Market power in California's water market

July 5, 2023

Abstract

We estimate market power in California's surface water market. Market power may distort the potential welfare gains from water marketing. We use a Nash-Cournot model and derive a closed-form solution for the extent of market power in a water market setting. We then use this solution to estimate market power in a newly assembled dataset on California's water economy. We show that, under the assumptions of the Nash-Cournot model, market power in this thin market is limited.

Keywords: Water markets, Market power, California, Nash-Cournot.

JEL classification: C72, D43, Q25.

1 Introduction

We study the extent and impact of market power in water markets. Such markets are not abundant globally but their prevalence has been increasing. In countries or river basins where rights to water use are allowed to be sold or leased, water marketing reallocates water from lower to higher value uses (Grafton et al., 2011). While reallocation is known to substantially increase the efficiency of water use (Vaux Jr. and Howitt, 1984; Jenkins et al., 2004; Bruno and Jessoe, 2021; Browne and Ji, 2023; Rafey, 2023), most water markets have relatively little trading, suggesting they remain underutilized. Why? Two leading explanations are high transaction costs (Carey et al., 2002; Regnacq et al., 2016; Leonard et al., 2019) and market power (Rosegrant and Binswanger, 1994; Easter et al., 1999; Jacoby et al., 2004; Holland, 2006; Chakravorty et al., 2009; Bruno and Sexton, 2020; Wheeler, 2022). Both sources of friction may distort the potential welfare gains from water marketing, but they may call for different policy remedies, depending on which is more important.

Our setting is California's statewide surface water market. California is one of the world's largest water markets by quantity and value of water traded. At the same time, water transactions are a small fraction of total water used, suggesting large potential gains from trade. Recent work has made progress in quantifying the role of transaction costs (Hagerty, 2019), but the contribution of market power is less clear. Existing conditions may support market power: Trading of water in California is dominated by large water districts that hold rights on behalf of farmers or households. Some of these districts are enormous; for example, the Metropolitan Water District of Southern California serves 19 million people. The belief that large districts, in particular those that are buyers, may have market power is held by many stakeholders and it is supported by previous literature (cf. Tomkins and Weber, 2010; Hansen et al., 2014; Hagerty, 2019).

Our contribution is twofold. The first contribution is methodological; we propose a novel measure of market power tailored to the setting of water markets. This measure is derived using a Nash-Cournot model of water transactions that is inspired by the model set-up of Ansink and Houba (2012). In Section 3 we motivate this model choice by comparing it with a Nash-in Nash bargaining model. Under two main assumptions, discussed below, this model allows us to derive a closed-form solution for the extent of market power in a water market setting. We write this solution in terms of willingness-to-pay and -accept. It is sufficiently general to be adapted and applied to other endowment economies, including permit markets. A methodological advantage of our model is that we do not rely on a conjectural variations approach that employs consistent conjectures (Bresnahan, 1989).

This approach is not compatible with standard notions of rational behavior since the game theory revolution (cf. Lindh, 1992).[1] Despite of this, our solution allows for an extension to conjectural variations that avoids prior selection of the side of the market that holds market power.

The second contribution is that we empirically test the extent of market power in California's surface water market. We apply our model to a newly assembled dataset on California's water economy by Hagerty (2019) and Hagerty (2021). The data that we use are 1993-2015 panel data on water transactions in California, with detailed information on quantities and prices by water district, combined with detailed spatial data on locations of buying and selling districts as well as geographical factors that may affect market power. The data allows us to control for main water uses of buying and selling districts as well as various types of associated transaction costs. The results of our estimation allow us, ultimately, to estimate measures of market power for California's water market.

Our model starts with two main assumptions, both of which are later relaxed. One assumption is that we fix the side of the market on which market power resides. Our starting point is buyer-side market power, as explained above. We then estimate the model in the other direction, allowing for seller-side market power, and find none. To check the relevance of this assumption, we also employ a model specification where we allow for market power on both sides at once; we find compelling empirical support for market power on the buyer side only. The second main assumption is that we use linear demand, originating from a quadratic benefit function of water use. This functional form is commonplace in the water economics literature and allows for a straightforward empirical strategy to derive our results. Constant linear demand across selling districts may not be realistic, however, and therefore we relax this assumption in an alternative specification where, instead, we impose a constant price elasticity. This alternative specification, with constant price elasticity, is presented as part of a larger class of model specifications featuring homogeneous demand, for which we present a closed-form solution as well.

Our results suggest that market power in California's water market is limited. Our main specification implies that buyer power yields an average markdown of 6% of the transaction price. This result is obtained for the linear model, but continues to hold for the non-linear specification and is robust to other model modifications. Our result is surprising in the sense that the thinness of water markets, including California's, is conventionally associated with higher possibilities of exploiting market power. Intuition for this result may

[1] In Hagerty (2019), the same dataset is analyzed, but the focus is on the impact of transaction costs in obstructing water markets. As a robustness check, the potential impact of market power as an alternative explanation for market friction is explored, using an approach that employs consistent conjectures. Other papers, including Bruno and Sexton (2020), use this same approach.

be found in the idea that an individual buyer rarely purchases a large fraction of any seller's endowment of water. Each seller not only has many potential buyers but also consumes water directly, leaving residual supply highly elastic to each buyer.

Our results suggest that market power is not a first-order concern for policy reforms aimed at improving the efficiency and flexibility of water allocation in California, and perhaps in other water-scarce contexts as well. Reform efforts can focus on other considerations instead, such as transaction costs and political economy factors. Proposals to, for example, break the control of water districts over water rights may or may not be desirable on other grounds, but market power in the statewide water market does not appear to be one of them.

Our empirical estimates are relevant to the specific context of surface water in California, and similar analyses in other settings may reveal different results. However, a broadly relevant lesson is that market power is less likely to be a major problem in settings in which water rights are the primary allocation mechanism and markets perform only secondary reallocation. The reason is simply that potential buyers or sellers must compete with the outside option of consuming the water one already owns. While such a configuration would not describe a situation in which, for example, all water rights are initially held by a single entity, it does describe the vast majority of settings in which water markets are active or under consideration.

We first present the data and background information on California's surface water market in Section 2. Next, we introduce the model and our main model specification in Section 3 and our empirical strategy in Section 4. Subsequently, we present model results in Section 5, focusing on our estimation of market power measures for California's water market. This main result is compared with a conjectural variations approach and checked for robustness in Section 6. In Section 7, we conclude.

2 Background and Data

Water in California is initially allocated each year on the basis of water rights and contracts. Water rights originate from historical claims and continue to be recognized by the state.[2] Contracts are long-term agreements held with federal and state agencies under three major water projects: the Central Valley Project, the State Water Project, and the Lower Colorado operations. The main agents that hold these water rights and contracts are mostly local water utilities we refer to as water districts. Water districts divert water from rivers, or

[2]Most water rights are based in the legal doctrine of prior appropriation; California also recognizes riparian rights, but these constitute a small fraction of total water use.

receive deliveries from canals, and distribute them among their constituent farmers and/or households.[3]

Following initial yearly allocations, water districts may transfer water to other parties, typically in return for monetary compensation. Physically, it is feasible to move water between nearly any pair of districts in the state, due to an extensive network of canals, reservoirs, and pumps that constitute one of the world's largest interconnected systems of surface water infrastructure. Legally, transfers are allowed, but they are subject to many different regulatory approval processes depending on the project or type of right the transferred water is based on, where the water originates, the intended destination, and the transfer pathway (California State Water Resources Control Board, 1999). Practically, transfers are completed on a bilateral basis, without a central clearinghouse or auctions; water managers (and sometimes independent brokers) reach out to each other and reach agreements with the help of attorneys (Brewer et al., 2008).

While water transactions are possible, they also face transaction costs arising from all of these frictions: the need to pump water uphill and account for evaporation losses (conveyance charges and delivery costs), the need to comply with regulatory requirements and delays (regulatory, verification, and monitoring costs), and the need to find potential counterparties and reach agreement (search, information, negotiation, and contracting costs). Transaction costs in the California water market are more thoroughly detailed by Regnacq et al. (2016), Scheer (2016), and Hagerty (2019). Our treatment of transaction costs is consistent with a frequently-cited definition provided by McCann et al. (2005): "the resources used to define, establish, maintain, and transfer property rights."

Many types of districts and other parties buy and sell water. Transactions may be agricultural to urban (i.e., a district that primarily supplies irrigation water to farms selling to a district that primarily supplies domestic water to residential households), agricultural to agricultural, urban to urban, or (rarely) urban to agricultural. They also may be agricultural or urban to environmental, in which an environmental nonprofit or government agency purchases water from farmers or municipalities for instream use (i.e., to be left in the river). Our analysis encompasses all these types of transactions, which are conducted bilaterally, at arms length, and at freely negotiated prices. We exclude longstanding arrangements between wholesale and retail water districts (e.g., the Metropolitan Water District and the City of Los Angeles), as as well as government programs that offer to buy or sell water at

[3]By water district, we mean any organization that supplies water to wholesale or retail customers. Most are local government agencies, though some are nonprofit or for-profit, and they carry various names, such as irrigation districts, water conservation districts, or mutual water companies. Some individual farms hold their own water rights without a district intermediary, but they constitute a small fraction of total water use in the state.

fixed prices (e.g., in cases of late season surplus or deficit), since observed prices are less informative in these situations.

There are a variety of reasons why water districts trade water. The key point is that trading allows water to be reallocated from lower- to higher-value uses. This reallocation allows selling districts to raise additional revenue streams while buying districts do so to adjust their water allocation strategy, minimizing losses when faced with water shortage. In such circumstances, agricultural districts with a large share of perennial crops like almonds and grapes will value water more than those with predominantly annual crops. Urban districts will want to secure water for household and industrial use, while buy-backs by the government may be crucial for environmental purposes particularly when water is scarce.

There are two main types of surface water transactions: leases and permanent sales. Leases are short-term transfers of water volumes with a high degree of certainty. Permanent sales are transfers of the underlying right either to divert water from rivers or to receive deliveries from the federal or state water projects. Permanent transfers bundle an expected annual water volume with year-to-year variability on the basis of weather conditions and seniority. We focus our analysis on leases because their prices are easier to interpret without strong assumptions about discount rates and risk preferences, and because they allow us to use year-to-year variation to identify parameters in panel data. Leases also constitute a majority of transactions, and their observed prices are much less noisy than for permanent sales.

The timing of water management and agricultural decisions in California follows an annual cycle known as the water year, which runs from October 1 through September 30. Nearly all precipitation in the state falls between October and March, and most water demand aligns with the main growing season, roughly April through September. By the end of winter, the quantity of water available for the upcoming summer is known with high certainty. For example, yearly allocations in the Central Valley Project are initially announced in the third week of February; they are often revised through April or later, but the revisions are small compared with the inter-annual variation (Stene, 1995; Hagerty, 2021). As we show below, most water leases are completed during April and May for delivery sometime during the rest of the water or calendar year. This timing comes after districts resolve most of their uncertainty over their water endowment for that year, and before they consume most of that endowment. Because both water transactions and market activity follow this kind of annual cycle, we aggregate our analysis to yearly observations of water use and transaction choices.

In focusing on surface water, our analysis also excludes transactions of groundwater rights or permits. Groundwater markets also exist in a few basins that have undergone a process called adjudication to clarify property rights (Ayres et al., 2021), though in most of the state extraction is largely unregulated and unmonitored. Legally, groundwater is fully distinct from surface water, and trade of groundwater rights is sufficiently different as to require distinct theoretical and empirical models. For example, groundwater transactions typically transfer the right to pump water, rather than the water itself, and groundwater pumping also can only be traded within a basin, whereas surface water can be transported across long distances and may be thought of as a statewide market.

Market power in California's water market is not unlikely. In the Cournot-Nash model that we develop in Section 3, such power originates from the number of players on either side of the market. If buyers are on the short side, then they can exploit their favorable position with markdowns on the price. Several aspects have been found to affect this relative position of buyers and sellers in California's water market. Hansen et al. (2014) point to the presence of some 'dominant' buying districts. Tomkins and Weber (2010) suggest informational asymmetries with selling districts being less aware of buying districts' options. Bruno and Sexton (2020) mention capacity constraints in conveyance infrastructure as well as an environment that is 'conducive to forming cartel-like coalitions'. Combined, these aspects may cause buyers to be able to exercise some degree of market power.

We use recently assembled data on California's water economy, first described by Hagerty (2019). We mainly use three datasets. The first is a proprietary dataset compiled by WestWater Research, LLC, listing prices, volumes, and other information related to water transactions in California. These raw transactions are supplemented by data on the geographical and institutional characteristics of water districts, assembled from geospatial files and other sources. This dataset identifies the locations of buying and selling districts and is used to estimate distances and identify other parameters related to transaction costs. Full details on these two datasets and their cleaning and processing are provided in Hagerty (2019, Section 4 and Appendix G). The third is a dataset on yearly surface water entitlements and deliveries in California, compiled from the archives of the California Department of Water Resources, the U.S. Bureau of Reclamation, and the State Water Resources Control Board. Details of this dataset are provided in Hagerty (2021).

The combined dataset provides panel data on water deliveries and transaction prices in California over the 23-year period 1993-2015. The panel data is unbalanced since districts can be involved in more than one transaction per year. Our unit of observation is the water district. The water district is the lowest possible level where (a) we can unambiguously match transactions to units, and (b) we have sufficient information on the units' entitlements and deliveries. Roughly 75% of all transactions in our transaction dataset can be matched to districts with complete information on entitlements and deliveries.[4]

The WestWater water transactions database includes a total of 6,309 transactions over the period 1990–2015. We impose several sample restrictions following the discussion above. First, to focus on the surface water market, we drop transactions of permits for small volumes of groundwater pumping within adjudicated basins in Southern California. Second, to limit our analysis to short-term leases, we drop transactions of permanent water rights. Third, to focus on transactions from which we can infer useful information, we exclude transactions whose price is missing or known to be set administratively (i.e., not at arms length).

Since we will assess transactions both from the sellers' and from the buyers' perspective, we duplicate each transaction and split the dataset into two; one for buyers and one for sellers. A minority of transactions involve more than one district on each side of the transaction. We split up such transactions (i.e. by dividing the transaction volume equally over the number of districts involved), such that each observation contains one selling and one buying district. Finally, we exclude transactions executed before 1993 (since data on water deliveries is only available from 1993 onward). We subsequently lose 28% of our remaining observations (slightly more for buyers than for sellers) when merging our transactions dataset with our dataset on districts' entitlements and deliveries. Our final dataset contains 1,131 observations, 592 for sellers and 539 for buyers.

