluctant to sell water because they fear voters will think they don't need it and may take away their property rights in the future (Carey and Sunding 2001).

Many types of transaction costs may be pure loss to society, such as administrative costs. Some types may represent transfers, in the case of attorney fees. Others may have important benefits in preventing negative externalities. Environmental reviews protect public goods like wildlife and ecosystem services, and determining consumptive use protects the property rights of downstream users. Political economy constraints may help prevent pecuniary externalities to the origin community.

However, the evidentiary standards for regulatory approval of water transactions are much higher than for the use of the water in the first place. Many stakeholders have proposed reforms that could accomplish many of the same regulatory goals while streamlining the process (Western

Water Company 2000; Culp et al. 2014; Gray et al. 2015; Association of California Water Agencies 2016). In my analysis, I focus on the costs of these regulations, which then can be weighed against their benefits. In addition to the costs incurred for transactions that do take place, transaction costs also may prevent many mutually beneficial trades from occurring.

3 Theoretical Framework

In this section I lay out a model of California's water market to guide my empirical analysis. I first present a simplified graphical model to build intuition. I then model an exchange economy in which consumers trade endowments of a homogeneous good while incurring transaction costs. Finally, I derive my four-step empirical procedure.

3.1 Simplified graphical model

Figure 3 plots inverse demand curves for two districts d and o (for destination and origin), which might be thought of as an urban district d and agricultural district o. They have initial endowments of water E_o and E_d and inverse demand curves $V_d(Q_d)$ and $V_o(Q_o)$, which give their marginal valuations of water as a function of quantity demanded. The axis is reversed for district o and lined up such that the overall width of the graph is equal to the overall resource constraint (the sum of endowments) in the style of an Edgeworth box. If the two districts could trade costlessly with each other, they would arrive at an equilibrium allocation (Q_d^*, Q_o^*) and price P^*.

If instead trading involves transaction costs, prices will generally not equal marginal valuations. Suppose the districts incur per-unit transaction costs τ_{od}^b (to buy) and τ_{od}^s (to sell) for cross-district trades from o to d. Then, the buyer's marginal willingness-to-pay is shifted down relative to its marginal valuation, and the seller's marginal willingness-to-accept is shifted up. The buyer requires a price discount equal to its transaction costs, and the seller a price premium, in order to be willing to transact. In equilibrium, the sum of these transaction costs is exactly equal to the price gap:

$$\underbrace{V_o + \tau_{od}^s}_{Seller\ MWTA} = P_{od} = \underbrace{V_d - \tau_{od}^b}_{Buyer\ MWTP} . \tag{1}$$

If a pair of districts trades, the price equalizes the seller's marginal willingness-to-accept with the buyer's marginal willingness-to-pay, net of transaction costs. This no-arbitrage condition is the foundation of my empirical analysis.

If we know the true demand curves, we can calculate the welfare gains from eliminating transaction costs. The change in surplus, $\Delta CS_d + \Delta CS_o$, is made up of (1) a rectangle representing the gains from reducing costs in transactions that already occur, and (2) a Harberger triangle representing the gains from transactions that become newly profitable.

3.2 Model of an exchange economy with transaction costs

I now extend the insights above to more than two districts and clarify the precise assumptions required. Consider N water districts indexed by n (or by o for origin and d for destination). Each district has initial endowment E_n of a single homogeneous good (water) that is allocated efficiently among a continuum of consumers. Consumers' preferences can be aggregated such that each district has an inverse demand function $V_n(Q_n)$, which gives marginal valuations as a function of quantity consumed Q_n. Inverse demand is decreasing and twice differentiable.

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 $V_o(Q_o)$) 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 $V_d(Q_d)$). Sellers and buyers meet at exchange points unique to each pair of districts, where prices P_{od} are determined. Each transaction i generates transaction costs for both sellers and buyers:

Assumption 1. Per-unit transaction costs. *To complete a transaction of q_{iod} units of water, sellers and buyers each incur constant marginal transaction costs. They are non-negative and may be specific to each origin-destination district pair and directionally asymmetric:*

$$\begin{aligned} &(Sellers) \qquad C^s_{od}(q_{iod}) \equiv \tau^s_{od} q_{iod} \\ &(Buyers) \qquad C^b_{od}(q_{iod}) \equiv \tau^b_{od} q_{iod}. \end{aligned} \tag{2}$$

These transaction costs capture not only physical transportation costs and legal contracting costs but also factors such as hassle costs, regulatory costs, and political pressures. Their flexible specification reflects that costs vary not only across buyers and sellers, but also by who they choose to transact with, and in which direction. Fixed costs of trading are ruled out, an assumption that is strong but crucial to make the empirical analysis tractable. It is an approximation that still captures many important aspects of California's water market: some percentage of water is lost in conveyance, larger transactions receive greater regulatory scrutiny, owners of canals and pipelines often charge per-unit fees ("wheeling" charges), and larger transactions may generate more political backlash.

Assumption 2. Perfect competition. *Each district has enough intermediaries that they behave as price takers. That is, the quantity sold or purchased by any one intermediary does not affect the equilibrium price for any district pair: $dP_{od}/dq_{iod} = 0$ for all o and d.*

This assumption is also strong but is less crucial. In Appendix A, I extend the model to allow for market power and conduct my empirical analysis allowing for markups or markdowns by either buyers or sellers. There, I find little evidence that market power explains the observable transaction costs that I estimate below. These results are consistent with a separate and contemporaneous analysis by Tomori et al. (2022).

Sellers and buyers choose non-negative quantities for each origin and destination pair od to maximize net profits:

$$\begin{aligned} &(Sellers) \quad \max_{q^s_{iod}} P_{od} q^s_{iod} - V_o(Q_o) q^s_{iod} - \tau^s_{od} q^s_{iod} \quad \text{s.t. } q^s_{iod} \geq 0 \\ &(Buyers) \quad \max_{q^b_{iod}} V_d(Q_d) q^b_{iod} - P_{od} q^b_{iod} - \tau^b_{od} q^b_{iod} \quad \text{s.t. } q^b_{iod} \geq 0 \end{aligned} \tag{3}$$

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: $V_d(Q_d) - \tau^b_{od} \geq V_o(Q_o) + \tau^s_{od}$. Second, if buyers and sellers do trade, first-order conditions yield the no-arbitrage condition in Equation 1. 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 transaction costs that each incurs.

Trading quantities traded are defined implicitly by combining the inverse demand curves $V_o(Q_o)$ and $V_d(Q_d)$ with the market clearing condition: the sum of quantities sold by sellers must equal the sum of quantities bought by buyers. Intuitively, in partial equilibrium (holding constant transactions in all other pairs of districts), sellers and buyers increase quantities until the marginal surplus is arbitraged away and the first-order conditions are met. In general equilibrium, quantities adjust so that each district's marginal valuation is equalized across all of the other districts it sells to or buys from.

3.3 From theory to estimation

I now derive my empirical procedure. Although neither marginal valuations nor transaction costs are directly observed, I use theoretically-grounded restrictions on prices to identify transaction costs and back out marginal valuations.[3]

Step 1: Estimate transaction costs from observable factors.

Although the overall levels of transaction costs τ^s_{od} and τ^b_{od} are unobserved, portions of these costs may be due to observable cost determinants such as conveyance distance or specific regulatory reviews. Since many of these determinants vary across transactions, I can estimate their costs using variation in prices across transactions within district and year. In equilibrium, a district selling to two buyers is indifferent between them at the margin, so any difference between them can be interpreted as marginal transaction costs. For example, suppose I observe Redding selling both to Sacramento at \$50 and to Los Angeles at \$70. If the sale to Los Angeles undergoes a regulatory

[3] In principle, marginal valuations could be observed from prices in intra-district water markets. Unfortunately, these internal prices are difficult to obtain, and most water districts do not keep systematic records. Observations in my data are cross-district transactions.

review that the sale to Sacramento does not, I can interpret the price difference of \$20 as the cost of this review to Redding.

To formalize this intuition, I make a further assumption on the functional form of transaction costs. At this point I introduce time subscripts t, allowing all variables and parameters to vary arbitrarily across time periods.[4]

Assumption 3. Determinant-specific transaction costs are additively linear and constant across transactions. *Transaction costs can be decomposed as* $\tau^s_{odt} = \tau^s \mathbf{B}_{od} + \tilde{\tau}^s_{odt}$ *(for sellers) and* $\tau^b_{odt} = \tau^b \mathbf{B}_{od} + \tilde{\tau}^b_{odt}$ *(for buyers), where* $\mathbf{B}_{od}$ *is a vector of observable transaction cost determinants and* τ^s *and* τ^b *are vectors of coefficients, and all parameter and cost determinant values are non-negative.*

This assumption carries three substantive restrictions: first, transaction costs arising from specific cost determinants $\mathbf{B}_{od}$ are constant across districts with the same values of $\mathbf{B}_{od}$; second, they are constant across time; and third, cost determinants do not interact with each other – when a transaction is subject to multiple cost determinants, the total costs equal the sum of their parts.[5]

Combining Assumption 3 with Equation 1 yields equations that can be estimated with regression:

$$\begin{aligned} (Sellers) \qquad P_{odt} &= \alpha_{ot} + \tau^s \mathbf{B}_{od} + \varepsilon^s_{odt} \\ (Buyers) \qquad P_{odt} &= \alpha_{dt} + \tau^b \mathbf{B}_{od} + \varepsilon^b_{odt}, \end{aligned} \tag{4}$$

where the α's collect terms fixed within district and year and the ε's collect terms that are mean-zero by construction, which I treat as econometric errors.[6]

Step 2: Adjust prices to recover equilibrium marginal valuations.

Next, I recover equilibrium marginal valuations by adjusting prices for both known and unknown determinants of transaction costs. This step is essentially a bounding exercise: Marginal valuations do not vary across transactions, so any price variation within district and year is due to transaction costs and can be eliminated. First, I adjust prices to remove transaction costs due to *observed* determinants, as estimated in Step 1. Second, I adjust prices to remove transaction costs due to *unobserved* determinants, by taking the minimum price paid (for sellers) or maximum price accepted (for buyers) in each year.

Since prices equal marginal willingness to pay or to accept (Equation 1), and all components of transaction costs are non-negative (Assumption 3), marginal valuations are bounded by observed prices after adjusting for observed cost determinants: $V_{ot} \leq P_{odt} - \tau^s \mathbf{B}_{od}$ for all sellers d, and $V_{dt} \geq P_{odt} + \tau^b \mathbf{B}_{od}$ for all buyers o. I then assume that these bounds are "good enough," which allows me to interpret marginal valuations as point estimates rather than bounds:

[4] In my empirical analysis I define a time period as one year; this is a reasonable approximation since almost two-thirds of observed transactions are arranged in the same three-month period of May through July.

[5] In principle I could relax either the first or the second restriction, but not both; for τ^s and τ^b to be identified, they must be constant in at least one dimension. The third restriction in principle could be relaxed nonparametrically.

[6] That is, $\alpha_{ot} \equiv V_{ot} + \mathbb{E}_d[\tilde{\tau}^s_{odt}]$; $\alpha_{dt} \equiv V_{dt} + \mathbb{E}_o[\tilde{\tau}^b_{odt}]$; $\varepsilon^s_{odt} \equiv \tilde{\tau}^s_{odt} - \mathbb{E}_d[\tilde{\tau}^s_{odt}]$; $\varepsilon^b_{odt} \equiv \tilde{\tau}^b_{odt} - \mathbb{E}_o[\tilde{\tau}^b_{odt}]$.

Assumption 4. Unobserved transaction costs are zero for at least one transaction. *Unobserved transaction costs $\tilde{\tau}^s_{odt}$ and $\tilde{\tau}^b_{odt}$ are zero for the transaction in each district and period with the lowest (for selling districts) or highest (for buying districts) adjusted price.*

This assumption ensures that the bounds hold with equality for at least one transaction per district and period:

$$
\begin{aligned}
(Sellers) \qquad V_{ot} &= \min_d \{P_{odt} - \tau^s \mathbf{B}_{od}\} \\
(Buyers) \qquad V_{dt} &= \max_o \{P_{odt} + \tau^b \mathbf{B}_{od}\}.
\end{aligned} \tag{5}
$$

To the extent this assumption is not true, I will understate the dispersion in marginal valuations. Because more dispersion means more mutually beneficial transactions, my estimates of the potential gains from trade would likely be a lower bound. However, this assumption may not be unreasonable, since my empirical estimates of τ^s are large enough to put some sellers' marginal valuations V_{ot} near zero.

Step 3: Estimate demand curves.

Having obtained equilibrium marginal valuations, I combine them with data on quantities to estimate price elasticities of demand. By combining marginal valuations, quantities, and price elasticities, I construct inverse demand functions $V_n(Q_n)$ for each district that allow me to extrapolate away from the observed equilibrium.

Step 4: Simulate counterfactual allocations.

Finally, I use the demand curves to simulate the efficient allocation without transaction costs, and calculate the resulting gain in surplus. I do so by solving the social planner's problem, a constrained optimization exercise in which I choose Q_n for each district n to maximize total surplus, subject to the resource constraint that the sum of final quantities is equal to the sum of initial endowments. I analyze multiple scenarios that add various constraints. Since purely physical costs of conveyance can never be eliminated in a real water market, most scenarios include constraints representing these costs as estimated from my hydrological network model.

4 New Data on California's Water Economy

4.1 Water transactions

I use a proprietary dataset on water transactions 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 impact analyses. No government agency or other institution maintains a centralized listing of all water transactions in California.[7]

I focus on (1) surface water transactions as opposed to groundwater, (2) the spot market (within-year leases) as opposed to permanent transfer of rights, (3) freely negotiated transactions as opposed to those where prices are set administratively, and (4) transactions involving at least one party that diverts water for consumption, as opposed to instream uses or storage. The WestWater dataset includes 6,264 transactions in total, but most are groundwater leases within adjudicated basins. Prices are available for 671 transactions that remain after applying these inclusion criteria. Table 1 shows summary statistics. Panel A shows that the distribution of volumes is highly dispersed and the mean price is $221 per acre-foot, in 2010 dollars. Panel B shows that the Sacramento River hydrologic region is the greatest net exporter and the Tulare Lake and South Coast regions are the greatest net importers. Figure 4 shows the distribution of prices on a base-10 logarithmic scale; transactions are centered around $100-300 but have substantial mass in the tails.

I identify the geographic location and sector (urban/municipal, agricultural, or environmental) of most buyers and sellers in the data, using several methods described in Appendix D.

4.2 Water quantities and allocation percentages

I assemble the universe of wholesale surface water deliveries in California, by user, sector, and year, from 1993 through 2016.[8] By deliveries I mean the volume of surface water actually received diverted. There are four sources of deliveries: the Central Valley Project (CVP), the State Water Project (SWP), the Lower Colorado Project, and surface water rights. For the CVP, SWP, and Lower Colorado, deliveries come from archives of the DWR and USBR. For deliveries of surface water rights, I use diversion reports collected by the SWRCB.[9]

Yearly allocation percentages from the CVP and SWP come from archives of the DWR and USBR. For surface water rights and Lower Colorado entitlements, allocation percentages are always 100 percent. To link users 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). To identify the locations, boundaries, and areas of water users, I combine several publicly available shapefiles into a single geospatial dataset.

Details of sources, cleaning, and processing of all these datasets are described in Hagerty (2022).

[7]One water transactions dataset has been assembled and made publicly available by Gary Libecap at UC Santa Barbara (e.g., Brewer et al. 2008); however, it has only a fraction of the transactions in California as the WestWater dataset. Another dataset has been assembled by Ellen Hanak at the Public Policy Institute of California, but it is not public. This dataset also appears to focus on transaction volume rather than prices (Hanak and Stryjewski 2012), whereas prices are mostly complete in the WestWater dataset.

[8]Data is publicly available at github.com/hagertynw/data-surface-water.

[9]I use time-averaged self-reported diversions rather than the face value of rights because the latter are frequently outdated for post-1914 appropriative rights and are not recorded for pre-1914 or riparian rights. This compilation uses recently available data made possible by a law that required all surface water rights holders to report their water use starting in 2010. It is reasonable to treat these reported diversions as the full value of present water rights due to the "use it or lose it" aspect of appropriative water rights.

