Title
Essays on Water Resource Economics and Agricultural Extension

Author
Buck, Steven Charles

Publication Date
2011

Peer reviewed|Thesis/dissertation

Essays on Water Resource Economics and Agricultural Extension

By

Steven Charles Buck

A dissertation submitted in partial satisfaction of the

requirements of the degree of

Doctor of Philosophy

in

Agricultural and Resource Economics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Maximilian Auffhammer, Co-Chair
Professor David Sunding, Co-Chair
Professor Steven Raphael
Professor David Zilberman

Fall 2011

Essays on Water Resource Economics and Agricultural Extension

by

Steven Charles Buck

Abstract

Essays on Water Resource Economics and Agricultural Extension

by

Steven Charles Buck

Doctor of Philosophy in Agricultural and Resource Economics

University of California, Berkeley

Professor Maximilian Auffhammer, Co-Chair

Professor David Sunding, Co-Chair

This dissertation discusses topics in the microeconomics of water resource economics and agricultural extension. In one chapter I use a hedonic model to explain the price of land transactions, and from this an implied value of irrigation water is inferred. In a separate chapter I develop measures of willingness-to-pay for water supply reliability measures, and estimate how consumers respond to changes in the price of residential water. My final chapter develops a model of a farmer's decision to invest in learning from agricultural technicians about a new integrated pest management technology that improves yields and reduces agricultural run-off of pesticides into surface waterways. Each chapter is grounded in microeconomic models of decision-making.

Chapter 2 presents a hedonic analysis of farmland values to recover a value of irrigation water. Irrigation accounts for most freshwater diversions worldwide. Water available to farmers may be reduced in the long run by climate change and other factors. We use differences in farmland prices to value the availability of irrigation water. Using panel data on a set of farm parcels, we estimate the value of irrigation water using a plot-level fixed effects design – an approach to hedonic analysis that effectively controls for unobservable factors and may be useful in numerous settings. Our results suggest that the value of water for irrigation is much larger than measured in previous studies based on cross-sectional analysis. This finding has important implications for the analysis of policies that influence water supply reliability, and for the assessment of climate impacts.

Chapter 3 describes a method for evaluating consumer willingness-to-pay for water supply reliability when fixed costs are a large portion of the marginal price of residential water. We show that the willingness-to-pay for water supply reliability is determined by consumer demand for water, how the utility covers its costs, and the source of unreliability. The framework is applied to the case of 52 large urban water utilities in California covering the greater San Diego, Los Angeles, and San Francisco Bay Areas. We estimate residential water demand using a panel dataset tracking annual price and consumption data from 1994-2009 for 82 California water retailers. Based on our estimated price elasticities along with observed rates and the marginal costs of supply, we calculate

willingness-to-pay (WTP) measures for reliability improvements for segments of both Northern and Southern California.

Chapter 4 uses a simple utility maximization model to examine how trust attitudes affect farmer learning during an agricultural training. Using data originally collected in 2005 we examine if trust conditions a farmer's decision to learn during an agricultural training. We present a model of farmer behavior during the agricultural training in order to link trust measures to behavior in the field. We find evidence that farmers who trust agricultural technicians relatively more than community farmers learn more during training. The results provide insight into the design of agricultural extension services in Ecuador.

To Kathleen.

Contents

1 Introduction **1**

2 The Economic Impact of Changes in the Availability of Irrigation Water **3**
- 2.1 Economic Model . . . 4
- 2.2 Data . . . 5
  - 2.2.1 Outcome Variable: Farmland Sale Price Per Acre . . . 5
  - 2.2.2 Explanatory Variable of Interest: Surface Water Delivery Right Per Acre . 6
  - 2.2.3 Control Variables . . . 6
  - 2.2.4 Sample Selection . . . 7
- 2.3 Empirical Research Design . . . 10
  - 2.3.1 Empirical Challenge . . . 10
  - 2.3.2 Estimating Equations . . . 11
- 2.4 Results and Discussion . . . 13
- 2.5 Conclusions and Future Work . . . 16

3 The Value of Supply Reliability in Urban Water Systems **18**
- 3.1 Loss Framework . . . 19
- 3.2 Residential Water Demand Estimation . . . 22
  - 3.2.1 Econometric Specification . . . 22
  - 3.2.2 Data . . . 23
  - 3.2.3 Estimation Results . . . 23
- 3.3 Economic Losses of a Water Supply Disruption . . . 25
  - 3.3.1 Parameterizing the Loss Function . . . 25
  - 3.3.2 Data for Calculation of Losses . . . 26
  - 3.3.3 Results of the Loss Analysis . . . 27
- 3.4 Conclusions . . . 28

4 Agricultural Extension, Trust and Learning **30**
- 4.1 Defining Trust . . . 32
- 4.2 The Trust Game . . . 33
- 4.3 Modeling Farmer Behavior in the Agricultural Training . . . 33
  - 4.3.1 Model of Learning from Community Farmers and Technicians . . . 34
  - 4.3.2 Relative Trust & Selection into Learning Regimes . . . 36
- 4.4 Data . . . 37
  - 4.4.1 Background . . . 37

4.4.2 Sample Summary 38
4.4.3 Summary Statistics of Trust Measures 39
4.4.4 Summary of Exam Scores 42
4.5 Empirical Framework 43
4.5.1 Estimation Strategy 43
4.5.2 Sensitivity analysis and sources of potential bias 44
4.5.3 Standard Error and Pivotal Statistic Adjustments 45
4.6 Results and Discussion 45
4.7 Conclusions 51

List of Figures

3.1 Economic Losses Under Volumetric Prices and Constant Unit Cost 20

4.1 Histograms of trust measures 42

List of Tables

2.1 Comparison of plot characteristic means across samples 9
2.2 Cross-sectional sample: Regress price/acre on federal water deliveries 14
2.3 Panel sample: Regress price/acre on federal water deliveries w/o plot-level fixed effects 14
2.4 Panel sample: Regress price/acre on federal water deliveries w/ plot-level fixed effects 15

3.1 Residential water demand estimation 24
3.2 Controlling for weather in the residential water demand estimation 25
3.3 Welfare losses due to shortages of 10%, 20% and 30% 27

4.1 Summary statistics for control and treatment groups 39
4.2 Description of trust measures 40
4.3 Regress exam score on trust measures 47
4.4 Regress exam score on relative trust measures 49

Acknowledgments

There are many people that have helped me along the way, and I thank you all. Several people have played outstanding roles; I try to acknowledge them here. I'll start with those who were there from the very beginning: Mom, Dad, Barb, Dave, Mike and Nancy—I am so lucky to have you all in my life. Jeffrey Alwang, whose patience as a co-author has been incredible, provided wonderful guidance on my first real research project and continues to be a sounding board. Arriving at Berkeley I met a brillant ARE cohort—they helped me survive my first year of the Ph.D. program, and made it lots of fun. Jenny, Kyriakos, Ben, Leslie, Santiago and the rest of you made me a better researcher and person. Erick Gong and Kelly Jones were the perfect officemates. They were always ready to answer a "quick" metrics questions, offer support in times of distress, and brainstorm research ideas–you guys are amazing. Maria Bowman is wise, loving and a meticulous editor; she is also fantastic friend, and I am so glad she joined me in Berkeley after studying at Virginia Tech together. Among the graduate students, I have to thank my soccer mates: "One, Two .. Che!". And I thank Kate and Neil for introducing me to board games and plenty of thoughtful conversation.

I thank everyone in the Development Workshop for teaching me econometrics, and special thanks to Elisabeth Sadoulet and Ethan Ligon. I learned a great deal from them on how to pursue research. Thanks to all the ARE faculty—especially, Brian Wright, Peter Berck, Alain de Janvry, Gordon Rausser and Leo Simon who were great listeners and were always ready to offer support if I knocked on their doors. I thank my dissertation committee. Steven Raphael was a perfect IGERT mentor, and outside committee member. David Zilberman made me appreciate the importance of using a model, even a simple one, to motivate a research agenda and related empirical analyses. I thank Max Auffhammer for his encouragement and practical advice. He is a natural teacher, and helped me resolve details that mattered in my dissertation–he pushed me to the very end. I thank David Sunding for introducing me to California water issues. He has taught me to not be distracted by too many details and has helped me cultivate a "big picture" perspective. He has also been a tremendous advocate for me; I could not ask for a more supportive advisor.

I am sure important people have been left out, but I am grateful to them as well. Finally, someone who I could never forget: Kathleen you have done so much to help me complete this dissertation, you receive the biggest thanks of all.

Chapter 1

Introduction

Both the allocation of water and management of its quality are controversial issues confronting societies worldwide, and they will likely continue to be the cause of conflict for a long time to come. Farmers in India overdraft their groundwater reserves; industrial activity in China pollutes their surface and groundwater; drought, agriculture and environmental flows generate unrest in Australia's Murray-Darling Basin. In Africa water hungery fertilizer and new agricultural technologies such as drought resistant crop varieties or irrigation systems are not widely used or are never adopted. This dissertation adds to the literature on water resource economics with two chapters on valuing water in California, and a third chapter that considers how Ecuadorian farmers learn about a new agricultural technology which has benefits for surface water quality. All three of these chapters use microeconomic models to generate hypotheses and motivate subsequent empirical analysis. In the remainder of this introduction I provide a brief overview of each chapter.

Chapter 2 presents a hedonic analysis of farmland values to recover a value of irrigation water. Irrigation accounts for most freshwater diversions worldwide. Water available to farmers may be reduced in the long run by climate change and other factors. We use differences in farmland prices to value the availability of irrigation water. Using panel data on a set of farm parcels, we estimate the value of irrigation water using a plot-level fixed effects design – an approach to hedonic analysis that effectively controls for unobservable factors and may be useful in numerous settings. Our results suggest that the value of water for irrigation is much larger than measured in previous studies based on cross-sectional analysis. This finding has important implications for the analysis of policies that influence water supply reliability, and for the assessment of climate impacts.

