A Comparison of Actual and Predicted Groundwater Trading Behavior: the role of transaction costs and trading restrictions in explaining deviations

Elizabeth M. Juchems and Karina Schoengold*

Abstract

Previous research on water trading has focused on surface water trading and theoretical approaches to analyzing groundwater trading. Empirical analysis of groundwater trading is a new area of research due to limited data on recorded usage, infrequent trades, and the lack of binding constraints on groundwater use by landowners. Groundwater trading can help move water from low-value to high-value areas of use for the benefit of the participants and the public. The paper uses data on groundwater trades to estimate the factors that affect utilization of groundwater trading. Specifically, the paper considers both formal and informal trading of groundwater used for crop irrigation and attempts to identify those characteristics that predict the probability of trade participation and whether an individual is a buyer or seller of groundwater rights. Results indicate a strong desire to participate in trades, but high transactions costs have limited the number of trades that have occurred. Utilizing empirical models improves the accuracy of predicting trade participation and direction, and therefore the accuracy of models of how introducing water transfers affects water supplies and stream flows. Such information is critical for policymakers who are considering introducing or expanding the use of groundwater trading.

Acknowledgements

The authors would like to thank Dirk Dinnel, Nathan Jenkins, and Jasper Fanning of the Upper Republican Natural Resource District for making the data on water trades available and answering countless questions. We would also like to acknowledge the USGS 104B Grant 2012NE224B and NSF/USDA Water Sustainability and Climate Grant USDA AG-2012-67003-19981 for funding this research.

*Corresponding author: kschoengold2@unl.edu, Associate Professor, Department of Agricultural Economics and School of Natural Resources, University of Nebraska at Lincoln. PRELIMINARY DRAFT. Please do not cite without permission of the authors.

1 Introduction

Economists have a long history of advocating market development and tradable permits to manage scarce resources. Markets have been widely used to manage environmental quality (see Fowlie and Perloff (2013) and Goulder (2013) for examples of permit markets) and some common property resources (see Newell, Sanchirico, and Kerr (2005) and Costello, Gaines, and Lynham (2008) for examples from the fishery industry). While the literature has discussed the potential benefits of using a market to manage groundwater resources (Thompson et al., 2009; Kuwayama and Brozović, 2013; Palazzo and Brozović, 2014), there is little empirical support for these results. This is in stark contrast to surface water markets, which have been widely developed and analyzed (Brown, 2006; Brewer et al., 2008). While surface water trading has occurred regularly throughout the western United States and the rest of the world for decades, the use of any type of groundwater trading has been very limited. However, groundwater is increasingly under stress from overuse and many areas are starting to regulate groundwater use. While the public benefits from efficient use of groundwater include adequate stream flow in hydrologically connected areas and future availability of groundwater supplies, there are significant private benefits to landowners especially in water short areas. Groundwater trading can help move water from low-value to high-value areas of use. Empirical analysis of groundwater trading is a new area of research due in part to the lack of information on three measures: recorded usage, trades, and binding constraints on groundwater use by landowners.

While groundwater trading may be a beneficial tool to move water to high value uses and encourage conservation, it requires that the water market operate efficiently. Dinar, Rosegrant, and Meinzen-Dick (1997) and Saliba (1987) agree that efficient water markets require: 1) many sellers and buyers with full knowledge of the market institutions and facing similar transaction costs; 2) participation decisions are made independent from other buyers and sellers; 3) outcomes are not affected by the decisions of other participants; 4) participants are assumed to be maximizing profits; 5) completely specified, enforceable, and transferable property rights. Market systems that have met these requirements will move resources from low value uses to high value uses, resulting in an economically efficient allocation of resources for both individuals and society, so long as the gains in value are large enough to offset the costs of completing the transaction. Because markets move water from low value to high value by allowing compensation for the water sold, they provide an incentive for more efficient use of water and reduce the stress on the water supply from high value uses. A well-designed water market requires the measurement and monitoring of water withdrawals, enforcement of withdrawal rules, and should consider any externalities or third party effects.

However, if high transaction costs persist, differences in marginal water values will continue to influence the market and the prices of water rights will vary between uses. Because detailed water rights are generally heterogeneous, transaction costs tend to be higher as buyers and sellers must engage in market searches that fulfill the institutional regulations on legal and hydrological characteristics of the water rights involved. High transaction costs reduce the level of profitable transactions but are not necessarily a sign of an inefficient market. Saliba (1987) concludes that water markets are in fact functioning well where the economic incentives for transfers outweigh the transaction costs involved and where the policies regarding markets are not costless, but are necessary.

The current paper begins to fill this gap in the literature. Specifically, we use data from two types of actual groundwater trades in a region where producers rely entirely on groundwater for agricultural production. We refer to these types of trades as *informal* and *formal*. There are two primary differences between these types of trades. Implementing an informal trade is almost costless, while formal trades require finding a trading partner. This can have high transaction costs. Informal trades can occur for any number of years, from a single year lease to a repeated transfer. Formal trades must be permanent transfers of an irrigation right. Both are restricted based on geographical constraints, where the buyer and seller must both be located within a six by six mile square.

2 Background and Description of the Study Region

Previous research on water markets have primarily focused on surface water trading, but interest in developing markets for groundwater trading has begun to increase in recent years. Although the majority of surface water rights use the appropriation doctrine of "first in time, first in right", groundwater rights in Nebraska utilize the correlative rights doctrine that is administered by the local Natural Resource District (NRD). The NRD system in Nebraska includes 23 districts that are responsible for management of natural resources such as soil and groundwater. The correlative rights doctrine ties all appropriated water rights together and assigns equal priority to all rights. This idea implies that when a shortage in groundwater within a NRD has been triggered, all groundwater-rights holders within that NRD have their allocation, or supply, decreased proportionally.

Our analysis uses information from the Upper Republican Natural Resource District (URNRD), a three-county region in southwestern Nebraska (see Figure 1). The majority of the land is used for grain production and cattle ranching. Land used for grain production is either irrigated to produce corn, soybeans or wheat or used for dryland production of wheat and edible beans. In 2012, the URNRD had a total of 3,179 active irrigation wells servicing 452,395 certified acres (Palazzo and Brozović, 2014). On average, the study area receives 15 to 20 inches of precipitation annually (U.S. Department of the Interior, 2005). This limited rainfall during the growing season in the western portion of Nebraska results in the high use of irrigation systems to produce grain crops like corn and wheat, which require about 22 inches of rainfall a growing season to reach high yielding maturity.[1] The primary source of

[1] See http://www.extension.org/pages/14080/corn-water-requirements for more information.

the water used for irrigation is groundwater from the High Plains Aquifer, much of which is hydrologically connected to rivers and watersheds located in the state.