4.3 Spatial aggregation

To estimate demand curves and simulate counterfactuals, I aggregate data to categories defined by geography and sector, which I call "units". Units are the unique combination of planning areas (46 geographical regions defined by the DWR) and sectors (urban or agricultural).

Aggregation serves two purposes. First, it ensures accurate matches between transactions and water quantities data; district-level analysis would instead yield noisy and nonsensical results.[10] Second, unit-level analysis simplifies the analysis by allowing me to avoid modeling selection into the water market. Individual district-level demand is difficult to estimate because I observe prices only for years in which a district chooses to trade. Instead of modeling both the intensive and extensive margin of trading, I estimate aggregate demand curves that subsume both margins. Unit-level marginal valuations can be viewed as the marginal valuation of the marginal district.

4.4 Hydrologic network model

I construct a model of California's hydrological network to calculate characteristics of transaction conveyance paths, and to allow for physical transaction costs in counterfactual simulations. This is a set of nodes and edges corresponding to all major water conveyance channels in California: rivers, canals, aqueducts, and pipelines. Routes 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), and conveyance losses (to percolation, evaporation, and required outflows from the Sacramento–San Joaquin Delta). I run a graph-theory algorithm to obtain least-cost delivery pathways for each unique pair of units. As compared with more detailed hydrologic models such as CALVIN (e.g., Howitt et al. 1999), this model lacks comprehensive information on capacity constraints.

5 Step 1: Estimating Transaction Costs from Observable Factors

The first step in my empirical approach is to estimate marginal transaction costs that arise from observable determinants.

5.1 Selecting cost determinants

To choose cost determinants for estimation, I start with the list compiled in Section 2.3. For some of these cost determinants, there are sellers and buyers for whom the determinant is incident on some counterparties but not others. From this list, I select the cost determinants that are both observable and heterogeneous across counterparties, and therefore have econometrically identified marginal

[10] This is often because a wholesale district receives water deliveries but does not appear in transactions data, while a geographically-overlapping retail district has no water deliveries yet is recorded as selling water. This is because wholesale and retail water districts often have long-term purchasing agreements and other complex linkages, which are difficult to track. Aggregating to unit simplifies such situations, since it treats districts that share jurisdictions together as one agent.

transaction costs. Cost determinants incurred equally in all of a district's transactions (such as the costs of writing a contract, or disutility from market participation) do not, because they are indistinguishable in price data from a level shift in marginal valuations. Table 2 lists and defines the cost determinants that meet these criteria.

5.2 Econometric specification

I stack the selected cost determinants into a vector $\mathbf{B}_{od}$, depending on seller (origin) o and buyer (destination) d. Most of these cost determinants are discrete so I use binary indicator variables; distances and elevation enter linearly. To identify marginal transaction costs nominally incident on sellers, I regress price P_{jodt} in transaction j in year t on this vector $\mathbf{B}_{od}$ and seller-by-year fixed effects, following Equation 4 in Section 3.3. To identify costs incident on buyers, I use separate regressions with buyer-by-year fixed effects.

The seller regression measures price gaps across transactions within seller and year. It isolates cases in which a seller completed transactions with two different buyers in the same year, of which one was subject to a particular cost determinant and the other was not. Each coefficient τ_h^s estimates the price premium paid to sellers to compensate for the hth cost determinant indexed in $\mathbf{B}_{od}$. The buyer regressions have exactly mirrored interpretations: τ_h^b measures differences in prices paid by a buyer to different sellers.

For unbiased estimates, I also must assume selection on observables: that unobserved determinants of prices are uncorrelated with the cost determinants $\mathbf{B}_{od}$. In the special case in which all unobserved components of transaction costs are equal across counterparties, this condition follows immediately from my model. In the more general case in which transaction costs vary arbitrarily across transactions, this assumption requires that, conditional on seller and year, the selected cost determinants are uncorrelated with other unobserved price determinants. In robustness checks, I partially test this assumption by adding covariates.

Because my sample of transactions is not large, I avoid overfitting the model by using the least absolute shrinkage selection operator (LASSO) to perform variable selection (Belloni et al. 2014). Many cost determinants are ex ante plausible, yet not all may necessarily be empirically important; others will be zeroed out. I cluster standard errors in all regressions at the level of subregion-by-year to allow for local spatial correlation. In some specifications, I use coarser fixed effects to explore robustness to specification.

5.3 Results

Table 3, Panel B shows the marginal transaction costs associated with factors selected by LASSO. Columns (4) and (8) are my preferred specifications, with seller-by-year and buyer-by-year fixed effects. Several of these factors are indeed costly. For example, prices received by sellers are \$76 per acre-foot higher for transactions that must cross the Sacramento–San Joaquin Delta, relative to other transactions by the same seller in the same year. Similarly, prices paid by buyers for transactions that export water from a federal or state water project are \$194 per acre-foot lower than other transactions. These price gaps are large, considering the mean price in the sample is \$221.

Other factors costly to sellers are distance conveyed in rivers and having an agricultural buyer. Factors costly to buyers are distance of virtual conveyance and review by the State Water Boards. Not all coefficients are statistically distinguishable from zero, but the fact they were selected by LASSO suggests they matter for the overall model. The importance of some of these factors is also supported by the specifications with coarser fixed effects (columns 2-3 and 6-7), in which many of the same factors are selected, and the Delta-crossing coefficient is larger and more precise. Across these specifications, all coefficients selected by LASSO are positive for sellers and negative for buyers, offering internally consistent evidence that these regressions are measuring the expected marginal transaction costs.

Results without variable selection are given in Panel A of Table 3. Columns 1 and 5 show coefficients from separate regressions in which the vector $\mathbf{B}_{od}$ includes only one known cost determinant at a time. Columns 2-4 show the results from regressions in which all known cost determinants are included together. Specifications are otherwise identical to the corresponding columns of Panel B. When including all proposed cost determinants, many coefficients are not statistically distinguishable from zero, and some have the opposite sign as expected. Still, the overall patterns hold up and remain broadly consistent across specifications. Buyers require price discounts, and sellers require price premiums, to choose transactions subject to these cost determinants.

6 Step 2: Recovering Marginal Valuations of Water

Second, I recover marginal valuations of water for each district and year by adjusting prices for both observed and unobserved determinants of transaction costs. To do so, I follow Equation 5, first adjusting observed prices for transaction costs estimated in Step 1 and then taking the minimum (for sellers) or maximum (for buyers) of adjusted prices within each district-year cell.[11] I classify districts as either buyers or sellers in each year on the basis of net transaction volume (the vast majority only buy or sell, not both).

Figure 5 plots the kernel density of the resulting marginal valuations, with the density of raw prices for comparison, for the district-years in which at least one price is observed. Marginal valuations have even greater dispersion than observed prices, even though prices are quite dispersed themselves. While marginal valuations are not yet directly comparable to each other,[12] the dispersion suggests that gains are available from reducing transaction costs.

7 Step 3: Estimating Demand Elasticities

Next, I estimate unit-level price elasticities of demand for surface water in the wholesale market.

[11] I impose free disposal by censoring at zero when occasionally necessary.

[12] The graph pools multiple years, districts do not appear in the data every year, and districts select into the observed market according to their marginal valuations

7.1 Econometric specification

I regress the log of equilibrium water quantities for unit k in year t on log prices of transactions j (made by district n with a counterparty in unit l). The simplest specification contains unit fixed effects; the preferred specification includes agent fixed effects, unit-pair fixed effects (to absorb differences in transaction costs across transactions within a unit), and unit-specific time trends (to control for changes in demand over time):

$$\ln Q_{kt} = \eta \ln P_{jnklt} + \psi_{nk} + \zeta_{kl} + \theta_k t + \nu_{jnklt}. \tag{6}$$

To overcome the joint determination of prices and quantities, I instrument for quantities using yearly water allocation percentages. These percentages are determined by precipitation in distant mountains during the previous winter and other regulatory considerations, making their year-to-year variation exogenous to demand. They are relevant instruments because transaction costs create inertia in initial endowments – only in a frictionless market would Coasian independence hold. Figure 7 plots allocation percentages aggregated to several regional categories.

My instruments are the full set of interactions between each allocation-percentage time series and unit indicator variables. These interactions allow the primary determinants of each unit's endowments to have separate effects on every other unit. For example, one unit's endowment may have a strong effect on its own outcomes, a moderate effect on the outcomes of a unit with which it frequently trades, and no effect on other units. These interactions yield a large number of instruments relative to the number of observations.[13] To avoid model overfitting and weak instrument problems, 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). IV Lasso uses data-driven penalization to choose an optimal subset of instruments.

I omit year fixed effects for two reasons. One is precision; most of the useful variation in prices and quantities is temporal. The other is that year effects introduce bias by comparing across units within the same water market. In this setting, 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). Appendix C.1 offers a proof of the bias from year effects within an interconnected market; the intuition is that they difference out statewide average prices and quantities, but these averages are themselves endogenous. Although I can no longer control for unobserved statewide shocks to demand, the most important time-varying factors are related to water availability, which is flexibly captured by the instruments.

The regression uses prices instead of marginal valuations in order to use all information available in the data, and to avoid estimation error from the previous steps. Because observed prices differ

[13] There are 13 series of allocation percentages (each corresponding to a different project contract type), 9 regions, and 62 units, giving a total of 923 potential instruments. The contract types are: SWP Agricultural, SWP Municipal, CVP North of Delta Agricultural, CVP North of Delta Urban, CVP North of Delta Settlement Contractors, CVP American River Urban, CVP In Delta (Contra Costa), CVP South of Delta Agricultural, CVP South of Delta Urban, CVP South of Delta Settlement Contractors, CVP Eastside Division, CVP Friant Class 1, and CVP Friant Class 2.

from marginal valuations only by the amount of transaction costs, it suffices to control for the determinants of transaction costs across transactions within a unit. (In robustness checks, I use marginal valuations instead of prices and find that the results are not significantly different.)

Although the level of observation is transaction-agent, the regression estimates unit-level elasticities, as quantities are measured at the unit level. Accordingly, I cluster standard errors by unit-year, the level of variation in both my instruments and outcome variable. Because only the endogenous variable varies within unit-year, omitting fixed effects would make this regression equivalent to one that uses unit-year observations and mean prices.

Identification assumptions This regression has two key identification assumptions. The first is conditional independence: conditional on covariates, changes in allocation percentages are not correlated with any other time-varying factors that independently affect prices or quantities. The second is the exclusion restriction: changes in allocation percentages affect quantities 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, one omitted variable might be local weather patterns, 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. Intuitively, 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.[14]

[14] A decrease in surface water endowments may lead a unit to extract more groundwater, reducing surface water quantity 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 surface water quantities to prices, net of any shifts in groundwater extraction.

7.2 Results

Table 4 reports the results of these regressions. As a first check that allocation percentages are indeed strong instruments, I estimate first stage and reduced form regressions by ordinary least squares using a single instrument: own-unit allocation percentage. The results are shown in Panel A for illustrative purposes; they are not used for further analysis because they do not account for the possible effects of other units' endowments. Columns 1 and 4 use the simplest specification with only unit fixed effects, columns 2 and 5 include additional terms, and columns 3 and 6 use the preferred specification (Equation 6).

Across all specifications, the instrument appears to be relevant and strong. Effects are estimated fairly precisely and point in the expected directions: higher allocations decrease prices and increase quantity consumed. F-statistics exceed the standard rule-of-thumb values, although not by a large amount in the first stage; including additional Lasso-selected instruments will help. The much larger F-statistics for the reduced form are also reassuring, since in just-identified two-stage least squares the roles of the first stage and reduced form can be swapped: the elasticity is estimated by one ratio, while the inverse elasticity can be estimated by its reciprocal.

Panel B shows results of instrumental variable regressions, in which the IV Lasso algorithm chooses instruments from the full set of potential instruments. In column 1, the simplest specification with only unit fixed effects, the Lasso selects 10 instruments, and the instrumented effect of log prices on log quantities is -0.14. This result implies that a 10% change in price results in a 1.4% change in quantity. Columns 2-6 report results from specifications that include additional sets of fixed effects and covariates. In the preferred specification (column 6), the Lasso selects 19 instruments, and the estimated elasticity is -0.20. Point estimates are relatively stable across specifications, despite the Lasso selecting as many as 34 or as few as 7 instruments. Although all first-stage F-statistics are large, I also report the results of the sup-score test for weak instruments recommended by Belloni et al. (2012). Two of the six specifications fail to reject the null hypothesis, indicating they do not pass the test. However, this test is known to be conservative, and the point estimates from these specifications are not significantly different from the others.

Panel C explores heterogeneity in demand elasticities. Each column estimates the preferred specification for a different subset of the sample. I find that urban units have a smaller elasticity (-0.10) than agricultural units (-0.23). Breaking up the agricultural sector into three regions, I find substantial geographic heterogeneity, with elasticities increasing from north to south. Elasticities are very small in the Sacramento Valley (-0.07), still somewhat small in the San Joaquin Valley (-0.11), and fairly large in the Tulare Lake Basin (-0.80). Appendix B explores robustness to alternative specifications.

These are the first estimates of price elasticities of demand in the wholesale surface water market in California. They tend to be smaller in magnitude than prior estimates for retail customers. For example, my elasticity of -0.10 for urban water demand is considerably smaller than an estimate for Riverside County (Baerenklau et al. 2014, -0.76) and the mean in a meta-analysis of urban water elasticities across the US (Dalhuisen et al. 2003, -0.41). My elasticity of -0.23 for agricultural water demand is considerably smaller than the mean in a meta-analysis across the US (Scheierling et al. 2006, −0.48).[15] In contrast, my estimates are more similar to prior wholesale-level estimates; for example, my elasticity of −0.10 for urban water demand is similar to a previous finding for urban agencies in California (Buck et al. 2016, −0.14).

8 Step 4: Simulating the Gains from Trade

Finally, I simulate counterfactual water market scenarios without transaction costs and calculate the resulting gains from trade.

8.1 Defining baseline scenarios

Because the gains from trade may be very different under different environmental conditions, I construct three baseline scenarios to carry forward into welfare analysis and counterfactual simulations: a wet year, a median year, and a dry year. For water quantities in these scenarios, I calculate each unit's mean within its least water-scarce, middle, and most water-scarce quintile of annual quantities between 1993 and 2015. For marginal valuations in each scenario, I compute fitted values from an inverted Equation 6, assigning each unit its corresponding elasticity from Table 4, Panel C.[16]

Figure 6 maps these fitted marginal valuations for the median-year scenario.[17] Despite wide dispersion in marginal valuations across individual water districts, there is less dispersion in marginal valuations across geographical regions in a typical year. There are few broad patterns: with the exception of lower marginal valuations in much of urban Southern California, marginal valuations fall within \$100-300 across most of the state. More familiar patterns emerge in the dry-year scenario (shown along with the wet-year scenario in Figure B1). Here, marginal valuations are high in much of urban Southern California, moderately high in the San Joaquin Valley and Tulare Basin regions, and generally low in the mostly-agricultural Sacramento Valley.

8.2 Surplus from the observed market

I first calculate the economic surplus achieved to date by observed transactions across units. This surplus for each unit k is calculated as the difference between the demand curve and the equilibrium marginal valuation (i.e., price net of transaction costs), integrated between initial endowment E_k and

[15] Two more recent estimates are more similar (Bruno 2017, −0.19; Hendricks and Peterson 2012, −0.10), but they are not as comparable because they are specific to groundwater demand.

[16] For the small number of agricultural units that fall outside of the three named regions, I apply the overall agricultural estimate.