Chapter 3 describes a method for evaluating consumer willingness-to-pay for water supply reliability when fixed costs are a large portion of the marginal price of residential water. We show that the willingness-to-pay for water supply reliability is determined by consumer demand for water, how the utility covers its costs, and the source of unreliability. The framework is applied to the case of 52 large urban water utilities in California covering the greater San Diego, Los Angeles, and San Francisco Bay Areas. We estimate residential water demand using a panel dataset tracking annual price and consumption data from 1994-2009 for 82 California water retailers. Based on our estimated price elasticities along with observed rates and the marginal costs of supply, we calculate willingness-to-pay (WTP) measures for reliability improvements for segments of both Northern and Southern California.

Chapter 4 uses a simple utility maximization model to examine how trust attitudes affect farmer

learning during an agricultural training. Using data originally collected in 2005 we examine if trust conditions a farmer's decision to learn during an agricultural training. We present a model of farmer behavior during the agricultural training in order to link trust measures to behavior in the field. We find evidence that farmers who trust agricultural technicians relatively more than community farmers learn more during training. The results provide insight into the design of agricultural extension services in Ecuador.

Chapter 2

The Economic Impact of Changes in the Availability of Irrigation Water

Agriculture accounts for most freshwater diversions worldwide (Foley, 2005). Irrigation is essential for crop production in arid parts of the developed and developing world. Changes in the availability of irrigation water can have large welfare consequences since water is an important input into agricultural production in these arid regions. Beyond natural fluctuations in precipitation, water supply to agriculture can be affected by numerous factors. Climate change has emerged as a significant threat to freshwater resources in many regions (McDonald, 2011). In California, Hayhoe (2004) find that under certain scenarios climate change will reduce snow pack in the Sierra Nevada mountains by 73-90%, thus, significantly reducing the main source of fresh water in the state. Beyond climate change, environmental regulations can also impact the availability of water for agriculture. For example, regulation under the Endangered Species Act has reduced the supply of irrigation water to farmers in California, Arizona, Texas and other locations (Moore, Mulville, and Weinberg, 1996). Reliable estimates of the economic impacts of expected reductions in irrigation water supplies are of crucial importance to the design of efficient water policy.

The welfare implications of changes in water availability are difficult to observe directly since water in agricultural settings is not allocated by a competitive price mechanism in the vast majority of agricultural areas. While there is usually not a competitive market for water that establishes a market-clearing price, there is a highly competitive market for farmland. This paper exploits differences in observed transaction-based farmland values among parcels with differential access to irrigation water deliveries to infer the value of water to producers. We base our analysis on a dataset of farmland transactions in California's San Joaquin Valley, one of the most important agricultural areas in the world. We contribute to the larger hedonics literature, by estimating a hedonic model using a plot-level fixed effects design to control for unobservable time-invariant factors.

With cross-sectional data it is difficult to adequately control for unobserved variables that may be correlated with the explanatory variable of interest and the outcome variable leading to biased coefficients. Panel data on the repeated sale of the same parcel of farmland over time would allow one to parse out any bias due to unobserved time-invariant characteristics. We make a first step in this direction by constructing a small panel dataset to estimate the value of water availability while controlling for unobservables at the parcel level.

Consistent with theory, the panel analysis suggests that the value of irrigation water in California agriculture is higher than estimates obtained from the pooled cross-sectional analysis. Indeed, the estimates from our panel analysis using a plot-level fixed effects design are several times the size of cross-sectional estimates. We provide evidence that this result is not an artifact of a selection process; we run the cross-sectional analysis on the same set of plots that appear in the panel analysis. We find results are similar to other studies and are consistent with the presence of important unobservable cross-sectional variation that biases cross-sectional estimates of the value of irrigation water supplies towards zero. Thus, farmers' willingness-to-pay for an additional acre-foot of water per acre appears to be much larger than prior estimates.

The hedonic approach is commonly used for evaluating environmental policies; examples include the Clean Air Act (Chong, Phipps, and Anselin, 2003; Chay and Greenstone, 2005) and the Superfund program for cleanup of hazardous waste sites (Greenstone and Gallagher, 2008). The approach is also used for valuing natural resources such as water quality (Leggett and Bockstael, 2000) and climate (Mendelsohn, Nordhaus, and Shaw, 1994; Schlenker, Hanemann, and Fisher, 2006). Other examples of the hedonic method applied to policy evaluation and resource valuation are numerous. Kuminoff, Parmeter, and Pope (2010) point out that in the presence of omitted variables the hedonic price function can produce severely biased estimates. Unobserved characteristics that affect the sale price of a good may also be correlated with the explanatory variable of interest; in this paper we make this point explicit. To do so, we evaluate the quality of hedonic valuations of irrigation water supplies using cross-sectional estimates compared to panel estimates using plot-level fixed effects.

Starting with Selby (1945) and most recently Petrie and Taylor (2007) empirical studies have analyzed how access to irrigation water is capitalized into farmland value. Other examples include Hartman and Anderson (1962), Crouter (1987), and Faux and Perry (1999). Most similar to this study is the work of Schlenker, Hanemann, and Fisher (2007)—SHF henceforth, who use cross-sectional data on self-reported plot-level land values from the June Agricultural Survey. The survey is geo-referenced which allows them to match farm-level data on the value of a farm with climate and soil databases. Their study takes place over a similar time period and location; therefore, their results can be compared to ours.

The remainder of this chapter is organized as follows. In the next section we present a simple economic model to motivate the empirical analysis. Section 2.2 describes the data sources. In section 2.3 we present the main estimating equations. The subsequent section shares the estimation results along with a discussion. The final section concludes and indicates areas for future work.

2.1 Economic Model

This section presents the economic logic of the hedonic price analysis for the valuation of irrigation water supplies. Following Crouter (1987) and others, we introduce a model of farmland prices applying Rosen's theory of hedonic price functions (1974). Rosen's theory allows one to infer the market value of a good by examining the prices of a composite good which include the good of interest. Applying the theory, we infer that there is an implicit market for irrigation water deliveries that works through the explicit farmland market. The main economic concept exploited in the hedonic price analysis is that the price of farmland equals the net-present value (NPV) of economic rents expected from the farmland while the price of the differentiated attribute is the shadow value of the attribute in terms of net present value.

A potential buyer of a plot of farmland observes that the land comes with a quantity of expected water availability and that this land and water may be combined with variable inputs to produce output in year t. We define output in year t in terms of a production function, $f(L,W,v_t)$ where L is land, W is water available to the plot and v_t is some optimal quantity of a variable input in time period t. The buyer assumes the price of the output is p_t and the cost of the variable input to be c_t in period t; the cost of the land with its associated water supply is $q(L,W)$, which is the hedonic price function. Based on these factors, the potential buyer considers the economic rents that may be derived in a year by choosing the optimal level of v_t given available land and water, and facing prices p_t, c_t, and $q(L,W)$. Said differently, she solves the following profit maximization problem:

$$\max_{(v_t)} \Pi = \sum_{t=0}^{\infty} \delta^t (p_t f(L,W,v_t) - c_t v_t) - q(L,W) \tag{2.1}$$

The resulting Π^* equals the expected economic rent to be obtained from the land. Next we consider how to assess the shadow value of a permanent change in yearly water availability. First, we note that a change in W will affect Π_t^*; this change in $\Pi_t^*$is the shadow value of additional water in a year of production, λ. The shadow value of a permanent change in yearly water availability is the sum of discounted shadow values of additional water for each year. The NPV of a permanent change in annual water supply can be written as: $\lambda_W = \sum_{t=0}^{\infty} \delta^t \cdot \lambda$. In this simple framework λ_W is the value of an additional unit of irrigation water deliveries in perpetuity.

To complete this analysis we tie λ_W to changes in land prices associated with a unit change in irrigation water deliveries. Because land markets are competitive, all economic rents will be bid away—that is, the price of land will reflect the NPV of production on the land. Therefore, any differences in land prices reflect differences in the NPV of the land. In summary, the partial derivative of the hedonic price function with respect to W is equivalent to the shadow value of a permanent increase in water availability:

$$\frac{\partial q(L,W)}{\partial W} = \lambda_W = \sum_{t=0}^{\infty} \delta^t \cdot \lambda \tag{2.2}$$

Importantly, the aim of the empirical analysis is to recover a consistent estimate of λ_W .

2.2 Data

We gather data on sale prices of all land transactions in eight California counties between 2001 and 2008, water deliveries and land acreage for irrigation districts, groundwater depth measurements, historical temperature and precipitation, soil quality measures, land classification codes, and measures of population density. All data resources are combined into a single dataset for analysis. This final dataset is comprised of 8 counties, 19 hydrological unit areas, and 28 irrigation districts.

2.2.1 Outcome Variable: Farmland Sale Price Per Acre

The farmland price data were purchased from DataQuick, a private firm which collects data on land sales, mostly from county court houses. Each observation is geo-referenced with latitude and longitude, and both the street address and lot size are observed, which allows for clear identification of repeat observations. In addition to land prices and the aforementioned variables, the transaction data include information on transaction characteristics such as whether the sale was a foreclosed property. Characteristics related to structures on the property are in the dataset including whether there is building on the property, total square footage of the building, the number of bedrooms, and the number of bathrooms, as well as an estimate of the percentage improvements made to the property. Other property characteristics include lot size and a approximate measure of the primary use of the land according to a county administrator.

2.2.2 Explanatory Variable of Interest: Surface Water Delivery Right Per Acre

In 1933 the State of California passed the Central Valley Project Act. This act authorized the government to begin fundraising for the construction of water infrastructure such as reservoirs, dams and canals to support irrigated agriculture. Due to the Great Depression, the federal government made several financial transfers to the State of California to complete the project. Today, the United States Bureau of Reclamation is responsible for the administration of the Central Valley Project (CVP) that delivers surface water to irrigation districts in California's Central Valley. Farmland within an irrigation district has a contractual right to buy a fixed amount of water in a given year from the Bureau of Reclamation. However, due to changing hydrologic conditions and other factors, the federal government does not guarantee that any specific amount of surface water deliveries will be made in any given year. Therefore, farmers may use the current year's surface water deliveries as a best estimate of water rights entitled to a plot of land.