Starting in 1978, the URNRD has been actively involved in the management of the groundwater resources within the district through the adoption and enforcement of rules and regulations. Along with monitoring pumping from each well, the URNRD has established correlative irrigation allocation rights. Since the first allocation limits were established in 1983, the annual irrigation allocation has decreased from 22 inches per year to the current allocation of 13 inches per year. These allocations are issued as an aggregate amount to be allocated over five years.

Unlike many other groundwater-dependent areas across the nation, the Upper Republican Natural Resource District (URNRD) in Nebraska has had metering and use restrictions in place for over 30 years. The URNRD has also developed some mechanisms to help producers use water most efficiently under allocation restrictions. Two of the tools available include creating pools and formally trading water. Formal water trading occurs when the irrigation rights are permanently transferred from one field to another field.[2] The URNRD also allows for the pooling of field-specific allocations into an aggregate allocation for groups of fields owned by the same person, partnership, or corporation, subject to a signed agreement. When multiple fields are combined to create a pool, a producer can shift a portion of the water allocation from one field to another field in the same pool. Within-pool transfers are conducted when fields under the same owner aggregate their total allocations into a pool, as approved by the URNRD board, and then redistribute the water to each field at the owner's discretion. These within-pool transfers have much lower transaction costs than the formal trades due to less time and money spent finding a trading partner, gaining board approval, and paying any associated broker fees. Therefore, the URNRD provides a great opportunity

[2]The data is almost exclusively made up of single well-single field pairs. Thus, a reference to a field is synonymous with a reference to the associated well.

to better understand the impacts of allowing some restricted trading, and how the rules associated with trading affect the decision to participate in trades.

3 Analytical Framework and Stylized Facts

We analyze two different types of trades in the analysis. Informal trades occur when a producer has received official approval to combine, or pool, the allocations from individual fields, and he uses that flexibility to shift water from one field to another. Fields that are not part of a pool are constrained by the allocation. Formal trades occur when a producer chooses to sell his allocation to another person. The rules implemented by the district require that these trades are permanent and for the entire allocation (i.e., the seller will shift to dryland farming or grazing, and the buyer can use the allocation on a new field).

3.1 Single Fields

A producer that is not part of a pool will choose water application to maximize expected profits. Specifically, we assume that a producer solves the following optimization problem:

$$\max_{w_i} \Pi_i(w_i, z_i) \text{ s.t. } w_i \leq \bar{w}, \tag{1}$$

where w_i denotes irrigation use on field i, Π_i denotes profit on field i, z_i denotes the field-level characteristics such as soil type and slope, and $\bar{w}$ denotes the allocation.

3.2 Pooled Fields

A producer that operates N fields in a pool solves the following optimization problem:

$$\max_{\{w_i\}} \sum_{i=1}^{N} \Pi_i(w_i, z_i) \text{ s.t. } \sum_{i=1}^{N} w_i \leq N\bar{w} \tag{2}$$

3.3 Formal Trades

Formal trades require a landowner to permanently sell his entire irrigation allocation and put the land into dryland production or grazing land. Thus, a producer that is considering a formal trade chooses between selling his allocation and earning the profit in Equation 1 or Equation 2.

$$\Pi_i = \max\{\Pi_i(0, z_i) + \tau\bar{w} - T(z_i), \Pi_i(w_i^*, z_i)\}, \tag{3}$$

where w_i^* is the profit maximizing level of water allocation determined in Equation 1 or Equation 2, τ is per-unit sale price for a water allocation, and $T(z_i)$ are the transaction costs associated with finding another landowner to trade with. We assume that those costs may depend on individual characteristics. For example, a landowner who manages his own operation may find it easier to find someone to trade with than a landowner who rents his land to a tenant and lives out of state.

3.4 Predictions Based on Stylized Facts

Using the stylized facts that describe the groundwater management framework in the URNRD and economic theory, we can make some predictions about trading behavior by producers. Specifically, we can show that we expect the following outcomes to hold:

- Pooled fields will equate the marginal abatement cost for all fields in the pool
- A field in a pool is less likely to participate in a formal trade, as the opportunity cost of not trading is higher for a single field than for a field in a pool

We can test these hypotheses using detailed data on formal and informal groundwater trades in the URNRD.

4 Description of the Data

The data used in the analysis includes publicly available well data from the Nebraska DNR and proprietary data collected and provided by the URNRD. The Nebraska DNR well database contains technical information on each well, including pump depth, pump rate, date of drilling, and the current status of the well (e.g, active, abandoned) and on the use (e.g., irrigation, livestock, domestic). Specific geographic location information and well ownership information is also included within the file. There are 4,604 wells in the URNRD region listed in the DNR database; however, only 3,274 are used for active irrigation after reconciling with the data provided by the URNRD.

Much of the data for the analysis came directly from the URNRD. The URNRD provided records of formal allocation transfers of groundwater in the 2006 to 2011 period. This information was used to identify 35 formal trades involving 100 unique field IDs and was verified using the allocation adjustment data set. Figure 2 shows a map of the wells that participate in formal trades. Other data provided by the URNRD includes the annual water usage of each irrigation well and the relevant groundwater pool (if applicable). using the historical data, we identify 1,974 fields that participated in pools during 2005-2007 and 1,914 during 2008-2012. Due to a change in the pool naming convention that made it difficult to track the history of the pools to those that existed before the name change, we only use the pools in the two recent allocation periods in this analysis. The URNRD also provided data on the following for each field ID: owner/operator IDs in 2011, the associated well ID, annual certified acres, crop planted, pool ID (if applicable), and water use by year. When reconciled, the final data contains information on 3,179 unique field IDs.