[17] Each polygon in the map represents a planning area, which contains two units, urban and agricultural. To generate the map, I average marginal valuations over the two units and weight by quantity consumed. (I also drop three units with very small quantities that may be mismeasured, by the criterion that their transaction volume is more than twice their quantity consumed. Including these units would not substantively change the results, since their quantities sum to 0.02 percent of total statewide quantity.)

baseline quantity Q_k^0:

$$Surplus_k(E_k, Q_k^0) \equiv \int_{E_k}^{Q_k^0} \Big[V_k(\varphi) - V_k(Q_k^0) \Big] d\varphi. \tag{7}$$

Initial endowments are calculated by subtracting mean observed net purchase volume (from the relevant weather scenario) from baseline quantity. The integral of the second term evaluates to $(Q_k^0 - E_k)V_k^0$, where $V_k^0 \equiv V_k(Q_k^0)$ is the fitted marginal valuation. For the first term, I assume an isoelastic functional form; the integral of inverse demand can then be evaluated analytically:

$$\int_A^B V_k(\varphi) d\varphi = \exp\Big(-\frac{\psi_k}{\eta_k}\Big)\Big(\frac{1}{\eta_k}+1\Big)^{-1}\Big(B^{\frac{1}{\eta_k}+1} - A^{\frac{1}{\eta_k}+1}\Big) \tag{8}$$

where η_k is the unit's price elasticity of demand and ψ_k is the intercept of the log-log demand model from Equation 6. Transaction costs are not explicitly included in the surplus calculation because marginal valuations already take them into account.[18]

8.3 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 relative to baseline represents the potential gains from trade. To capture inescapable physical transaction costs, I calculate pair- and direction-specific conveyance costs in my hydrological network model and include them as costs in the objective function.

The social planner chooses the vector of bilateral transaction quantities q_{kl} (net volumes delivered from unit k to unit l, for all pairs of units) that maximizes aggregate valuation of water net of physical transaction costs, subject to the resource constraint.[19] Aggregate valuation is the area under each unit's demand curve between baseline quantities Q_k^0 and the final quantity Q_k^f, summed over units. Each unit's final quantity is its baseline quantity plus net transaction quantities. Physical transaction costs are directionally asymmetric: $c_{kl}q_{kl}$ for a positive transaction quantity (net delivery from k to l) and $c_{lk}q_{kl}$ for a negative transaction quantity (net delivery from l to k), where c_{kl} is the physical marginal cost of delivering water from k to l. The resource constraint requires that final quantities be nonnegative, meaning that no unit can sell more than its baseline quantity. Together, the full optimization problem is:

$$\max_{\{q_{kl}\}_{k,l>k}} \sum_k \int_{Q_k^0}^{Q_k^f} V_k(\varphi) d\varphi - \sum_k \sum_{l>k} \Big[1(q_{kl}>0)c_{kl} - 1(q_{kl}<0)c_{lk} \Big] q_{kl} \tag{9}$$

[18] Surplus can equivalently be expressed as $\int_{E_d}^{Q_d} V_d(\varphi)d\varphi - \sum_o P_{od}q_{od} - \sum_o \tau_{od}^b q_{od}$ for buyers and $\int_{E_o}^{Q_o} V_o(\varphi)d\varphi + \sum_d P_{od}q_{od} - \sum_d \tau_{od}^s q_{od}$ for sellers. These expressions are equal to Equation 7 because in equilibrium each district's marginal values are equalized across transactions: $P_{od} + \tau_{od}^b = V_d(Q_d)$ and $P_{od} - \tau_{od}^s = V_o(Q_o)$.

[19] Prices are irrelevant because they are merely transfers between units, canceling out in the sum.

subject to

$$\text{(definition of final quantities)} \qquad Q_k^f = Q_k^0 - \sum_{l>k} q_{kl} + \sum_{l<k} q_{lk} \; \forall k$$

$$\text{(resource constraint)} \qquad Q_k^f \geq 0 \; \forall k.$$

I solve this problem using the `patternsearch` solver in Matlab, using baseline quantities as the initial conditions and evaluating the integral with Equation 8. Appendix C.2 proves that the solution to this planner's problem also satisfies the conditions of an efficient market.[20]

I calculate unit-specific gains relative to baseline in the same way as Equation 7:

$$H_k \equiv \int_{Q_k^0}^{Q_k^f} \left[V_k(\varphi) - V_k(Q_k^f)\right] d\varphi. \tag{10}$$

This expression uses the post-trading marginal valuation instead of marginal transaction costs, which avoids the need to directly calculate the incidence of transaction costs for each pair of buyers and sellers. Because marginal valuations equalize up to transaction costs, the sum of H_k over all units is equivalent to the maximand in Equation 9; a proof is in Appendix C.3.

In addition to gains from new transactions, an efficient market also reduces the costs of transactions that were already taking place. These gains are the difference between endowments and baseline quantities multiplied by the difference between baseline marginal valuation and post-trading marginal valuation:[21]

$$R_k \equiv \left(Q_k^0 - E_k\right)\left(V_k^0 - V_k(Q_k^f)\right). \tag{11}$$

8.4 Main results

Scenario 1: Gains from observed spot-market transactions.

I first calculate the gains achieved to date by spot-market transactions, the same observations I use to estimate demand. Table 5, Panel A reports that the total volume of water traded across units in short-term transactions is 156,000 acre-feet for a typical year in my dataset (corresponding to the median quintile of water quantities consumed). Figure 8 maps the geographic patterns of net water sales among observed transactions. Sellers tend to be in the Sacramento Valley, northern San Joaquin Valley, and along the Colorado River; buyers tend to be in urban Southern California.

These observed transactions result in economic surplus of $13.4 million per year in the median-year scenario – a figure that is small relative to annual water-related expenditures in California.[22]

[20] Intuitively, marginal valuations must be equalized up to physical costs in both the planner's solution and market equilibrium; if they are not, then buying from one unit and selling to another would be profitable for an arbitrageur, while the same reallocation would increase welfare for the social planner.

[21] I restrict R_k to be non-negative; negative values occasionally arise if the simulation predicts net trading in the opposite direction of observed trading.

[22] 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.

In wet years (highest quintile), transaction volume and total gains are similar. In dry years (lowest quintile), transaction volumes and marginal valuations are more than twice as large. Total gains are larger (but still relatively small) at $87.7 million per year.

Scenario 2: Gains achieved by all observed transactions, including permanent transfers.

Next, I consider the gains achieved by all observed transactions across units, including both spot-market and permanent transfers of water rights and entitlements. Figure 8 shows that the geographic pattern of trade is similar when including permanent transactions, with additional buyers in the urban areas of both Northern and Southern California.

When including permanent transfers in the median-year scenario, gains increase with the transaction volume, but they are still relatively small, at $35.9 million per year. Gains are much larger in the dry-year scenario: $822 million per year. This pattern of results suggests that permanent transfers have already allowed many parts of California to avoid substantial welfare losses from reduced surface water availability during droughts.

Scenario 3: Gains from an efficient market, eliminating all non-physical transaction costs.

Results from the main simulation are shown in Table 5, Panel B. All quantities and gains are relative to the baseline scenarios, so they are additional to the gains already achieved by the existing market (Scenarios 1 and 2).

Total quantity traded increases dramatically, to over one million acre-feet per year. Total gains range from $85.8 million (in a median year) to $278 million (in a dry year). Although eliminating transaction costs does result in lower costs for transactions that already take place ($11.3 million in a median year), most of the gains arise from new transactions that would not have otherwise occurred ($74.5 million in a median year). These results suggest that reducing transaction costs could increase the gains achieved by the existing short-term market, but the gains are still not large.

Geographic patterns of trading for the median-year and dry-year scenarios are mapped in Figure 8. In a median year, some water is sold within the Sacramento Valley, larger amounts are sold from the San Joaquin Valley to the lower Tulare Lake Basin, and some water moves from urban Southern California to other parts of the state. In a dry year, more water is sold from the Sacramento Valley to the San Joaquin Valley and points further south. Despite previous findings of costly water shortages in urban Southern California (Buck et al. 2016), these results suggest that districts in this region would not actually buy more water in drought years if transaction costs were lower, and in fact they might even sell water (or buy less) in non-drought years.

8.5 Extensions and sensitivity checks

Environmental constraints. Because my simulation does not take into account capacity constraints, a natural question is how important these might be. Even if additional pumping and conveyance infrastructure could be built, there may be points in the system at which additional flow would cause unacceptable environmental harm (for example, as defined by federal environmental regulations). I explore the impact of such environmental constraints by examining the case of the Sacramento–San Joaquin Delta, the single most ecologically sensitive juncture in California's hydrology. Specifically, I simulate an efficient market with the additional constraint that outflow from the Sacramento River may be no more than it is at present. This restriction addresses a common concern over the prospect of expanding water transactions in California: if additional water flows out of the Sacramento River and is pumped into canals, it could affect salinity levels and flow direction, harming water quality and the migration ability of endangered fish species.

Scenario 4 implements this simulation by separating the market into two segments, one north of the Sacramento–San Joaquin Delta (i.e., the mouth of the Sacramento River) and one south of it, and then simulating an efficient market in each segment as in Scenario 3. This allows water to be efficiently reallocated both within the Sacramento Valley and south of the Delta, but with no change in the outflow of the Sacramento River.

Table 5, Panel C shows that in this scenario, both trading volumes and economic gains would be virtually the same as in Scenario 3 – the environmental constraint barely seems to bind. This result is consistent with the geographic patterns of trading from Scenario 3, which show both purchases and sales within all regions of the state (the Sacramento Valley, the San Joaquin Valley, and Southern California). It suggests the gains from a more efficient surface water market could be achieved almost entirely by trading within each side of the Sacramento–San Joaquin Delta, and they do not require relaxing current environmental regulations.

Functional form sensitivity. Next, I explore how sensitive my results are to the functional form for demand. In simulating counterfactual market equilibria and evaluating the resulting surplus, I rely on an assumption of isoelastic demand. This assumption may be innocuous if the change in quantity is small for all units, but otherwise it may be an important factor driving the results.

To test sensitivity to functional form, **Scenario 5** instead assumes that demand is linear. This scenario uses the same estimated elasticities and marginal valuations as in the previous scenarios but evaluates the integral of inverse demand using the corresponding expression for linear demand.[23] Results are very similar, with trading volumes and economic gains within 10 percent of the corresponding figures for Scenario 3. This suggests that no major aspect of my results is driven by the choice of functional form.

Physical transaction costs. I also investigate the extent to which my results are driven by the physical transaction costs. Are the gains from an efficient market small because it is too expensive to arbitrage across regions, or simply because regions are not very heterogeneous? In **Scenario 6**, I simulate an efficient market with all transaction costs c_{kl} set to zero – implying, implausibly, that all water conveyance and pumping is costless.

The results are, once again, similar to the main results from Scenario 3. Gains in median and

[23] $\int_A^B V_k(\varphi)d\varphi = (1-1/\eta_k)(B-A)V_k^0 + (2\eta_k)^{-1}(B^2-A^2)(V_k^0/Q_k^0)$.

wet years are only slightly larger without physical transaction costs, and gains in a dry year are only 17 percent larger. This suggests that purely physical transaction costs are not a major obstacle to additional gains from trade.

Individual versus unit-level elasticities. My analysis uses the inferred preferences of retail water districts rather than the preferences of individual consumers and producers, who usually cannot directly access California's wholesale water market. However, individual-level preferences may differ from district-level behavior. In particular, the demand elasticities I estimated in Step 3 are smaller (less elastic) than the individual-level elasticities in much of the prior literature.

In **Scenario 7**, I simulate an efficient market using larger elasticities that are representative of the individual-level literature: -0.48 for agricultural units (from Scheierling et al. 2006),[24] and -0.143 for urban units (from Buck et al. 2016),[25] with all other parameters the same as in Scenario 3. I find that the gains from trade are much larger. Surplus nearly doubles in median and wet years, and in dry years it reaches $711 million. These results suggest that the relatively small elasticities I estimate, particularly for agricultural water districts, are a major factor driving the surprisingly small results.

Dispersion of marginal valuations. Another way in which individual-level preferences may differ from district-level behavior is in the initial level of the marginal valuations. In particular, the district-level marginal valuations I estimate may be less dispersed than the prices (or shadow prices) faced by districts' retail customers, even after adjusting for the marginal costs of treatment and distribution.

To investigate the sensitivity of my results to the dispersion of marginal valuations, **Scenario 8** applies a transformation to estimated marginal valuations before calculating the fitted values in Step 3. Specifically, I double the variance of the logarithmic distribution of marginal valuations, by doubling the difference between each log marginal valuation and the grand mean of the log distribution and adding the grand mean to find each new marginal valuation. Doubling the log distribution is an arbitrary choice simply meant to illustrate the sensitivity.

The gains are again larger when using these more-extreme marginal valuations, with surplus in median and wet years approximately triple the surplus achieved in Scenario 3. In **Scenario 9**, I apply the modifications of both Scenarios 7 and 8, using individual-level elasticities as well as more-extreme marginal valuations. Here, once again, the gains increase, ranging from $392 to $610 million per year depending on weather.

[24] This is the mean price elasticity in a meta-analysis of irrigation water demand in the United States. I apply it to all agricultural units except for those in the Tulare Lake Basin, which was already using a larger-magnitude elasticity.

[25] This is the most credible analysis of price elasticities that includes a large set of municipal districts in California; I use the central instrumental variables estimate.

8.6 Discussion

My main results suggest that reducing transaction costs in California's surface water market would achieve between \$86 and \$278 million dollars per year, depending on weather. These benefits are meaningful but significantly lower than often thought. Why are the potential benefits so small, given ample anecdotal evidence suggesting considerable misallocation of water throughout the state? Three explanations appear to be supported by evidence.

One explanation is that statewide water allocation is simply not as inefficient as anecdotal evidence suggests. In most non-drought years, marginal valuations are not dramatically heterogeneous across regions and sectors, varying by factors of two or three, not orders of magnitude. Results from Scenario 2 indicate that the existing water market has already reduced the cost of drought substantially; marginal valuations would otherwise be much more dispersed.

Another explanation is that many of the potential benefits from water markets may be local rather than statewide. The gains I estimate are from simulating trading among units, i.e., across regions and sectors. However, as an intermediate step I also calculate marginal valuations for individual water districts within units. The variance of marginal valuations within units is similar to the variance across units,[26] suggesting that many of the potential gains are in fact local.

A third explanation is that water districts do not accurately represent the preferences of their retail customers when participating in this wholesale market. My approach infers and extrapolates from the objective functions of water districts, which may differ substantially from the aggregate preferences of retail customers for a range of reasons. My estimates imply that these districts behave conservatively in the wholesale surface water market, showing more inelastic demand than prior individual-level estimates in the literature. Simulations using these alternative elasticities find substantially larger gains, suggesting that the internal governance and pricing policies of these water districts may be a major part of the explanation.

Overall, my analysis suggests that if there do exist large potential gains from reallocating water within California, they might not be achieved by reforms to reduce transaction costs in the wholesale water market without also addressing the internal workings of retail water districts. Policy reforms to improve local allocation and make water districts' behavior more responsive to their customers' preferences might include: (1) simplifying rate schedules, (2) raising per-unit rates to reflect wholesale opportunity costs in addition to other marginal costs of provision, (3) establishing buy-back programs at wholesale market prices, (4) adjusting rates from year to year to reflect changing conditions in the wholesale market, (5) creating ways for customers to participate directly in the wholesale market, and (6) merging and consolidating water districts.

[26] The within-unit standard deviation (i.e., after partialing out both year and unit) is 194.3, while the between-unit standard deviation of marginal valuations (i.e., after taking the mean within unit-year and partialing out year) is 183.5.

9 Conclusion

This paper studies the role of transaction costs and the potential gains from trade in California's statewide surface water market. I develop a revealed-preference approach to analyze welfare in thin markets with transaction costs, and I apply it to comprehensive new data on California's water economy. First, I estimate marginal transaction costs arising from observable factors by measuring price gaps across transactions within district and year. I find that several cost determinants give rise to large transaction costs, up to several hundred dollars per acre-foot. Second, I recover equilibrium marginal valuations of water by taking the maximum or minimum price after adjusting for estimated marginal transaction costs. I find that marginal valuations are substantially more dispersed than observed prices, implying that welfare gains are available from reallocation. Third, I estimate demand elasticities using weather-driven variation in yearly surface water allocations. I find that both agricultural and urban sectors have fairly inelastic demand, though there is substantial heterogeneity within agriculture. Fourth, I combine marginal valuations and demand elasticities to conduct welfare analysis and simulate counterfactual equilibria. I find that observed trading across regions and sectors achieves gains of \$36 to \$822 million per year. Reducing transaction costs in an efficient market would result in additional gains of \$86 to \$278 million per year.