Surface water deliveries from the Central Valley Project account for a significant share of agricultural water use in California. Data on federal surface water deliveries come from the Bureau of Reclamation which records annual deliveries at the irrigation district level. We make the strong assumption that the irrigation district level delivery is evenly divided among all land in the district. Based on conversations with irrigation district managers this is a valid approach given that this is how water rights are allocated. Hence, to compute a per-acre measure of surface water delivery rights we divide the total quantity of surface water deliveries by the total amount of land in acres within an irrigation district.

2.2.3 Control Variables

Other sources of surface water deliveries are the State Water Project and private projects which are usually administered by local governments; both are observed in our dataset. The last major source of water for agriculture is groundwater, which varies significantly across the State and within our sample. We gather data on groundwater availability from the California Department of Water Resources. They record approximated measures of groundwater levels in over 15,000 wells for several decades. Many wells are not measured regularly or during the sample period, and the wells used for groundwater measurements are not evenly distributed across agricultural land. Therefore, we use regional averages of groundwater availability to create plot-level measures of groundwater availability; these are likely measured with error. Data on groundwater quality and other hydrological characteristics such as groundwater flow direction are not easily collected and so are unobserved in this analysis. However, hydrologists have surveyed California and defined areas with similar hydrological characteristics, we reference these as 'hydrological units'. These are contiguous areas smaller than a county, yet crossing county lines; the land transactions in the sample are spread across 19 hydrological units.

We collect data on soil quality from the United States Department of Agriculture's Natural Resources Conservation Service which maintains both STATSGO and SSURGO2 soil databases. Neither of these soil databases are ideal for plot-level analysis because, like the groundwater measures, they are a weighted average of soil type over large swaths of land; this implies significant variation within each soil survey unit and non-trivial measurement error. This poses a problem if the unmeasured component of soil quality is correlated with both deliveries and land price. For example, there may be a particularly high quality piece of land within a low average soil quality area. If the high quality land requires less irrigation water to effectively water plants, then this source of measurement error would cause a downward bias in our point estimate on deliveries.

In addition to irrigation water availability and soil quality, climate is likely to be another important determinant of farmland value. We use the same high-resolution temperature and precipitation climate data that has been employed by others (Schlenker, Hanemann, and Fisher, 2006). These climate data were organized by the Spatial Climate Analysis Service at Oregon State University for the National Oceanic and Atmospheric Administration. Plot-level measures of climate are interpolated using the PRISM model also developed by researchers at Oregon State. We use thirty-year historical annual rainfall and both maximum and minimum temperatures to control for climate in our analysis. There is evidence that alternative climate measures such as degree days and rainfall timing within a year as well as diurnal temperatures are more informative climate measures than historical means. However, we only use climate measures in the cross-sectional analysis—the impact of climate on farmland values will ultimately be parsed out in the panel data analysis because climate does not vary in our sample time period.

Another important factor to control for in any analysis of farmland values is urban development potential. Our approach to this problem is to create a five mile buffer zone around population centers. Visual inspection of these buffer zones using ArcGIS confirms these are more densely populated areas.

2.2.4 Sample Selection

We proceed with the criteria employed in sample selection. For inclusion in the final sample we only use observations which satisfy several criteria. First, we drop all plots that also receive water from the State Water Project (SWP). Deliveries from the SWP may be negatively correlated with CVP deliveries and positively correlated with price. We do not observe the actual quantity delivered from the SWP; in order to avoid omitted variable bias we exclude these observations. Second, variation in urban development potential may still be significant even after controlling for plots near population centers. Land prices over $20,000 per acre in year 1998 prices likely represent plots of farmland with urban development potential and so were removed from the sample. For similar reasons we dropped land classified as rural residential by the California Department of Conservation. We also drop plots sold more than three times in our sample period or sold twice within the same year. The concern is that property with higher turnover may also signal urban development potential, or other characteristics unique to properties which are sold more than three times in a decade. Third, plots with houses and other significant infrastructure will greatly affect the price per acre. Similar to urban development potential, plot attributes such as housing may dwarf variation in prices due to surface water deliveries. To address this issue we drop all observations with bedrooms, bathrooms or with buildings larger than 1500 square feet (this allows for service sheds). We also had to drop sales whose recorded price included non-farmland parcels because information on the non-farmland portion of the property were not reported in our data. Finally, we drop plots of land that have relatively little use for irrigation water; farmland plots used for timber, poultry or dairy production are removed. The admittedly small panel sample has 304 observations representing 146 parcels of farmland. In order to evaluate the external validity of our panel estimates we compare these plots of farmland for which we have repeated observations to all plots satisfying the described selection criteria; we reference the latter as the non-panel sample. Table 1 reports the means and sample standard deviations for observable covariates by the non-panel and panel samples. The samples are comparable although statistically significant differences exist between them. The plots of land in the panel sample receive more rain, have higher average maximum daily temperatures, and are at a higher elevation than plots of land in the non-panel sample.

Table 2.1: Comparison of plot characteristic means across samples

			All	Panel Only	
	Minimum	Maximum	Mean	Mean	P-value
Price per acre (dollars)	177	19921	7612	6195	0.000
			(4540)	(3848)	
Deliveries in year of sale (acre-feet)	0	4.205	0.453	0.449	0.961
			(.819)	(.785)	
Lot size (in thousands of acres)	0.003	0.640	0.05	0.04	0.852
			(.06)	(.07)	
Building structure/storage shed (d)	0	1	0.10	0.08	0.422
			(.29)	(.26)	
Square footage of building structure	0	14210	214.74	214.14	0.994
			(916.3)	(1252.25)	
Private water deliveries access (d)	0	1	0.48	0.55	0.125
			(.5)	(.5)	
No groundwater (d)	0	1	0.42	0.42	0.998
			(.49)	(.49)	
Mean depth to groundwater (ft)	0	446	37.47	34.90	0.538
			(48.35)	(45.38)	
Elevation (meters)	6.18	215.27	85.78	91.23	0.093
			(37.53)	(36.24)	
Historical mean rainfall	0.016	0.048	0.03	0.03	0.051
			(.01)	(.01)	
Historical mean max temp (°C)	2.204	2.605	2.48	2.49	0.009
			(.07)	(.06)	
Historical mean min temp (°C)	0.823	1.195	0.92	0.93	0.021
			(.04)	(.03)	
Storie Index for soil quality	0.000	0.098	0.06	0.06	0.613
			(.03)	(.03)	
Orchards (d)	0	1	0.35	0.36	0.845
			(.48)	(.48)	
Vinyards (d)	0	1	0.19	0.17	0.519
			(.39)	(.38)	
Distance to freeway	0.000	0.035	0.01	0.01	0.538
			(.01)	(.01)	
Rural residential buffer (d)	0	1	0.33	0.34	0.782
			(.47)	(.48)	
Partial property sale (d)	0	1	0.03	0.01	0.343
			(.16)	(.12)	
Total Observations			1537	146	

Note: Sample standard deviations presented in parantheses.

2.3 Empirical Research Design

The main empirical concern is that an unobservable plot characteristic is correlated with both water deliveries and land prices. We describe this empirical challenge in more detail and justify the need for panel data to address the potential problem of omitted variable bias. Then we indicate how we will proceed with the empirical analysis and present the main estimating equations.

2.3.1 Empirical Challenge

To begin, we assume land and water are the only two factors of production and that land quality is homogeneous. Therefore, the hedonic price equation may be estimated using:

$$\text{price/acre}_{it} = \beta_0 + \beta_1 \text{deliveries}_{it} + \varepsilon_{it} \tag{2.3}$$

In this equation, i denotes the plot of land, the outcome variable is the price per acre of the land sale, and 'deliveries' captures the expected stream of annual surface water to be delivered in perpetuity to plot i. The parameter β_1 can be interpreted as the shadow value of a permanent shock to water deliveries supply (λ_W); β_1 is expected to be positive. If one observes land price and volumetric water deliveries, then one can estimate β_1 consistently if the following identifying assumption holds:

$$\mathbb{E}[\varepsilon_{it}|\text{deliveries}_{it}] = 0 \tag{2.4}$$

Assuming equation (2.4) holds then β_1 can be estimated using Ordinary Least Squares to recover the capitalized value of surface water deliveries. However, the identifying condition can be violated in many ways. The most basic way is that there are unobservable environmental quality characteristics of the plot that are correlated with deliveries so that:

$$\varepsilon_{it} = \gamma \cdot EQ_i + \eta_{it} \tag{2.5}$$

where EQ_i is the underlying environmental quality of plot i. If EQ_i is observable then one can consistently estimate the coefficient of interest if the following condition holds:

$$\mathbb{E}[\eta_{it}|\text{deliveries}_{it}, EQ_i] = 0 \tag{2.6}$$

This is the assumption maintained in pooled cross-sectional analyses such as in (SHF, 2007). This identifying assumption is violated if there are plot specific unobservable characteristics, which are correlated with either the other environmental quality measures or deliveries. That is, the error term from equation (2.3) may take the form:

$$\varepsilon_{it} = \theta_i + \gamma \cdot EQ_i + \nu_{it} \tag{2.7}$$

where,

$$\mathbb{E}[\theta_i|EQ_i, \text{deliveries}_{it}] \neq 0. \tag{2.8}$$

Relaxing the assumption of homogeneous land quality and assuming the expressions in equations (2.7) and (2.8) hold, then the only way to estimate β_1 consistently is using a fixed effects estimator, which requires repeat sales of the same plot.