Groundwater price and quantity values used for the marginal abatement cost (MAC) indicators were generated by Palazzo and Brozović (2014). The MAC values are determined using the Nebraska DNR well data and results from the Water Optimizer program (Martin et al., 2007). Each point in the MAC curve is determined by calculating the change in the profit associated with increasing the stringency of the pumping constraint, starting from the current allocation of 13 inches per year. These values are based on well and field level characteristics (e.g., depth to groundwater, soil type), average values for weather variables, average prices of inputs including energy for irrigation pumps, and average market prices for harvested crops. We use the MAC curves to determine the predicted buyer and seller in an actual trade and to measure expected gains from trade. Figure 3 shows an example from one of the formal trades in the data. In Figure 3, Well B has an expected MAC of approximately \$8 with a reduction of two acre-inches of irrigation, while Well A has an expected MAC that is close to zero (about \$0.18). This is not unrealistic in the study region, as some fields do not need the full allocation of irrigation water in an average year to maximize profit. Based on the differences in MAC curves, we expect that Well A is the seller if two inches of water are transferred between these two fields. When multiple fields are combined into a pool or used in a formal trade, we use the average MAC value for all fields in the pool to determine expected buyers and sellers. The MAC curves for Well A and B do not intersect in this example (i.e., the MAC for Well B is above or equal to the MAC for Well A throughout the abatement distribution). However, that is not consistently true, and in cases where the MAC curves intersect, we need to use the size of the transfer to determine the expected buyer and seller.

We determine the expected buyers and sellers differently for pools than for formal trades. Formal trades require that the full allocation is transferred and the seller switch from irrigated production to dryland crop production or grazing. Thus, we determine the expected buyer and seller based on the total abatement cost (TAC), which is equivalent to the area under the MAC curve. Informal trades occur with much smaller quantities. Figure 4 shows the distribution of the size of informal trades. The dashed lines indicate the two and three acre-inch thresholds. Analyzing the distribution shows that over 70 percent of informal trades are less than two acre-inches, and over 80 percent are less than three acre-inches. Thus, we use the abatement cost at two acre-inches to predict buyers and sellers.

5 Empirical Methodology

The primary goal of our analysis is to estimate actual trading behavior, compare it to predicted behavior, and determine if transaction costs and trading constraints are limiting the occurrence of otherwise economically beneficial trades. We recognize that the placement of fields into pools, and the choice to participate in a formal trade, are not randomly assigned. Thus, we first test if there are systematic differences in the groups. The existence of systematic differences provides evidence that we need to incorporate a selection equation in the estimation of trade direction. Table 1 compares the differences in means and proportions between three different groups of fields: those that are members of a pool, those that participate in a formal trade, and those that are constrained by the URNRD's current water allocation. This last group is determined by measuring if average historical use for the field is above 12.9 inches. The first two groups form the basis of our empirical investigation and the third is included for informational purposes.

Table 1 divides the variables into three types of well characteristics: physical (hydrological, geographical, and agronomic), technical, and operational. Physical characteristics include the *pumping water level*, which is the depth to groundwater at the time the well was installed; the *SDF* or stream depletion factor, which is a measure of the impact of groundwater pumping on surface water flow; and soil type and geographical indicators. Technical characteristics include the MAC (measured at two acre-inches), the *GPM* or gallons per minute, and the year that the well was completed. Operational characteristics include the number of times the field grows multiple crops in a year (*Multicrop Years*), the average amount of irrigation water applied (*Average Use*), the size of the operation (*Operation Size*), whether the field well is managed by the landowner (*Owner Operator*), and whether or not the well is constrained at the current allocation of 13 inches per year, based on historical use.

The group comparison shows that some important differences exist between those wells/fields that trade water (either informally or formally) and those that do not. Fields that join pools are those that have a greater benefit and need from sharing a water allocation. On average, those fields use more irrigation water and are more likely to be constrained by the URNRD allocation. The URNRD also requires that a pool only include fields owned by a single landowner, thus fields in pools are in larger operations than those fields operated independently. A comparison between fields that participate in formal trades and those that do not shows that participants differ in several physical characteristics, including the pumping water level and soil type. Both of these will affect the relative benefit of trading water. Fields that participate in formal trades are more likely to multicrop the land, and, this is a potential reason that it is water buyers find it beneficial to purchase an allocation. Those wells that are constrained differ from the unconstrained set in almost every characteristic, although many of the differences are intuitive. For example, wells with a greater capacity (measured in gallons per minute) or those with coarse soil (relative to fine or medium) are more likely to be constrained.

5.1 Empirical Framework

Given the nonrandom nature of whether a field is in a pool, or if it participates in formal trades, we use a probit model with sample selection (Van de Ven and Van Praag, 1981) to estimate if a field is a buyer or seller. The model assumes that there is an underlying relationship that measure the profitability of buying a water allocation,

$$y_j^* = x_j\beta + u_{1j}, \tag{4}$$

but that we only observe a binary outcome for individual j, based on the following,

$$y_j = \begin{cases} 0 & \text{if } y_i^* \leq 0 \\ 1 & \text{if } y_i^* > 0 \end{cases} \tag{5}$$

In addition, we only observe $y_j \in 0, 1$ if the field selects into a pool or a formal trade. Formally,

$$y_j^{observed} = z_j \gamma + u_{2j}, \tag{6}$$

where $u_1 \sim N(0,1)$, $u_2 \sim N(0,1)$, and $corr(u_1, u_2) = \rho$.

The general structure of the model, given by Equations 4 - 6, is the same for estimating trading behavior within pools and formal trades. However, the choice of variables will differ for a couple reasons. If there are no transaction costs, we expect the same variables to predict participation and buying/selling decisions in both subsets. If transaction costs are higher for formal trades (as we hypothesize based on anecdotal evidence), we need to include variables to incorporate this difference. We also need to include variables that reflect the differences in marginal abatement costs under a partial-allocation transfer and a full-allocation transfer. Finally, the variables we include to predict whether an individual is a buyer or seller are different than those variables which determine trade participation. In determining if a field is a buyer or seller, it is the field characteristic relative to the trading partners that matters. However, when a landowner decides to form a pool or enter a formal trade, the actual characteristic affects the transaction costs and benefits of doing so.

5.2 Variables included in Participation Decision

The first stage of the estimation (the participation stage) is where we predict if a landowner chooses to participate in a formal trade and/or to put his field into a pool. We expect that important variables fall into one of two categories: those that affect an individual's transaction costs of participating, or those that affect an individual's relative benefits.