These results carry at least three important limitations. First, my approach does not account for all types of transaction costs; it cannot adjust for cost determinants that are both unobserved and constant within user, such as search or contracting. As a result, my results may understate the true potential gains from reducing transaction costs. Second, my approach focuses on the statewide water market across regions and sectors; it does not account for gains from reallocation among retail customers within each district (nor among districts of the same sector within each planning area). Third, my approach accounts for only the potential benefits of reducing transaction costs in water markets. Regulatory sources of transaction costs may have benefits (such as achieving ecological goals and avoiding hydrological externalities) while reducing other types of transaction costs may require significant new investments (such as creating expanded water-use monitoring systems or a centralized trading exchange). A complete policy analysis should account for both the benefits and costs.

References

Adamopoulos, Tasso, Loren Brandt, Jessica Leight, and Diego Restuccia, "Misallocation, Selection, and Productivity: A Quantitative Analysis with Panel Data from China," 2017.

Ahrens, A., C.B. Hansen, and M.E. Schaffer, "pdslasso and ivlasso: Programs for post-selection and post-regularization OLS or IV estimation and inference," 2018.

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

Atkin, David and Dave Donaldson, "Who's Getting Globalized? The Size and Nature of Intranational Trade Costs," 2015.

Ayres, Andrew B., Eric C. Edwards, and Gary D. Libecap, "How Transaction Costs Obstruct Collective Action to Limit Common-Pool Losses: Evidence from California's Groundwater," 2017.

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

Belloni, Alexandre, Daniel Chen, Victor Chernozhukov, and Christian Hansen, "Sparse Models and Methods for Optimal Instruments with an Application to Eminent Domain," *Econometrica*, 2012, *80* (6), 2369–2429.

_ , 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.

Bergquist, Lauren Falcao, "Pass-through, Competition, and Entry in Agricultural Markets: Experimental Evidence from Kenya," 2016.

Brewer, Jedidiah, Robert Glennon, Alan Ker, and Gary Libecap, "2006 Presidential Address: Water Markets in the West: Prices, Trading, and Contractual forms," *Economic Inquiry*, 2008, *46* (2), 91–112.

Bruno, Ellen M, "Agricultural Groundwater Markets: Understanding the Gains from Trade and the Role of Market Power," 2017.

Bryan, Gharad and Melanie Morten, "Economic Development and the Spatial Allocation of Labor: Evidence From Indonesia," 2015.

Buck, Steven, Maximilian Auffhammer, and David Sunding, "Land markets and the value of water: Hedonic analysis using repeat sales of farmland," *American Journal of Agricultural Economics*, 2014, *96* (4), 953–969.

_ , _ , 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.

California Department of Water Resources and U.S. Bureau of Reclamation, "Technical Information for Preparing Water Transfer Proposals (Water Transfer White Paper)," 2015.

California State Water Resources Control Board, "A Guide to Water Transfers," 1999.

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.

Carey, Janis M. and David L. Sunding, "Emerging Markets in Water: A Comparative Institutional Analysis of the Central Valley and Colorado-Big Thompson Projects," *Natural Resources Journal*, 2001, (Spring).

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.

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.

Culp, Peter W., Robert Glennon, Gary Libecap, George A. Akerlof, Roger C. Altman, Gordon S. Rentschler Memorial, and Laura D. Andrea Tyson, "Shopping for Water: How the Market Can Mitigate Water Shortages in the American West," 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.

Donaldson, Dave, "Railroads of the Raj: Estimating the Impact of Transportation Infrastructure," 2012.

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

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

Gray, Brian, Ellen Hanak, Richard Frank, Richard Howitt, Jay Lund, Leon Szeptycki, and Barton "Buzz" Thompson, "Allocating California's Water: Directions for Reform," 2015.

Hagerty, Nick, "Adaptation to Surface Water Scarcity in Irrigated Agriculture," 2022.

Hahn, Robert W., "Market Power and Transferable Property Rights," *The Quarterly Journal of Economics*, 1984, *99* (4), 753–765.

Hanak, Ellen and Elizabeth Stryjewski, "California's Water Market, By the Numbers: Update 2012," 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," *Public Policy Institute of California*, 2014.

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.

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).

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.

Kyle, Albert S., "Informed Speculation with Imperfect Competition," *The Review of Economic Studies*, 1989, *56* (3), 317–355.

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.

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

__ **and Juan Pablo Montero**, "Market Power in an Exhaustible Resource Market: The Case of Storable Pollution Permits," *Economic Journal*, 2011, *121* (551), 116–144.

Malueg, David A. and Andrew J. Yates, "Bilateral oligopoly, private information, and pollution permit markets," *Environmental and Resource Economics*, 2009, *43* (4), 553–572.

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.

Mukherjee, M. and K. Schwabe, "Irrigated Agricultural Adaptation to Water and Climate Variability: The Economic Value of a Water Portfolio," *American Journal of Agricultural Economics*, 2014, *97* (2009), 809–832.

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.

Restuccia, Diego and Raul Santaeulalia-Llopis, "Land Misallocation and Productivity," 2017.

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.

Rostek, Marzena and Marek Weretka, "Price Inference in Small Markets," *Econometrica*, 2012, *80* (2), 687–711.

_ and _ , "Dynamic Thin Markets," *Review of Financial Studies*, 2015, *28* (10), 2946–2992.

Scheer, Jennifer Beth, "Participation, Price, and Competition in the Sacramento Valley Water Market," 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.

Schlenker, Wolfram, W. Michael Hanemann, and Anthony C. Fisher, "Water availability, degree days, and the potential impact of climate change on irrigated agriculture in California," *Climatic Change*, 2007, *81* (1), 19–38.

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.

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," 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.

Tomori, Françeska, Erik Ansink, Harold E. D. Houba, Nick Hagerty, and Charles S. Bos, "Market power in California's water market," 2022.

UNDP, "Human Development Report 2006 - Beyond scarcity: Power, poverty and the global water crisis," 2006.

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

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

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

Figures

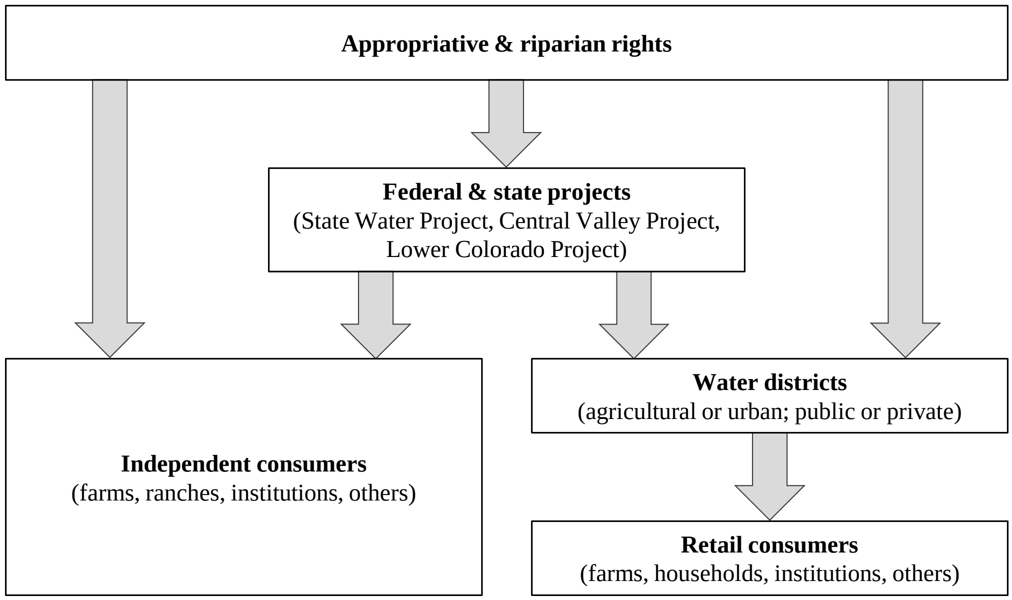

Figure 1: Structure of water use in California. All water use is governed by a legal system of appropriative and riparian rights. Some water districts and independent consumers directly hold their own water rights; others hold long-term contracts with the federal and state water projects, which in turn hold water rights. Retail consumers purchase water from water districts. [Back]

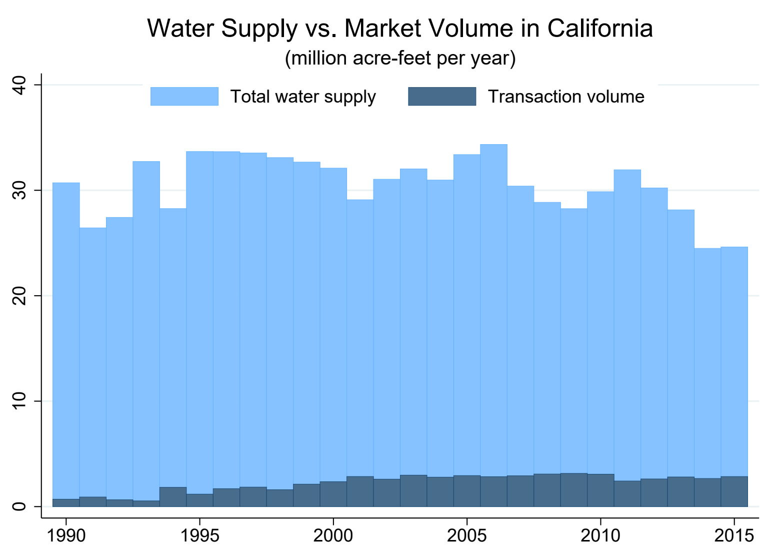

Figure 2: Observed volume of market transactions compared with total water supply. Market transactions include spot market transactions (within-year leases), longer-term leases, and permanent transfers of rights (which are counted in the transaction year and every year afterward). Water supply includes surface water rights, allocations from the federal and state water projects, and average annual groundwater supply. [Back]

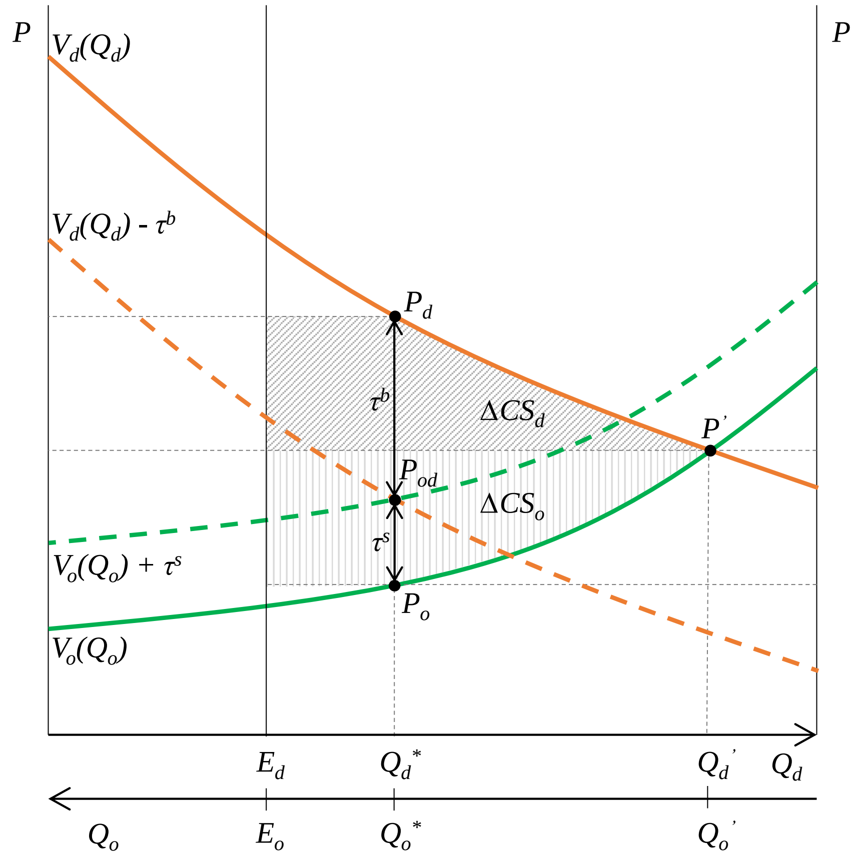

Figure 3: Trade between two districts, buying district d (for destination) and selling district o (for origin). Districts have endowments E_n and inverse demand $V_n(Q_n)$; demand for o is on a reversed axis so that the total width of the graph equals the sum of endowments. With costless trading, the competitive equilibrium among consumers in the two districts would result in price P'. If cross-district transactions incur per-unit transaction costs τ^s (for sellers) and τ^b (for buyers), buyers' marginal willingness to pay shifts down relative to demand, and sellers' marginal willingness to accept shifts up. Competitive equilibrium in this case would result in three distinct prices: a cross-district price P_{od} and within-market prices P_o and P_d, with $P_d - \tau^b = P_{od} = P_o + \tau^s$. Consumer surplus for buyers is shaded with diagonal lines, and surplus for sellers (analogous to producer surplus) is shaded with vertical lines. [Back]

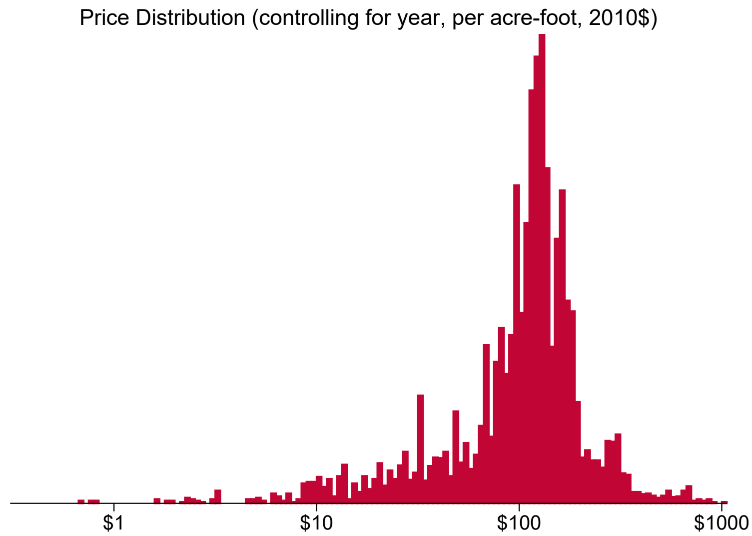

Figure 4: Distribution of prices, controlling for year, logarithmic scale. Graph plots a histogram of observed prices on California's wholesale surface water market, 1980-2015, converted to 2010 dollars using the CPI. I control for year by regressing log price on year fixed effects, taking the residual, and adding the grand mean. [Back]

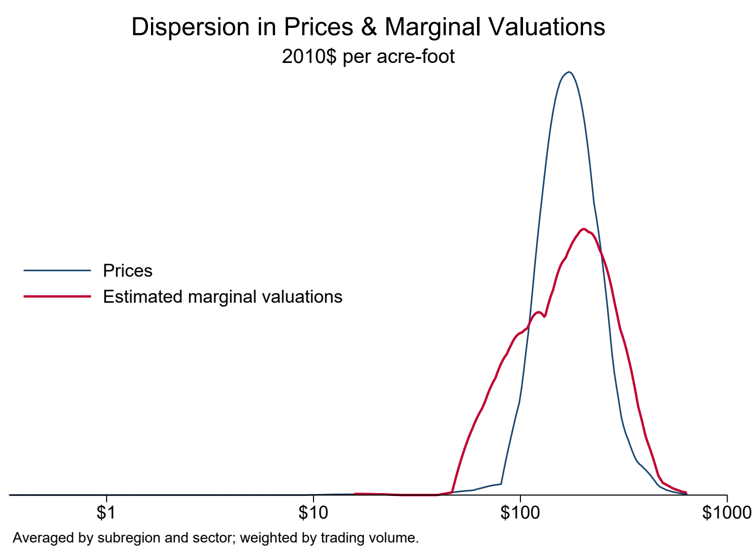

Figure 5: Distributions of prices and estimated equilibrium marginal valuations, logarithmic scale. Marginal valuations are constructed by taking the highest price paid (for buyers) or the lowest price accepted (for sellers) in each year, after adjusting prices for estimated transaction costs from known determinants. Distributions are estimated by kernel density, where an observation is one unit (defined as the intersection of subregion and sector), and units are weighted by trading volume. Within unit, prices and marginal valuations are averaged first across all users observed trading in a particular year and then across years. Unit-level valuations are direct inputs to the counterfactual simulations in Step 4. [Back]