2.3.2 Estimating Equations

The empirical analysis can be divided into two parts. In the first part, following previous literature we estimate the hedonic price function using a cross-sectional specification without plot-level fixed effects. We also consider several cross-sectional specifications with various controls, and estimate a model using a random effects specification clustering at the hydrological unit level as this is, in general, a more efficient estimator than Ordinary Least Squares (OLS) when using clustered data. Second, we estimate a model using plot-level fixed effects, and then run a Hausman test to evaluate the consistency of the random effects estimator clustered at the plot-level. For all specifications we measure expectations about the stream of future water deliveries using deliveries in the year of sale[1].

In the sub-sections that follow, we present the estimating equations and discuss advantages and disadvantages of each specification.

Pooled Cross-Sectional Analysis

We first estimate the linear relationship between price per acre and acre-feet of water deliveries per acre using OLS including controls for lot size, whether there is building on the property, the building size, alternative surface water supply availability, groundwater depth, historical rainfall and temperature, soil quality, crop production (orchards, vineyards and row crops/pasture land as the excluded group) population density and transaction characteristics such as whether the sale was a partial property sale. For a full list of the control variables see table 2.1.

In another specification we estimate a random effects model with spatial clustering at the hydrological unit level. As mentioned earlier, a hydrological unit is a region of land (smaller than a county and larger than an irrigation district) with similar underlying hydrological characteristics. Because hydrological characteristics are correlated with soil, land within the same hydrological unit will likely have more similar soil than land across hydrological units. It is also true that plots within the same hydrological units are likely to experience similar climates. If this is the case, then the random effects model will produce inconsistent estimates. In a third specification, we include hydrological unit fixed effects to account for groundwater availability, water quality and flow direction common to all land within in the hydrological unit. In fact, the hydrological unit fixed effects capture all characteristics common to all land within the hydrological unit. For these reasons, hydrological unit fixed effects may be a better way to control for these important production factors when using cross-sectional data. Perhaps the most compelling argument for this spatial fixed effects estimator is that hydrological units cross over irrigation districts lines so that one can compare a plot in an irrigation district (i.e., a plot that receives federal water deliveries) to a plot that lies just outside the irrigation district but within the same hydrological unit. In this sense, the estimation is similar to a border discontinuity design; however, we cannot explicitly employ a regression discontinuity approach since we do not observe data delineating borders.

There is still the potential that we could be omitting variables that vary across time at the hydrological unit level such as changing urban development potential that, in turn, affect farmland prices. To partially control for this possibility, we estimate a fourth specification in which we

[1] We argue that current deliveries is the correct measures of future deliveries on the basis that water policy in California is in a constant state of evolution, and that the recent past has been so unusual that it is a poor indicator of water supply allocations.

interact the hydrological unit fixed effects with a linear time trend. In a fifth specification we add the set of hydrological unit fixed effects interacted with a quadratic time trend. This final estimating equation is given by:

$$\text{price/acre}_{ijt} = \beta_1 \text{deliveries}_{it} + \beta_2 X_i + \beta_3 Z_{it} + \mu_j + \tau_t + \mu_j \cdot t + \mu_j \cdot t^2 + \varepsilon_{ijt} \tag{2.9}$$

where i indicates the plot; j indicates the hydrological unit; t indicates the year; deliveries$_{it}$ represents the water deliveries; X_i are time-invariant observables like historical mean temperature, rainfall, and soil quality; Z_{it} are time-varying transaction characteristics; μ_j is the hydrological unit fixed effect; τ_t is the year fixed effect; $\mu_t \cdot t + \mu_j \cdot t^2$ is the hydrological unit fixed effect interacted with time trends. It is worth noting that the interpretation of β_1 is that of λ_W from equation (2.1), the capitalized value of an acre-foot of water in perpetuity.

The problem with OLS and the spatial fixed estimators such as the model described in equation (2.9) is that they all ultimately rely on cross-sectional variation to identify the effect of irrigation water availability on land price. However, none of these estimators do an adequate job of addressing the main empirical challenge: accounting for unobservable cross-sectional variation. Controlling for soil quality using SSURGO2 measures or through the use of spatial fixed effects is not sufficient as soil quality measures are based on weighted averages over large areas and so, like spatial fixed effects, assume homogeneous soil quality for large regions of land. Furthermore, hydrological units are large areas that encompass significant variation in environmental quality. Thus, the use of hydrological unit spatial fixed effects does not satisfy the homogeneity assumption. As already described, to overcome this one could collect plot-level measures of soil quality, although if there are remaining time-invariant omitted variables then the estimates may still be biased. For this reason, an estimator employing panel data is attractive.

Panel Analysis

Environmental quality is slow to change over time so one may consider it time-invariant over the sample period; a plot-level fixed effects estimator will parse out all variation due to time invariant plot-level characteristics. The estimating equation for the plot-level fixed effects analysis is given by:

$$\text{price/acre}_{ijt} = \delta_1 \text{deliveries}_{it} + \delta_2 Z_{it} + \theta_i + \tau_t + \mu_j \cdot t + \mu \cdot t^2 + \varepsilon_{ijt} \tag{2.10}$$

where i, j and t are as before; θ_i is the plot-level fixed effect and Z_{it} are time-varying transaction characteristics; δ_1 is interpreted as λ_W from equation (2.1), the capitalized value of an acre-foot of water in perpetuity. This specification will control for variation in the outcome variable due to time-invariant characteristics including underlying environmental quality of the land. Similar to the cross-sectional specification, we estimate a hydrological unit specific quadratic time trend to account for time variant omitted variables which accounts for factors exhibiting a quadratic time trend common within a hydrological unit.

Standard Error and Pivotal Statistic Adjustments

Statistical inference is complicated due to the clustered structure of the data. To address clustering and within-cluster heteroskedasticity we compute robust standard errors clustered at the county level (eight clusters). Bertrand, Duflo, and Mullainathan (2004) show that when there are a small number of clusters (ten or less), the performance of statistical inference using cluster-robust standard errors is unreliable. Cameron, Gelbach, and Miller (2008) find the same result and consider a variety of standard error adjustments for clustered data and then evaluate which adjustments offer reliable performance for statistical inference. They suggest using the wild cluster-bootstrap method to obtain pivotal t-statistics as this offers asymptotic refinement. We follow their advice, and compute a wild cluster-bootstrap pivotal t-statistic for each regression coefficient; then the analytically computed cluster-robust standard error is used with the adjusted pivotal t-statistic to perform hypothesis testing.

2.4 Results and Discussion

We begin with the results presented in table 2.2, which presents the results of the cross-sectional analysis using the pooled non-panel sample. As we move across columns we see point estimates from different regression specifications. In column (1) we present the simple OLS estimate of $219 for the capitalized value of an acre-foot of water. In the second specification we present the random effects model clustered at the hydrological unit level. Surprisingly, the random effects point estimate is identical to the OLS estimate. This suggests that there is no dependence between error terms within a hydrological unit (the level at which the random effects model is clustered). Notably, this random effects specification is most similar to the specification employed by SHF (2007). The point estimate of $219 is less than the point estimate of $656 obtained in their study; however, differences between the samples may account for the observed difference in point estimates. Furthermore, the standard errors on point estimates from SHF and our cross-sectional analysis are not small enough to reject the null hypothesis that they are identical.

In column (3) of table 2.2 we present the results from the hydrological unit fixed effects estimation. The point estimate is slightly larger, although we cannot reject a Hausman test that the point estimates from the random effects and fixed effects models are different. This evidence against clustering is consistent with the findings of SHF's results in which they favor a random effects model over a fixed effects model, and with the result that the OLS and random effects point estimates are identical. In subsequent columns we test if hydrological unit specific time trends affect the point estimates and we find no significant differences.

Table 2.2: Cross-sectional sample: Regress price/acre on federal water deliveries

Dependent variable: Price per acre-foot/acre of water (Mean: 9,521, S.D.: 5,802)					
	(1)	(2)	(3)	(4)	(5)
Capitalized value of one acre-foot/acre	219	219	368	411	415
federal surface water of (λ_W)	(170)	(170)	(185)	(175)	(201)
# of observations	1,702	1,702	1,702	1,702	1,702
R2	0.302	0.302	0.3365	0.359	0.377
Year fixed effects	Yes	Yes	Yes	Yes	Yes
Time variant controls	Yes	Yes	Yes	Yes	Yes
Hydrological unit fixed effects	No	No	Yes	Yes	Yes
Hydro. unit fixed effects*Linear Trend	No	No	No	Yes	Yes
Hydro. unit fixed effects*Quadratic Trend	No	No	No	No	Yes

Robust standard errors clustered at the county level in parentheses.
*** p<0.01, ** p<0.05, * p<0.1.

In table 2.3 we see the results of the cross-sectional analysis using the panel sample. The regression specifications in each column are the same as those in the corresponding columns of table 2.2. The point estimates in the first two columns are identical, this is consistent with what we observed in the previous table—there is little spatial dependence between observations within a hydrological unit. The subsequent columns represent specifications including the hydrological fixed effects. These models produce larger point estimates for the value of irrigation water, all estimates are greater than \$1,000. However, we cannot reject the Hausman test's null hypothesis that the random effects estimate is consistent. As before, this result corroborates previous evidence that there is little spatial dependence across observations within a hydrological unit. Based on this cross-sectional analysis, we have no evidence that the value of irrigation water is different than the \$656 estimate obtained by SHF (2007).

Table 2.3: Panel sample: Regress price/acre on federal water deliveries w/o plot-level fixed effects

Dependent variable: Price per acre-foot/acre of water (Mean: 9,339, S.D.: 5,779)					
	(1)	(2)	(3)	(4)	(5)
Capitalized value of one acre-foot/acre	226	226	1133	1137	1210
of federal surface water (λ_W)	(129)	(129)	(423)	(573)	(585)
# of observations	304	304	304	304	304
R2	0.382	0.382	0.44	0.485	0.545
Year fixed effects	Yes	Yes	Yes	Yes	Yes
Time variant controls	Yes	Yes	Yes	Yes	Yes
Hydrological unit fixed effects	No	No	Yes	Yes	Yes
Hydro. unit fixed effects*Linear Trend	No	No	No	Yes	Yes
Hydro. unit fixed effects*Quadratic Trend	No	No	No	No	Yes

Robust standard errors clustered at the county level in parentheses.
*** p<0.01, ** p<0.05, * p<0.1.