The land tenure indicator (*Owner Operator*) is a binary variable that equals one when the owner of the field is also the operator, based on the 2012 URNRD data. The sign of this variable in the participation models is expected to be positive as a landowner is likely to find it easier to find a trading partner for formal trades. With informal trades, he is also more likely to put forward the effort to create a pool. Doing so is a fixed cost, but a landowner is more likely to continue to operate the field in the future and receive the benefits from having the field in a pool. The operation size (*Operation Size*) is based on URNRD owner-operator data from 2012. This is expected to be a positive determinant of participation in pools and formal trades. With pools, the reason is obvious. The URNRD has a preference for a pool of fields that all have the same owner, although it still requires the pooled fields to be within a set distance from each other. A larger operation is simply more likely to have multiple fields that satisfy those requirements. A larger operation is also more likely to participate in formal trades, as the operator will have more resources for finding a trading partner (i.e., reduce the transaction costs).

The average use pre-2008 or pre-trade is included to explain pool or formal trade participation. We expect that fields with a higher level of historical irrigation water use will have a greater benefit from participating in a pool or trade. The soil type indicators (*medium* and *coarse*) are binary variables that measure the soil's ability to retain water after an irrigation cycle. Relative to *fine* soil (the omitted category), we expect that landowners with medium and coarse soil are more likely to participate in formal trades or pools. These fields are more likely to be constrained with the allocation limits and create a higher benefit of the increased flexibility from a pool, or from purchasing a water allocation. The pool participation variable is a binary indicator of whether the field was in a pool in the 2005-2007 period.[3] As

[3] The data includes pools from 2005-2007, and from 2008-2012. Because of a change in naming convention, it is difficult to determine if the pools remain constant between the two periods.

predicted based on our stylized facts, we expect that fields that are already in a pool will have a lower marginal benefit of participating in a formal trade relative to those fields that do not have the flexibility from an existing pool. The county indicator variables (*Perkins* and *Dundy*) are included to capture unobservable differences between the three counties in the URNRD.

5.3 Variables included in Behavior Decision

We only include a small number of variables in the estimation of trade direction. A critical difference between the estimation of trade participation and trade direction is that variables that determine trade direction are relative values, not absolute values. To determine whether a field is predicted to be a buyer or seller, we compare the MAC from the field to the average MAC value of all fields in the pool or trade. A field with a higher (lower) MAC than the pool/trade average is predicted to be the buyer (seller). Specifically, Field j is a predicted buyer if $TAC_{2j} > \sum_{i=1}^{N} TAC_{2i}/N$ where N is the number of fields in the pool or trade. Since the MAC values already incorporate many of the physical characteristics that differentiate fields (e.g., absolute soil quality, pumping water level, gallons per minute), we do not include those as additional covariates.

We include two additional explanatory variables: constrained status and relative soil quality. The *Constrained* variables are calculated differently for a pool than a formal trade. With a pool, the variable measures whether the overall pool is constrained at the current URNRD water allocation. Due to the convexity of the MAC over small reductions in water allocation, we expect more fields overall are sellers, and thus those fields in a constrained pool are more likely to be a seller than a buyer. Since we expect behavioral differences between constrained and unconstrained pools, we include an interaction term with the predicted buyer variable. A measure of *Relative Soil Quality* indicates the value of soil quality for the individual field relative to others in the pool or trade. Specifically, we create an indicator variable *Soil quality*, which equals 1, 2, and 3 for fine, medium, and coarse soil, respectively. Field j has soil that is more coarse than the average field in the pool or trade if $\textit{Soil quality}_j > \sum_{i=1}^{N} \textit{Soil quality}_i / N$ where N is the number of fields in the pool or trade. Since coarse soil requires more applied water for the same level of output, we expect a field with positive relative soil quality to be a net buyer of water.

Finally, we include a variable to measure if the *relative stream depletion* factor is a determinant of the direction of a groundwater trade. The stream depletion factor (SDF) is a number between zero and one that indicates the impact of groundwater extraction at the well location on stream flow in the Republican River. The relative measure indicates whether the field has a higher or lower SDF than the pool or trade average. This is not a variable that affects the marginal profit of additional irrigation water. So, a priori, there is no reason to expect producers to base trading decisions on a field's SDF. However, the URNRD is under pressure to reduce the impact of groundwater pumping on stream flow in order to meet interstate compact requirement. The URNRD is considering adding restrictions on trade that reduce an allocation proportionally if it is moved from a field with a low SDF to a high SDF. Thus, producers that want to move water to a high SDF field may have had an incentive to do so before these regulations were added.

6 Results and Discussion

We present the results in Table 2. We estimate pool participation and direction using the Heckman selection model, as described in Equations 4 - 6 and find strong evidence that the two pool equations are correlated. However, the test statistic shows no evidence of correlation in the participation and direction equations for formal trades. Thus, we report the results from two separate probit estimations. Overall, results are stronger and more consistent with our expectations with the estimation of pool participation and direction than with the formal trade estimation. Given the small number of formal trades and associated fields/wells, this is not surprising. The estimation does a reasonable job of predicting pool participation and trade direction, with 68.5, 70.8, and 67.0 percent of outcomes correctly predicted using a prediction threshold of 0.5. A threshold of 0.5 does not work well for predicting formal trade participation. A small proportion of all fields participate in formal trades (about 2.6 percent), and thus, the model predicts no trade participation with a probability threshold of 0.5. Lower threshold values perform better.

The participation regression results are in the bottom panel of Table 2. We expect that *Operation size* is an indicator of the relative importance of the transaction costs associated with trading. A larger operation will have greater familiarity with the area and more resources to find trading partners. The results support this hypothesis, as larger operations are more likely to participate in pools and in formal trades. This is an expected result for pool participation based on URNRD regulations about field ownership/management and pool formation. However, those regulations do not apply to formal trades, and the variable maintains statistical significance. Despite the statistical significance, the magnitude of the marginal effect is small (0.0003), indicating that an operation with 10 additional fields is only 0.3 percent more likely to participate in a formal trade.

We find that a history of a high level of irrigation use has mixed results. Fields with a higher average irrigation water use before 2008 (for pool participation) are more likely to participate in a pool, a result that may reflect a larger marginal benefit of the additional flexibility the pool provides. We find the opposite result for formal trade participation. This result is driven by the large number of low water users that participate in formal trades and sell their water allocation. Specifically, the pre-trade average use is 8.78 inches per year for sellers and 12.66 inches per year for buyers, relative to a district-wide average of 11.70 inches per year. While this result is not surprising, it does create concern about the potential for a trade to result in an increase in overall groundwater depletion if the URNRD does not limit the water transfer to the historical use, and instead allows a buyer to use the full allocation.