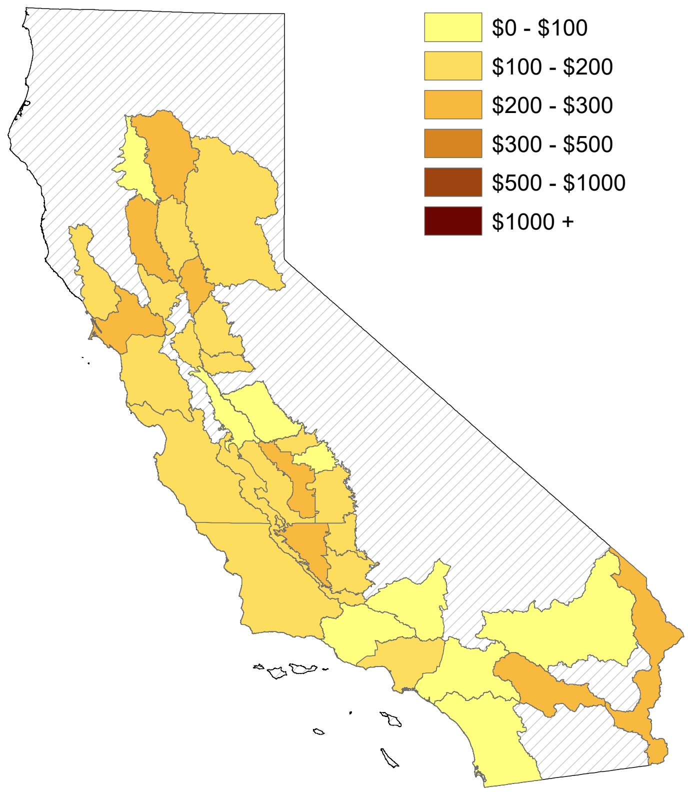

Figure 6: Estimated marginal valuations (per acre-foot) by geography for a median-year scenario. Geographic polygons correspond to subregions (i.e., planning areas as defined by the California Department of Water Resources); areas with diagonal shading have no observed transactions. Values shown are the fitted values from inverting the unit-specific demand models estimated in Step 3 and plugging in the median quantity consumed for each unit across years. (Unit is defined as the intersection of planning area with sector: urban or agricultural). The map shows the weighted average (by quantity) of unit-level marginal valuations across the two units within each planning area. Dry-year and wet-year scenarios are shown in Appendix B. Unit-level valuations in these three scenarios are direct inputs to the counterfactual simulations in Step 4. [Back]

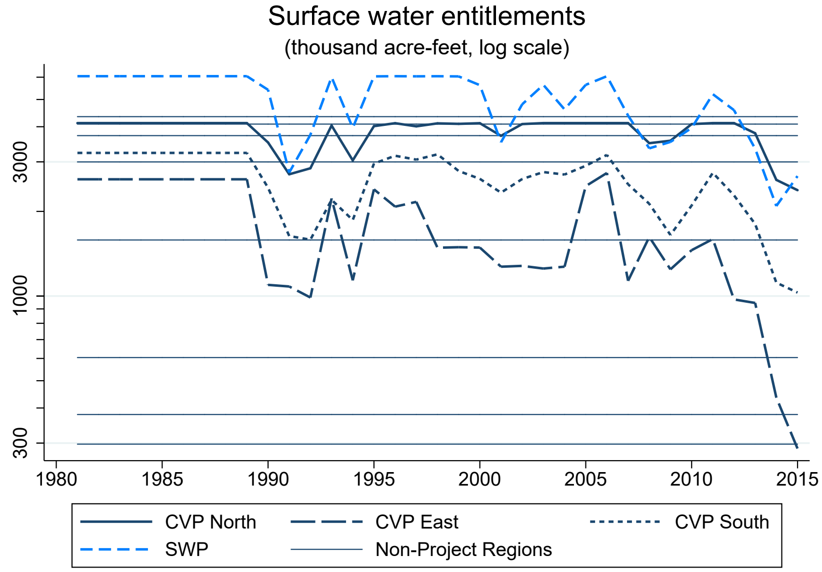

Figure 7: Entitlements over time for several categories of water endowments. This variation is used to estimate demand elasticities. Surface water entitlements are the sum of water rights and allocations from the federal and state water projects (CVP = Central Valley Project; SWP = State Water Project). Water rights are time-invariant, while project allocations vary year to year on the basis of weather conditions. This variability is set separately for each of 14 contract types. For clarity of illustration, each time series on the graph represents users aggregated by project and region. [Back]

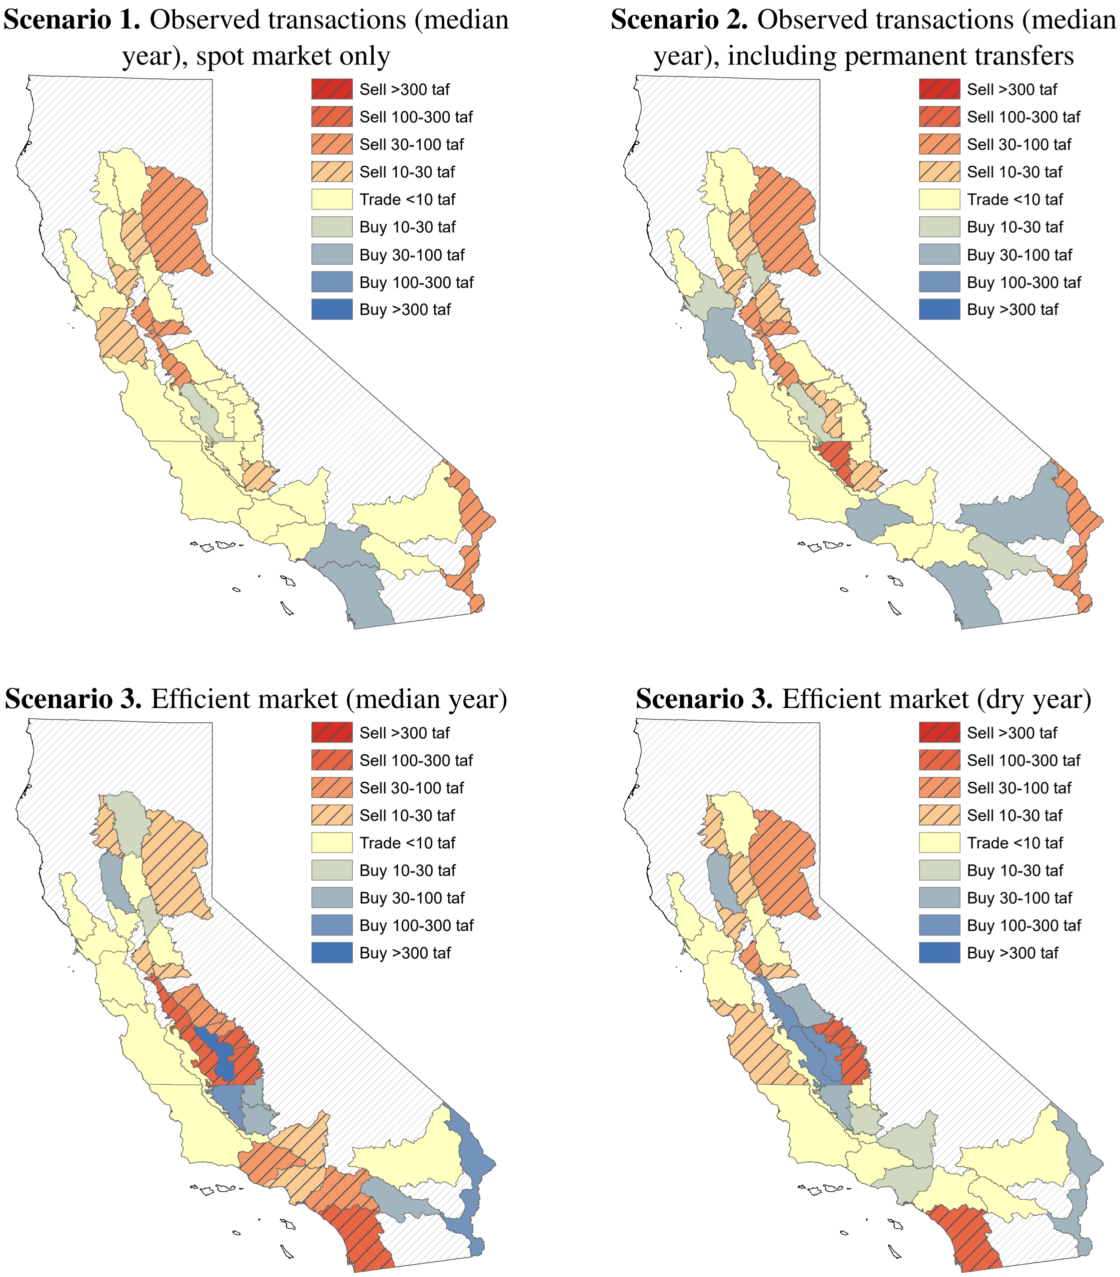

Figure 8: Net annual surface water quantities traded ("taf" = thousand acre-feet) in the existing market (Scenarios 1-2) and additional quantities traded in counterfactual simulations (Scenario 3), by geographic subregion. Median- and dry-year scenarios show trading relative to median and lowest-quintile water quantities consumed, respectively. Trading shown is limited to wholesale market transactions across units (defined by subregion and sector). Subregions are planning areas defined by the California Department of Water Resources; areas shaded with thin diagonal lines have no observed transactions. Quantities shown on map sum across sectors (i.e., units) sharing a subregion. For Scenario 3, transactions incur physical conveyance costs calibrated from a hydrological network model, but all other transaction costs are eliminated. [Back]

Tables

Table 1: Summary statistics of transactions data

Panel A: Summary statistics of transactions data

	Mean	Standard deviation	Observations
Volume (acre-feet)	13,629	31,213	671
Price (2010$/acre-foot)	221.3	297.4	647
Distance, Euclidean (km)	111.8	168.5	654
Distance conveyed (km)	149.3	234.9	665
Distance conveyed in rivers (km)	61.2	99.9	665
Distance conveyed in canals (km)	88.2	172.5	665
Distance of virtual conveyance (km)	72.4	117.0	665
Elevation lift (feet)	274.9	795.9	665
Crosses the Sacramento-San Joaquin Delta	18.3%	0.385	665
Reviewed by the State Water Boards	37.8%	0.483	671
Within project	40.8%	0.491	671

Panel B: Transactions by sector and region

	As Origin (Seller)		As Destination (Buyer)		
	Count	Volume (TAF)	Count	Volume (TAF)	Net sales (TAF)
By Sector					
Agricultural	597	7,606.3	359	2,195.4	5,410.9
Urban	101	833.0	217	2,661.0	-1,828.0
By Hydrologic Region					
North Coast	2	3.7	2	0.7	2.9
North Lahontan	0	0.0	0	0.0	0.0
Sacramento River	273	4,231.5	40	100.4	4,131.1
San Francisco Bay	17	189.7	47	346.1	-156.3
Central Coast	28	23.5	35	22.9	0.5
San Joaquin River	246	2,462.8	139	1,467.5	995.3
Tulare Lake	99	534.4	201	1,215.0	-680.6
South Coast	11	154.3	72	1,568.8	-1,414.4
South Lahontan	2	14.9	31	34.6	-19.7
Colorado River	20	845.2	9	100.5	744.8

TAF = thousand acre-feet. Panel A reports statistics for observed spot market (within-year) transactions of surface water in California that are freely negotiated and involve at least one party that diverts water for agricultural or urban consumption (i.e., not a predominantly environmental user nor a federal or state government agency). Variables are defined in Table 2. Panel B reports the count and total volume of transactions which begin (left) and end (right) in each sector or hydrologic region. Column sums do not exactly agree because a small number of transactions that involve more than two parties count in more than one hydrologic region. Net sales within category do not sum to zero because inclusion criteria are applied separately to each side of a transaction (i.e., many agricultural sales are purchased by environmental or government entities, which do not appear in the table as buyers). [Back]

Table 2: Cost determinants for which marginal transaction costs are econometrically identified in transactions data

	Cost determinant	Definition	Potentially costly because:
1	**Elevation gain (pump lift)**	Total vertical distance lifted.*	Energy is required to pump water uphill.
2	**Distance conveyed in rivers**	Total distance conveyed along a river.*	Some water is lost to evaporation and percolation.
3	**Distance conveyed in canals**	Total distance conveyed along a canal, aqueduct, or pipeline.*	Some water is lost to evaporation. Regulators and canal operators charge fees.
4	**Virtual distance**	Total distance of virtual movement, in which water is transferred against the direction of flow in a river or canal. Transaction quantity is diverted earlier, reducing channel flow.	No conveyance losses but may capture other non-physical costs that increase with distance between buyer and seller.
5	**Crosses the Sacramento–San Joaquin Delta**	Water must pass through the Delta and be pumped into a southward canal.	Fraction of each transaction must flow to the ocean to maintain salinity levels. Ecological requirements lead to delivery risks and heightened regulatory scrutiny.
6	**Import into federal or state water project**	Water is transferred into a project, instead of between contractors within the same project.**	Transactions within a project are subject to less regulatory scrutiny and may be easier to set up.
7	**Export from federal or state water project**	Water is transferred out of a project, instead of between contractors within the same project.**	Transactions within a project are subject to less regulatory scrutiny and may be easier to set up.
8	**State Water Boards (SWRCB) review**	Transaction involves a change in the place of use of a post-1914 appropriative water right.	Requires an extensive regulatory review.
9	**Agricultural counterparty**	Other party (buyer, for sellers; seller, for buyers) primarily uses water in the agricultural sector, rather than municipal/industrial.	Transactions involving agricultural users may incur a greater burden of ongoing monitoring and/or political challenges, or may be more difficult due to factors such as less market experience.