We now turn to the results of the panel analysis using plot-level fixed effects. Table 2.4 presents the results of these regressions. The first column presents a random effects specification with the cluster at the plot-level. The point estimate is -$369, and has a large standard error, especially given the large bootstrapped critical t statistic. Subsequent columns include plot-level fixed effects. The second column reports a point estimate of $2,655 and is significant at the 99% level; we strongly reject the Hausman test. We obtain a similar result in column (3) when we control for the percentage of capital improvements on the land. In the subsequent columns we add a hydrological unit linear and quadratic time trends, respectively. Adding these region specific time trends substantially increases the point estimate on deliveries to over $4,000. All of these estimates are significantly different than zero at the 99% level when using the robust standard errors clustered at the county level and our wild cluster bootstrapped critical t statistics for hypothesis testing. Based on SHF (2007) we also consider the null hypothesis that the capitalized value of an acre-foot of water is $656; we are able to reject the null at the one percent level for columns (2) and (3), and at the five percent level for columns (4) and (5). Based on our own data, the largest cross-sectional estimate for the capitalized value of an acre-foot of water is $1,210. If we test our point estimates from the plot-level fixed effects models against this as the null hypothesis we still reject the null at the one percent level for the specifications in columns (2) and (3). This indicates that if one believes the plot-level fixed effects specification consistently estimates the shadow value of expected future deliveries then the estimates from the cross-sectional analysis exhibit a large downward bias.

Table 2.4: Panel sample: Regress price/acre on federal water deliveries w/ plot-level fixed effects

Dependent variable: Price per acre-foot/acre of water (Mean: 9,339, S.D.: 5,779)					
	(1)	(2)	(3)	(4)	(5)
Capitalized value of one acre-foot/acre	-336	2655***	2603***	4611***	4286***
of federal surface water (λ_W)	(369)	(380)	(379)	(1680)	(2029)
	[4.59]	[1.82]	[1.74]	[1.81]	[1.69]
# of Observations	304	304	304	304	304
R2		0.759	0.759	0.793	0.838
Plot-level fixed effects (146 plots)	No	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes
Controls	No	No	Yes	Yes	Yes
Hydro. unit fixed effects*Linear Trend	No	No	No	Yes	Yes
Hydro. unit fixed effects*Quadratic Trend	No	No	No	No	Yes

Robust standard errors clustered at the county level in parentheses.
Wild cluster bootstrapped critical t-statistic for 99th percentile.
*** p<0.01, ** p<0.05, * p<0.1.

In terms of magnitude, our estimates of $2,603 to $4,611 per capitalized acre-foot of water are two and a half to four times the size of the cross-sectional estimates using the same data. The estimates are roughly the same degree of magnitude larger than those of SHF who put the capitalized value of irrigation water at \$656 per acre-foot. Second, the average price of an acre of land in our sample is \$9,339 with the average plot receiving 0.444 acre-feet per acre. Our estimate of \$4,286 suggests that, on average, irrigation water rights account for approximately 20% of the sale price. Taken together our results suggest 1) that farmers have a much higher willingness-to-pay for water than previously thought, and 2) that this result plays a non-trivial role in the determination of irrigated farmland prices.

One may wonder how climate change will affect landowners with rights to surface water. To address this, we consider the work of Chung (2005) and Vicuna (2006) who use the CALSIM 2 simulation model to make predictions of the impact of climate change on CVP surface water deliveries south of the Delta between 2035 and 2064, and between 2070 and 2099. Their analyses are based on the A2 emission scenario modeled using the Geophysical Fluid Dynamics Laboratory global climate model, details can be found in Hanemann et al. (2006) and Cayan et al. (2008). Under the GFDL A2 emission scenario, Chung (2005) find that median expected CVP surface water deliveries are 2,435 thousand acre-feet (TAF) between 2035 and 2064, 14.5% less than base deliveries during the same time period. Using the same GFDL A2 emission scenario, Vicuna (2006) finds median expected deliveries are 1,944 TAF, 31.4% less than base deliveries between 2070 and 2099. As an exercise, we assume these impacts are not currently capitalized into farmland values and generate an estimate of the net present value (NPV) of climate change's aggregate impact via water reductions on landowners with rights to CVP water. We calculate:

$$\triangle \text{Welfare} = \sum_{t=24}^{58} \delta^t \cdot \lambda \cdot (\text{Base TAF}_1 - \text{Reduced TAF}_1) + \sum_{t=59}^{88} \delta^t \cdot \lambda \cdot (\text{Base TAF}_2 - \text{Reduced TAF}_2) \tag{2.11}$$

In this welfare loss calculation δ is the discount factor and is based on a five percent discount rate; λ is the implied value of an acre-foot of water. The change in welfare is the sum of losses from the period 2035-2069 and the period 2070-2099; t represents the number of years from 2011 that the loss will be realized. Based on a capitalized value for one acre-foot of \$4,286, then assuming a 5% discount rate, the implied value of an acre-foot of water, λ, is \$204. The base allocation in the period 2035-2069 is assumed to be 2,716 TAF while the reduced allocation is 2,435 TAF. The base allocation in the period 2069-2099 is assumed to be 2,833 TAF while the reduced allocation is 1,944 TAF. A substitution of the relevant numbers into equation (2.11) indicates the climate change aggregate welfare loss to California farmers due to reduced water deliveries is \$460 million dollars. Notably, this figure is based on the assumption that climate change has no impact on deliveries until 2035. Nor does the calculation consider water reductions beyond 2099 although these may be considerable. One final caveat is that these predictions are the outcome of one simulated climate model and a particular discount rate, other scenarios will lead to different predictions.

2.5 Conclusions and Future Work

Using a small sample of repeated sales of agricultural land in California, we find evidence that existing cross-sectional estimates of the value of water deliveries to agriculture are significantly downward biased. Using a plot-level fixed effects estimation we account for unobserved factors correlated with deliveries and farmland prices, and obtain novel estimates of the value of water deliveries in California agriculture. The estimated capitalized value of one acre-foot of water is two and half to four times larger than the estimate obtained in the cross-sectional analysis. Based on this result, changes in water availability are likely to induce larger changes in producer welfare than previously thought. Finally, these findings inform policy analysis on issues related to water infrastructure projects, the protection of habitat for endangered species and climate change.

There are several ways to extend the current work. One could explicitly incorporate groundwater availability and crop choice into an agro-economic model to make predictions about how the shadow value of surface water varies with these other sources of heterogeneity. It would also be worthwhile to examine how buyers form expectations about future water deliveries—insight would give us more confidence about the appropriateness of using deliveries in the year of sale a measure of expectations of future water deliveries. Also, as Kuminoff, Parmeter, and Pope (2010) recognize, the hedonic price equilibria may vary across time. If this is the case, then the implied shadow value of water deliveries may not remain constant. One way to examine time-variant implicit shadow values is to interact the explanatory variables with the year fixed effects; however, with a small dataset this approach does not permit precise estimation. Future work with more observations would allow for such tests of time-varying hedonic equilibria.

Chapter 3

The Value of Supply Reliability in Urban Water Systems

Water systems in many regions of the planet are vulnerable to supply fluctuations. Natural variation in precipitation and runoff leads to contemporaneous changes in water supply that pose challenges to urban water purveyors who seek to meet target levels of demand. In recent years, environmental protections for endangered fish and aquatic habitats have increasingly impinged on the ability of water managers to reliably divert water towards urban uses, for instance the pumping restrictions placed on diversions to the California Aqueduct to protect the delta smelt under the "Wanger decision" in December, 2010. Yet, despite these challenges to urban water supplies, urban water demand in many parts of the United States continues to grow, and is projected to do so for decades to come.

Taken together, the natural variation in water supply coupled with continued and often unabated growth in water demand highlight the importance of understanding the economic losses that result from periodic water shortages. This paper considers the economic dimensions of urban water shortages by considering water supply disruptions of various levels. We consider water supply disruptions mediated through the residential segment of the market, which accords with the typical policy response of many water planners to water shortages in the range of 10-30 percent below baseline levels of use. Our methodology extends readily to other potential urban water management responses, for instance the rationing of commercial and industrial users.

We frame our analysis around regional residential water demand estimates in California. Our demand estimates, which take into account time-invariant characteristics, are refined to the local retailer level and increment on the methods employed by Renwick and Green (2000). The analysis we produce is based on administrative data reporting market prices and consumption levels. Our loss assessment has the advantage of being based on actual valuations of water units by residential consumers as opposed to stated preferences or the results of hypothetical optimization scenarios. In addition, the standard approach to measure welfare losses during shortage, the value of supply reliability, only considers the loss of the consumer surplus triangle. However, demand-side information does not provide sufficient information when prices correspond to average cost rather than marginal costs and fixed costs are large.

Our approach to measuring the economic value of water supply reliability innovates on previous methods by addressing the financial structure of water agencies. In many urban areas, water supply infrastructure involves a substantial fixed cost component. This distinction is essential for measuring welfare losses from water supply disruptions, because most water agencies are regulated as natural monopolists under a profit (or net revenue) constraint that limits their ability to set water prices. In many cases, a substantial share –often the largest share– of residential water rates corresponds with fixed cost recovery rather than marginal cost pricing. Fixed costs, which are sunk at the time of a water supply disruption, have no bearing on welfare outcomes. The welfare loss from a supply disruption can be many times greater than the loss evaluated using standard measures of compensating variation or equivalent variation in consumer utility functions.