Summary statistics (mean, standard deviation, and number of observations) on transactions for both buyers and sellers are shown in Table 1. In addition to transaction volumes and prices, this table lists statistics on six different factors that were found by Hagerty (2019) to be costly to buyers or sellers and thereby generate transaction costs. The first three are costly to sellers: (S1) transactions that cross the Sacramento-San Joaquin Delta, (S2) transactions where the buyer is primarily using water for agricultural purposes, and (S3) the total distance if water is conveyed along a river. The next three are costly to buyers: (B1) the virtual distance between buyer and seller if water is transferred against the direction of flow, (B2) transactions that are subject to a State Water Boards review, and (B3) transactions that export water from a federal or state water project. Two factors cause differences in the data between buyers and sellers. One is that, in merging transactions with entitlements, we lose more observations for buyers than for sellers and this difference is apparently not a random draw. The second factor is that the buyer observations include

[4]The alternative to districts as units of observation would be to either use planning areas or DAU-county areas (both are hydro-geographical areas defined by the California Department of Water Resources). Doing so would facilitate the matching with entitlements and deliveries. The downside, however, is that it would severely reduce the number of observations in our final dataset since transactions would be lumped into fewer units.

a substantial share, 24%, where water is acquired for instream use, while for sellers this is only 1%. Such transactions tend to have much lower prices, roughly 50% lower than transactions where buyers are purchasing water for consumptive use. We will check whether the inclusion of these transactions affects our results in Section 5.

Table 1: Summary statistics on transactions by sellers/buyers.

	Sellers			Buyers		
	Mean	SD	Obs	Mean	SD	Obs
Price (2010$/AF)	237.49	296.79	592	185.50	173.75	539
Volume (AF)	8.74	24.70	583	8.93	26.67	530
S1: Delta crossing (1=yes)	0.33	0.47	568	.	.	0
S2: Agricultural buyer (1=yes)	0.47	0.50	592	.	.	0
S3: River distance (km)	0.09	0.10	568	.	.	0
B1: Virtual distance (km)	.	.	0	0.08	0.11	534
B2: State Water Boards review (1=yes)	.	.	0	0.42	0.49	539
B3: Export from project (1=yes)	.	.	0	0.05	0.22	539

Transactions mostly occur in a limited number of hydrologic regions. Sellers are mostly located in the Sacramento River and San Joaquin River regions, while buyers are mostly located in the Tulare Lake, San Joaquin River, and South Coast regions. We find only few instances of districts that both sell and buy, suggesting that we can assume fixed roles for districts as sellers or buyers in our model that we present in the next section. Examining transactions for which we have the exact date of transaction, we find that they are largely spread out over months 3–7 with more than 50% of transactions occurring in April and May, the start of the summer growing season. Transactions in our database cover a total of 161 districts, which implies a mean number of $592/161 = 3.7$ transactions per district over our 23-year period from the sellers' perspective and $539/161 = 3.3$ for buyers. This low number suggests that California's water market is thin.[5]

3 Model

3.1 A model of market power in water markets

Models that are able to quantify market power range from the Cournot-Nash model to the recent Nash-in-Nash bargaining model (Collard-Wexler et al., 2019). The former model provides a succinct framework that centers around competition in quantities and has been

[5]One could argue that our data suffers from selection bias since we only observe realized transactions and these are typically from pairs of trading districts with low transaction costs. Note, however, that we only observe equilibrium transactions and any non-observed transaction price would be 'out-of-equilibrium'.

adapted to accommodate water markets (Ansink and Houba, 2012). The Nash-in-Nash model bargaining model shares certain similarities with the Cournot-Nash model. A major advantage is that it incorporates bilaterally negotiated contracts, which would capture many realistic aspects of our setting.

There are, however, two disadvantages to the use of Nash-in-Nash bargaining. The first is practical. Nash-in-Nash has yet to find effective applications in non-standard settings such as the one given by water markets. There is no closely fitting version of Nash-in-Nash that accommodates our setting of an endowment economy with endogenous trading networks and transaction costs. The second disadvantage is more conceptual. Nash-in-Nash bargaining recognizes the asymmetry of bargaining power as an additional source of market power and is not able to distinguish the two. Collard-Wexler et al. (2019), for example, maintain a general notion of contracts that leaves the relationship between quantities and prices unspecified, thereby equating market power to bargaining power. More recently, Alviarez et al. (2023) provide a formulation of Nash-in-Nash that employs negotiated bilateral quantities and inverse demand functions to determine prices. In general, application of Nash-in-Nash requires specifying one side of the market operating along either the demand or supply curve. While this formulation bears similarities to the Cournot-Nash model, it also implies a concept of market power that combines both bargaining power and strategic quantity manipulation.

For these reasons, we resort to the established Nash-Cournot model as the structural model in our investigation of market power in the Californian water market and leave application of the Nash-in-Nash approach for future research. Note, though, that both models share an important commonality. By relying on the Nash equilibrium concept, it is presumed that an invisible hand steers individual expectations toward equilibrium even in the absence of market institutions such as auctions or clearinghouses. In the Californian water market, this makes sense. Water district managers and water brokers have both a good understanding of the market and good information on aspects that may affect market trading such as weather fluctuations, institutional rules, and different types of transaction costs.[6]

We develop a Nash-Cournot model of water transactions in order to derive a novel measure for the extent of market power in a water market setting. Consider a water market with water transactions between sellers at origins $o = 1, 2, \ldots, N_o$ and buyers at destinations $d = 1, 2, \ldots, N_d$. Water is a homogeneous good and purchases from different

[6]The thinness of the market and lack of a central clearinghouse will create additional transaction costs. Such costs are consistent with our model as is explained in Section 3.2 and shown in Appendix C where we present a model version that includes transaction costs.

sellers are (physically) perfect substitutes. Both sellers and buyers have endowments of water, denoted either $e_o > 0$ or $e_d \geq 0$, depending upon their role. Although variation in rainfall and snow-melt may cause endowments to change over time, we suppress time subscripts in this section to keep notation simple. The amount of water sold by seller o to buyer d is denoted $q_{od} \geq 0$. Obviously, sellers cannot sell more water than their endowments, i.e. $\sum_{d=1}^{N_d} q_{od} \leq e_o$.

Water use by buyers consists of their endowment plus purchased water: $Q_d \equiv e_d + \sum_{o=1}^{N_o} q_{od}$. Buyer d's benefit from using this total sum of water equals $f_d(Q_d)$, which is increasing in the neighborhood of e_d (buyers are unsatiated at e_d), strictly concave, and twice continuously differentiable in Q_d. For later reference, we introduce buyers' willingness-to-pay, denoted WTP, which is defined as the partial derivative of net benefits with respect to water use. Formally,

$$\text{WTP}_d(Q_d) = f_d'(Q_d). \tag{1}$$

In any bilateral trade, buyer d does not pay more than their WTP_d through the transaction-specific price $p_{od} \leq f_d'(Q_d)$.

Water use by sellers consists of their endowment minus sold water: $Q_o \equiv e_o - \sum_{d=1}^{N_d} q_{od}$. Seller o's benefit from using the unsold amount of water equals $f_o(Q_o)$, which is increasing in the neighborhood of e_o (sellers are unsatiated at e_o), strictly concave, and twice continuously differentiable in Q_o. Seller o's net benefits of water use are now given by $f_o(Q_o)$ plus revenues from selling water, introduced below. For later reference, we introduce sellers' willingness-to-accept, denoted WTA, which is defined as the partial derivative of net benefits with respect to water use. Formally,

$$\text{WTA}_o(Q_o) = f_o'(Q_o). \tag{2}$$

In any bilateral trade, sellers must be financially compensated for these opportunity costs through the transaction-specific price and this imposes $p_{od} \geq f_o'(Q_o)$.

Recall that we consider the case where buyers hold all market power. In the standard oligopoly model in economics, consumers react to the quantities chosen by the competing producers and this is modeled by setting the market price equal to the inverse demand or price function. The price function can be interpreted as either the market-clearing price that equates market demand to the chosen market supply, or as the willingness to pay for the chosen market supply. In our case, seller's WTA takes over the role of the price function. Then WTAs can be seen as either market-clearing prices that equate market supply to the market demand chosen by the buyers, or as the willingness to accept for the chosen market demand. Formally,

$$p_{od} = \mathrm{WTA}_o\left(\mathrm{Q}_o\right). \tag{3}$$

Buyer d's expenditure on buying water from seller o is then given by $q_{od} \cdot \mathrm{WTA}_o(\mathrm{Q}_o)$. Similar to standard Cournot oligopoly models, our model is static and all buyers take their quantity decisions simultaneously. Therefore, each buyer maximizes over all potential sellers to purchase water. Formally,

$$\max_{q_{1d},\ldots,q_{\mathrm{N}_o d}} f_d\left(\mathrm{Q}_d\right) - \sum_{o=1}^{\mathrm{N}_o} q_{od} \cdot \mathrm{WTA}_o\left(\mathrm{Q}_o\right). \tag{4}$$

Using the positive relation between Q_d and q_{od} as well as the negative relation between Q_o and q_{od}, a buyer's first-order condition with respect to q_{od} for an interior solution is given by[7]

$$f_d'\left(\mathrm{Q}_d\right) - \mathrm{WTA}_o\left(\mathrm{Q}_o\right) + q_{od} \cdot \mathrm{WTA}_o'\left(\mathrm{Q}_o\right) = 0. \tag{5}$$

Substituting (1) into (5) and rewriting yields

$$\begin{aligned} \mathrm{WTP}_d\left(\mathrm{Q}_d\right) &= \mathrm{WTA}_o\left(\mathrm{Q}_o\right) - q_{od} \cdot \mathrm{WTA}_o'\left(\mathrm{Q}_o\right) \\ &\geq \mathrm{WTA}_o\left(\mathrm{Q}_o\right). \end{aligned} \tag{6}$$

Note that in a partial equilibrium of perfect competition the inequality in (6) would hold with equality for every pair of trading districts, see e.g. Hagerty (2019). Under buyers' market power (and $-q_{od} \cdot \mathrm{WTA}_o'(\mathrm{Q}_o) > 0$), the willingness to pay will be larger than the willingness to accept for every pair of trading districts while such power keeps markets away from equality. The wedge between buyers' WTP and sellers' WTA reflects the possible price range for each transaction.

As the last step we substitute (3) into (6) and rewrite to obtain the following system of equilibrium conditions that we will use in Section 4 in our estimations:

$$p_{od} = \mathrm{WTA}_o\left(\mathrm{Q}_o\right), \tag{7a}$$

$$p_{od} = \mathrm{WTP}_d\left(\mathrm{Q}_d\right) + q_{od} \cdot \mathrm{WTA}_o'\left(\mathrm{Q}_o\right). \tag{7b}$$

[7] For simplicity we omit the first-order condition for pairs of districts that do not trade. To derive this first-order condition for pairs of trading districts we use the chain rule to obtain $\frac{\partial}{\partial q_{od}} f_d(\mathrm{Q}_d) = \frac{\partial}{\partial \mathrm{Q}_d} f_d(\mathrm{Q}_d) \cdot \frac{\partial \mathrm{Q}_d}{\partial q_{od}} = f_d'(\mathrm{Q}_d)$. Likewise, we obtain $\frac{\partial}{\partial q_{od}} \mathrm{WTA}_o(\mathrm{Q}_o) = \frac{\partial}{\partial \mathrm{Q}_o} \mathrm{WTA}_o(\mathrm{Q}_o) \cdot \frac{\partial \mathrm{Q}_o}{\partial q_{od}} = -\mathrm{WTA}_o'(\mathrm{Q}_o)$.

Recall that $\text{WTA}'_o(Q_o) < 0$ such that the last term of (7b) is negative. This markdown makes the market price lower than the buyer's willingness to pay.

In many economic contexts, it is standard to construct measures of market power based on relative markdowns or markups.[8] We, therefore, rewrite our pair-specific relative markdowns to construct a novel pair-specific measure of heterogeneous market power given by

$$\frac{\text{WTP}_d(Q_d) - \text{WTA}_o(Q_o)}{\text{WTA}_o(Q_o)} = -\frac{q_{od}}{Q_o} \cdot \frac{Q_o \text{WTA}'_o(Q_o)}{\text{WTA}_o(Q_o)}. \tag{8}$$

It consists of the quantity elasticity of the seller's WTA weighted by the ratio of transaction volume to this seller's total water use. This elasticity is the reciprocal of seller o's individual price elasticity of demand. Our estimate of the system of equations (7) provides us with an estimate of $\text{WTA}'_o(Q_o)$ and $\text{WTA}_o(Q_o)$ from which we calculate our measure of market power (8) to address such power in California's water market.

Our model is illustrated in Figure 1. With two types of districts (buyers and sellers) and one good (water), whose supply is given, our model is an endowment economy and so we can visualize it in a chart with a secondary mirrored vertical axis, while the total available water is on the horizontal axis. Demand for water is displayed using the $\text{WTA}_o(Q_o)$ curve for sellers and the $\text{WTP}_d(Q_d)$ curve for buyers. Starting from water endowments e_o and e_d in Figure 1, water transactions increase buyers' water consumption and decrease sellers' water consumption, while closing the wedge between buyers' WTP and sellers' WTA. Compared with the competitive equilibrium, buyer power implies a lower transaction volume, which leaves a positive wedge, as discussed in this section and as illustrated in the figure.

The necessary and sufficient condition for a positive bilateral trade in the Nash-Cournot equilibrium of Figure 1 is that $\text{WTP}_d(e_d) > \text{WTA}_o(e_o)$, or $f'_d(e_d) > f'_o(e_o)$. The interpretation is that the marginal benefit of water use at the initial endowment of each destination exceeds the marginal benefit of water use at the initial endowment of each origin. In other words, trade must be (marginally) beneficial at the initial endowment levels.

A relevant indicator for water authorities is the (relative) damage caused by market power. Similar to the Herfindahl-Hirschman Index used by antitrust authorities, we investigate – as an easy approximation of individual damage – the markdown times the quantity sold: $\big(\text{WTP}_d(Q_d) - \text{WTA}_o(Q_o)\big) \cdot q_{od}$. This damage consists of the rectangle bordered by

[8] Closely related measures are the (markup) Lerner Index used in Industrial Organization and the markdown of monopsony power used in Labor Economics. For homogeneous static labor supply ℓ, revenue $R(\ell)$ and endogenous wage $w(\ell)$, the latter states that $\frac{R'(\ell) - w(\ell)}{w(\ell)} = \frac{\ell \cdot w'(\ell)}{w(\ell)}$ (see e.g. Ashenfelter et al., 2010). Note the difference with (8) where we add the term $-\frac{q_{od}}{Q_o}$ to adjust for heterogeneity between pairs of trading districts.