* Along the least-cost distance, as calculated by the hydrological network model. ** "Project" refers to either the federal Central Valley Project or the State Water Project. [Back]

Table 3: Marginal transaction costs from specific cost determinants

Dependent variable: **Price** (2010$/acre-foot)	**SELLERS** (positive coefficients are costly)				**BUYERS** (negative coefficients are costly)			
Panel A: All known cost factors								
	Each alone	*All together*			*Each alone*	*All together*		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
River distance (km)	0.25 *	-0.34 *	-0.17	0.13	0.13 *	0.20	0.23 **	0.19 **
	(0.14)	(0.19)	(0.20)	(0.18)	(0.07)	(0.13)	(0.10)	(0.10)
Canal distance (km)	0.03	-0.30 *	-0.09	0.08	-0.10	-0.26	-0.21	-0.24
	(0.05)	(0.16)	(0.14)	(0.13)	(0.25)	(0.26)	(0.30)	(0.37)
Virtual distance (km)	-0.03	-0.11	0.02	-0.15	-0.32 **	0.00	-0.19 *	-0.13
	(0.09)	(0.13)	(0.12)	(0.10)	(0.16)	(0.16)	(0.11)	(0.11)
Pumping lift (ft)	0.00	0.04	0.02	-0.02	0.03	0.09 **	0.07	0.08
	(0.01)	(0.03)	(0.03)	(0.02)	(0.02)	(0.04)	(0.05)	(0.06)
Delta crossing	109.6 **	243.0 ***	139.9 **	74.3	27.8	-42.3	-47.3	-30.8
	(49.2)	(59.5)	(59.3)	(61.1)	(25.3)	(37.2)	(37.3)	(40.2)
State Water Boards review	8.9	253.8 ***	170.1	-29.8	-47.0 **	-82.2 **	-67.8 *	-68.1 *
	(16.1)	(95.5)	(119.3)	(43.1)	(23.1)	(38.8)	(34.5)	(36.6)
Import into project	10.9	23.3	69.1	0.0	-40.3	74.8	44.6	21.5
	(20.4)	(46.9)	(73.4)	(27.9)	(47.9)	(52.4)	(55.2)	(59.2)
Export from project	14.0	51.9	7.9	44.9	-207.7 **	-247.3 ***	-213.1 ***	-202.7 ***
	(16.5)	(50.4)	(63.9)	(43.9)	(99.9)	(74.5)	(65.1)	(67.5)
Agricultural counterparty	46.3 **	1.1	26.4	31.5 *	5.1	55.3	24.6	22.3
	(20.5)	(28.2)	(35.6)	(17.7)	(26.5)	(48.2)	(40.3)	(42.8)
Panel B: Post-LASSO (Cost factors selected by LASSO)								
		(2)	(3)	(4)		(6)	(7)	(8)
River distance (km)				0.08				
				(0.15)				
Virtual distance (km)							-0.18	-0.16
							(0.11)	(0.11)
Delta crossing		106.6 ***	103.0 ***	75.7				
		(19.16)	(29.89)	(50.67)				
State Water Boards review		290.5 ***	162.7 *				-44.3 **	-50.3 **
		(90.85)	(83.21)				(20.05)	(21.90)
Import into project			68.2					
			(62.01)					
Export from project						-240.2 **	-211.7 ***	-193.7 **
						(97.9)	(79.5)	(78.0)
Agricultural counterparty				30.1 *				
				(17.0)				
Seller FE		✓	✓					
Seller's region by year FE		✓						
Seller's subregion by year FE			✓					
Seller by year FE	✓			✓				
Buyer FE						✓	✓	
Buyer's region by year FE						✓		
Buyer's subregion by year FE							✓	
Buyer by year FE					✓			✓
Observations	532	532	532	532	431	431	431	431
Clusters	507	507	507	507	304	304	304	304

Regressions of transaction price on cost determinants from Table 2; unitless variables are binary indicators. The left side (columns 1-4) include seller-side fixed effects, comparing prices across buyers. Positive coefficients here reflect a price premium, indicating that the cost determinant is costly to the seller. The right side (columns 5-8) include buyer-side fixed effects, comparing prices across sellers. Negative coefficients here reflect a price discount, indicating that the cost determinant is costly to the buyer. Fixed effects and counts at the bottom describe both Panels A and B. In Panel A, columns 1 and 5 report the results from separate regressions in which the row variable is the only regressor; columns 2-4 and 6-8 each show results from one regression containing all variables as regressors. Variables omitted from Panel B are not selected by LASSO in any of the specifications. Mean price in data is $221. Standard errors are shown in parentheses and clustered by transaction. Region and subregion are hydrological region and planning area as defined by the California Department of Water Resources. * p<.1, ** p<.05, *** p<.01. [Back]

Table 4: Demand elasticities

Panel A: First Stage & Reduced Form with One Instrument (OLS)

	First stage: Ln(Price)			Reduced form: Ln(Quantity consumed)		
	(1)	(2)	(3)	(4)	(5)	(6)
Ln(Own allocation percentage)	-0.81 ***	-0.33 ***	-0.37 ***	0.36 ***	0.34 ***	0.34 ***
	(0.18)	(0.10)	(0.10)	(0.03)	(0.04)	(0.05)
Unit FE	✓	✓	✓	✓	✓	✓
Region-specific time trends		✓	✓		✓	✓
Buyer/seller unit pair FE		✓	✓		✓	✓
Cost factors X Direction			✓			✓
Agent FE			✓			✓
Observations	1031	1031	1031	1031	1031	1031
Clusters	331	331	331	331	331	331
F-statistic	19.8	11.2	14.4	109.9	59.4	51.7

Panel B: Price Elasticities of Demand (IV Lasso)

	Ln(Quantity consumed)					
	(1)	(2)	(3)	(4)	(5)	(6)
Ln(Price)	-0.14 ***	-0.14 ***	-0.21 ***	-0.16 ***	-0.16 ***	-0.20 ***
	(0.02)	(0.03)	(0.04)	(0.03)	(0.04)	(0.05)
Lasso-selected instruments	10	7	34	10	10	19
Unit FE	✓	✓	✓	✓	✓	✓
Region-specific time trends		✓	✓	✓	✓	✓
Buyer/seller unit pair FE			✓			✓
Cost factors X Direction				✓		✓
Agent FE					✓	✓
Observations	1053	1053	1053	1053	1053	1053
Clusters	345	345	345	345	345	345
First stage F-statistic	167.1	101.5	70.0	89.6	110.1	65.8
Sup-score weak-ID test	reject	reject	fail to reject	reject	reject	fail to reject

Panel C: Heterogeneous Price Elasticities of Demand (IV Lasso)

	Ln(Quantity consumed)				
			Agricultural regions		
	Urban	Agricultural	Sacramento Valley	San Joaquin Valley	Tulare Lake Basin
	(1)	(2)	(3)	(4)	(5)
Ln(Price)	-0.10 *	-0.23 ***	-0.07 ***	-0.11 ***	-0.80 ***
	(0.06)	(0.05)	(0.02)	(0.04)	(0.12)
Lasso-selected instruments	2	17	9	4	3
Unit FE	✓	✓	✓	✓	✓
Region-specific time trends	✓	✓	✓	✓	✓
Buyer/seller unit pair FE	✓	✓	✓	✓	✓
Cost factors X Direction	✓	✓	✓	✓	✓
Agent FE	✓	✓	✓	✓	✓
Observations	259	767	235	287	235
Clusters	88	243	79	70	84

Regressions estimating unit-level price elasticities of surface water demand by instrumental variables. Units aggregate all observed water users and market participants falling within the same geographic subregion (i.e., planning area as defined by the California Department of Water Resources) and sector (agricultural or urban). Observations are unit-transaction since multiple transaction prices are observed in the same unit-year, but quantities are aggregated within unit-year, and standard errors (in parentheses) are clustered by unit-year. Instruments are allocation percentages (of maximum contract amounts) determined yearly for each of 14 contract types within the Central Valley Project and State Water Project (set on the basis of precipitation and snowmelt), interacted with region and unit indicators. Panel A shows the effects on prices and quantities of a single instrument, own allocation percentage. These results are illustrative and not directly used in further analysis because they do not account for the possible effects of others' endowments. Panels B and C reports instrumental variables regressions in which the instruments are selected by Lasso from the full set of potential instruments: each contract type's allocation percentage interacted with unit indicators. Cost factors are the cost determinants listed in Table 2; they are interacted with transaction direction (i.e., whether buying or selling). Agent is the particular identity of the transactor. Region and subregion are hydrological region and planning area as defined by the California Department of Water Resources. * p<.1, ** p<.05, *** p<.01. [Back]

Table 5: Annual economic benefits from wholesale surface water markets in several scenarios

Panel A: Observed transactions in existing market

	Scenario		Volume traded (acre-feet)	Average marginal valuation	Total gains (millions)
1	**Observed market**	Med	156,000	$ 151	$ 13.4
	(within-year spot market	Wet	141,000	$ 121	$ 9.9
	only)	Dry	364,000	$ 595	$ 87.7
2	**Observed market**	Med	293,000	$ 151	$ 35.9
	(including permanent	Wet	251,000	$ 121	$ 23.5
	transfers)	Dry	569,000	$ 595	$ 822.2

Panel B: Simulated efficient market

	Scenario		Additional volume traded (acre-feet)	Average equilibrium price	Gains (millions): From lower costs of observed transactions	Gains (millions): From new transactions	Gains (millions): Total
3	**Efficient market** (only	Med	1,085,000	$ 141	$ 11.3	$ 74.5	$ 85.8
	physical transaction costs;	Wet	1,370,000	$ 111	$ 4.7	$ 96.6	$ 101.3
	no capacity constraints)	Dry	1,009,000	$ 254	$ 52.9	$ 225.5	$ 278.3

Panel C: Extensions and sensitivity checks

	Scenario		Additional volume traded (acre-feet)	Average equilibrium price	From lower costs of observed transactions	From new transactions	Total
4	**Key environmental**	Med	1,079,000	$ 145	$ 11.4	$ 73.9	$ 85.2
	constraint (Sacramento	Wet	1,374,000	$ 102	$ 4.8	$ 90.3	$ 95.1
	River outflow held fixed)	Dry	1,009,000	$ 254	$ 53.3	$ 225.5	$ 278.8
5	**Linear demand**	Med	1,026,000	$ 121	$ 6.4	$ 78.1	$ 84.5
	(instead of isoelastic)	Wet	1,250,000	$ 86	$ 3.4	$ 98.3	$ 101.7
		Dry	1,096,000	$ 197	$ 55.7	$ 195.9	$ 251.7
6	**Zero transaction costs**	Med	1,120,000	$ 139	$ 12.3	$ 76.7	$ 89.0
	(as if conveyance were	Wet	1,427,000	$ 111	$ 5.6	$ 102.6	$ 108.1
	costless)	Dry	1,130,000	$ 237	$ 74.5	$ 250.2	$ 324.8
7	**Larger elasticities** (using	Med	1,869,000	$ 141	$ 10.3	$ 132.9	$ 143.2
	individual-level estimates	Wet	2,537,000	$ 103	$ 3.8	$ 180.1	$ 183.8
	from the literature)	Dry	2,530,000	$ 277	$ 56.9	$ 653.9	$ 710.7
8	**More-extreme marginal**	Med	1,900,000	$ 170	$ 19.9	$ 259.4	$ 279.3
	values (using a twice-as-	Wet	2,156,000	$ 151	$ 13.2	$ 260.2	$ 273.4
	dispersed log distribution)	Dry	1,671,000	$ 246	$ 81.8	$ 325.8	$ 407.6
9	**Both larger elasticities**	Med	3,139,000	$ 161	$ 14.5	$ 394.9	$ 409.4
	and more-extreme	Wet	3,397,000	$ 146	$ 9.9	$ 382.5	$ 392.3
	marginal values	Dry	3,130,000	$ 202	$ 75.0	$ 536.0	$ 610.9

Per-year welfare analysis of the existing market (Panel A) and counterfactual simulations (Panels B and C). Simulation trading volumes and prices are found by solving constrained optimization problems as described in Step 4, and trading volumes are relative to observed post-trading allocations. Each scenario is run in three versions: median ("med"), wet, and dry; these use baseline quantities corresponding to the middle, first, and last quintiles of quantities observed over time. For Scenarios 3-5 and 7-9, transactions incur physical conveyance costs calibrated from the hydrological network model. Scenarios 4-9 introduce extensions to the main simulation (Scenario 3) that are fully explained in the text. All dollar figures are in 2010 USD; gains are per year. [Back]

A Market Power as an Alternative Explanation for Price Gaps

In the main sections of this paper, I document price gaps associated with regulatory barriers and other cost determinants, and I interpret these price gaps as marginal transaction costs. In this appendix, I investigate whether these price gaps may instead be explained by market power.

Because my model of bilateral negotiations rules out the successful exercise of market power, I first develop an alternative model of spatial trade that allows both buyers and sellers to exploit market power. Then, I derive and perform an empirical procedure that adjusts raw prices for market power, using estimated pass-through rates as sufficient statistics. Both the model and empirical procedure are based on, and extend the approach of, Atkin and Donaldson (2015). Finally, using these adjusted prices, I re-estimate marginal transaction costs as in Section 5. I find little change in the overall pattern of marginal transaction costs, and so I conclude that the issue of market power is small relative to transaction costs.

Even if markups are non-negligible, however, the approach in the main sections of this paper are still interpretable. The price gaps I estimate in 5 would instead measure both marginal transaction costs and differences in markups (and markdowns) that are correlated with the selected cost determinants. Because both transaction costs and markups drive wedges between observed prices and marginal valuations, the bounds I estimate for marginal valuations using these estimated price gaps are still valid.[27] Then, as long as these differential markups are endogenous to the transaction costs (as opposed to arising from exogenous aspects of the market structure), eliminating transaction costs would also eliminate markups, and counterfactual price and quantity analysis proceeds through the same way. What markups do change is the interpretation of surplus in welfare analysis. The deadweight loss triangles still represent true potential economic gains from removing the cost determinants, but now some of the price gap rectangle may simply represent transfers between parties without lost efficiency.

A.1 Model of spatial trade with market power

As in the main model, N water districts indexed by n (or by o for origin and d for destination) are given initial endowment E_n that is allocated efficiently among a continuum of consumers. Consumers' preferences can be aggregated such that each district has an inverse demand function $V_n(Q_n)$, which gives marginal valuations as a function of quantity consumed Q_n. Trade across districts is conducted by two layers of intermediaries. In each district, selling intermediaries ("sellers") can buy units of water from consumers (at their marginal valuations $V_o(Q_o)$) 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 $V_d(Q_d)$). Sellers and buyers meet at exchange points unique to each pair of districts, where prices P_{od} are determined.

Each transaction i generates constant marginal transaction costs for both sellers and buyers, as

[27] Under an analogous assumption for markups as Assumption 3 for transaction costs: that differences in markups associated with specific cost determinants are constant and additively linear.

in Assumption 1. However, sellers and buyers may also incur fixed costs: $C^s_{od}(q_{iod}) = \tau^s_{od} q_{iod} + F^s_{od}$ and $C^b_{od}(q_{iod}) = \tau^b_{od} q_{iod} + F^b_{od}$.

Also unlike in the main model, strategic interactions are possible among sellers at specific destinations, and among buyers at specific origins. One seller's quantity decisions may affect the profits of other sellers to the same district through the aggregate quantity sold, and similarly for buyers. I summarize this strategic interaction using a competitiveness index, following Atkin and Donaldson (2015).

Assumption 5. Market structure. *Buying and selling intermediaries choose quantities q_{iod} to maximize profits subject to the expected response of other intermediaries. The competitiveness index $\phi^s_{od} \equiv \left(\frac{dQ_d}{dq_{iod}}\right)^{-1}\left(\frac{Q_d}{q_{iod}}\right)$ for sellers and $\phi^b_{od} \equiv \left(\frac{dQ_o}{dq_{iod}}\right)^{-1}\left(\frac{Q_o}{q_{iod}}\right)$ for buyers is fixed for each origin-destination district pair.*

This approach nests many specific models of market structure. For Cournot oligopoly with identical intermediaries, the competitiveness index ϕ^s_{od} equals Q_d/q_{iod}, or the number of intermediaries. For perfect collusion or a pure monopoly, $\phi^s_{od} = 1$, and for perfect competition, $\phi^s_{od} \longrightarrow \infty$. Note that both buyers and sellers cannot both successfully exercise market power in the same origin/destination pair. Market power comes from keeping quantities lower or higher than efficient; these are mutually exclusive. However, by permitting both in the same model, I allow the market power to appear on either side (which may even change from place to place).

Next, I define markups (also referred to as markdowns, for buyers) as any difference between prices and willingness to pay (or to accept):

$$\begin{aligned} &(Sellers) \qquad \mu^s_{od}(Q_o) \equiv P_{od}(Q_o) - V_o(Q_o) - \tau^s_{od} \\ &(Buyers) \qquad \mu^b_{od}(Q_d) \equiv V_d(Q_d) - \tau^b_{od} - P_{od}(Q_d). \end{aligned} \tag{12}$$

I then assume there are no strategic interactions *across* sellers and buyers.

Assumption 6. No interactions across sides of the market. *Markups set by sellers and buyers are not affected by the overall quantities consumed in their own districts: $\frac{\partial \mu^s_{od}}{\partial Q_o} = 0$ and $\frac{\partial \mu^b_{od}}{\partial Q_d} = 0$.*

This says that in choosing quantities, intermediaries consider only the other side's demand response, not any expected strategic change in markups. A sufficient condition for this assumption is that market power exists on only one side of the market: for any origin-destination pair, one side may exert market power, while the other consists of price-takers.

With this setup in place, the first-order conditions for buyers and sellers yield expressions for markups:

$$\begin{aligned} &(Sellers) \qquad \mu^s_{od} = -\left(\frac{\partial V_d}{\partial Q_d}\right)\frac{Q_d}{\phi^s_{od}} \\ &(Buyers) \qquad \mu^b_{od} = +\left(\frac{\partial V_o}{\partial Q_o}\right)\frac{Q_o}{\phi^b_{od}}. \end{aligned} \tag{13}$$

Note that if markups are zero (such as under perfect competition), Equation 12 becomes identical to the result for price determination in the main model.