Soil characteristics affect the marginal benefit of additional flexibility in water allocation, and we find that fields with medium and coarse soil are more likely to participate in pools. Soil quality has a small amount of power in explaining formal trades, but is much more limited. Based on the stylized facts presented earlier, we hypothesize that a landowner with a field in a pool is less likely to participate in a formal trade, since he already has some benefit from the flexibility a pool provides. The pool participation variable is negative and significant, with a field in a pool 1.2 percent less likely to participate in a formal trade.

The trade direction estimation results are in the top panel of Table 2. Results from the trade direction equations are important for two reasons. First, existing literature suggests that there could be large gains from creating groundwater markets and facilitating transfers, and that creating markets will provide incentives to shift water from low-value uses to high-value uses (Thompson et al., 2009; Kuwayama and Brozović, 2013; Palazzo and Brozović, 2014). However, limited data has not allowed those theories to be tested using actual transaction data. Second, regulators have been cautious about expanding the use of transfers due to concern about increasing overall water use and affecting hydrologically connected streams and associated ecosystems. Analyzing actual transfers will allow regulators to learn more about how to design new regulation, and which concerns need to be incorporated into regulatory design.

The estimation of informal trade direction shows that the predicted buyer measure is significant in predicting actual buying choices, although the marginal effect is small. A field that is predicted to be a buyer (relative to the rest of the field) is 2.9 percent more likely to actually purchase water. The pool constraint indicator shows that more fields will sell when the overall pool is constrained. This is consistent with our predictions based on the shape of the MAC curves. Many fields are able to reduce water use from current allocation levels with no effect on profitability in an average year. Thus, a landowner can shift a water allocation from several unconstrained fields to a single constrained field and increase profitability. The ability to do so may be reduced in the future if allocations are reduced from the current level.

The relative stream depletion factor is negative (insignificant) with informal (formal) trading. This is reassuring, as there is no indication that producers are shifting water to fields with a greater impact on stream flow because of an expectation that those trades will not be allowed in the future. In fact, despite the fact that stream depletion was not considered in the approval process during the study period, the overall effect of trading has been to reduce stream depletion.

The impact of a binding allocation level (reflected in a constrained pool or field) is the opposite with formal trades than with pools. The definition of the variable is a little different. With pools, the variable measures if the overall pool is constrained at the current allocation. With formal trades, the variable measures if the field is constrained based on its historical use. Those fields that have been constrained historically are more likely to purchase water. The marginal effect of 30.7 percent has both statistical and economic significance.

We are especially interested in whether we can correctly predict which fields purchase water. Doing so could help support the URNRD or other locations that are interested in creating groundwater markets. Accurately predicting buyers and sellers will allow regulators to predict the outcomes of new groundwater markets better. In cases where the spatial distribution of water use matters, this is important for understanding impacts on neighboring wells, hydrologically connected streams, and associated ecosystems. While we find marginal statistical significance with informal trades, we do not find evidence that the predicted buyer indicator helps to predict formal trade purchases. This may be for a couple of reasons. First, the transaction costs of finding a trading partner may overwhelm any expected gains from the trade. Also, the predictions for a buyer or seller are based on average values for input and output prices. Since formal trades require a permanent transfer of the water allocation, a change in those prices could switch the predicted buyer and seller. Finally, the relatively small number of formal trades make it difficult to accurately assess the determinants of those decisions.

7 Conclusions

The water trading literature states that there are significant economic gains to be achieved by moving water from areas of low efficiency to areas of higher efficiency within a region. This study utilized a unique set of data from a region where groundwater is constrained and monitored to analyze the factors that predict the probability of participating in formal and informal trades, as well as the direction of trade among participants. The participation models use field-level variables and provide insight into the participation decision process. The trade direction models use those variables as well as other factors that are unique to the field and trade. The results of the model support our hypotheses and can be used to guide ex-ante evaluation of groundwater trading in other regions. However, results also show that the presence of high transaction costs and restrictive trading regulations are limiting trades that would otherwise be economically beneficial.

A major constraint to the formal trade estimation is limited data on formal participation, resulting in poor statistical significance of empirical results. However, the large participation in informal trades is an indicator that there is, in fact, substantial interest in trading water but that there are currently barriers preventing more formal trading. Once a pool is formed, the marginal cost of trading water is effectively zero, whereas the marginal cost of formal trading under the current process is significantly higher. The large participation in informal trading is a sign of potential economic gains for the URNRD from reducing the transaction costs associated with formal trading. If the marginal cost of participating in a formal trade is reduced (for example, through the aid of an online trading platform to find potential trading partners) participation in formal trades is expected to increase substantially.

Expanding the formal trading market to include annual use trades (leases) in addition to the permanent trades also has the potential to open the market and allow for more observations for the formal participation model. The ranking of the MAC curves are not stationary under different precipitation and crop price scenarios, indicating that buyers and sellers may switch roles under different production situations. The ability for buyers and sellers to switch roles, as predicted by their MAC relationship, indicates that annual leases would provide an additional risk management tool for producers by generating flexibility in the field's annual allocation.

References

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* 46:91–112.

Brown, T.C. 2006. "Trends in water market activity and price in the western United States." *Water Resources Research* 42.

Costello, C., S.D. Gaines, and J. Lynham. 2008. "Can catch shares prevent fisheries collapse?" *Science* 321:1678–1681.

Dinar, A., M.W. Rosegrant, and R.S. Meinzen-Dick. 1997. *Water allocation mechanisms: principles and examples*. 1779, World Bank Publications.

Fowlie, M., and J.M. Perloff. 2013. "Distributing pollution rights in cap-and-trade programs: are outcomes independent of allocation?" *Review of Economics and Statistics* 95:1640–1652.

Goulder, L.H. 2013. "Markets for Pollution Allowances: What Are the (New) Lessons?" *The Journal of Economic Perspectives* 27:87–102.

Kuwayama, Y., and N. Brozović. 2013. "The regulation of a spatially heterogeneous externality: Tradable groundwater permits to protect streams." *Journal of Environmental Economics and Management* 66(2):364–382.

Martin, D., R. Supalla, B. McMullen, and S. Nedved. 2007. "Water Optimizer: a decision support tool for producers with limited water." *University of Nebraska-Lincoln Departments of Biological Systems Engineering and Agricultural Economics*, pp. .