The goal of this research is to measure the losses that will actually occur in urban parts of California under various levels of shortage, given existing institutions and rules for allocating water. The California water system is highly fragmented with water provided to customers by over 400 major urban utilities. Partly as a result of this fragmentation, there is no single number for the economic loss resulting from a water shortage that can be applied to all regions of California. Indeed, we find substantial variation in the economic losses across regions depending on prevailing residential water prices, the financial structure of the urban utility, and the price elasticity of demand for residential water. We calculate regional losses at the level of regional water purveyors using regional price elasticities from our demand estimation model, and then aggregate regional losses into measures of total economic loss and average loss per unit of supply disruption ($/AF) for both Northern and Southern California.

The remainder of the paper is structured as follows. The next section lays out the conceptual model underlying the loss analysis. We follow with a discussion of the empirical model for urban residential water demand, a description of the data, and a presentation of our estimation results. The penultimate section uses point estimates from our empirical model to calculate economic losses from various levels of water supply disruptions for Northern and Southern California, and the final section concludes.

3.1 Loss Framework

Our loss framework considers the economic impact of an urban water supply disruption. Water shortages following a supply disruption have the potential to adversely affect economic outcomes among several types of water users, including rural and agricultural users; however, we confine our analysis to the case of urban water management. Specifically, we consider a drought response framework for urban water suppliers in which water supply reductions of a given magnitude are mediated through the residential segment of the market. A drought response framework that primarily targets residential water use, for instance by imposing restrictions on outdoor irrigation, corresponds to the typical urban water management response to supply disruptions of relatively small magnitude and short duration.

We measure residential losses for the markets served by regional water purveyors in California as the difference between the area under residential demand and the avoided cost of reduced water service; these are the marginal costs of service delivery. Given that a substantial portion of the cost of water service represents fixed cost, the vast majority of which is sunk, our methodology represents a significant innovation over the more commonly used calculation of the consumer surplus triangle as a measure of the economic loss of water supply disruptions. Our measure of the economic loss reconciles with the consumer surplus measure when the avoided cost of reduced water service is described by a constant unit cost function at the prevailing residential water rate, but differs, often substantially, from the consumer surplus measure under circumstances in which residential water rates exist primarily for fixed cost recovery. Hence, our approach underscores the novel and important insight that the method of financing by a water supplier, for instance the extent to which fixed costs are recovered through fixed charges per household meter rather than through market prices for water, is essential for calculating economic losses.

Figure 1 provides a graphical depiction of the economic loss in the case of constant unit cost of water and volumetric pricing. The essential framework we employ for measuring economic losses can be seen by considering a water supply disruption of $dQ = Q^R - Q^*$ units. The water supply reduction rations price along the demand curve in the figure by $dP = P^R - P^*$ units. Under a zero profit constraint on water district pricing, which approximates the situation facing water purveyors in California, the equilibrium price of water prior to the water shortage is $P^* = AC$, which satisfies $P^* > MC$ when fixed costs exist, so that the economic loss, shown as the shaded region in the figure, is generally larger than the loss of consumer surplus.

Figure 3.1: Economic Losses Under Volumetric Prices and Constant Unit Cost

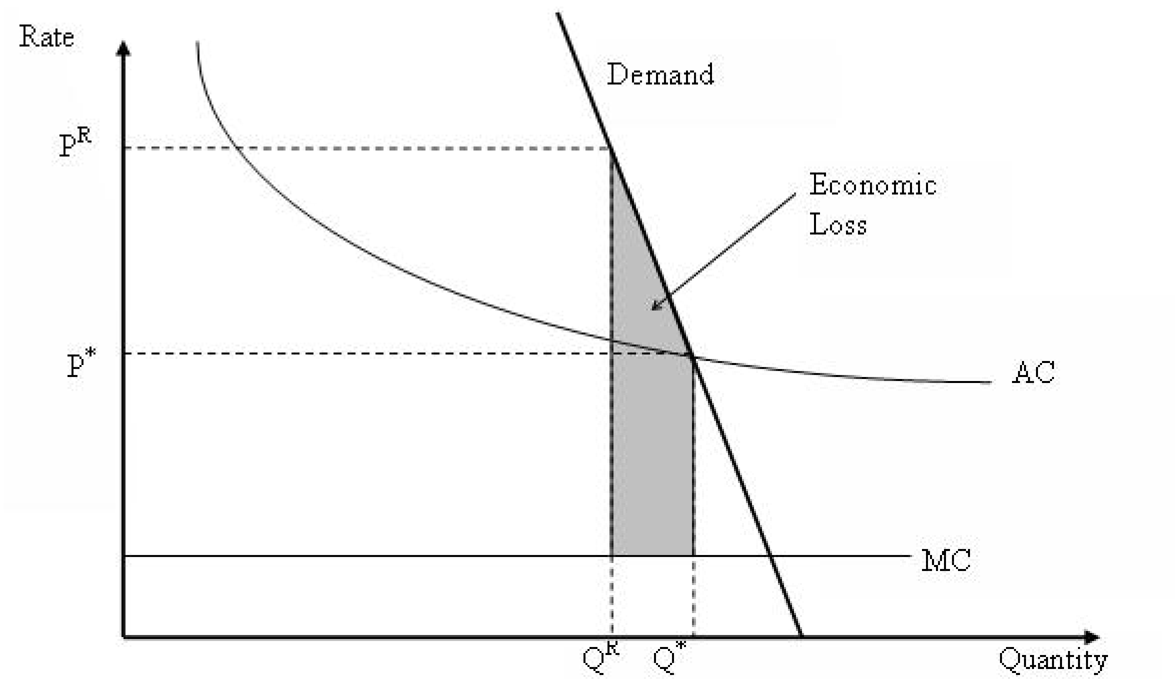

Economic losses are determined by both the magnitude and duration of the water supply disruption. Following Brozovic, Sunding, and Zilberman (2007), we define the severity of the water supply interruption in region i at time t as $z_{it} \in [0, 1]$, where $z_{it} = 0$ corresponds to a complete outage and $z_{it} = 1$ corresponds to the baseline level of service.

Let $f_{it}(z_{it})$ denote the probability density function of residential water disruption z_{it} in region i at time t and let $W_i(z_{it})$ denote consumer willingness-to-pay to avoid a supply disruption z_{it} in region i at time t. For a period of duration T until baseline water service is reestablished, consumer willingness-to-pay to avoid a cumulative service disruption across I regions and T periods is given by

$$W^R = \sum_{t=1}^{T} \sum_{i=1}^{I} \int_0^1 W_i(x) f_{it}(x) dx. \tag{3.1}$$

For a given region and time, the computation of $W_i(z_{it})$ that underlies the loss framework in equation (3.1) involves integrating the area under a demand curve for a shortage level of z_{it}. Specifically, willingness-to-pay to avoid a supply disruption of magnitude z_{it} in region i at time t can be defined as:

$$W_i(z_{it}) = \int_{Q_i(z_{it})}^{Q_i^*} P_i(x) dx \tag{3.2}$$

where $P_i(Q_i)$ is the (inverse) demand function for residential water in region i, $Q_i^* = Q_i(z_{it} = 1)$ is the baseline quantity of water delivered to residences in region i prior to a supply disruption, and $Q_i(z_{it})$ is the quantity of supply available after a water supply disruption in region i at time t.

Consumer willingness-to-pay to avoid a (contemporaneous) water supply disruption of a given magnitude in equation 3.2 is calculated for each agency by constructing an aggregate demand curve to represent the residential water segment. For agencies in which residential customers pay volumetric water rates, $P_i^* = P_i(Q_i^*)$ is the volumetric rate paid by residential consumers under baseline conditions prior to the water supply disruption in region i. For agencies with an inclining block pricing structure, $P_i^* = P_i(Q_i^*)$ is the marginal rate paid by a representative residential consumer on the highest tier of consumption under baseline conditions in region i. Integrating losses in equation 3.2 over the probability distribution of outcomes and the duration of the water supply disruption results in the consumer surplus loss measure described by equation 3.1.

Economic losses that result from water shortage in a given market are mitigated to the extent that delivering a smaller quantity of water reduces the system-wide cost of water service. The economic loss that occurs in region i following a water supply disruption is therefore the difference between the consumer surplus loss in equation 3.2 and the avoided cost of service. If water service is characterized by constant unit cost at the prevailing baseline price level, P_i^*, then the avoided cost of service is $P_i^*(Q_i^* - Q(z_{ij}))$, and the net economic loss following a water supply disruption of a given magnitude reduces to the consumer surplus triangle.

For most urban water suppliers, the cost of water service includes a large fixed cost component for elements of the water distribution network related to infrastructure, repair and maintenance of water lines, metering service, and public administration. To the extent that fixed costs facing water purveyors are sunk, the avoided cost of water service corresponds to the reduction in variable cost in region i for the duration of the water supply disruption. The variable cost of water supply include the energy and chemical costs of treating water and conveyance and distribution costs. The avoided cost of water service encompasses the reduction in these variable cost components in response to reduced water deliveries.

The avoided cost of service depends on the water portfolio of a given region and on the distribution of the water supply disruption across water sources within the portfolio. A disruption of relatively high-cost imported water supplies would entail a larger avoided cost of service component than a disruption of local water supplies by more greatly reducing procurement costs of water in the system. Given the variation in the water portfolios of regional purveyors, the avoided cost of service for units of water no longer delivered in response to a residential supply disruption can vary considerably across water agencies.

Let $c_i(z_{it})$ denote the avoided unit cost of service in region i at time t. Accordingly, the contemporaneous economic loss in region i of a given magnitude water supply disruption is

$$L_i(z_{it}) = \int_{Q_i(z_{it})}^{Q_i^*} (P_i(x) - c_i(x))dx. \tag{3.3}$$

Notice that the contemporaneous welfare loss in equation 3.3 corresponds with a consumer surplus measure in the case where $c_{it}(z_{it}) = P_i^*$. In this case, equation 3.3 reduces to

$$L_i(z_{it}) = \int_{Q_i(z_{it})}^{Q_i^*} P_i(x)dx - P_i^*(Q_i^* - Q(z_{ij})). \tag{3.4}$$

.