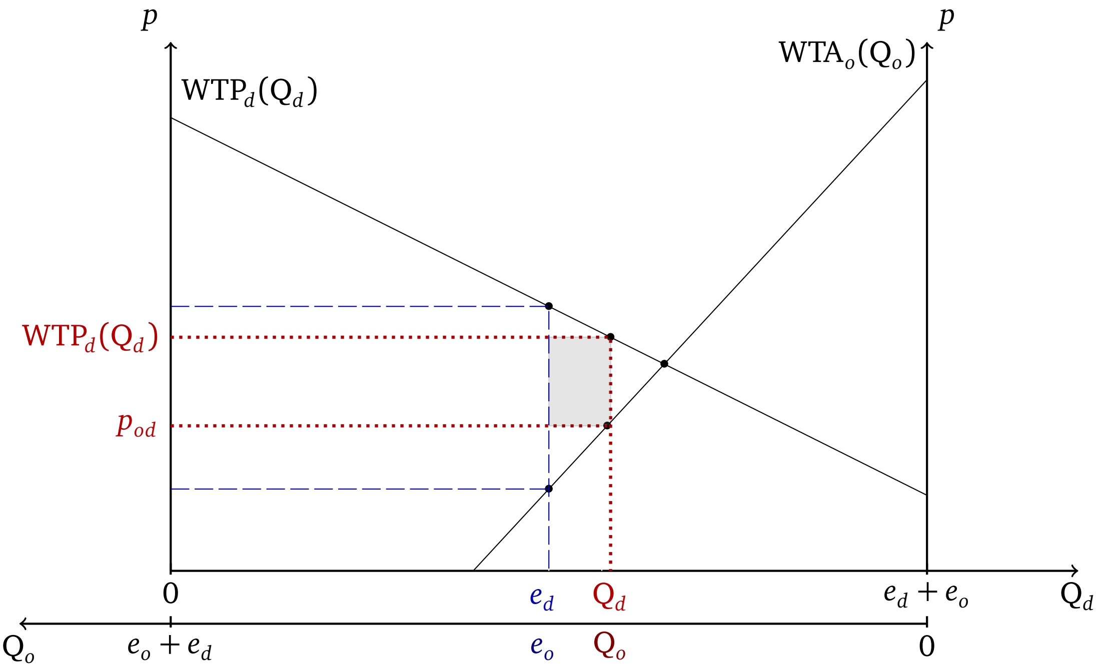

Figure 1: Stylized visualization of endowments (dashed, blue) and Nash-Cournot equilibrium (dotted, red), where $\mathrm{WTP}_d(Q_d) - p_{od}$ equals the markdown, see (7b). The shaded area is an approximation of damage due to market power.

three red dotted lines and the blue dotted line in Figure 1. Next, we relate this damage to the monetary value of own water use by individual sellers by dividing by $p_{od} \cdot Q_o$. Doing so and making use of (3) and the measure of market power (8), yields as approximated individual relative damage:[9]

$$-\left(\frac{q_{od}}{Q_o}\right)^2 \cdot \frac{Q_o \mathrm{WTA}'_o(Q_o)}{\mathrm{WTA}_o(Q_o)}. \tag{9}$$

Aggregate relative damage follows after summing over all sellers in the market. The latter can be seen as a modified Herfindahl-Hirschman Index for water markets. Further investigation of this index is left for future research.

[9]Our measure resembles the individual contribution to industry's damage relative to industry's revenue in deriving the Herfindahl-Hirschman Index for Cournot oligopolies in Hirschman (1945). In this index, individual market shares (in quantities) are squared and divided by the price elasticity of market demand and then summed over all firms. In the Cournot model, both the elasticity and aggregate production are the same across firms and then the sum of individual damage can be rewritten as a simple formula. In water markets, heterogeneity in both Q_os and elasticities rule out a simple formula.

3.2 Main specification

The preferred specification of our model uses quadratic benefit functions for both buyers and sellers. This specification allows us to estimate a linear model, as explained in Section 4. Our proposed benefit functions allow for heterogeneity across buyers and sellers as well as over time, which is why we add time subscripts from here on.

For each origin we have $f_{ot}(Q_{ot}) = Q_{ot}(\alpha_{ot} - \frac{1}{2}\delta Q_{ot})$, where $\alpha_{ot} = \phi_o + \beta_t + \nu_{ot}$ captures heterogeneity in productivity between different sellers and time periods, while parameter δ is kept constant. This benefit function implies that $f'_{ot}(Q_{ot}) = \alpha_{ot} - \delta Q_{ot}$, which is the sellers' WTA in (2). Similarly, for each destination we have $f_{dt}(Q_{dt}) = Q_{dt}(a_{dt} - \frac{1}{2}\gamma Q_{dt})$, with $a_{dt} = \psi_d + \beta_t + u_{dt}$ and therefore $f'_{dt}(Q_{dt}) = a_{dt} - \gamma Q_{dt}$, which is the buyers' WTP in (1). Note that in Appendix C we generalize our main model specification to also allow heterogeneous pair-specific transaction costs. We do so after presenting the solution to our main model specification under homogeneity in Appendix A and under heterogeneity in productivity in Appendix B.

Recall that the necessary and sufficient condition for positive quantities in the Nash-Cournot equilibrium is that $f'_{dt}(e_{dt}) > f'_{ot}(e_{ot})$, which implies

$$a_{dt} - \gamma e_{dt} > \alpha_{ot} - \delta e_{ot}. \tag{10}$$

The interpretation is that the marginal benefit of water use at the initial endowment of each destination exceeds the marginal benefit of water use at the initial endowment of each origin. In other words, trade is (marginally) beneficial at the initial endowment levels.

4 Empirical strategy

The objective of our empirical exercise is to measure market power in California's water market. We do so using our measure for market power (8). Calculation of this index requires an estimate of $\text{WTA}'_{ot}(Q_{ot})$. For the linear model specification introduced in Section 3.2, we have $\text{WTA}'_{ot}(Q_{ot}) = f''_{ot}(Q_{ot}) = -\delta$, which we will estimate using the system of equations (7). Note that this parameter δ is the only estimate that we need according to our markdown measure of market power (8). To see this, note that our linear model specification with buyer power allows us to write this index in terms of δ as well as transaction prices and quantities p_{odt} and q_{odt}, which are present in our transaction data:

$$\begin{aligned}\frac{\mathrm{WTP}_{dt}(Q_{dt}) - \mathrm{WTA}_{ot}(Q_{ot})}{\mathrm{WTA}_{ot}(Q_{ot})} &= -\frac{q_{odt}}{Q_{ot}} \cdot \frac{Q_{ot}\mathrm{WTA}'_{ot}(Q_t)}{\mathrm{WTA}_{ot}(Q_{ot})} \\ &= \delta \cdot \frac{q_{odt}}{p_{odt}}. \end{aligned} \tag{11}$$

Below, we present our empirical strategy to estimate parameter δ.

Given our panel data on transaction prices and quantities, we construct a fixed effects model, which exploits variation in observed transaction prices, WTA, and WTP across trading districts and across time. This approach rests on two requirements. The first is that we have sufficient variation in WTA and WTP over time. In our data, such variation over time is caused by variation in water endowments over time, which imply movements along the benefit function of water use, thereby changing districts' marginal benefits of water use. Water endowments are determined by the interaction of weather fluctuations with historically-determined allocation rules, which are markedly different across regions of California. The second requirement is that WTA and WTP are exogenous, conditional on unobserved district characteristics (as captured by the fixed effects). We meet this requirement by assumption, since our model dictates that WTA (and, implicitly, WTP) determines transaction prices.

There are two possible sources of endogeneity in our data, one of which is that omitted variables may cause biases. Ideally, we would control for these using both year fixed effects as well as time-invariant district-by-counterparty fixed effects (the counterparty being the other district involved in the transaction). The latter would capture any variation in prices caused by unobserved heterogeneity across pairs of trading districts. Unfortunately, we do not have sufficient observations per trading district-pair to estimate such fixed effects. We resort to separate seller- and buyer fixed effects instead. The second possible source of endogeneity is reverse causality, which we discuss at the end of this section.

We substitute the linear specification of our model into the system of equations (7):

$$p_{odt} = -\delta Q_{ot} + \phi_o + \beta_t + v_{ot}, \tag{12a}$$

$$p_{odt} = -\gamma Q_{dt} - \delta q_{odt} + \psi_d + \beta_t + u_{dt}. \tag{12b}$$

An implicit assumption underlying the regression of *individual* transaction prices on (some function of) *total* water use levels is that districts do not hedge against the risk of water shortage within each year. One example would be that districts buy 'too much' water and will try to re-sell later that same year. We find, however, that only a handful of districts in our dataset have ever been active on both sides of the market within one year. Hence, this assumption of no risk hedging seems warranted. It is also consistent with the situation in many Western US watersheds, where predictions on water availability in early spring provide ‘reasonably accurate forecasts’ of actual availability (Draper, 2001). Although districts face some degree of uncertainty before the start of the growing season as water allocation percentages are being updated, this uncertainty does not appear to affect their trading behavior.

Without risk hedging, price differences across transactions for a particular district and year should not occur, except in the case of transaction costs, or market power. In model variations, we, therefore, control for various types of transaction costs, as introduced in Section 2. Transaction costs are pair-specific and time-invariant, and they apply to either the seller or the buyer in a specific transaction as summarized in Table 1. In the regressions below, transaction costs are included as $T_{odr} = \tau_r C_{odr} + \tau_o + \tau_d + \varepsilon_{odr}$, where vector C_{odr} includes seller-, buyer-, and pair-specific transaction costs, with units (mostly dummies) as presented in Table 1, while τ_o and τ_d represent district fixed costs. Since each transaction is assessed twice in this system, we add r to indicate whether the transaction is assessed from the seller's perspective ($r = 0$) or the buyer's perspective ($r = 1$). We expect $\tau_r \geq 0$ if $r = 0$ and $\tau_r \leq 0$ if $r = 1$. That is, transaction costs enter sellers' WTA positively, because these have to be compensated for sellers on top of the sellers' net benefits, while transaction costs enter the buyers' WTP negatively, because these decrease sellers' net benefits.

We add transaction costs to (12) and re-order and re-label terms:

$$\begin{aligned} p_{odrt} &= -\delta Q_{ot} + (\phi_o + \tau_o) + \tau_d + \beta_t + \tau_r C_{odr} + (\nu_{ot} + \varepsilon_{odr}) \\ &= -\delta Q_{ot} + \phi_o + \psi_d + \beta_t + \tau_r C_{odr} + \epsilon_{odrt} \end{aligned} \quad (13a)$$

$$\begin{aligned} p_{odrt} &= -\gamma Q_{dt} - \delta q_{odt} + \tau_o + (\psi_d + \tau_d) + \beta_t + \tau_r C_{odr} + (u_{dt} + \varepsilon_{odr}) \\ &= -\gamma Q_{dt} - \delta q_{odt} + \phi_o + \psi_d + \beta_t + \tau_r C_{odr} + \epsilon_{odrt} \end{aligned} \quad (13b)$$

Note that coefficient δ appears in both equations. We estimate both equations simultaneously by constructing two variables, R^o_{odtk} and R^d_{odtk}, that combine the coefficients on water use from (13). We also add a counter k, since there can be multiple transactions between one origin o and one destination d within one year t:

$$R^o_{odtk} = \begin{cases} Q_{ot} & \text{if } r = 0 \\ q_{odtk} & \text{if } r = 1, \end{cases} \quad \text{and} \quad R^d_{odtk} = \begin{cases} 0 & \text{if } r = 0 \\ Q_{dt} & \text{if } r = 1. \end{cases}$$

The combined regression equation, which also suppresses the intercept, is:

$$p_{odrtk} = -\delta \mathrm{R}^{o}_{odtk} - \gamma \mathrm{R}^{d}_{odtk} + \phi_o + \psi_d + \beta_t + \tau_r \mathrm{C}_{odr} + \epsilon_{odrtk}. \tag{14}$$

In the next section, we will estimate variations of (14) using linear regression.

Unlike standard models of supply and demand, we are estimating a system with two demand functions (with slopes given by parameters γ and δ), while the annual supply of water is determined by rainfall and snow-melt. With hydrological variation between years, the total amount of water in the system changes exogenously each year. Summed over all districts, annual supply cannot respond to changes in price.

Despite this exogeneity in supply, individual districts may still respond to price changes by changing the volume of water bought or sold. We therefore also estimate (14) while instrumenting for water use with districts' water allocations in each year. Water allocations are the product of a time-invariant maximum entitlement for each district and a year-varying allocation percentage. Allocation percentages are determined bureaucratically on the basis of environmental conditions (i.e., rainfall and snowmelt in the mountains during the preceding winter) and do not respond to demand-driven factors. They are determined separately for each of 13 categories of water entitlements (for more details see Hagerty (2019) and Hagerty (2021)). With district and year fixed effects, identifying variation comes from fluctuations in each district's annual water endowment, relative to other districts in the same year. In line with Hagerty (2019), we argue that it is reasonable to assume changes in water allocations reflect pure supply shocks and are uncorrelated with changes in demand. [10]

5 Results

The estimates of regression equation (14) are shown in Table 2. Recall that the aim of this regression is to estimate the impact of market power on transaction prices via the markdown $\mathrm{WTP}_{dt}(\mathrm{Q}_{dt}) - \mathrm{WTA}_{ot}(\mathrm{Q}_{ot})$. Applying a model with quadratic benefit functions implies that *Seller water use* (i.e., $\mathrm{R}^{\mathrm{O}}_{odtk}$) is one of the independent variables, whose coefficient gives the slope of the sellers' benefit function, parameter δ. Multiplied by transaction volume, this parameter gives the markdown for each transaction.

In model (1), estimated using OLS, we model water use as our only explanatory variable, combined with seller-, buyer-, role-, and year fixed effects. The coefficient on *Seller water*

[10] Hagerty (2021) shows that the effect of surface water supplies on crop choice is precisely unaffected by the inclusion or exclusion of flexible controls for weather in the same or preceding year.