Newell, R.G., J.N. Sanchirico, and S. Kerr. 2005. "Fishing quota markets." *Journal of Environmental Economics and Management* 49(3):437–462.

Palazzo, A., and N. Brozović. 2014. "The role of groundwater trading in spatial water management." *Agricultural Water Management*, pp. .

Saliba, B.C. 1987. "Do water markets work? market transfers and trade-offs in the southwestern states." *Water Resources Research* 23:1113–1122.

Thompson, C.L., R.J. Supalla, D.L. Martin, and B.P. McMullen. 2009. "Evidence supporting cap and trade as a groundwater policy option for reducing irrigation consumptive use." *Journal of the American Water Resources Association* 45:15081518, doi: 10.1111/j.1752-1688.2009.00384.x.

Van de Ven, W.P., and B.M. Van Praag. 1981. "The demand for deductibles in private health insurance: A probit model with sample selection." *Journal of econometrics* 17:229–252.

Table 1: Comparison of Characteristics by Group

Variable	Pool Participation			Formal Trade			Constrained		
	Yes	No	Test Statistic	Yes	No	Test Statistic	Yes	No	Test Statistic
Physical Characteristics									
Pumping Water Level (feet)	125.47	153.70	10.69***	112.97	136.94	3.20***	105.51	154.40	19.11***
Stream Depletion Factor	0.57	0.57	0.10	0.58	0.57	-0.48	0.54	0.59	5.61***
Fine$^{\Psi}$	0.28	0.42	8.36***	0.31	0.33	0.39	0.22	0.40	9.97***
Medium$^{\Psi}$	0.35	0.32	-1.68*	0.44	0.33	-2.32**	0.33	0.34	0.83
Coarse$^{\Psi}$	0.36	0.24	-7.07***	0.21	0.32	2.28**	0.44	0.24	-11.68***
Perkins$^{\Psi}$	0.23	0.39	9.60***	0.11	0.29	3.94***	0.11	0.39	16.53***
Dundy$^{\Psi}$	0.30	0.27	-1.65*	0.25	0.29	0.74	0.41	0.21	-11.77***
Technical Characteristics									
MAC (at 2 inches)	23.02	18.78	-10.72***	21.48	21.42	-0.05	26.59	18.34	-21.93***
Gallons per Minute	1545.03	1460.13	-3.07***	1486.59	1513.61	0.35	1657.80	1428.22	-8.33***
Completed Year	1975.29	1975.51	0.65	1972.90	1975.45	2.60**	1974.41	1975.94	4.33***
Operational Characteristics									
Multicrop Years	2.58	2.88	2.08**	3.82	2.74	-2.53**	2.47	2.88	2.78***
Average Use (acre-inches)	12.57	10.53	-18.63***	11.09	11.91	2.71***			
Operation Size (fields)	15.83	9.43	-9.58***	14.11	13.38	-0.39	19.12	10.04	-13.71***
Owner Operator$^{\Psi}$	0.56	0.58	0.75	0.52	0.57	1.12	0.51	0.61	5.40***
Constrained$^{\Psi}$	0.46	0.24	-13.19***	0.38	0.37	-0.16			
N	1969	1399		99	3075		1173	1997	

Ψ indicates a binary indicator variable.

The test statistic is a t-test of mean differences for continuous variables and a z-test test of proportional differences for indicator variables. Significance at the 10%, 5%, and 1% levels are denoted by *, **, and ***, respectively.

Table 2: Determinants of Trade Participation and Direction: Regression Results

	Informal					Formal			
	Coefficient	T-Stat	Marg Effect	T-Stat		Coefficient	T-Stat	Marg Effect	T-Stat
					Trade Direction = Buyer				
Constrained Pool$^{\Psi}$	-0.057	-0.89	-0.029	-1.75*	Constrained Field$^{\Psi}$	0.679	1.66*	0.307	3.14***
Pred Buyer$^{\Psi}$	0.113	1.89*	0.029	1.50μ		-0.460	-1.07	-0.117	-1.05
Constrained x Pred Buyer	-0.052	-0.56				0.288	0.51		
Relative Soil Quality	-0.047	-1.39	-0.017	-1.40		0.504	0.95	0.175	0.97
Relative Stream Depletion	-0.351	-1.46μ	-0.125	-1.46μ		-2.126	-1.34	-0.741	-1.38
Constant	0.560	6.98***				-0.046	-0.16		
					R-squared	0.109			
					N	92			
Correct Predictions	70.8%				Correct Predictions	67.0%			
	Informal					Formal			
	Coefficient	T-Stat	Marg Effect	T-Stat		Coefficient	T-Stat	Marg Effect	T-Stat
					Participation				
Owner Operator$^{\Psi}$	0.035	0.75	0.030	1.55μ		-0.065	-0.65	-0.004	-0.65
Operation Size	0.006	4.01***	0.003	5.35***		0.006	2.05**	0.0003	2.03**
Average Use Pre-2008	0.121	15.36***	0.038	11.85***	Average Use Pre-Trade	-0.051	-3.36***	-0.003	-3.26***
Soil = medium$^{\Psi}$	0.308	5.55***	0.106	4.84***		0.176	1.54μ	0.011	1.53
Soil = coarse$^{\Psi}$	0.374	5.94***	0.134	5.77***		-0.049	-0.35	-0.003	-0.35
Pool Participation$^{\Psi}$						-0.192	-1.87*	-0.012	-1.85*
Perkins$^{\Psi}$	-0.304	-5.19***	-0.179	-7.58***		-0.625	-4.21***	-0.038	-4.01***
Dundy$^{\Psi}$	-0.334	-5.68***	-0.171	-7.03***		-0.329	-2.66***	-0.020	-2.61***
Constant	-1.288	-12.48***				-1.136	-5.90***		
Correlation coefficient (ρ)	-0.888				R-squared	0.058			
Chi2	77.54				N	3148			
Chi2 test of $\rho = 0$	0.00								
Correct Predictions	68.5%				Correct Predictions	97.4%$^{\phi}$			

Ψ indicates a binary indicator variable. The informal trade participation model has 3071 observations, with 1199 censored and 1872 uncensored.
Significance at the 15%, 10%, 5%, and 1% levels are denoted by μ, *, **, and ***, respectively.
ϕ: This number is artificially high because only about 2.6 percent of all fields participate in formal trades. Thus, predicting no participation is almost always correct. With a threshold value of 0.5 the model does not predict that any fields participate in formal trades. Lower threshold values increase the proportion that are determined correctly.