A.2 Deriving the estimation procedure

Following Atkin and Donaldson (2015), I use pass-through rates as sufficient statistics for the competitive structure of the market. I first define pass-through rates as the absolute rate at which costs (for sellers) or revenues (for buyers) are passed through to market equilibrium prices: $\rho^s_{od} \equiv \partial P_{od}/\partial MWTP_{od}$ and $\rho^b_{od} \equiv \partial P_{od}/\partial MWTA_{od}$, where $MWTP_{od} = V_o(Q_o) + \tau^s_{od}$ and $MWTA_{od} = V_d(Q_d) + \tau^b_{od}$. Then, I assume demand takes a particularly convenient functional form.

Assumption 7. Bulow-Pfeiderer demand. *Consumer preferences are time-invariant and take the form*

$$V_n(Q_n) = \begin{cases} a_n - b_n(Q_n)^{\delta_n} & \text{if } \delta_n > 0, a_n > 0, b_n > 0, 0 < Q_n < (a_n/b_n)^{1/\delta_n} \\ a_n - b_n \ln(Q_n) & \text{if } \delta_n = 0, a_n > 0, b_n > 0, 0 < Q_n < \mathrm{e}^{(a_n/b_n)} \\ a_n - b_n(Q_n)^{\delta_n} & \text{if } \delta_n < 0, a_n \geq 0, b_n < 0, 0 < Q_n < \infty. \end{cases}$$

Under this assumption, pass-through rates can be expressed as $\rho^s_{od} = \left(1 + \frac{\delta_d}{\phi^s_{od}}\right)^{-1}$ and $\rho^b_{od} = \left(1 + \frac{\delta_o}{\phi^b_{od}}\right)^{-1}$, and each is fixed within seller-buyer pair. Inserting this functional form for demand into buyers' and sellers' first-order conditions, and adding time subscripts, yields:

$$\begin{aligned} (Sellers) \qquad & P_{odt} = \gamma^s_{od} V_{ot} + \rho^s_{od}\tau^s_{odt} - (1-\rho^s_{od})(\tau^b_{odt} - a_d) + (1-\rho^s_{od})(\frac{1}{\rho^b_{od}} - 1)a_o \\ (Buyers) \qquad & P_{odt} = \gamma^b_{od} V_{dt} - \rho^b_{od}\tau^b_{odt} + (1-\rho^b_{od})(\tau^s_{odt} + a_o) + (1-\rho^b_{od})(\frac{1}{\rho^s_{od}} - 1)a_d, \end{aligned} \tag{14}$$

where $\gamma^s_{od} \equiv (\rho^s_{od} + \rho^b_{od} - 1)/\rho^b_{od}$ and $\gamma^b_{od} \equiv (\rho^s_{od} + \rho^b_{od} - 1)/\rho^s_{od}$.

From here, I can obtain estimates of transaction costs – accounting for markups – via a two-step procedure.

Step 1: Estimate pair-specific pass-through rates for buyers and sellers. First, to transform 14 into a regression model, I again adopt Assumption 3 (additively linear determinant-specific transaction costs). I absorb demand levels (terms depending on a_o and a_d) and time-invariant portions of transaction costs with buyer-seller pair fixed effects. Then, I regress prices on previously-estimated marginal valuations (also restoring Assumption 4 to ensure marginal valuations are point estimates rather than bounds):

$$\begin{aligned} (Sellers) \qquad & P_{iodt} = \gamma^s_{od}\hat{V}_{ot} + \lambda^s_{od} + \varepsilon^s_{iodt} \\ (Buyers) \qquad & P_{iodt} = \gamma^b_{od}\hat{V}_{dt} + \lambda^b_{od} + \varepsilon^b_{iodt}, \end{aligned} \tag{15}$$

where I assume movement in marginal valuations within buyer-seller pair is uncorrelated with other, unobserved determinants of prices. (The error terms arise from idiosyncratic time variation in transaction costs, $\tilde{\tau}^s_{odt}$ and $\tilde{\tau}^b_{odt}$.) Upon obtaining estimates $\hat{\gamma}^s_{od}$ and $\hat{\gamma}^b_{od}$, I can back out $\hat{\rho}^s_{od}$ and $\hat{\rho}^b_{od}$ by solving the system of two equations defining γ^s_{od} and γ^b_{od}: $\rho^s_{od} = \gamma^b_{od}/(\gamma^s_{od} + \gamma^b_{od} - \gamma^s_{od}\gamma^b_{od})$ and $\rho^s_{od} = \gamma^s_{od}/(\gamma^s_{od} + \gamma^b_{od} - \gamma^s_{od}\gamma^b_{od})$.

Step 2: Adjust prices for pass-through and re-estimate transaction costs from gaps in adjusted prices. Second, by inserting Assumption 3 into Equation 14 and rearranging, I obtain

$$\begin{aligned}
&(Sellers) \qquad \frac{1}{\hat{\rho}^s_{od}}\left(\hat{\rho}^s_{od}P_{iodt} + (1-\hat{\rho}^s_{od})\hat{V}_{dt}\right) = \omega_{ot} + \tau^s\mathbf{B}_{od} + v^s_{iodt} \\
&(Buyers) \qquad \frac{1}{\hat{\rho}^b_{od}}\left(\hat{\rho}^b_{od}P_{iodt} + (1-\hat{\rho}^b_{od})\hat{V}_{ot}\right) = \omega_{dt} + \tau^b\mathbf{B}_{od} + v^b_{iodt}.
\end{aligned} \tag{16}$$

These can be estimated as regressions, yielding new unbiased estimates of determinant-specific transaction costs $\hat{\tau}^s$ and $\hat{\tau}^b$. Here I have to assume that inverse demand levels a_n (absorbed by the error terms v_{iodt}) are uncorrelated with the cost determinants $\mathbf{B}_{od}$, which may appear to be a strong assumption but is also implicitly assumed in the main sections of this paper.

Finally, note that the full price gap as estimated in Section 5 (i.e., without adjusting prices for imperfect pass-through) identifies the sum of determinant-specific transaction costs and determinant-specific markups. Therefore, the difference between these two estimates identifies the determinant-specific markups.

A.3 Empirical implementation

To bring the regressions from Step 1 and Step 2 to the data, I make several more choices.

First, I coarsen the level of origin and destination *od* in order to conserve statistical power. Instead of estimating parameters at the consumer level, I use the level of groups defined by the interaction of hydrologic region and major cost determinants (the ones for which I measure transaction costs). Regression 15 is more demanding of the data than other regressions in this paper, as it requires estimating both an intercept and a slope (i.e., a fixed effect and a coefficient) for each origin-destination pair. If I continued to use consumer-level pairs, very few parameters would be identified, since most consumer pairs do not have repeated transactions over time. This coarser grouping of origins and destinations still allows markups to vary discontinuously across the cost determinants as well as across geography. Of 1,190 possible group-pair cells, 48 have at least 3 observations (the minimum necessary to identify both point estimates and standard errors).

Second, I instrument for marginal valuations using the leave-out mean of marginal valuations within each of these cells, estimating Equation 15 via two-stage least squares. This addresses the potential concern of mechanical correlation between prices and marginal valuations, since marginal valuations are constructed using within-consumer price data. It does not necessarily solve any other potential endogeneity concerns. Unfortunately, I cannot use the same instrument as in estimating demand elasticities, namely fluctuations in water entitlements. Here, I need to estimate cell-specific parameters, but water entitlements have variation within only a subset of the cells.

Third, I impute 1 for any parameters $\hat{\gamma}^s_{od}$ or $\hat{\gamma}^b_{od}$ that are unidentified, in order to still calculate values for $\hat{\rho}^s_{od}$ and $\hat{\rho}^b_{od}$. That is, I assume complete pass-through, as in the case of perfect competition. This affects origin-destination cells that have only 1 or 2 observations, which comprise 15% of my dataset.

Finally, in the second step, I drop observations for which the pass-through rate ($\hat{\rho}^s_{od}$ or $\hat{\rho}^b_{od}$) is below 0.2 or greater than 2, following Atkin and Donaldson (2015). This ensures that the estimates are not driven by noisy outliers in estimated pass-through rates, which may be amplified because they appear in the denominator on the left-hand side.

A.4 Results

In Step 1, I estimate Equation 15 by two-stage least squares regression, obtain the estimates $\hat{\gamma}^s_{od}$ and $\hat{\gamma}^b_{od}$, and algebraically solve their definitions to obtain $\hat{\rho}^s_{od}$ and $\hat{\rho}^b_{od}$. These are pass-through rates for sellers and buyers, estimated using time-series variation in regional marginal valuations.

Figure A1 plots histograms of these estimated pass-through rates. Sellers' pass-through rates cluster tightly around 1, which is consistent with models in which sellers are not able to successfully exercise market power, including the model of perfect competition in the main sections of this paper. Buyers' pass-through rates are more dispersed, with a mode below 1, suggesting that buyers may have some ability to exercise market power. (Note that depending on market structure, market power may be consistent with pass-through rates either below or above 1).

In Step 2, I adjust raw prices for possible markups and markdowns and then re-estimate marginal transaction costs associated with the cost determinants selected in Section 5. Under the assumptions of the model in this appendix, the adjusted prices reflect what prices would be in a partial-equilibrium counterfactual without markups or markdowns. The prior estimates of transaction costs (in Section 5) then may be biased, since they capture both marginal transaction costs as well as any differential markups or markdowns that vary with the cost determinants (and presumably arise from the cost determinants, since transaction costs reduce competition). In contrast, the estimates here account for markups and markdowns, allowing me to better isolate the marginal transaction costs.

Table A1 shows the results of these marginal transaction cost regressions, from Equation 16. This table is identical to Table 3 from Section 5, except that the dependent variable is adjusted prices rather than raw prices. Overall, the results are noisier than in the main analysis, but the broad pattern of results is unchanged. Many individual point estimates, especially when using less restrictive fixed effects, are quite similar to their counterparts in Table 3. Generally, known cost factors are associated with price premiums for sellers (positive coefficients) and price discounts for buyers (negative coefficients).

The results of this exercise suggest that market power does not explain the large marginal transaction costs I estimate in Section 5. After adjusting for possible markups, many marginal transaction costs remain large and statistically distinguishable from zero. I conclude that the assumption of no exercise of market power, maintained throughout the main sections of this paper, is reasonable, and that I can interpret marginal transaction costs as true deadweight loss rather than transfers between consumers.

Figure A1: Estimated pass-through rates

Histograms of estimated pass-through rates for sellers and buyers. Each pass-through rate corresponds to a different origin-destination region pair, where regions are defined by the interaction of hydrologic region and major cost determinants. Pass-through rates are derived from transaction-level regressions of prices on marginal valuations (as estimated in Step 2 of the main part of this paper) in which both coefficients and fixed effects are estimated separately for each origin-destination pair. Each marginal valuation is instrumented with the leave-out regional mean of marginal valuations, meaning the coefficients are identified using time-series variation in regional marginal valuations.

Table A1: Marginal transaction costs, adjusting for incomplete pass-through

Dependent variable: **Price** (2010$/acre-foot)	**SELLERS** (positive coefficients are costly)				**BUYERS** (negative coefficients are costly)			
Panel A: All known cost factors								
	Each alone	*All together*			*Each alone*	*All together*		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
River distance (km)	0.02	-0.53 **	-0.25	0.25	-0.34	-0.14	0.07	-0.03
	(0.10)	(0.24)	(0.29)	(0.38)	(0.22)	(0.21)	(0.17)	(0.18)
Canal distance (km)	0.02	-0.32 *	-0.06	0.14	-0.68	-1.16	-0.89	-0.96
	(0.08)	(0.19)	(0.16)	(0.18)	(0.60)	(0.98)	(1.10)	(1.11)
Virtual distance (km)	-0.04	-0.10	0.05	-0.14	0.41 *	0.45	0.42	0.57 *
	(0.09)	(0.17)	(0.14)	(0.17)	(0.21)	(0.35)	(0.30)	(0.32)
Pumping lift (ft)	-0.01	0.05	0.03	-0.03	-0.07	0.27 *	0.24	0.25
	(0.01)	(0.03)	(0.03)	(0.03)	(0.13)	(0.16)	(0.17)	(0.17)
Delta crossing	37.7	261.5 ***	98.5	-48.5	-173.5 **	-51.8	-118.7	-69.2
	(33.7)	(80.6)	(88.1)	(99.6)	(74.7)	(89.6)	(86.8)	(96.6)
State Water Boards review	-103.1	273.2 **	178.1	-25.7	-13.3	-36.4	-1.8	-8.5
	(122.7)	(120.5)	(148.0)	(87.2)	(34.5)	(67.3)	(68.3)	(67.8)
Import into project	-17.2	23.4	90.7	-11.2	-163.7	-71.4	-161.0	-171.9
	(19.6)	(68.0)	(104.3)	(22.8)	(123.5)	(145.0)	(146.0)	(153.4)
Export from project	-103.7	32.7	-50.3	-91.8	20.5	-302.1 *	-224.8	-208.4
	(115.1)	(87.8)	(103.9)	(82.0)	(24.6)	(175.1)	(188.5)	(196.5)
Agricultural counterparty	47.8 **	-13.2	14.5	26.0	6.4	147.1	93.1	101.5
	(18.8)	(44.1)	(48.5)	(30.8)	(73.4)	(103.3)	(92.0)	(89.8)
Panel B: Post-LASSO (Cost factors selected by LASSO)								
		(2)	(3)	(4)		(6)	(7)	(8)
River distance (km)								
Virtual distance (km)								
Delta crossing						-169.2 **	-174.0 **	-184.0 **
						(67.1)	(68.2)	(79.0)
State Water Boards review		326.9 **	120.1	-34.2				
		(137.5)	(145.5)	(89.8)				
Import into project								-147.6
								(108.9)
Export from project				-73.7				
				(52.9)				
Agricultural counterparty								
Seller FE		✓	✓					
Seller's region by year FE		✓						
Seller's subregion by year FE			✓					
Seller by year FE	✓			✓				
Buyer FE						✓	✓	
Buyer's region by year FE						✓		
Buyer's subregion by year FE							✓	
Buyer by year FE					✓			✓
Observations	445	445	445	445	283	283	283	283
Clusters	428	428	428	428	192	192	192	192

Regressions of transaction price on cost determinants from Table 2, adjusting prices for incomplete pass-through (and therefore for some possible forms of market power). Table is otherwise identical to Table 3. * p<.1, ** p<.05, *** p<.01.

B Appendix Tables and Figures

Table B1 explores robustness to alternative regression specifications in the estimation of demand elasticities. As compared with the main results in Table 4: (1) Precision is similar when clustering more conservatively by unit (Panel A, columns 1-3). (2) Magnitudes are similar when using estimated marginal valuations as the endogenous variable instead of prices (Panel A, columns 4-6); although standard errors appear to be smaller, they do not account for additional estimation error from the prior stage of estimating the marginal valuations. (3) Results are also similar when providing the Lasso algorithm different sets of potential instruments: allocation percentages interacted with only region indicators (Panel B, columns 1-3) or with only unit indicators (Panel B, columns 3-6). (4) Panel C confirms that including year effects produces results that are unstable across specifications, are imprecise, have the wrong sign, and/or have weak first stages, both when using IV Lasso (columns 1-3) or a only single instrument (columns 4-6).