For a period of duration T until baseline water service is reestablished, the economic loss resulting from a cumulative service disruption across I regions and T periods is given by

$$L^R = \sum_{t=1}^{T}\sum_{i=1}^{I} \int_0^1 L_i(x) f_{it}(x)dx \tag{3.5}$$

where $L_i(z_{it})$ is defined in equation 3.3.

In the next section, we describe our econometric methods for estimating regional demand functions. The economic loss from a water supply disruption is then calculated using a parameterized loss measure similar to equation 3.5 for one period and regional estimates of demand conditions facing various water purveyors in California.

3.2 Residential Water Demand Estimation

An important factor determining economic losses in the residential sector is the elasticity of demand. The elasticity of demand is a parameter that summarizes how the willingness-to-pay of consumers to avoid a water service disruption changes with the level of water consumption. Specifically, the price elasticity of demand for water at any price P and consumption quantity Q is given by $\varepsilon = (dQ/DP)(P/Q)$, which represents the percentage change in willingness-to-pay for a percentage change in water deliveries. Our annual consumption data is aggregated at the retailer level. The choice of price index for our demand analysis is an important one. If all retailers used a uniform price then the choice of a price index would be a simple one. However, many retailers used a tiered rate structure, usually increasing rate block schedules. There are theoretical arguments for one price index versus another, but we have to rely on the data available for our analysis. From our data sources we have a measure of the average price paid by consumers in each retailer across years, and for this reason we use consumer response to average price as the relevant price elasticity.

3.2.1 Econometric Specification

Our basic estimating equation is:

$$\ln(q_{it}) = \alpha_1 \ln(\bar{p}_{it}) + \mu_i + \tau_t + \varepsilon_{it} \tag{3.6}$$

where q_{it} is the single family residence average monthly consumption for retailer i and year t; $\bar{p}_{it}$ is the average price per ccf of residential water; μ_i is a retailer fixed effect; τ_t is a year fixed effect; and ε_{it} captures all unobserved factors affecting the dependent variable. This specification will capture all unobserved time-invariant retailer specific characteristics.

However, the specification will fail to capture any time-variant retailer specific characteristics. To partially control for this we estimate retailer specific time trends, which will capture characteristics that display a common linear or quadratic trend across time within a retailer. For example, changes in technology or social norms might be captured by these retailer-specific time trends. Finally, for the welfare analysis we would like to account for spatial heterogeneity in price elasticities, especially between disparate regions of California like the San Francisco Bay Area and Southern California. To address this, we capture heterogeneity of the price elasticity across space by interacting our price variables with an indicator variable for Southern California. Our preferred specification can be written as:

$$\ln(q_{it}) = \beta_1 \ln(\bar{p}_{it}) + \beta_2 \ln(\bar{p}_{it}) \cdot \text{SoCal}_i + \mu_i + \mu_i \cdot t + \mu_i \cdot t^2 + \tau_t + \varepsilon_{it} \tag{3.7}$$

3.2.2 Data

The price elasticity of demand for water is estimated using single-family residential annual average price and consumption data from urban water retailers in the San Francisco Bay Area and in the Metropolitan Water Disctrict of Southern California (MWD). Average monthly consumption, average price, and customer data for water retailers in the San Francisco Bay Area were recorded from the Bay Area & Water Supply Conservation Agency Annual Surveys from 1994-2009; similar data for water retailers in the MWD service area was obtained by directly contacting water retailers. The final sample of retailers used to estimate our econometric model includes 26 retailers from in the San Francisco Bay Area and 56 retailers from within the MWD service area. In total, the sample contains 1050 observations spanning 82 water retailers.

Average annual single-family household consumption levels for each water retailer was calculated by dividing the total single-family residential consumption level for the fiscal year by the number of single-family residential accounts for that fiscal year. A monthly average consumption level was created by dividing the yearly average by twelve. The equilibrium price of water for the typical user in each region was assumed to be equal to the average price charged ($/ccf) to a residential customer, and in practice this measure is likely to be very close the marginal price facing residential consumers.

3.2.3 Estimation Results

Table 3.1 presents the results for three variations of the empirical model. In column (1) we present the results of our basic estimating equation, 3.6, which includes both retailer and year fixed effects. Under this specification we obtain a precisely estimated elasticity of -0.091; a result slightly smaller than previous studies. In column (2) we add retailer specific linear and quadratic time trends. The point estimate changes to -0.238, this suggests a trending omitted variable which is correlated with both consumption and price. The result is generally consistent with other estimates in the academic literature, although somewhat more elastic. Column (3) is identical to the specification in column (2) except we have introduced some heterogeneity into the price elasticity by interacting price with an indicator denoting Southern California. With a price elasticity of -0.149 water consumers in Northern California are more price inelastic than their Southern Californian counterparts who have a significantly lower price elasticity of -.308. These results make intuitive sense: residential consumers in the Southern California service areas we consider have more outdoor water use than those considered in Northern California.

Table 3.1: Residential water demand estimation

	(1)	(2)	(3)
ln(price)	-0.091***	-0.238***	-0.149**
	(0.019)	(0.056)	(0.050)
ln(price)*Southern California			-0.159**
			(0.057)
Observations	1050	1050	1050
R-squared	0.9525	0.9818	0.9819
Retailer fixed effects (m=82)	Yes	Yes	Yes
Year fixed effects (n=15)	Yes	Yes	Yes
Retailer specific quadratic time trend	No	Yes	Yes

Standard errors clustered at county level reported in parentheses.
*** p<0.01, ** p<0.05, * p<0.1

One concern with these specifications is that we do not account for deviations from typical weather patterns within a retail service area. However, in our sample we believe it is rare for retailers to adjust prices to contemporaneous weather patterns. To examine if excluding weather biases our price elasticities we consider a subset of retailers for which we have annual weather data. We re-run the specifications presented in table 3.1 with the addition of temperature and precipitation weather controls. The results are presented in table 3.2. Although the coefficient on price interacted with Southern California is different from table 3.1, we observe that the inclusion of weather does not significantly change the estimated price coefficients in columns (2) and (3) relative to columns (1) and (2). This is consistent with our prior thought that retailers rarely respond to contemporaneous weather patterns by adjusting price. We have evidence that the price elasticities shown in table 3.1 are not biased due to omitted weather control variables.

Table 3.2: Controlling for weather in the residential water demand estimation

	(1)	(2)	(3)	(4)
ln(price)	-0.159**	-0.158**	-0.144**	-0.143**
	(0.043)	(0.042)	(0.050)	(0.048)
ln(price)*Southern California			-0.080	-0.081
			(0.138)	(0.136)
Average max. summer temperature		0.015		0.016
		(0.013)		(0.013)
Total annual precipitation		0.007		0.006
		(0.044)		(0.041)
Observations	488	488	488	488
R-squared	0.9872	0.9872	0.9872	0.9872
Retailer fixed effects (m=82)	Yes	Yes	Yes	Yes
Year fixed effects (n=15)	Yes	Yes	Yes	Yes
Retailer specific quadratic time trend	Yes	Yes	Yes	Yes

Standard errors clustered at county level reported in parentheses
*** p<0.01, ** p<0.05, * p<0.1

3.3 Economic Losses of a Water Supply Disruption

In this section, we estimate welfare losses for Northern and Southern California in their respective residential sector during times of shortage.

3.3.1 Parameterizing the Loss Function

We assume a constant elasticity of demand specification:

$$P_i = A_i Q_i^{\frac{1}{\varepsilon_i}} \tag{3.8}$$

for $i = 1,..., n$, where ε_i is the elasticity of water demand in region i and A_i is a constant. Let P_i and Q_i^* respectively denote the retail water price and quantity of water consumed by residential households in region i under baseline conditions. For a given water shortage with an available level of water given by $Q_i(z_{it}) < Q_i^*$, it is helpful to define the relationship between these quantities in terms of the percentage of water that is rationed in region i at time t, r_{it}, as:

$$Q_i(z_{it}) = (1 - r_{it})Q_i^*. \tag{3.9}$$

Making use of equations 3.8 and 3.9, the economic loss following a supply disruption of magnitude z_{it} in region i at time t can be calculated from equation (3.4) as:

$$L_i(z_{it}) = \frac{\varepsilon_i}{1+\varepsilon_i} P_i^* Q_i^* [1 - (1 - r_{it})^{\frac{1+\varepsilon_i}{\varepsilon_i}}] - \int_{Q_i(z_{it})}^{Q_i^*} c_i(x)dx \tag{3.10}$$

3.3.2 Data for Calculation of Losses

Using equation (3.10) losses are calculated for each agency with estimated residential water demand parameters from our econometric model. The resulting agency-level values are used to calculate the total economic losses in Nothern and Southern California by summing across agencies.

The loss of being short an additional unit of water is the difference between consumer's willingness-to-pay for the unit, and the marginal cost of the unit. The former is equivalent to the first term in equation 3.10, $\frac{\varepsilon_i}{1+\varepsilon_i}P_i^*Q_i^*[1-(1-r_{it})^{\frac{1+\varepsilon_i}{\varepsilon_i}}]$, and is calculated based on observed consumption, prices, shortage scenarios and estimated price elasticities. The avoided marginal costs of water delivery is based on reports from water district engineers.

We develop loss estimates due to water shortages for SFPUC, BAWSCA, Santa Clara Valley WD, and East Bay MUD in Northern California; and MWDSC in Southern California. As previously noted, losses also depend on regional prices, baseline demand, demand shortages, and prices. Estimates for all of the information required for the loss calculation are publicly available except for the price elasticities; these were generated based on the residential water demand estimation.