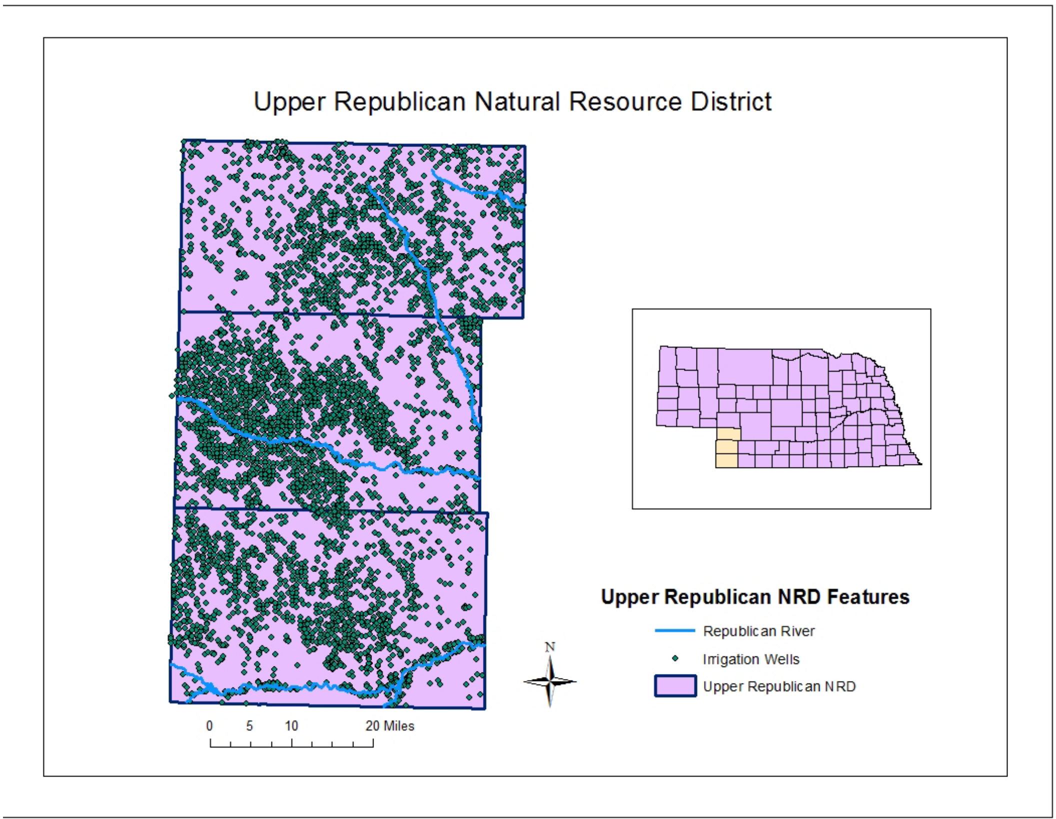

Figure 1: Map of All Irrigation Wells in the Upper Republican Natural Resource District