Table 2: Estimating WTA and WTP: Linear model

Price (2010$/AF)	(1)	(2)	(3)	(4)
	OLS	OLS	IV	IV
Seller water use (1,000 AF)	-0.0183^{**}	-0.0280^{**}	-0.580^{***}	-0.932^{**}
(coefficient $-\delta$)	(0.00821)	(0.0130)	(0.208)	(0.445)
Buyer water use (1,000 AF)	-0.00757^{**}	-0.0144^{**}	-0.309^{**}	-0.311^{**}
(coefficient $-\gamma$)	(0.00336)	(0.00612)	(0.132)	(0.156)
Seller fixed effects	✓	✓	✓	✓
Buyer fixed effects	✓	✓	✓	✓
Year fixed effects	✓			
Quadratic time trend		✓	✓	✓
Transaction costs				✓
# Observations	1034	1034	879	877
# Clusters	543	337	308	307
# FE dummies	212	190	164	163
Cragg-Donald F-statistic			9.936	8.681
First-stage (Seller water use):				
Seller entitlements (%)			128.3^{***}	225.6^{**}
			(26.75)	(75.20)
Buyer entitlements (%)			−76.92	−31.46
			(60.13)	(54.81)
First-stage (Buyer water use):				
Seller entitlements (%)			6.255	-227.4^{*}
			(47.77)	(97.74)
Buyer entitlements (%)			531.7^{***}	519.7^{***}
			(134.6)	(142.8)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.
Coefficient estimates from fixed effects models using OLS and IV. Standard errors in parentheses, clustered by seller, buyer, and year (but only by seller and buyer in models (2)–(4) where year fixed effects are replaced by a quadratic time trend).

use implies that $\delta = 0.0183$, which is more than double the size of $\gamma = 0.00757$, implied by the coefficient on *Buyer water use*. The difference indicates that selling districts have steeper demand curves than buying districts. In model (2) we attempt to improve efficiency of these estimates. Given the large number of clusters compared to observations, we replace year fixed effects by a time trend. No comparable simplification was found feasible for the other fixed effects. Particularly, there is no obvious possibility to replace seller- and buyer fixed effects with a coarser set of dummy variables. As a result of replacing the year fixed effect by the time trend, the number of clusters decreases sharply. Compared to model (1), the model (2) estimates for both δ and γ increase significantly.

In model (3) we instrument water use by districts' water entitlements. The resulting estimates of δ and γ increase sharply, in absolute terms, compared to those of models (1) and (2). The fact that the IV estimates are so much greater in magnitude than the OLS estimates suggests that the OLS estimates are biased by selection into trading. In the OLS estimates, prices appear to respond minimally to changes in water use – but this water use is observed after water users have already conducted endogenous water transactions. In contrast, the IV estimates isolate variation in water use driven by exogenous changes in water supply, so it more accurately identifies the slope of the benefit function. Finally, in model (4), we add seller- and buyer-specific transaction costs, which do not appear to improve the model results, decreasing the F-statistic and increasing the standard error of our main coefficient of interest. The Cragg-Donald F-statistics for the first stages of the IV models are both higher than their critical values as reported by Stock and Yogo (2005), suggesting that models (3) and (4) do not suffer from weak instruments. Variations of models (3) and (4) featuring year-fixed effects rather than a time trend would yield a large number of clusters compared to observations, similar to model (1). Such model variations would deteriorate efficiency and would turn the estimates of δ and γ insignificant.

Based on these model results and interpretation, our preferred model is model (3) and we use the main coefficient of interest from this specification, $\delta = 0.580$, in the remainder of this section.[11] The interpretation of δ is that sellers' WTA, which equals the water price in our model, increases by \$0.58/AF for each 1,000 AF sold.[12] More important for our analysis, however, is that δ is used to calculate the markdown $\text{WTP}_{dt}(Q_{dt}) - \text{WTA}_{ot}(Q_{ot}) = \delta q_{od}$. Doing so we find that the average markdown, after removing one outlier, equals \$4.60/AF (SD=8.66). This markdown corresponds to about 6.4% of the transaction price, on average, with markedly higher markdowns (both in absolute and relative terms) for transactions with low prices. We use transaction-specific markdowns to compute our measure of market power (11) and plot these in Figure 2. This figure shows that our measure is relatively low. It is markedly higher, though, for a small set of transactions with low prices, which also tend to have the highest transaction volumes; a substantial share of these transactions is water transferred from agriculture to environmental use. All in all, we find that market power is relatively low in California's water market.

[11] Another reason to prefer one of the larger coefficients is that we are less likely to underestimate the impact of market power. Since our main result is that market power is rather limited, this is a conservative choice.

[12] AF: acre-foot. One acre-foot equals 1,233 m^3.

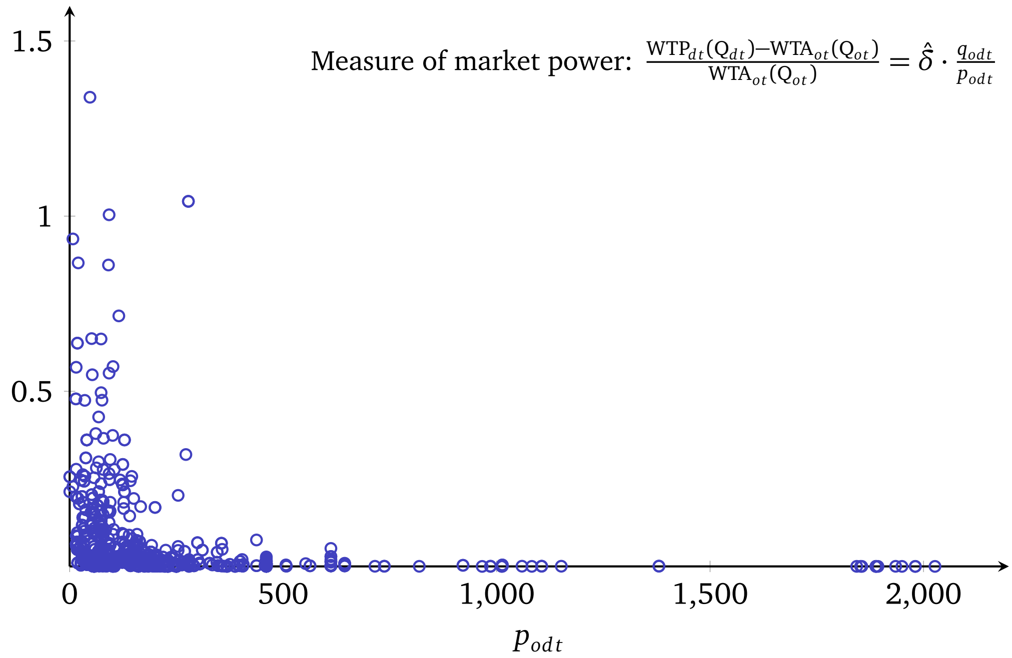

Figure 2: Scatter-plot of transaction prices and the measure of market power as given by (11) (one outlier removed).

6 Robustness

In this section, we report on six robustness checks. First, we use a conjectural variations approach. Second, we check robustness when we focus on relevant sub-samples of the data. Third, we apply an alternative model specification featuring non-linear benefit functions. Fourth, we alter the calculation of districts' water use to account for the timing of transactions within one year. Fifth, we check whether selling and buying districts can be reasonably assumed to have similar benefit functions. Finally, despite the results of our conjectural variations approach, we estimate a model with seller power.

Note that this list of robustness checks is not exhaustive. Importantly, we also checked for differential levels of market power. One such example would be differential market power occurring in wet versus dry years. In wet years, one could imagine that buyers have better opportunities to exercise market power. Using the Sacramento Valley Water Year Hydrological Classification Index to classify years, we fail to find such differences. Another option is differential market power depending on the location of buyers and sellers. The argument would be that buyers that are more central would have more opportunities to switch to another seller and could therefore achieve higher markdowns. This argument ignores, however, that the Californian water market features an almost complete hydrological network enabling water transfers between nearly any two districts.

As a result, while central buyers would probably face lower transaction costs, they do not have increased opportunities to exercise market power compared to buyers at the periphery.

6.1 A conjectural variations approach

We proceed to compare our results to those obtained using a conjectural variations approach in order to verify whether our assumption of buyer power is warranted. In this approach, the term expressing market power is multiplied by some weight that dampens this term. A recent example that employs this approach and analyzes Californian groundwater is Bruno and Sexton (2020). Accordingly, we introduce conjectural variations using parameter $\theta \in [0,1]$ that measures the degree of buyer power, while $\xi \in [0,1]$ measures the degree of seller power. Conceptually, at most one of these parameters can be positive and the other is equal to zero. Nevertheless, in order to avoid repetition of arguments, we include both parameters simultaneously in the derivation below. We rewrite (7) to include these market power weights:

$$p_{odt} = \mathrm{WTA}_{ot}\left(\mathrm{Q}_{ot}\right) - \xi \cdot q_{odt} \cdot \mathrm{WTP}'_{dt}\left(\mathrm{Q}_{dt}\right), \tag{15a}$$

$$p_{odt} = \mathrm{WTP}_{dt}\left(\mathrm{Q}_{dt}\right) + \theta \cdot q_{odt} \cdot \mathrm{WTA}'_{ot}\left(\mathrm{Q}_{ot}\right). \tag{15b}$$

The new terms capture districts' expectations about other districts' reactions to a change in transaction quantities. These expectations may tend to the expected reactions under perfect competition versus settings with buyer or seller power. As a result, the maximum possible markups or markdowns are dampened by, respectively, θ or ξ. In our analysis so far we have assumed $(\theta, \xi) = (1, 0)$, i.e. only buyer power. Two other special cases of the model are seller power – which would imply $(\theta, \xi) = (0, 1)$ – and perfect competition, which would imply $(\theta, \xi) = (0, 0)$.

We proceed to estimate this system of equations. The resulting values of θ and ξ will verify whether our assumption of buyer power is warranted using this conjectural variations approach. Taking similar steps as before, we first substitute the linear model specification:

$$\begin{aligned} p_{odrt} = \alpha - \delta \mathrm{Q}_{ot} + \gamma q_{odt} - (1-\xi)\gamma q_{odt} & \\ + \phi_o + \psi_d + \beta_t + \tau_r \mathrm{C}_{odr} + \epsilon_{odrt}, & \end{aligned} \tag{16a}$$

$$\begin{aligned} p_{odrt} = a - \delta q_{odt} - \gamma \mathrm{Q}_{dt} + (1-\theta)\delta q_{odt} & \\ + \phi_o + \psi_d + \beta_t + \tau_r \mathrm{C}_{odr} + \epsilon_{odrt}. & \end{aligned} \tag{16b}$$

The combined regression equation becomes:

$$p_{odrtk} = -\delta \mathrm{R}^{o}_{odtk} - \gamma \tilde{\mathrm{R}}^{d}_{odtk} + (1-\theta)\delta \hat{\mathrm{R}}^{o}_{odtk} - (1-\xi)\gamma \hat{\mathrm{R}}^{d}_{odtk} + \phi_o + \psi_d + \beta_t + \tau_r \mathrm{C}_{odr} + \epsilon_{odrtk} \quad (17)$$

with R^{o}_{odtk} as defined in Section 4, while $\tilde{\mathrm{R}}^{d}_{odtk}$, $\hat{\mathrm{R}}^{o}_{odtk}$ and $\hat{\mathrm{R}}^{d}_{odtk}$ are defined as follows:

$$\tilde{\mathrm{R}}^{d}_{odtk} = \begin{cases} -q_{odtk} & \text{if } r = 0 \\ \mathrm{Q}_{dt} & \text{if } r = 1, \end{cases} , \quad \hat{\mathrm{R}}^{o}_{odtk} = \begin{cases} 0 & \text{if } r = 0 \\ q_{odtk} & \text{if } r = 1, \end{cases} , \quad \hat{\mathrm{R}}^{d}_{odtk} = \begin{cases} q_{odtk} & \text{if } r = 0 \\ 0 & \text{if } r = 1. \end{cases}$$

In order to get a clear view on the parameters of interest, we apply extremum estimation of the IV criterion function, transformed such that we optimize our parameters δ, γ, θ and ξ. Table 3 reports the results for the case where the zero-one interval restrictions on the parameters in the IV criterion function were dropped in order to obtain precise estimates. We present three models. In model (1) we allow only buyer power and in model (2) only seller power. Both models are motivated by standard practice with an ex-ante subjective choice on which side holds market power. Ambiguity would arise if both models were to indicate some degree of market power. For this reason, we also include the novel and mathematically motivated "let the data speak" model (3) in which both market power parameters are allowed to take a non-zero value. Ideally, the estimation results would suggest which side holds market power by forcing the other side's parameter to a value indicating no market power. In addition to the coefficients on seller and buyer water use, $-\delta$ and $-\gamma$, we report coefficients on both market power weights, θ and ξ, while suppressing the coefficients on the terms $\hat{\mathrm{R}}^{o}_{odtk}$ and $\hat{\mathrm{R}}^{d}_{odtk}$, since these coefficients are combinations of the four parameters that are already reported.

The results in Table 3 support our assumption of buyer power. Starting with models (1) and (2), estimates for $-\delta$ and $-\gamma$ are very close to those obtained in our preferred model (3) of Table 2. The unrestricted estimates for seller and buyer power weight, ξ and θ, are found to lie outside the bounds of $[0,1]$. We find that the buyer power weight θ is larger than 1 in model (1) while the seller power weight ξ is negative in model (2). These results are both in a direction that is consistent with buyer market power. Imposing the restrictions that either $[\xi \in [0,1], \theta = 0]$, or $[\xi = 0, \theta \in [0,1]]$ leads to the expected result: estimates at the closest boundaries of these intervals. Given that the amount of observations in the dataset is relatively low compared to the number of parameters and fixed effects, and hence both estimates and standard deviations cannot be extracted with too great precision, we do not report these results but take them as an indication that buyer power is the most reasonable assumption.

The estimation results for model (3), which allows both market power parameters to

Table 3: Estimating WTA and WTP: Conjectural variations

Price (2010$/AF)	(1)	(2)	(3)
	IV Buyer power ($\xi = 0$)	IV Seller power ($\theta = 0$)	IV Both
Seller total water use (1,000 AF) (coefficient $-\delta$)	-0.577^{***} (0.162)	-0.610^{***} (0.180)	-0.573^{***} (0.162)
Buyer total water use (1,000 AF) (coefficient $-\gamma$)	-0.242^{***} (0.070)	-0.282^{***} (0.082)	-0.241^{***} (0.070)
Seller power weight (coefficient ξ)		-8.581^{**} (3.570)	1.251 (2.422)
Buyer power weight (coefficient θ)	5.664^{***} (0.658)		6.155^{***} (1.237)
Seller fixed effects	✓	✓	✓
Buyer fixed effects	✓	✓	✓
Year fixed effects			
Quadratic time trend	✓	✓	✓
Transaction costs			
# Observations	879	879	879
# FE dummies	164	164	164

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.
Coefficient estimates from fixed effects models using extremum estimation. The covariance matrix is computed as a robust sandwich covariance matrix, following the theory of extremum estimation (Cameron and Trivedi, 2005, Section 6.3.4). Standard errors in parentheses. Models (1)–(3) correspond to model (3) of Table 2, but using the conjectural variations approach. First-stage regressions are omitted for brevity.

take a non-zero value, support this result. Estimates for $-\delta$ and $-\gamma$ are, again, very close to those obtained in our preferred model (3) of Table 2. In addition, while the unrestricted estimate for seller power weight ξ is statistically not different from zero, the unrestricted estimate for buyer power weight θ is close to the buyer power estimate of model (1). All three models combined, we take these results as an indication that buyer power is the most reasonable assumption.