Table B1: Demand elasticities: alternative specifications

Panel A: Price Elasticities of Demand -- Robustness Checks (I)

	Ln(Quantity consumed)					
	Errors clustered by unit			Using marginal valuations		
	(1)	(2)	(3)	(4)	(5)	(6)
Ln(Price)	-0.14 ***	-0.15 ***	-0.20 **			
	(0.05)	(0.05)	(0.10)			
Ln(Marginal valuation)				-0.13 ***	-0.09 ***	-0.15 ***
				(0.03)	(0.02)	(0.03)
Lasso-selected instruments	10	9	18	12	9	17
Unit FE	✓	✓	✓	✓	✓	✓
Region-specific time trends		✓	✓		✓	✓
Cost factors X Direction		✓	✓		✓	✓
Buyer/seller unit pair FE			✓			✓
Agent FE			✓			✓
Observations	1031	1031	1031	574	574	574
Clusters	36	36	36	318	318	318
First stage F-statistic	165.8	76.7	65.4	71.3	39.2	29.7
Sup-score weak-ID test	fail to reject	reject	fail to reject	fail to reject	fail to reject	fail to reject

Panel B: Price Elasticities of Demand -- Robustness Checks (II)

	Ln(Quantity consumed)					
	Region instruments only			Unit instruments only		
	(1)	(2)	(3)	(4)	(5)	(6)
Ln(Price)	-0.14 ***	-0.19 ***	-0.21 ***	-0.21 ***	-0.07 **	-0.31 ***
	(0.03)	(0.04)	(0.05)	(0.05)	(0.03)	(0.09)
Lasso-selected instruments	6	5	8	12	6	13
Unit FE	✓	✓	✓	✓	✓	✓
Region-specific time trends		✓	✓		✓	✓
Cost factors X Direction		✓	✓		✓	✓
Buyer/seller unit pair FE			✓			✓
Agent FE			✓			✓
Observations	1031	1031	1031	1031	1031	1031
Clusters	331	331	331	331	331	331
First stage F-statistic	27.6	19.9	14.1	19.5	4.0	13.7
Sup-score weak-ID test	reject	reject	fail to reject	fail to reject	fail to reject	fail to reject

Panel C: Price Elasticities of Demand -- With Year Fixed Effects (Biased)

	Ln(Quantity consumed)					
	IV Lasso			Single instrument (own allocation)		
	(1)	(2)	(3)	(4)	(5)	(6)
Ln(Price)	0.23 **	-0.08	-1.41 ***	5.17	11.95	-7.47
	(0.11)	(0.10)	(0.10)	(7.41)	(36.07)	(8.06)
Number of instruments	2	2	3	1	1	1
Unit FE	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓
Cost factors X Direction		✓	✓		✓	✓
Buyer/seller unit pair FE			✓			✓
Agent FE			✓			✓
Observations	1031	1031	1031	1031	1031	1031
Clusters	331	331	331	331	331	331
First stage F-statistic	8.6	0.7	191.7	0.5	0.1	0.7
Sup-score weak-ID test	reject	reject	fail to reject	-	-	-

Regressions estimating unit-level price elasticities of surface water demand by instrumental variables. Units aggregate all observed water users and market participants falling within the same geographic subregion and sector. Observations are unit-transaction but quantities are aggregated within unit-year. Regressions are similar to those in Table 4 but with modifications. Panel A shows results from clustering standard errors by unit instead of unit-year, and from using marginal valuations as the endogenous variable instead of prices. Panel B shows results from using subsets of candidate instruments: allocation percentages interacted with only region indicators or only unit indicators. Panel C shows results from biased regressions that include year fixed effects. * p<.1, ** p<.05, *** p<.01. [Back]

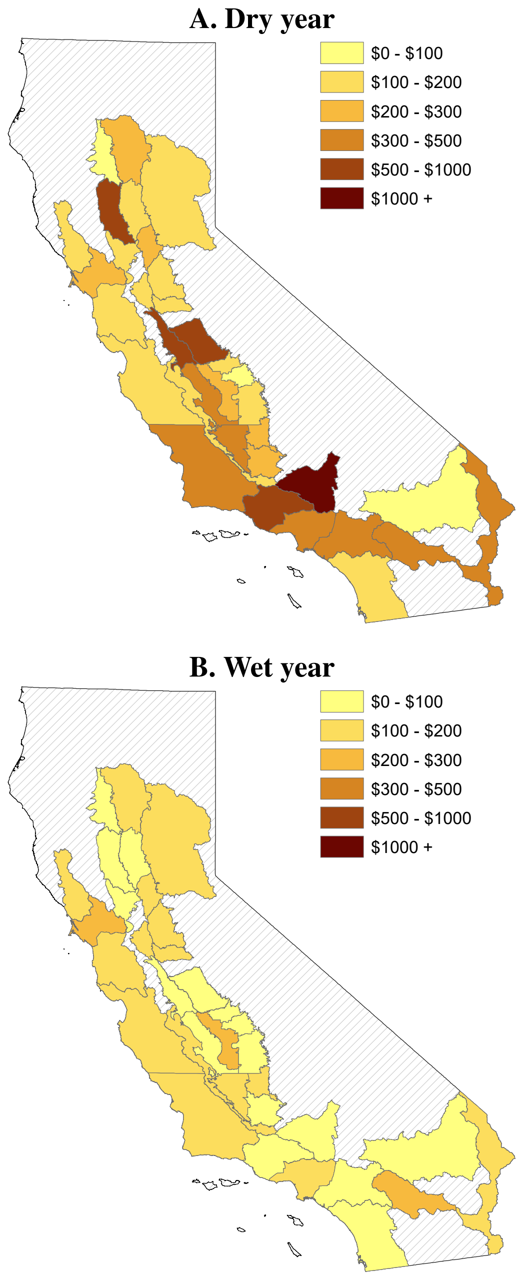

Figure B1: Estimated marginal valuations (per acre-foot) by geography for dry-year (A) and wet-year (B) scenarios. Geographic polygons correspond to subregions (i.e., planning areas as defined by the California Department of Water Resources); areas with diagonal shading have no observed transactions. Values shown are the fitted values from inverting the unit-specific demand models estimated in Step 3 and plugging in the median quantity consumed for each unit across years. (Unit is defined as the intersection of planning area with sector: urban or agricultural). The map shows the weighted average (by quantity) of unit-level marginal valuations across the two units within each planning area. The median-year scenario is shown in Table 6. Unit-level valuations in these three scenarios are direct inputs to the counterfactual simulations in Step 4. [Back]

C Proofs

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

Here I show that including year effects can bias elasticity estimates when prices and quantities are equilibrium outcomes of a common market. To keep the proof as simple as possible, I focus on the reduced form, 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 entitlement 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 entitlements. Total quantities are:

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

This model captures a market with inertia; $\beta = 1$ would correspond to autarky while $\beta = 0.5$ would suggest no inertia, since entitlements 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 entitlements (each agent's own entitlement and the other agent's entitlement) 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 entitlement must also control for the other agent's entitlement. An estimate of β from a regression containing only each agent's own entitlement would suffer from omitted variables bias unless the other agent's entitlement z_{-it} is uncorrelated with own entitlement 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:

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

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 entitlement on the agent's own quantities (β), but also the effect of the entitlement 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 entitlements 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:

$$\begin{aligned}
\sum_k \int_{Q_k^0}^{Q_k^f} V_k(\varphi) d\varphi &= \sum_k \exp\Big(-\frac{\psi_k}{\eta_k}\Big)\Big(\frac{1}{\eta_k}+1\Big)^{-1}\big((Q_k^f)^{\frac{1}{\eta_k}+1} - (Q_k^0)^{\frac{1}{\eta_k}+1}\big) \\
&= \sum_k \exp\Big(-\frac{\psi_k}{\eta_k}\Big)\Big(\frac{1}{\eta_k}+1\Big)^{-1}(Q_k^0)^{\frac{1}{\eta_k}+1}\left[\Big(\frac{Q_k^f}{Q_k^0}\Big)^{\frac{1}{\eta_k}+1} - 1\right] \\
&= \sum_k \exp\Big(-\frac{\psi_k}{\eta_k}\Big)\Big(\frac{1}{\eta_k}+1\Big)^{-1}(Q_k^0)^{\frac{1}{\eta_k}+1}\left[\Big(1 + \frac{-\sum_{l>k} q_{kl} + \sum_{l<k} q_{lk}}{Q_k^0}\Big)^{\frac{1}{\eta_k}+1} - 1\right].
\end{aligned}$$

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

$$
\begin{aligned}
0 &= \frac{d}{dq_{od}}\left\{\exp\left(-\frac{\psi_o}{\eta_o}\right)\left(\frac{1}{\eta_o}+1\right)^{-1}(Q_o^0)^{\frac{1}{\eta_o}+1}\left[\left(1+\frac{-\sum_{l>o}q_{ol}+\sum_{l<o}q_{lo}}{Q_o^0}\right)^{\frac{1}{\eta_o}+1}-1\right]\right\} \\
&\quad +\frac{d}{dq_{od}}\left\{\exp\left(-\frac{\psi_d}{\eta_d}\right)\left(\frac{1}{\eta_d}+1\right)^{-1}(Q_d^0)^{\frac{1}{\eta_d}+1}\left[\left(1+\frac{-\sum_{l>d}q_{dl}+\sum_{l<d}q_{ld}}{Q_d^0}\right)^{\frac{1}{\eta_d}+1}-1\right]\right\} \\
&\quad -\frac{d}{dq_{od}}\left\{\sum_{l>o}c_{ol}q_{ol}\right\} \\
&= \exp\left(-\frac{\psi_o}{\eta_o}\right)\left(\frac{1}{\eta_o}+1\right)^{-1}(Q_o^0)^{\frac{1}{\eta_o}+1}\frac{d}{dq_{od}}\left(1+\frac{-\sum_{l>o,l\neq d}q_{ol}+\sum_{l<o}q_{lo}-q_{od}}{Q_o^0}\right)^{\frac{1}{\eta_o}+1} \\
&\quad +\exp\left(-\frac{\psi_d}{\eta_d}\right)\left(\frac{1}{\eta_d}+1\right)^{-1}(Q_d^0)^{\frac{1}{\eta_d}+1}\frac{d}{dq_{od}}\left(1+\frac{-\sum_{l>d}q_{dl}+\sum_{l<d,l\neq o}q_{ld}+q_{od}}{Q_d^0}\right)^{\frac{1}{\eta_d}+1} \\
&\quad -c_{od}\frac{d}{dq_{od}}q_{od} \\
&= \exp\left(-\frac{\psi_o}{\eta_o}\right)\left(\frac{1}{\eta_o}+1\right)^{-1}(Q_o^0)^{\frac{1}{\eta_o}+1}\left(\frac{1}{\eta_o}+1\right)\left(\frac{Q_o^f}{Q_o^0}\right)^{\frac{1}{\eta_o}}\frac{-1}{Q_o^0} \\
&\quad +\exp\left(-\frac{\psi_d}{\eta_d}\right)\left(\frac{1}{\eta_d}+1\right)^{-1}(Q_d^0)^{\frac{1}{\eta_d}+1}\left(\frac{1}{\eta_d}+1\right)\left(\frac{Q_d^f}{Q_d^0}\right)^{\frac{1}{\eta_d}}\frac{1}{Q_d^0}-c_{od} \\
&= -\exp\left(-\frac{\psi_o}{\eta_o}\right)\left(Q_o^f\right)^{\frac{1}{\eta_o}}+\exp\left(-\frac{\psi_d}{\eta_d}\right)\left(Q_d^f\right)^{\frac{1}{\eta_d}}-c_{od}.
\end{aligned}
$$

Next, rearrange the demand model $Q_k^f = \left(V_k^f\right)^{\eta_k}e^{\psi_k}$ and substitute it for the parameters ψ_o and ψ_d:

$$
\begin{aligned}
0 &= -V_o^f\left(Q_o^f\right)^{-\frac{1}{\eta_o}}\left(Q_o^f\right)^{\frac{1}{\eta_o}}+V_d^f\left(Q_d^f\right)^{-\frac{1}{\eta_d}}\left(Q_d^f\right)^{\frac{1}{\eta_d}}-c_{od} \\
&= -V_o^f+V_d^f-c_{od}.
\end{aligned}
$$

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

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

These are identical to the first-order conditions for the market equilibrium in Equation 1.

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

I need to prove that the sum of unit-specific gains in Equation 10 equals the maximand in Equation 9. 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 10 summed over all units k, I rearrange, switch indices twice, and expand:

$$\begin{aligned}
-\sum_k \int_{Q_k^0}^{Q_k^f} V(Q_k^f) d\varphi &= -\sum_k (Q_k^f - Q_k^0) V(Q_k^f) \\
&= -\sum_k (-\sum_{l>k} q_{kl} + \sum_{l<k} q_{lk}) V_k^f \\
&= \sum_k \sum_{l>k} q_{kl} V_k^f - \sum_k \sum_{l<k} q_{lk} V_k^f \\
&= \sum_k \sum_{l>k} q_{kl} V_k^f - \sum_l \sum_{k<l} q_{kl} V_l^f \\
&= \sum_k \sum_{l>k} q_{kl} V_k^f - \sum_k \sum_{l>k} q_{kl} V_l^f \\
&= \sum_k \sum_{l>k} q_{kl} (V_k^f - V_l^f) \\
&= \sum_k \sum_{l>k} q_{kl} \left[1(q_{kl} > 0) + 1(q_{kl} < 0) \right] (V_k^f - V_l^f) \\
& \quad \sum_k \sum_{l>k} q_{kl} \left[1(q_{kl} > 0)(V_k^f - V_l^f) + 1(q_{kl} < 0)(V_k^f - V_l^f) \right]
\end{aligned}$$

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

$$\begin{aligned}
-\sum_k \int_{Q_k^0}^{Q_k^f} V(Q_k^f) d\varphi &= \sum_k \sum_{l>k} q_{kl} \left[1(q_{kl} > 0)(-c_{kl}) + 1(q_{kl} < 0) c_{lk} \right] \\
&= -\sum_k \sum_{l>k} q_{kl} \left[1(q_{kl} > 0) c_{kl} - 1(q_{kl} < 0) c_{lk} \right]
\end{aligned}$$

which is the second term of Equation 9.

D Processing of Water Transactions Data

WestWater Research, LLC provided a dataset containing 6,263 water transactions in California between 1990 and 2015. Variables include transaction date, volume, price, and duration; and name, latitude and longitude, and water use category of both origin and destination parties. I focus on the spot market, which I define as transactions that are initiated, delivered, and completed within one year. I drop multi-year leases and permanent transfers, leaving 4,906 spot market transactions. Of these, prices are available for 4,415.

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 13,328 observations. For transactions with multiple buyers or multiple sellers, if more specific information is not available, I divide transaction volume equally across parties.

Location. I classify all observations into one of 10 hydrologic regions (defined by California's Department of Water Resources [DWR]). When possible, I also generate latitude and longitude coordinates. I first attempt to use centroids from my user location file, matching parties to users via my crosswalk file. This matches 6,221 transaction-by-party observations. Second, I manually geolocate 65 users not appearing in the user location file but which are common in either transactions or entitlements data. For these users, I generate coordinates based on addresses, towns, or maps found via user websites and other publicly available documents; they match 180 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 3,721 additional observations. This process leaves 3,206 observations for which 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. Total historical allocations (21% of observations). If the party appears in the allocations dataset, I assign to agriculture or municipal depending on which sector receives a majority of total historical allocations.

2. Project agencies (3%). I classify the DWR and USBR themselves as environmental, because they either devote the water to environmental flows or act as intermediaries. This essentially excludes them from sector-wise analysis.

3. Keywords (41%). I classify users based on the following keywords in their name. Agriculture: almonds, citrus, contractors, dairies, dairy, famers, family, farm, farmers, farming,

grower, irrigating, irrigation, nurseries, nursery, orchard, ranch, river interests, trust. Municipal: archbishop, automobile, cement, cemetery, chemical, Chevron, church, city of, cold storage, college, communities, community services, companies, company, container, corporations, country club, developer, development, electric, energy, estate, foods, gardens, golf, gravel, homeowners, homes, housing, inc., Indians, industries, investment, K.O.A., landscaper, leasing, LLC, LP, military, mobile home, monastery, motor, municipal, mutual water, non-ag, oil, owners, park, paving, power authority, properties, property, railway, real estate, realty, recycled, refining, retail, rock, sanitation, school, speedway, Texaco, town of, tribe, university, ventures. Environmental: conservancy, duck hunting, ducks, fish & wildlife, forest service, forestry & fire prevention, water bank.

4. Original use categories (23%). I apply WestWater's original water use categories, based on agriculture, irrigation, and environmental, counting all other categories as urban.

5. Individual names (10%). I assume remaining names of individual people are farmers and therefore agricultural.

6. Remainder (2%). I assume all remaining observations are urban.