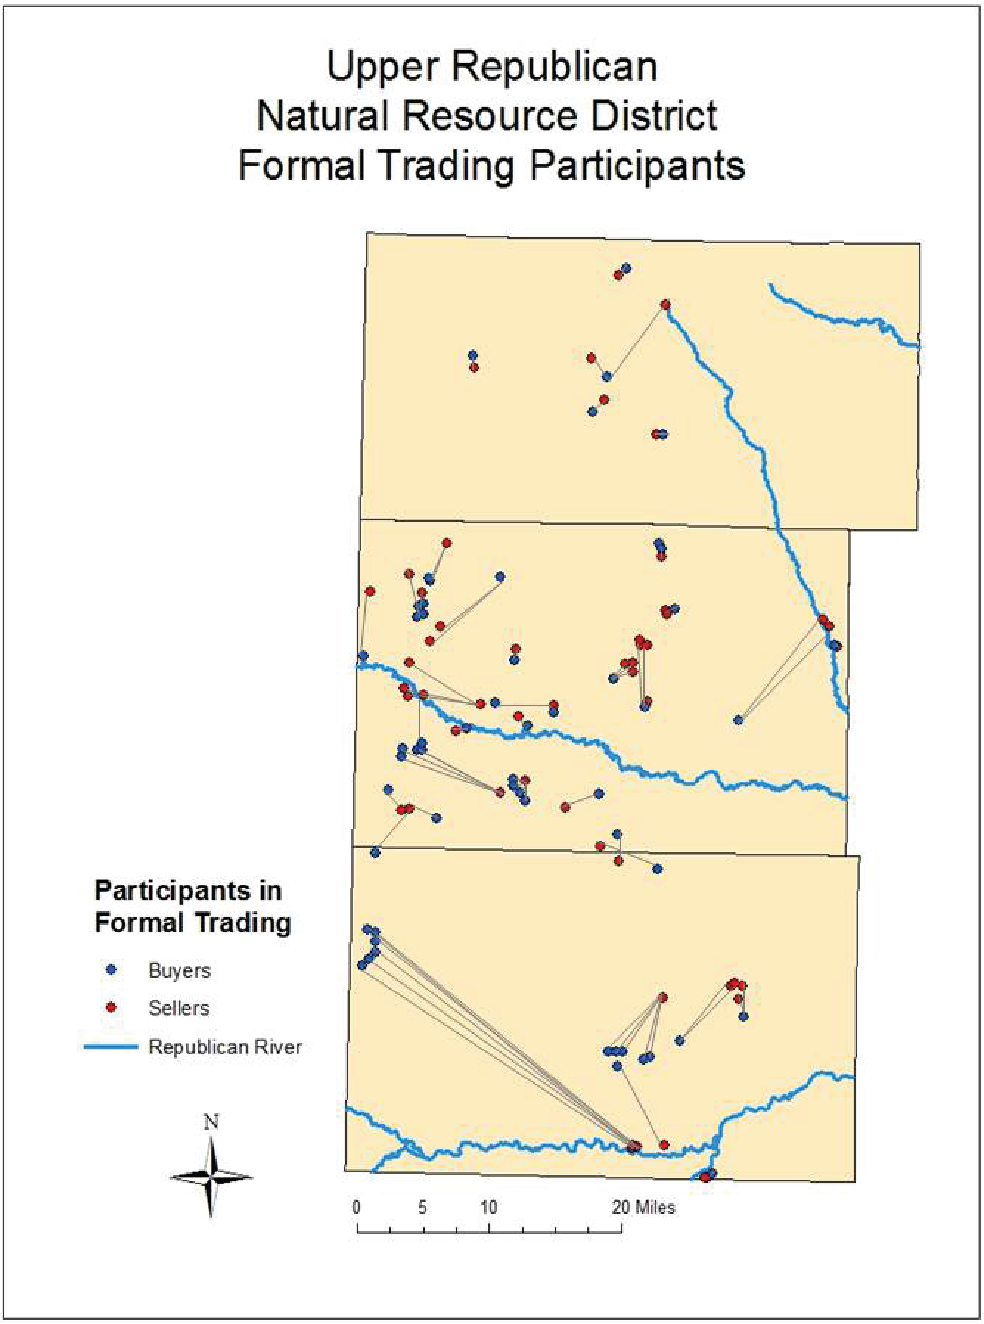

Figure 2: Location of Buyers and Sellers for Formal Trades

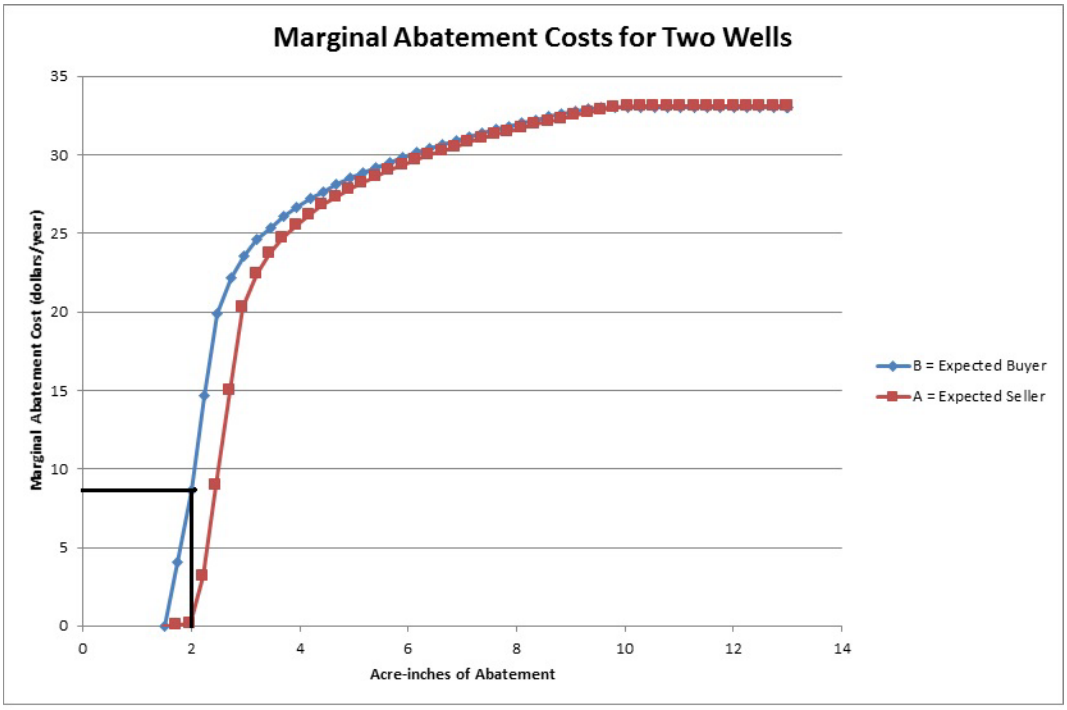

Figure 3: Determining Predicted Buyers and Sellers

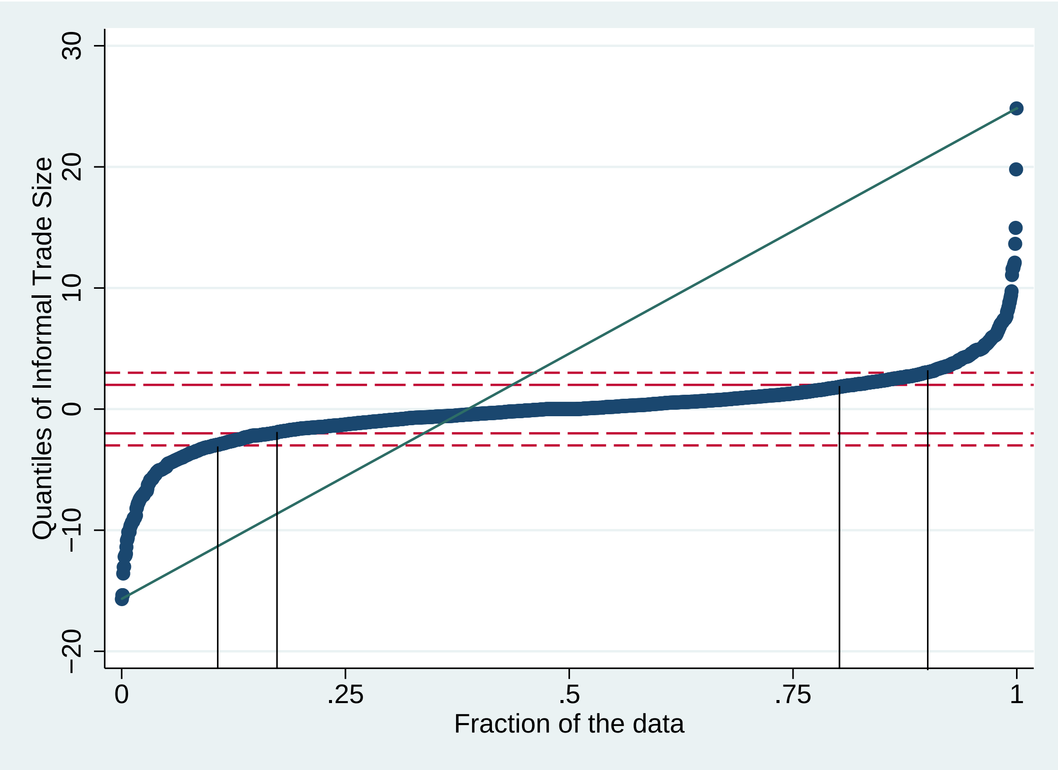

Figure 4: Distribution of Informal Trade Size (Acre-Inches)