6.2 Sub-sample analysis

We repeat our preferred model (3) of Table 2 for three sub-samples of interest. Table 4 shows the results of these additional regressions. For reference, we include the preferred model as model (1). In model (2) we drop all observations that involve water for environmental use, for instance, buy-backs by the government. Arguably, such transactions are markedly different from transactions between districts that intend to use the water for consumptive purposes. In model (3) we include only transactions where agricultural districts are selling, which seem to represent the smaller, weaker actors in the market. Unfortunately, our sample size does not allow us to focus only on water sales from agricultural to urban districts (slightly more than 100 transactions), which seem to represent the larger, stronger actors capable of exercising market power (cf. Isaaks and Colby, 2020). By focusing on all sales from agricultural districts, we may still capture the fact that agricultural districts may have less market power than the other types of districts. Note that half of these sales are to other agricultural districts, while the other half is shared roughly equally between buying urban districts and environmental projects. In model (4) we drop outlier transactions. We exclude the 5% transactions with lowest and 5% transactions with highest transaction prices and similarly for transaction volumes.

Table 4: Estimating WTA and WTP: Sub-samples

Price (2010$/AF)	(1)	(2)	(3)	(4)
	IV preferred	IV no env	IV ag sellers	IV no outliers
Seller water use (1,000 AF)	-0.580^{***}	-0.520^{**}	-0.570^{***}	-0.442^{***}
(coefficient $-\delta$)	(0.208)	(0.201)	(0.179)	(0.142)
Buyer water use (1,000 AF)	-0.309^{**}	-0.278^{**}	-0.284^{***}	-0.232^{**}
(coefficient $-\gamma$)	(0.132)	(0.127)	(0.108)	(0.0906)
Seller fixed effects	✓	✓	✓	✓
Buyer fixed effects	✓	✓	✓	✓
Quadratic time trend	✓	✓	✓	✓
# Observations	879	728	778	737
# Clusters	308	248	261	250
# FE dummies	164	149	133	132
Cragg-Donald F-statistic	9.936	7.057	14.01	10.82

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.
Coefficient estimates from fixed effects models using IV. Standard errors in parentheses, clustered by seller and buyer. Model (1) corresponds to our preferred model (3) of Table 2. In model (2) we drop transactions from or to environmental use. In model (3) we keep only transactions where agricultural districts are selling. In model (4) we drop transactions that are outliers in terms of price or volume. First-stage regressions are omitted for brevity.

Coefficients of sub-sample Models (2)–(4) are not statistically different from those of the preferred model. Model (2), which discards 17% of the observations, performs similarly in terms of precision and slightly worse in terms of the F-statistic. Unexpectedly, Model (3) does not show a higher coefficient (in absolute terms). Hence, there is no indication of more buyer power when buying from an agricultural district. Model (4) suggests that some of the market power we find is driven by outlier transactions in terms of price or volume, as one could expect. Combined, these additional regressions show that our main results are robust to including only specific sub-samples of interest.

6.3 Constant price elasticity

In the main specification of our model, we have imposed a constant slope of the benefit functions and a variable price elasticity. In this section, we impose instead that these functions have a constant price elasticity and, consequently, a variable slope. In particular, we consider the class of non-linear WTA functions that are homogeneous[13]. Given our earlier assumption of differentiability, we have that Euler's Homogeneous Function Theorem[14] applies to the equilibrium conditions and our measure of market power.

For arbitrary homogeneous $\mathrm{WTA}_{ot}(\mathrm{Q}_{ot})$ of order $-\kappa_o$, we can rewrite the markdown $\mathrm{WTP}_{dt}(\mathrm{Q}_{dt}) - \mathrm{WTA}_{ot}(\mathrm{Q}_{ot})$ as

$$-q_{odt} \cdot \mathrm{WTA}'_{ot}(\mathrm{Q}_{ot}) = -\frac{q_{odt}}{\mathrm{Q}_{ot}} \cdot \left[\mathrm{Q}_{ot} \cdot \mathrm{WTA}'_{ot}(\mathrm{Q}_{ot})\right] = \frac{q_{odt}}{\mathrm{Q}_{ot}} \cdot \left[\kappa_o \cdot \mathrm{WTA}_{ot}(\mathrm{Q}_{ot})\right]. \tag{18}$$

This implies that our measure of market power (8) can be updated to

$$\frac{\mathrm{WTP}_{dt}(\mathrm{Q}_{dt}) - \mathrm{WTA}_{ot}(\mathrm{Q}_{ot})}{\mathrm{WTA}_{ot}(\mathrm{Q}_{ot})} = \frac{q_{odt}}{\mathrm{Q}_{ot}} \cdot \kappa_o. \tag{19}$$

The empirical strategy to estimate κ_o has many similarities to the empirical strategy proposed in Section 4, and we refer to Appendix D for details. The resulting regression equation becomes:

$$\ln p_{odrtk} = -\kappa_o \overline{\mathrm{R}}^o_{odtk} - \kappa_d \overline{\mathrm{R}}^d_{odrtk} + \phi_o + \psi_d + \beta_t + \ln \tau_r \mathrm{C}_{odr} + \epsilon_{odrtk}, \tag{20}$$

where $\overline{\mathrm{R}}^o_{odtk}$ and $\overline{\mathrm{R}}^d_{odtk}$ are modifications of R^o_{odtk}, respectively, R^d_{odtk} that are defined in Appendix D.

We estimate variations of (20) using linear regression, similar to Table 2 for our main model specification. Table 5 shows the estimates of four models that are similar to models (1)–(4) of Table 2. Despite allowing for non-linear benefit functions, the models of Table 5 do not perform better than those of our linear model specification in Table 2 in terms of the Cragg-Donald F-statistic nor the precision of our coefficient of interest, the coefficient on *Seller water use*.

Again, we use model (3) to derive the main coefficient of interest for this model specification, $\kappa_o = 0.370$. Similar as before, we use this coefficient to calculate the markdown $\mathrm{WTP}_{dt}(\mathrm{Q}_{dt}) - \mathrm{WTA}_{ot}(\mathrm{Q}_{ot}) = \frac{q_{odt}}{\mathrm{Q}_{ot}} \cdot [\kappa_o \cdot \mathrm{WTA}_{ot}(\mathrm{Q}_{ot})]$, which now also depends on the ratio $\frac{q_{odt}}{\mathrm{Q}_{ot}}$. We find that the mean value of this ratio is heavily skewed by 15 districts

[13]The function $f : \mathrm{R} \to \mathrm{R}$ is homogeneous of order $\kappa \in \mathrm{R}$ if $f(\mu x) = \mu^k f(x)$ for all x and $\mu > 0$.

[14]Let the function $f : \mathrm{R} \to \mathrm{R}$ be homogeneous of order $\kappa \in \mathrm{R}$. Euler's Homogeneous Function Theorem states that $x \cdot f'(x) = \kappa\, f(x)$.

Table 5: Estimating WTA and WTP: Constant price elasticity

Log price (2010$/AF)	(1)	(2)	(3)	(4)
	OLS	OLS	IV	IV
Log (seller water use, 1,000 AF)	−0.0000147	−0.000660	−0.370**	−0.623*
(coefficient $-\kappa_o$)	(0.000224)	(0.000670)	(0.173)	(0.357)
Log (buyer water use, 1,000 AF)	−0.000576	−0.00185*	−0.402**	−0.458
(coefficient $-\kappa_d$)	(0.000568)	(0.00106)	(0.194)	(0.279)
Seller fixed effects	✓	✓	✓	✓
Buyer fixed effects	✓	✓	✓	✓
Year fixed effects	✓			
Quadratic time trend		✓	✓	✓
Transaction costs				✓
# Observations	942	942	827	825
# Clusters	465	292	274	273
# FE dummies	188	166	148	147
Cragg-Donald F-statistic			7.808	5.954

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.
Coefficient estimates from fixed effects models using OLS and IV. Standard errors in parentheses, clustered by seller and buyer, and year (but only by seller and buyer in models (2)–(4) where year fixed effects are replaced by a quadratic time trend). Models (1)–(4) correspond to models (1)–(4) of Table 2, but now with a non-linear model specification. First-stage regressions are omitted for brevity.

that sell the majority of their endowments at least once. After removing these outlier observations, we have $\frac{q_{odt}}{Q_{ot}} = 0.10$ and the corresponding average markdown equals $ 6.97/AF (SD=14.05), which is about 50% larger than the average markdown found for the linear model specification, but still small in percentage terms.

Note that we do not attach much weight to the results from this specification, both because of its sensitivity to removing outliers and also since the functional form of regression equation (20) depends on the specific implementation of a first-order Taylor expansion (see Appendix D for details), which may not be warranted. With these caveats in mind, the results of a model specification with constant price elasticity are largely consistent with those from the linear model specification.

6.4 Transaction timing

So far we have ignored information on the timing of transactions. As a result, in case of multiple transactions per district per year, each district's water use—as captured by variables R^i_{odtk}, $i = o, d$, in (14)—is identical for each of these transactions within one year. This approach is consistent with the assumption of no hedging against the risk of water shortage, such that districts can foresee how much water they are going to sell or buy within a year. In this section, we take the alternative approach and update Q_{dt} and

Q_{ot} after each transaction. This implies that we use counter z to calculate water use (just) after transaction $j = 1, 2, \ldots$ as $Q_{otj} = e_{ot} - \sum_{z=1}^{j} q_{od(j)tz}$ and $Q_{dtj} = e_{dt} + \sum_{z=1}^{j} q_{o(j)dtz}$, where $d(j)$ is the j^{th} counterparty of o and $o(j)$ is the j^{th} counterparty of d. When multiple transactions happen to occur within the same month, we order them by transaction volume such that smaller transactions go first. In an alternative specification, we reverse this order.

Table 6 shows the results. For reference, we include the preferred model from Table 2

Table 6: Estimating WTA and WTP: Dynamic updating

Price (2010$/AF)	(1)	(2)	(3)
	IV preferred	IV dynamic	IV dynamic reversed
Seller water use (1,000 AF)	-0.580^{***}	-0.530^{***}	-0.551^{***}
(coefficient $-\delta$)	(0.208)	(0.186)	(0.195)
Buyer water use (1,000 AF)	-0.309^{**}	-0.306^{**}	-0.305^{**}
(coefficient $-\gamma$)	(0.132)	(0.127)	(0.128)
Seller fixed effects	✓	✓	✓
Buyer fixed effects	✓	✓	✓
Quadratic time trend	✓	✓	✓
# Observations	879	879	879
# Clusters	308	308	308
# FE dummies	164	164	164
Cragg-Donald F-statistic	9.936	11.15	10.58

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.
Coefficient estimates from fixed effects models using IV. Standard errors in parentheses, clustered by seller and buyer. Model (1) corresponds to our preferred model (3) of Table 2. In models (2) and (3), water use is updated dynamically in case of multiple transactions per district per year. In model (2), multiple transactions in one month are ordered from small to large volume, in model (3) this is reversed. First-stage regressions are omitted for brevity.

as model (1). In models (2) and (3), we repeat this model using our dynamically updating measure of water use. The results show that the effect of transaction timing on prices is negligible and that model (1) still offers a conservative estimate of market power.

6.5 Sellers and buyers on one demand curve

So far we have estimated buyers' and sellers' demand curves separately rather than estimating a combined curve. We reject this possibility with multiple arguments. First, we test for equivalence of coefficients using our preferred model (3) of Table 2. Based on a Wald test (F(1,356)=4.92, $p = 0.027$), we reject equality of these coefficients. Second, we use theory and data to argue that selling and buying water districts differ in key characteristics, implying that buying districts cannot be on the same demand curve as selling districts, and hence our approach of modeling two distinct curves is correct.

Table 7 compares selling and buying districts in terms of their main type of water use (urban, agriculture, environment), levels of water entitlements, and water use, as well as whether or not a district trades with more than one counterparty in any given year. Clearly,

Table 7: Key differences between sellers and buyers.

	Sellers		Buyers	
	Mean	SD	Mean	SD
District: urban (share)	0.08	0.28	0.29	0.45
District: agriculture (share)	0.91	0.29	0.47	0.50
District: environment (share)	0.01	0.09	0.24	0.43
Water entitlements (1,000 AF)	193.72	298.63	207.97	570.51
Total water use (1,000 AF)	176.22	280.61	251.65	578.29
More than one counterparty (yes=1)	0.36	0.48	0.55	0.50

selling and buying districts differ in their types of water use. Sellers are more likely to use water for agriculture, while buyers are more likely to use water for urban or environmental uses. The key variable that underlines our argument that sellers and buyers are on different demand curves for water is *Water endowments*. Table 7 shows that buying districts have higher water entitlements than selling districts, and by purchasing water they end up with even higher levels of water use compared with selling districts. If selling and buying districts would have identical demand curves for water, then districts with higher water use would be selling water, rather than buying. In Figure 1 this implies that e_d would be located to the right of the competitive Q_d. This location implies that $WTP_d(Q_d) < WTA_o(Q_o)$, which is inconsistent with the occurrence of observed water transactions. It follows that sellers and buyers cannot be on the same demand curve.

A final difference between selling and buying districts is related to the dummy variable that measures whether a district has *More than one counterparty*. Comparison indicates that buyers have 53% more transactions with multiple counterparties than sellers do. This statistic points to buyer power, with sellers being on the long side of the market.

6.6 Seller power

Our main result is that buyer power is relatively low. Going against previous literature, stakeholder beliefs, and the results of our conjectural variations approach of Section 6.1, we now reverse our model to estimate seller power. This allows us to check if, rather counter-intuitively, a model with seller power would better explain our data than our model with buyer power. We start by adapting (7), as follows:

$$p_{od} = \mathrm{WTA}_o(\mathrm{Q}_o) - q_{od} \cdot \mathrm{WTP}'(\mathrm{Q}_d), \tag{21a}$$

$$p_{od} = \mathrm{WTP}_d(\mathrm{Q}_d). \tag{21b}$$

Taking similar steps as in Section 4, the resulting regression equation becomes:

$$p_{odrtk} = -\delta \tilde{\mathrm{R}}^o_{odtk} - \gamma \tilde{\mathrm{R}}^d_{odtk} + \tau_r \mathrm{C}_{odr} + \phi_o + \psi_d + \beta_t + \varepsilon_{odrtk}, \tag{22}$$

where $\tilde{\mathrm{R}}^o_{odtk}$ and $\tilde{\mathrm{R}}^d_{odtk}$ are modifications of R^o_{odtk}, respectively, R^d_{odtk} that are defined as follows:

$$\tilde{\mathrm{R}}^o_{odtk} = \begin{cases} \mathrm{Q}_{ot} & \text{if } r = 0 \\ 0 & \text{if } r = 1, \end{cases} \quad \text{and} \quad \tilde{\mathrm{R}}^d_{odtk} = \begin{cases} -q_{odtk} & \text{if } r = 0 \\ \mathrm{Q}_{dt} & \text{if } r = 1. \end{cases}$$

Table 8: Estimating WTA and WTP: Seller power

Price (2010$/AF)	(1)	(2)	(3)	(4)
	OLS	OLS	IV	IV
Seller water use (1,000 AF)	−0.0162**	−0.0251**	−0.576***	−0.867**
(coefficient $-\delta$)	(0.00786)	(0.0126)	(0.212)	(0.435)
Buyer water use (1,000 AF)	−0.00764**	−0.0145**	−0.329**	−0.324**
(coefficient $-\gamma$)	(0.00341)	(0.00620)	(0.147)	(0.159)
Seller fixed effects	✓	✓	✓	✓
Buyer fixed effects	✓	✓	✓	✓
Year fixed effects	✓			
Quadratic time trend		✓	✓	✓
Transaction costs				✓
# Observations	1034	1034	879	877
# Clusters	543	337	308	307
# FE dummies	212	190	164	163
Cragg-Donald F-statistic			8.507	8.231

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.
Coefficient estimates from fixed effects models using OLS and IV. Standard errors in parentheses, clustered by seller, buyer, and year (but only by seller and buyer in models (2)–(4) where year fixed effects are replaced by a quadratic time trend). Models (1)–(4) correspond to models (1)–(4) of Table 2, but now with seller power. First-stage regressions are omitted for brevity.

Results of this regression are displayed in Table 8. The resulting coefficients are very similar to those of models (1)–(4) of our main specification with buyer power in Table 2. Importantly, with seller power our measure of market power is now based on the coefficient on *Buyer water use*, i.e. γ rather than δ. Restricting the comparison to our preferred model (3), we find that model (3) of Table 8 does not perform better than model (3) of

Table 2 when comparing either the Cragg-Donald F-statistic or the precision of our coefficient of interest. In case one would still assume seller power, we obtain from model (3) of Table 8, that $\gamma = 0.329$, which is lower than $\delta = 0.580$ from model (3) of Table 2. This difference would imply measures of market power to be lower under seller power than under buyer power.

7 Conclusion

Using a Nash-Cournot model, we derive a closed-form solution for the extent of market power in a water market setting and we construct related measures for market power. Applying our model to surface water transactions in California over the period 1993-2015, we find only limited market power in California's water market, despite the thinness of this market. Our main specification implies that buyer power yields an average markdown of 6% of the transaction price. This result is important in the context of current discussions on Californian water market reform (cf. Maples et al., 2018) which, perhaps, should focus on other distorting factors, most notably transaction costs (Carey et al., 2002; Regnacq et al., 2016; Hagerty, 2019; Leonard et al., 2019).

Our model has three main assets: (1) it features a closed-form solution, (2) it does not rely on conjectural variations, and (3) it is sufficiently flexible that it can be applied to other types of endowment economies, including permit markets. On the downside, our model requires choosing a specific functional form for WTP and WTA that may not be warranted. In addition, while our current application is quite clear in terms of the side of the market where market power resides, this may not be the case in other applications.

One explanation for the limited extent of market power in California is that transaction quantities are, generally, small. These quantities enter our measure of market power linearly such that small quantities imply low markdowns. By the same line of reasoning, high prices also imply low markdowns. This effect was illustrated clearly in Figure 2. Another explanation for the limited extent of market power is that, although California's water market is 'thin' in trades, it is 'thick' in possibilities to trade. Recall from Section 6 that California features an almost complete hydrological network such that nearly any two districts can trade water. The fact that many do not does not imply that such trades are not feasible. Rather, it implies that such districts have high pair-specific transaction costs, which causes a relatively low WTP or a relatively high WTA. The threat of a counterparty switching to a competing district limits the possibility to exercise market power (Funaki et al., 2020). The extent to which such threats affect equilibrium outcomes is an avenue for future research.

References

Alviarez, V. I., M. Fioretti, K. Kikkawa, and M. Morlacco (2023). Two-sided market power in firm-to-firm trade. NBER Working Paper 31253, National Bureau of Economic Research.

Ansink, E. and H. Houba (2012). Market power in water markets. *Journal of Environmental Economics and Management 64*(2), 237–252.

Ashenfelter, O., H. Farber, and M. Ransom (2010). Labor market monopsony. *Journal of Labor Economics 28*(2), 203–210.

Ayres, A., K. Meng, and A. Plantinga (2021). Do environmental markets improve on open access? Evidence from California groundwater rights. *Journal of Political Economy 129*(10), 2817–2860.

Bresnahan, T. (1989). Empirical studies of industries with market power. In R. Schmalensee and R. Willig (Eds.), *Handbook of Industrial Organization, Volume 2*, pp. 1011–1057. North-Holland: Elsevier.

Brewer, J., R. Glennon, A. Ker, and G. Libecap (2008). 2006 Presidential Address: Water markets in the West: Prices, trading, and contractual forms. *Economic Inquiry 46*(2), 91–112.

Browne, O. R. and X. J. Ji (2023). The economic value of clarifying property rights: Evidence from water in Idaho's Snake River Basin. *Journal of Environmental Economics and Management 119*, 102799.

Bruno, E. and K. Jessoe (2021). Using price elasticities of water demand to inform policy. *Annual Review of Resource Economics 13*, 427–441.

Bruno, E. and R. Sexton (2020). The gains from agricultural groundwater trade and the potential for market power: Theory and application. *American Journal of Agricultural Economics 102*(3), 884–910.

California State Water Resources Control Board (1999). A Guide to Water Transfers.

Cameron, A. C. and P. K. Trivedi (2005). *Microeconometrics: Methods and Applications*. Cambridge: Cambridge University Press.

Carey, J., D. Sunding, and D. Zilberman (2002). Transaction costs and trading behavior in an immature water market. *Environment and Development Economics 7*(4), 733–750.

Chakravorty, U., E. Hochman, C. Umetsu, and D. Zilberman (2009). Water allocation under distribution losses: Comparing alternative institutions. *Journal of Economic Dynamics and Control 33*(2), 463–476.

Collard-Wexler, A., G. Gowrisankaran, and R. S. Lee (2019). "Nash-in-Nash" bargaining: A microfoundation for applied work. *Journal of Political Economy 127*(1), 163–195.

Draper, A. (2001). *Implicit stochastic optimization with limited foresight for reservoir systems*. Ph. D. thesis, University of California, Davis.

Easter, K. W., M. Rosegrant, and A. Dinar (1999). Formal and informal markets for water: Institutions, performance, and constraints. *World Bank Research Observer 14*(1), 99–116.

Funaki, Y., H. Houba, and E. Motchenkova (2020). Market power in bilateral oligopoly markets with non-expandable infrastructures. *International Journal of Game Theory 49*(2), 525–546.

Grafton, R. Q., G. Libecap, S. McGlennon, C. Landry, and B. O'Brien (2011). An integrated assessment of water markets: A cross-country comparison. *Review of Environmental Economics and Policy 5*(2), 219–239.

Hagerty, N. (2019). Liquid constrained in California: Estimating the potential gains from water markets. Working Paper, Department of Agricultural and Resource Economics, UC Berkeley.

Hagerty, N. (2021). Adaptation to surface water scarcity in irrigated agriculture. Working Paper, Montana State University.

Hansen, K., J. Kaplan, and S. Kroll (2014). Valuing options in water markets: A laboratory investigation. *Environmental and Resource Economics 57*(1), 59–80.

Hirschman, A. (1945). *National Power and the Structure of Foreign Trade*. Berkeley: University of California Press.

Holland, S. (2006). Privatization of water-resource development. *Environmental and Resource Economics 34*(2), 291–315.

Isaaks, R. and B. Colby (2020). Empirical application of Rubinstein bargaining model in Western U.S. water transactions. *Water Economics and Policy 6*(2), 1950010.

Jacoby, H., R. Murgai, and S. Rehman (2004). Monopoly power and distribution in fragmented markets: The case of groundwater. *Review of Economic Studies 71*(3), 783–808.

Jenkins, M., J. Lund, R. Howitt, A. Draper, S. Msangi, S. Tanaka, R., and G. Marques (2004). Optimization of California's water supply system: Results and insights. *Journal of Water Resources Planning and Management 130*(4), 271–280.

Leonard, B., C. Costello, and G. Libecap (2019). Expanding water markets in the Western United States: Barriers and lessons from other natural resource markets. *Review of Environmental Economics and Policy 13*(1), 43–61.

Lindh, T. (1992). The inconsistency of consistent conjectures: Coming back to Cournot. *Journal of Economic Behavior and Organization 18*(1), 69–90.

Maples, S., E. Bruno, A. Kraus-Polk, S. Roberts, and L. Foster (2018). Leveraging hydrologic accounting and water markets for improved water management: The case for a central clearinghouse. *Water 10*(12), 1720.

McCann, L., B. Colby, K. W. Easter, A. Kasterine, and K. Kuperan (2005). Transaction cost measurement for evaluating environmental policies. *Ecological Economics 52*(4), 527–542.

Rafey, W. (2023). Droughts, deluges, and (river) diversions: Valuing market-based water reallocation. *American Economic Review 113*(2), 430–471.

Regnacq, C., A. Dinar, and E. Hanak (2016). The gravity of water: Water trade frictions in California. *American Journal of Agricultural Economics 98*(5), 1273–1294.

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

Scheer, J. B. (2016). Participation, price, and competition in the Sacramento Valley water market. MSc thesis, University of California, Davis.

Stene, E. A. (1995). Central Valley Project Overview. Bureau of Reclamation.

Stock, J. and M. Yogo (2005). Testing for weak instruments in linear IV regression. In *Identification and Inference for Econometric Models*. New York: Cambridge University Press.

Tomkins, C. and T. Weber (2010). Option contracting in the California water market. *Journal of Regulatory Economics 37*(2), 107–141.

Vaux Jr., H. and R. Howitt (1984). Managing water scarcity: An evaluation of interregional transfers. *Water Resources Research 20*(7), 785–792.

Wheeler, S. (2022). Debunking Murray-Darling basin water trade myths. *Australian Journal of Agricultural and Resource Economics 66*(4), 797–821.

Appendices

A Thick markets under homogeneity

In this appendix, we illustrate the Nash-Cournot equilibrium concept for a tractable and symmetric version of our main model specification assuming fully connected water markets, homogeneity on each side of the market, and buyer market power. We derive general formulas for equilibrium quantities, prices, etc. We conclude with an example to illustrate that there are unrealistically many equilibrium trades and that homogeneity on each side of the market is causing this. In Appendix B and C, we discuss why heterogeneity in the main specification of our model leads to thin markets as observed in the Californian data set.

Formally, we assume arbitrary numbers of buyers (N_d) and sellers (N_o). Homogeneity on each side of the market imposes $\alpha_{ot} = \alpha$ for all o and $a_{dt} = a$ for all d. We maintain the necessary and sufficient condition (10) for positive trades, which can now be written as $a - \gamma e_d > \alpha - \delta e_o$. For each individual buyer, we can write the maximand of equation (4), i.e. the buyer's profit function, for all origins with which this buyer trades as

$$\begin{aligned}\pi_d &= f_d(\mathrm{Q}_d) - \sum_{o=1}^{\mathrm{N}_o} q_{od} \cdot \left[\mathrm{WTA}(\mathrm{Q}_o)\right] \\ &= \left(e_d + \sum_{o=1}^{\mathrm{N}_o} q_{od}\right)\left(a - \tfrac{1}{2}\gamma\left(e_d + \sum_{o=1}^{\mathrm{N}_o} q_{od}\right)\right) - \sum_{o=1}^{\mathrm{N}_o} q_{od} \cdot \left[\alpha - \tfrac{1}{2}\delta\left(e_o - \sum_{d=1}^{\mathrm{N}_d} q_{od}\right)\right].\end{aligned} \tag{A.1}$$

Applying (5), we take the derivative of the buyer's profit function (A.1) with respect to q_{od} and, by symmetry, we simplify the resulting condition by writing $q_{od} = q$:

$$a - \gamma(e_d + \mathrm{N}_o q) - \alpha + \delta(e_o - \mathrm{N}_d q) - \delta q = 0. \tag{A.2}$$

This condition implies $a - \gamma e_d - \alpha + \delta e_o = [\mathrm{N}_o \gamma + (\mathrm{N}_d + 1)\delta]q > 0$. Thus, the equilibrium quantity from seller to buyer q_{od} equals

$$q^* = q^*_{od} = \frac{a - \gamma e_d - \alpha + \delta e_o}{\mathrm{N}_o \gamma + (\mathrm{N}_d + 1)\delta}, \tag{A.3}$$

which is positive for all o and d by (10). We obtain a thick market in which every seller trades with every buyer.

In case the number of available buyers N_d and/or the number of available sellers N_o

increases, then each buyer would buy less water from each individual seller. The quantity in equilibrium can be expressed differently by substituting $\mathrm{S} = \alpha - \delta e_o$ and $\mathrm{B} = a - \gamma e_d$. Thus, $q^* = (\mathrm{B} - \mathrm{S})/(\mathrm{N}_o\gamma + (\mathrm{N}_d + 1)\delta)$. The numerator of this expression consists of the marginal benefits of water use at the initial endowments. If the buyers' marginal benefit B increases, trade will increase. In contrast, if the sellers' marginal benefit S increases, trade will decrease. The effects on trade of parameters a, α, γ, δ and initial endowments e_d and e_o follow immediately through their effects on either B or S. For example, an increase in the initial endowment e_d of individual buyers implies that individual buyers buy less. Similarly, an increase of the initial endowment e_o of individual sellers implies that individual sellers sell more.

We use equilibrium quantities as derived in (A.3) to derive the sellers' and buyers' equilibrium (marginal) benefits as well as prices. Using (2), we have that $\mathrm{WTP}(\mathrm{Q}_d) = a - \gamma \mathrm{Q}_d$ and $\mathrm{WTA}(\mathrm{Q}_o) = \alpha - \delta \mathrm{Q}_o$. By symmetry, we can therefore write the marginal benefit for, respectively, each buyer and each seller in equilibrium:

$$\mathrm{WTP}(\mathrm{Q}_d^*) = a - \gamma\left(e_d + \mathrm{N}_o q_{od}^*\right) = \frac{(\mathrm{N}_d + 1)\delta\mathrm{B} + \mathrm{N}_o\gamma\mathrm{S}}{\mathrm{N}_o\gamma + (\mathrm{N}_d + 1)\delta}, \tag{A.4a}$$

$$\mathrm{WTA}(\mathrm{Q}_o^*) = \alpha - \delta(e_o - \mathrm{N}_d q_{od}^*) = \frac{(\mathrm{N}_o\gamma + \delta)\mathrm{S} + \mathrm{N}_d\delta\mathrm{B}}{\mathrm{N}_o\gamma + (\mathrm{N}_d + 1)\delta}. \tag{A.4b}$$

From the WTP function we directly obtain $\mathrm{Q}_d = (a - \mathrm{WTP}(\mathrm{Q}_d))/\gamma$. The other component of benefit function f_d is $(a - \frac{1}{2}\gamma\mathrm{Q}_d)$ and it can also be expressed in terms of this WTP: $\left(a - \frac{1}{2}\gamma\mathrm{Q}_d\right) = \frac{1}{2}(a + a - \gamma\mathrm{Q}_d) = \frac{1}{2}(a + \mathrm{WTP}(\mathrm{Q}_d))$. Combining these expressions yields the buyers' benefit function:

$$\begin{aligned} f_d(\mathrm{Q}_d^*) &= \mathrm{Q}_d^*(a - \frac{1}{2}\gamma\mathrm{Q}_d^*) = \frac{1}{2\gamma}\left[a^2 - (\mathrm{WTP}(\mathrm{Q}_d^*))^2\right] \\ &= \frac{1}{2\gamma}\left[a^2 - \left(\frac{(\mathrm{N}_d + 1)\delta\mathrm{B} + \mathrm{N}_o\gamma\mathrm{S}}{\mathrm{N}_o\gamma + (\mathrm{N}_d + 1)\delta}\right)^2\right]. \end{aligned} \tag{A.5a}$$

Similar steps are applied to obtain the sellers' benefit function:

$$\begin{aligned} f_o(\mathrm{Q}_o^*) &= \mathrm{Q}_o^*\left(\alpha - \frac{1}{2}\delta\mathrm{Q}_o^*\right) = \frac{1}{2\delta}\left[\alpha^2 - \left(\mathrm{WTA}(\mathrm{Q}_o^*)\right)^2\right] \\ &= \frac{1}{2\delta}\left[\alpha^2 - \left(\frac{(\mathrm{N}_o\gamma + \delta)\mathrm{S} + \mathrm{N}_d\delta\mathrm{B}}{\mathrm{N}_o\gamma + (\mathrm{N}_d + 1)\delta}\right)^2\right]. \end{aligned} \tag{A.5b}$$

Given buyer power, equilibrium price equals the willingness to accept. Using (A.4b), we have

$$p^* = p_{od}^* = \mathrm{WTA}(\mathrm{Q}_o^*) = \frac{(\mathrm{N}_o\gamma + \delta)\mathrm{S} + \mathrm{N}_d\delta\mathrm{B}}{\mathrm{N}_o\gamma + (\mathrm{N}_d + 1)\delta}, \tag{A.6}$$

which is positive for all o and d. The importance of this last result is that the Law of Uniform Price holds in a fully connected water market with homogeneity on each side and buyer market power. It is a consequence of every seller trading with every buyer and every buyer trading with every seller. In doing so, each seller equates its willingness to accept across all the links it trades (which are all links to this agent's buyers) and each buyer accomplishes the same markdown across all the links it trades (which are all sellers in the market).

For completeness, we also verify that equation (6) holds. This equation states that the difference between WTP and WTA equals $-q_{od}\mathrm{WTA}_o'(\mathrm{Q}_o) = \delta q_{od}^* > 0$. Substitution of our equilibrium expressions (A.4b) and (A.4a) gives

$$\begin{aligned}\mathrm{WTP}(\mathrm{Q}_d^*) - \mathrm{WTA}(\mathrm{Q}_o^*) &= \frac{(\mathrm{N}_d + 1)\delta\mathrm{B} + \mathrm{N}_o\gamma\mathrm{S}}{\mathrm{N}_o\gamma + (\mathrm{N}_d + 1)\delta} - \frac{(\mathrm{N}_o\gamma + \delta)\mathrm{S} + \mathrm{N}_d\delta\mathrm{B}}{\mathrm{N}_o\gamma + (\mathrm{N}_d + 1)\delta} \\ &= \delta\frac{\mathrm{B} - \mathrm{S}}{\mathrm{N}_o\gamma + (\mathrm{N}_d + 1)\delta} \\ &= \delta q_{od}^*. \end{aligned} \tag{A.7}$$

Therefore, equation (6) holds and the difference is positive, as it should. This completes the derivation of the symmetric version of the main specification of our model.

Figure 3 illustrates the case of a thick market with nine equilibrium trades among three homogeneous sellers and three homogeneous buyers. Each arrow represents a trade. In each trade a quantity $q_{od}^* = \frac{\mathrm{B}-\mathrm{S}}{3\gamma+4\delta}$ of water is exchanged at the uniform price $p^* = p_{od}^* = \frac{(3\gamma+\delta)\mathrm{S}+3\delta\mathrm{B}}{3\gamma+4\delta}$. This thick market result is due to homogeneity on each side of the market. A thick market would also arise for relatively low homogeneous transaction costs per unit of water. These can be accommodated by subtracting such costs from the coefficients a and α. As long as $\mathrm{B} - \mathrm{S} > 0$, the quantities of (A.3) remain positive and imply that a thick market arises in equilibrium. In order to explain thin markets, we need heterogeneity on each side of the market. Such heterogeneity might be due to asymmetries in the productivity of water (see Appendix B) or to substantial and heterogeneous transaction costs (see Appendix C).

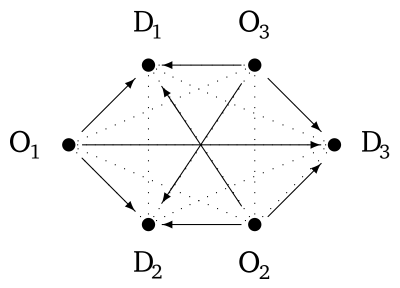

Figure 3: A fully connected market consisting of six agents who may buy or sell water from each other. The (marginal) productivity of the three identical origins O_1, O_2, O_3 is lower than that of the three identical destinations D_1, D_2, D_3. The vectors illustrate the direction of Nash-Cournot equilibrium trades. It is a thick market in which every agent trades with every agent on the other side and nine of the available fifteen links are utilized.

B Thin markets under heterogeneity: productivity

California's water market is thin and the insights of Appendix A point to heterogeneity on each side of the market as its main cause. Heterogeneity arises with respect to substantial differences in productivity across districts and substantial pair-specific transaction costs as identified in Table 1. In this appendix, we generalize the linear model specification introduced in Section 3.2 to allow for heterogeneity in productivity and we assess transaction costs in Appendix C.

The extended analysis of convex nonlinear program (4) that takes zero-trade boundary solutions into account would incorporate non-negativity conditions $-q_{od} \leq 0$ (in standard form) with shadow price $\lambda_{od} \geq 0$ and employ the Karush-Kuhn-Tucker conditions given by[15]

$$f_d'(Q_d) - \mathrm{WTA}_o(Q_o) + q_{od} \cdot \mathrm{WTA}_o'(Q_o) + \lambda_{od} = 0, \tag{B.1a}$$

$$\lambda_{od} \cdot q_{od} = 0. \tag{B.1b}$$

As before, a positive trade $q_{od} > 0$ implies a shadow price $\lambda_{od} = 0$ and we obtain first-order conditions (5) as stated in the main text. Furthermore, zero trade implies $q_{od} = 0$ and λ_{od} solves (B.1a). Because the shadow price is non-negative and $q_{od} = 0$, we obtain the weak inequality $\leq$ in (5) and its simplification $f_d'(Q_d) \leq \mathrm{WTA}_o(Q_o)$, i.e. no marginal benefits of

[15]The full analysis would also incorporate feasibility at each origin o given by $\sum_{d'} q_{od'} \leq e_o$ with shadow price $\mu_{od} \geq 0$. The modification of the Karush-Kuhn-Tucker conditions includes an extra $-\mu_{od}$ term in (B.1a) as well as the extra condition $\mu_{od} \cdot \left(\sum_{d'} q_{od'} - e_o\right) = 0$. In our database, no seller sells his entire endowment, implying a strict inequality in the feasibility constraint and $\mu = 0$.

trade at total consumption levels Q_d and Q_0.

Consider a setting where all buyers and all sellers are heterogeneous and no transaction costs. We update our benefit functions to allow for heterogeneity in terms of benefit parameters a, α, γ, and δ, while suppressing time subscripts to keep notation simple. For each destination we now have $f_d(Q_d) = Q_d(a_d - \frac{1}{2}\gamma_d Q_d)$ and for each origin we now have $f_o(Q_o) = Q_o(\alpha_o - \frac{1}{2}\delta_o Q_o)$. Therefore $f_d'(Q_d) = a_d - \gamma_d Q_d$, which is the WTP in (1), while $f_o'(Q_o) = \alpha_o - \delta_o Q_o$, which is the WTA in (2). We number sellers as $o = 1, 2, 3, \ldots$ and buyers as $d = -1, -2, -3\ldots$. Subscript $od = 2-1$ implies that seller 2 delivers to buyer 1.

Almost all contracts in our database consist of trades between a single seller and a single buyer who only trade with each other and non-traders in the background as potential alternative trading partners. Without loss of generality, consider one such non-trader on each side of the market. If we number seller 1 and buyer −1 as the trading parties with $q_{1-1} > 0$, then seller 2 and buyer −2 do not trade, i.e. $q_{1-2} = q_{2-1} = q_{2-2} = 0$. The equilibrium conditions are derived from buyer −1 who maximizes over quantities q_{1-1} and q_{2-1} and from buyer −2 who maximizes over quantities q_{1-2} and q_{2-2}. After adding subscripts d and o, we take the derivative of the buyer's profit function (A.1) with respect to q_{od} and obtain:

$$a_d - \gamma_d(e_d + q_{1d} + q_{2d}) - \alpha_o + \delta_o(e_o - q_{o-1} - q_{o-2}) + \delta_o q_{od} \leq 0. \tag{B.2}$$

For $o = 1, 2$ and $d = -1, -2$, we have $q_{1-1} > 0$ and $q_{1-2} = q_{2-1} = q_{2-2} = 0$, so we obtain four equilibrium conditions:

$$a_{-1} - \gamma_{-1}(e_{-1} + q_{1-1}) - \alpha_1 + \delta_1(e_1 - q_{1-1}) + \delta_1 \cdot q_{1-1} = 0, \tag{B.3a}$$
$$a_{-1} - \gamma_{-1}(e_{-1} + q_{1-1}) - \alpha_2 + \delta_2 e_2 \leq 0, \tag{B.3b}$$
$$a_{-2} - \gamma_{-2}e_{-2} - \alpha_1 + \delta_1(e_1 - q_{1-1}) \leq 0, \tag{B.3c}$$
$$a_{-2} - \gamma_{-2}e_{-2} - \alpha_2 + \delta_2 e_2 \leq 0. \tag{B.3d}$$

Before solving, we combine and rewrite these four equilibrium conditions in terms of equilibrium WTP or WTA. In doing so, note that because $q_{1-1} \geq 0$, we can rewrite the first condition as a weak inequality: $a_{-1} - \gamma_{-1}(e_{-1} + q_{1-1}) - \alpha_1 + \delta_1(e_1 - q_{1-1}) = -\delta_1 q_{1-1} \leq 0$.

We obtain

$$
\begin{aligned}
\alpha_1 - \delta_1(e_1 - q_{1-1}) &\geq \max\{a_{-1} - \gamma_{-1}(e_{-1} + q_{1-1}), a_{-2} - \gamma_{-2}e_{-2}\},\\
\alpha_2 - \delta_2 e_2 &\geq \max\{a_{-1} - \gamma_{-1}(e_{-1} + q_{1-1}), a_{-2} - \gamma_{-2}e_{-2}\},\\
a_{-1} - \gamma_{-1}(e_{-1} + q_{1-1}) &\leq \min\{\alpha_1 - \delta_1(e_1 - q_{1-1}), \alpha_2 - \delta_2 e_2\},\\
a_{-2} - \gamma_{-2}e_{-2} &\leq \min\{\alpha_1 - \delta_1(e_1 - q_{1-1}), \alpha_2 - \delta_2 e_2\}.
\end{aligned}
$$

The first two conditions indicate that, *in equilibrium*, the seller's WTA must be equal to or larger than the highest WTP for all buyers in the market, independent whether these sellers trade or not. The last two lines indicate that, *in equilibrium*, the buyers' WTP must be equal to or lower than the highest WTA from sellers in the market, independent whether these buyers trade or not. These insights generalize to any market with N_o sellers and N_d buyers independent whether these trade or not. These conditions imply that none of the pairs of trading districts has incentives to expand equilibrium trade.

We now check each of the equilibrium conditions in (B.3). Solving condition (B.3a) gives equilibrium trade between seller $o = 1$ and buyer $d = -1$. It is a special case of (A.3) and we obtain

$$
q^*_{1-1} = \frac{a_{-1} - \gamma_{-1}e_{-1} - \alpha_1 + \delta_1 e_1}{\gamma_{-1}}. \tag{B.4}
$$

Under $a_{-1} - \gamma_{-1}e_{-1} > \alpha_1 - \delta_1 e_1$, which is a straightforward modification of (10), this quantity is positive. Substitution of q^*_{1-1} into condition (B.3b) yields $\alpha_1 - \delta_1 e_1 \leq \alpha_2 - \delta_2 e_2$. Evaluated at the initial endowments, seller 1's WTA is lower than than that of seller 2, making seller 1 more efficient in supplying water. Rewriting after substitution of q^*_{1-1} into condition (B.3c) yields

$$
a_{-2} - \gamma_{-2}e_{-2} \leq \left(1 - \frac{\delta_1}{\gamma_{-1}}\right)(\alpha_1 - \delta_1 e_1) + \frac{\delta_1}{\gamma_{-1}}(a_{-1} - \gamma_{-1}e_{-1}). \tag{B.5}
$$

For $\frac{\delta_1}{\gamma_{-1}} \in [0, 1]$, the right-hand side is the convex combination of seller 1's WTA and buyer 1's WTP, both evaluated at the initial endowments. For the boundary case $\delta_1 = \gamma_{-1}$, the right-hand side simplifies to $a_{-1} - \gamma_{-1}e_{-1}$. Evaluated at the initial endowments, buyer −1's WTP is higher than that of buyer −2, making buyer −1 more efficient in purchasing water. If the gap in WTP between the two buyers is positive, then condition (B.5) also holds for δ_1 almost equal to γ_{-1}. Finally, condition (B.3d) specifies the condition that non-trading seller $o = 2$ and non-trading buyer $d = -2$ do not want to trade with each other. If rewritten as $a_{-2} - \gamma_{-2}e_{-2} \leq \alpha_2 - \delta_2 e_2$, it is the complement of modified condition (10).

To summarize, the configuration in which seller 1 exclusively trades with buyer −1 arises naturally in case seller 1 has a substantially lower WTA than competing seller 2, while buyer −1 has a substantially higher WTP than competing buyer −2. Also, the WTA of seller 2 is larger than the WTP of buyer −1. By the preceding discussion of the equilibrium conditions in (B.3), a sufficient condition for water trade between seller 1 and buyer −1 is the following:

$$\alpha_2 - \delta_2 e_2 > a_{-1} - \gamma_{-1} e_{-1} > \alpha_1 - \delta_1 e_1 > a_{-2} - \gamma_{-2} e_{-2}. \tag{B.6}$$

Equilibrium (marginal) benefits and prices for the heterogeneous case can be determined similarly as was done in Appendix–A.

So far, we have covered trades between a single seller and a single buyer who only trade with each other and non-traders in the background as potential alternative trading partners. Two other types of transactions occur frequently in the data, as introduced in Section 2. These are transactions with either one seller with multiple buyers or one buyer with multiple sellers. Such transactions can be analyzed in a similar way. First, expressions for the traded quantities need to be derived from all equations for which inequalities hold, which would be special cases of (A.3) when trading buyers are homogeneous and trading sellers are homogeneous. Second, the expressions for the quantities need to be substituted in all equations for which weak inequalities hold, which would yield a system of inequalities similar to (B.3a)-(B.3d), but with more inequalities. Deriving this system would be cumbersome and would involve a lot of repetition of the case with a single seller and single buyer without generating more insight.

The main insight is intuitive. The lowest WTA of non-trading sellers is higher than the highest WTP of trading buyers. Similarly, the highest WTP of non-trading buyers is lower than the lowest WTA of trading sellers. Combined, the most productive buyers trade with the least productive sellers, which illustrates the role of heterogeneous productivity in water markets.

C Thin markets under heterogeneity: transaction costs

After assessing heterogeneity in productivity across districts in Appendix B, we now consider heterogeneity due to substantial pair-specific transaction costs. Our starting point is that marginal transaction costs are constant (Hagerty, 2019). Although conceptually similar to heterogeneous productivity, the analysis requires some additional notation and subtle modifications of willingness to accept and willingness to pay. The arguments in this appendix justify our empirical strategy as explained in Section 4.

Formally, T_{odr} denotes agent-dependent, pair-specific, and constant marginal transaction costs for seller o and buyer d. Index $r = o, d$ indicates to whom these costs apply to, either the seller or the buyer. So, T_{odo} are seller o's marginal transaction costs of trading with buyer d while T_{odd} are d's marginal transaction costs of trading with o. Seller o's aggregate transaction costs over all possible trades are denoted $\sum_{\hat{d}=1}^{N_d} T_{o\hat{d}o} \cdot q_{o\hat{d}}$ and those of buyer d are denoted $\sum_{\hat{o}=1}^{N_o} T_{\hat{o}dd} \cdot q_{\hat{o}d}$.

The Nash-Cournot equilibrium with buyer market power assumes that sellers are price takers. Seller o's net benefits of water use become $f_o(Q_o)$ plus revenues from selling water minus transaction costs. Due to the negative relation between Q_o and q_{od}, we obtain that the partial derivative with respect to q_{od} is given by $-f_o'(Q_o) + p_{od} - T_{odo}$. Seller o's first-order condition for trading with buyer d as a price taker can be rewritten as

$$p_{od} = f_o'(Q_o) + T_{odo}. \tag{C.1}$$

The market-clearing price equals the opportunity costs of further marginal changes in selling water plus full compensation for seller o's transaction costs involved. Since pair-specific transaction costs are heterogeneous, the Law of Uniform Price ceases to hold. In a Nash-Cournot equilibrium, buyer d's expenditure on buying water from seller o is then given by $q_{od} \cdot \left[\mathrm{WTA}_o(Q_o) + T_{odo}\right]$. Buyer d's optimization problem is given by

$$\max_{q_{1d},\ldots,q_{N_o d}} f_d(Q_d) - \sum_{o=1}^{N_o} q_{od} \cdot \left[f_o'(Q_o) + T_{odo}\right] - \sum_{o=1}^{N_o} T_{odd} \cdot q_{od}. \tag{C.2}$$

Proceeding as in the main text, buyer d's first-order condition with respect to q_{od} for an interior solution is given by

$$f_d'(Q_d) - T_{odd} - f_o'(Q_o) - T_{odo} + q_{od} \cdot f_o''(Q_o) = 0. \tag{C.3}$$

This last equation can be rewritten in terms of a pair-specific willingness to accept and a pair-specific willingness to pay by modifying definitions (1) and (2) to include transaction costs. To achieve this, define $\mathrm{WTA}_{od}(Q_o) = f_o'(Q_o) + T_{odo}$ and $\mathrm{WTP}_{od}(Q_d) = f_d'(Q_d) - T_{odd}$.[16] That is, transaction costs enter sellers' WTA positively, because these have to be compensated for sellers on top of the sellers' net benefits, while transaction costs enter the buyers' WTP negatively, because they decrease sellers' net benefits. These modifications justify our

[16] Note that $\mathrm{WTA}_{od}'(Q_o) = \frac{d}{dQ_o}\left(f_o'(Q_o) + T_{odo}\right) = f_o''(Q_o)$.

empirical strategy in Section 4. With these modified definitions, after rewriting we obtain

$$\mathrm{WTP}_{od}\left(\mathrm{Q}_{d}\right)=\mathrm{WTA}_{od}\left(\mathrm{Q}_{o}\right)-q_{od}\cdot\mathrm{WTA}_{od}'\left(\mathrm{Q}_{o}\right). \tag{C.4}$$

Our modified pair-specific measure of heterogeneous market power is now equal to

$$\frac{\mathrm{WTP}_{od}\left(\mathrm{Q}_{d}\right)-\mathrm{WTA}_{od}\left(\mathrm{Q}_{o}\right)}{\mathrm{WTA}_{od}\left(\mathrm{Q}_{o}\right)}=-\frac{q_{od}}{\mathrm{Q}_{o}}\cdot\frac{\mathrm{Q}_{o}\mathrm{WTA}_{od}'\left(\mathrm{Q}_{o}\right)}{\mathrm{WTA}_{od}\left(\mathrm{Q}_{o}\right)}. \tag{C.5}$$

The necessary condition for trade in our main specification with heterogeneous productivity is stated in 10. It can be easily modified for transaction costs by applying the main insight of this appendix: including these costs in the willingness to accept and willingness to pay. We obtain $\left(a_{d}-\mathrm{T}_{odd}\right)-\gamma_{d}e_{d}>\left(\alpha_{o}+\mathrm{T}_{odo}\right)-\delta_{o}e_{o}$ *as if* the model has modified coefficients $a_{odd}=a_{d}-\mathrm{T}_{odd}$ and $\alpha_{odo}=\alpha_{o}+\mathrm{T}_{odo}$, which is how to read the system of equations (13a) and (13b). Positive trade requires sufficiently small pair-specific transaction costs depending on differences in productivity between the seller and buyer.

We conclude this appendix with two illustrations of thin markets by adjusting Figure 3 such that transaction costs of some pairs are sufficiently high as to prevent trade. To keep both analyses tractable, we ignore heterogeneity as discussed in Appendix B and we simply impose condition (10). As a result, we have $\alpha_{odo}=\alpha$ and $a_{odd}=a$ for all o and d; and $\mathrm{S}=a-\gamma e_{d}>\alpha-\delta e_{o}=\mathrm{B}$. Recall that S can be seen as willingness to accept at e_{o} and B as willingness to pay at e_{d}, both in the absence of transaction costs.

The left panel of Figure 4 illustrates a thin market with three identical sellers, denoted $\mathrm{O}_1,\mathrm{O}_2,\mathrm{O}_3$, and three identical buyers, each denoted $\mathrm{D}_1,\mathrm{D}_2,\mathrm{D}_3$. The pairs $(\mathrm{O}_1,\mathrm{D}_1)$, $(\mathrm{O}_2,\mathrm{D}_2)$ and $(\mathrm{O}_3,\mathrm{D}_3)$ are involved in 1 seller-1 buyer trades. This equilibrium may arise if transaction costs within each trading pair are zero and elsewhere these are larger than $\mathrm{B}-\mathrm{S}$. So, the results of Appendix A apply for $\mathrm{N}_o=\mathrm{N}_d=1$ in each trade and we obtain as the expressions for the equilibrium quantity and price within each pair:

$$q_{od}^{*}=\frac{\mathrm{B}-\mathrm{S}}{\gamma+2\delta}>0\quad\text{and}\quad p^{*}=\frac{(\gamma+\delta)\mathrm{S}+\delta\mathrm{B}}{\gamma+2\delta}\in[\mathrm{S},\mathrm{B}]. \tag{C.6}$$

The right panel of Figure 4 illustrates a thin market with four identical sellers, denoted $\mathrm{O}_1,\mathrm{O}_2,\mathrm{O}_3,\mathrm{O}_4$, and two identical buyers, denoted $\mathrm{D}_1,\mathrm{D}_2$. The trios $(\mathrm{O}_1,\mathrm{O}_3,\mathrm{D}_1)$ and $(\mathrm{O}_2,\mathrm{O}_4,\mathrm{D}_2)$ are involved in 2 seller-1 buyer trades. This equilibrium may arise if transaction costs within each trading trio are zero and elsewhere these are larger than $\mathrm{B}-\mathrm{S}$. So, the results of Appendix A apply for $\mathrm{N}_o=2$ and $\mathrm{N}_d=1$ in each trade and we obtain as the

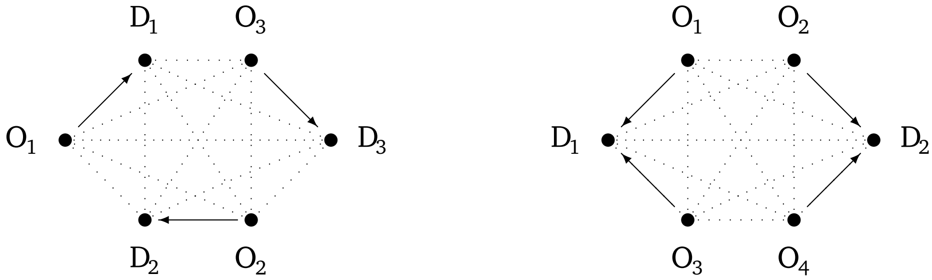

Figure 4: Fully connected water markets similar to Figure 3. Left panel: three identical origins and three identical destinations. Transaction costs lead a thin market in which only three of the available fifteen links are utilized. Right panel: four identical origins and two identical destinations. Transaction costs lead a thin market in which only four of the available fifteen links are utilized.

expressions for the equilibrium quantity and price within each pair:

$$q^*_{od} = \frac{B - S}{2(\gamma + \delta)} > 0 \quad \text{and} \quad p^* = \frac{(2\gamma + \delta)S + \delta B}{2(\gamma + \delta)} \in [S, B]. \tag{C.7}$$

D Non-linear demand

In this appendix, we present our empirical strategy for the model of Section 6.3 featuring a non-linear WTA function that is homogenous. Our aim is to estimate κ_o so that we can measure market power for this model specification.

The strategy is largely similar to that of Section 4 for the linear model specification. We start with the following system of regression equations, based on (7), and substitute (18) to obtain

$$p_{od} = \text{WTA}(Q_o), \tag{D.1a}$$
$$p_{od} = \text{WTP}(Q_d) - \kappa_o \cdot \frac{q_{od}}{Q_o} \cdot \text{WTA}(Q_o). \tag{D.1b}$$

Substituting p_{od} for $\text{WTA}_o(Q_o)$, we solve the last equation for p_{od}, which yields the non-linear system

$$p_{od} = \text{WTA}(Q_o), \tag{D.2a}$$
$$p_{od} = \left(1 + \frac{q_{od}}{Q_o} \cdot \kappa_o\right)^{-1} \text{WTP}_d(Q_d). \tag{D.2b}$$

This system can be written in logarithmic form as

$$\ln p_{od} = \ln \mathrm{WTA}(\mathrm{Q}_o), \tag{D.3a}$$
$$\ln p_{od} = \ln \mathrm{WTP}(\mathrm{Q}_d) - \ln\left(1 + \frac{q_{od}}{\mathrm{Q}_o} \cdot \kappa_o\right). \tag{D.3b}$$

To extract parameter κ out of the last term, we approximate it by the first-order Taylor expansion of the logarithmic function around 1.[17] This yields the following non-linear system:

$$\ln p_{od} = \ln \mathrm{WTA}(\mathrm{Q}_o), \tag{D.4a}$$
$$\ln p_{od} = \ln \mathrm{WTP}(\mathrm{Q}_d) - \kappa_o \cdot \frac{q_{od}}{\mathrm{Q}_o}. \tag{D.4b}$$

We proceed to estimate (D.4) for the specification $\mathrm{A}_i(\mathrm{Q}_i)^{-\kappa_i}$, $i = o, d$ and $\kappa_i > 0$, that features constant price elasticity equal to $-1/\kappa_i$. Substitution, rewriting and including multiplicative transaction costs in the factor A_i, $i = o, d$, as well as seller-, buyer-, and year fixed effects, yields

$$\ln p_{odrtk} = -\kappa_o \ln \mathrm{Q}_{ot} + \ln \tau_r \mathrm{C}_{odr} + \phi_o + \psi_d + \beta_t + \varepsilon_{odrtk}, \tag{D.5a}$$
$$\ln p_{odrtk} = -\kappa_o \frac{q_{odtk}}{\mathrm{Q}_{ot}} - \kappa_d \ln \mathrm{Q}_{dt} + \ln \tau_r \mathrm{C}_{odr} + \phi_o + \psi_d + \beta_t + \varepsilon_{odrtk}. \tag{D.5b}$$

Similar to the procedure used in deriving the regression equation for our linear model specification, we combine both equations. This combination requires the construction of two new variables that are defined by

$$\overline{\mathrm{R}}^o_{odtk} = \begin{cases} \ln \mathrm{Q}_{ot} & \text{if } r = 0 \\ q_{odtk}/\mathrm{Q}_{ot} & \text{if } r = 1, \end{cases} \quad \text{and} \quad \overline{\mathrm{R}}^d_{odtk} = \begin{cases} 0 & \text{if } r = 0 \\ \ln \mathrm{Q}_{dt} & \text{if } r = 1. \end{cases}$$

The combined regression equation is:

$$\ln p_{odrtk} = -\kappa_o \overline{\mathrm{R}}^o_{odtk} - \kappa_d \overline{\mathrm{R}}^d_{odrtk} + \ln \tau_r \mathrm{C}_{odr} + \phi_o + \psi_d + \beta_t + \varepsilon_{odtk}. \tag{D.6}$$

Results of the estimation of this regression equation are presented in Table 5 and discussed in Section 6.3.

[17] The first-order Taylor expansion of $\ln(1+x)$ around $x_0 = 0$ is given by $\ln(1+x_0) + \frac{1}{1+x_0}(x - x_0) = x$. In our case $x = \frac{q_{od}}{\mathrm{Q}_o} \cdot \kappa$.