Baseline demand data is recorded from FY 2009-2010 Urban Water Management Plans for each water agency included in the analysis. The price data is acquired from separate sources for Northern and Southern California. The Northern California price data is the 2009 median rate reported in the BAWSCA survey, or as recorded from agency websites. From these we calculate a quantity-weighted average for the price index to be used in the loss analysis. The calculation of the Southern California data is slightly more complicated because not all MWD member agencies are retailers, some of them are wholesalers who sell to retailers. We use FY 2009-2010 Urban Water Management Plans to calculate median rates for all retailers within each member agency. The MWD member agency average price we use in the loss analysis is a quantity-weighted average of the average prices for retailers residing in an MWD member agency's service area.

The demand estimation suggests that the price elasticity of water is significantly different between Southern and Northern California; however, we do not have evidence that there is statistically significant variation within these two regions. For the loss analysis, we assume one price elasticity for each of these two regions based on the results of column (3) in table 3.1. We assume all regions in Northern California have a single family residential price elasticity of demand equal to -.149 MWDSC we assume a price elasticity of -.308. In terms of the residential loss functions, the impact of regional demand elasticity values is to raise economic losses from water service disruptions when demand is relatively inelastic and to lower economic losses in regions where demand is more elastic. Predictions of residential water shortages in California reflect considerable uncertainty, and have a wide range. Thus, we calculate losses due to 10, 20, and 30 percent shortages.

The final piece of information required to calculate losses due to a shortage is the marginal cost of service delivery. The marginal cost of supply depends on the supply source. For simplicity, in Northern California we assume Hetch Hetchy reservoir is the relevant supply source. Communication with water managers in the San Francisco Bay Area suggest the marginal cost is \$50/AF. For the MWD member agencies, the Delta water from the State Water Project is the relevant supply source. Communication with water managers in the MWD service area suggest the marginal cost is \$225/AF.

3.3.3 Results of the Loss Analysis

Table 3.3 reports the results of the loss analysis for shortage scenarios of 10%, 20% and 30% for the Northern and Southern Californian service areas considered in our study. We present welfare losses under different assumptions about the marginal costs of service delivery; welfare losses are equivalent to willingness-to-pay for water supply reliability. Row one in each panel presents results under the assumption that marginal costs are zero. In this instance, the loss experienced is exactly consumer willingness-to-pay for the water. In the second row of each panel we present the results under the assumption that marginal costs are $50/AF for Northern California regions and $225/AF for Southern California regions. The third row presents losses under the assumption of marginal cost pricing—losses are equivalent to the consumer surplus triangle. The latter is often used as a measure of WTP for water supply reliability.

Table 3.3: Welfare losses due to shortages of 10%, 20% and 30%

Panel A: Northern California			
Price Index: $1,501/AF; Total M&I Demand: 911,969; Elasticity: -0.149			
	10% disruption	20% disruption	30% disruption
WTP for water not delivered (MC = 0)			
Total	$197,801,573	$617,566,589	$1,598,202,006
Average	$2,169	$3,386	$5,842
Estimated Losses due to a shortage (0 < MC <p)*			
Total	$193,241,728	$608,446,899	$1,584,522,471
Average	$2,119	$3,336	$5,792
Consumer surplus triangle (MC = p)*			
Total Loss	$60,918,733	$343,800,909	$1,187,553,487
Average	$668	$1,885	$4,341
Panel B: Southern California			
Price Index: $1,267/AF; Total M&I Demand: 3,796,681; Elasticity: -0.308			
	10% disruption	20% disruption	30% disruption
WTP for water not delivered (MC = 0)			
Total	$571,840,667	$1,393,703,383	$2,630,423,148
Average	$1,506	$1,835	$2,309
Estimated Losses due to a shortage (0 < MC <p)*			
Total	$486,415,344	$1,222,852,738	$2,374,147,181
Average	$1,281	$1,610	$2,084
Consumer surplus triangle (MC = p)*			
Total Loss	$90,801,184	$431,624,417	$1,187,304,700
Average	$239	$568	$1,042

For each magnitude of supply disruption, the total economic loss is the average loss ($/AF) multiplied by the decrease in total M&I deliveries. The lower bound for economic losses ranges from $239/AF for a 10% supply disruption in Southern California with marginal cost pricing to

$5,842/AF for a 30% supply disruption with zero marginal costs in Northern California. At all levels of supply disruption, the difference between the lower bound value and upper bound value of economic losses in each region is given by the price index.

Notice that the average economic loss is rising in the magnitude of the supply disruption. The reason is that consumer willingness-to-pay for an incremental unit of water is decreasing in the level of water consumption as the value of water applied to various end-uses declines with the total level of supply. In response to a water supply disruption, residential consumers select the least costly conservation activities before dispensing with higher valued uses of water.

Basing the loss calculation on the typical water use level of residential households in the service territory of each member agency implicitly assumes that a water supply disruption of a given magnitude is met with proportionate rationing on each tier of the pricing structure. For member agencies that meet a supply reduction by raising water rates predominantly on the higher tiers of the pricing structure, the actual economic losses would be larger than the losses described above.

There are at least two important points to be taken away from table 2. First, the source of unreliability can greatly affect the size of the loss according to the marginal costs of service delivery from the supply source. For example, notice that losses (the value of supply reliability) in Southern California under a 10% shortage when marginal costs are close to zero is roughly six times the size of losses under marginal cost pricing. Failure to distinguish between the marginal costs of service delivery from different supply sources in Northern and Southern California would distort the value of supply reliability. Second, heterogeneity in the price elasticity has a huge impact on losses across regions. Under a 30% shortage the losses per acre foot in Northern California are two to four times the size of average losses in Southern California. In other words, if we applied the same price elasticity of -0.15 to Southern California as we did in Northern California, then this could quadruple the value of water supply reliability. This is not inconsequential as such a difference in value could influence the economic viability of large infrastructure projects designed to secure water supply reliability.

3.4 Conclusions

We estimate welfare losses to residential water consumers during water supply shortages, and in doing so generate a value of water supply reliability. In the economics of water utilities a simple, yet often overlooked fact, is that utilities load a large share of the high infrastructure costs onto the unit price of water they charge their customers. Therefore, the consumer surplus triangle which is often used to determine welfare losses due to supply shortages is theoretically incorrect. Said differently, fixed costs are sunk at the time of a water supply disruption and so have no bearing on welfare outcomes. The welfare loss from a supply disruption, accordingly, can be many times greater than the loss evaluated using standard measures of compensating variation or equivalent variation in consumer utility functions. Consistent with this, our loss analysis shows large differences in welfare losses depending on the marginal cost of service delivery.

Germane to the discussion is the fact that there is regional variation in the marginal cost of service delivery; hence, regional data on these marginal costs should be an object of interest for water managers and policy-makers alike. Further, marginal costs relate explicitly to the source of supply. If a region has multiple supply sources then water managers and policy-makers must consider the supply source facing unreliability when evaluating the benefits of investing in a reliable alternative supply.

A final point this paper makes relates to the estimation of price elasticities. Our work empirically demonstrates that the price elasticity of residential water demand displays statistically significant spatial variation. Accounting for this variation is essential for accurately assessing welfare losses, and evaluating reliability infrastructure investments. Understanding spatial variation in price elasticities will also aid water managers when allocating water shortages across space.

Our analysis offers useful guidance for those interested in more accurate assessments for the value of water supply reliability. However, there are several fronts for future research. To our knowledge, water managers in California or their offices do not keep detailed records on the marginal costs of service delivery. Our research suggests they would be wise to do so, and if collected, the data could inform subsequent research on the value of water supply reliability. Data limitations exist; for example, our demand estimation relies on consumption data aggregated to the retailer level. Administrative records can provide household level measures of consumption, and would significantly improve the demand estimation. Absent such micro-level data, information on the distribution of consumption within a retailer could be combined with retailer specific tiered rate structures to compute the theoretically correct price index for estimating price elasticities of demand using retailer-level household consumption averages. It may also be possible to introduce heterogeneity into price elasticities at a more disaggregate level than Northern and Southern California. Finally, while a panel of price and consumption data is preferable to cross-sectional data, our demand specification may still suffer from omitted variable bias or simultaneity issues. For instance, conservation practices such as ordinances to curtail outdoor water use or conservation campaigns may vary across time within a district and so would not be accounted for using retailer fixed effects. These retailer-level conservation efforts may be correlated with both price and consumption, and so would bias our estimation results. Data collection efforts for retailer level information on conservation efforts would allow one to address this form of omitted variable bias. Likewise, if retailers are easily able to update prices to reflect changes in demand or anticipated demand then our estimates would be biased. A more complete picture of how prices are set would be informative when choosing a specification to estimate demand.

Chapter 4

Agricultural Extension, Trust and Learning

Technical change in agriculture is an important engine of poverty-reducing growth in developing countries. Transmission of knowledge about technical alternatives is, however, costly and governments are increasingly seeking cost-effective means of spreading information (Godtland et al., 2004). An important determinant of cost-effectiveness is farmer learning. Adegbola and Gardebroek (2007) and Moser and Barrett (2006) study how farmers learn about agricultural technologies from agricultural technicians (extension agents) and community farmers. They find differences in adoption depending on the information source, although there is scant evidence about factors that account for such differences. An important, yet understudied, determinant of uptake of new knowledge may be the recipient's trust in the motives of the information source. Because they may question underlying motives, farmers do not value all sources of agricultural information equally. Some farmers may trust neighboring farmers while others, if given the opportunity, may be more inclined to trust agricultural technicians.

Few studies have tried to address the causal effect of trust on behavior in the field environment; notable exceptions are that of Bouma, Bulte, and Van Soest (2008) and Tu and Bulte (2010). We study how trust attitudes affect learning in the context of agricultural extension training for naranjilla farmers in the Amazon region of Ecuador. Naranjilla is a solanaceous fruit used mainly for juices and its farmers face severe pest and disease problems, such as fruit worms, borer beetles, and fungi (Ochoa et al., 2001). Anecdotal evidence from technicians and farmers indicates the most common management tool is the application of agrochemicals; Andrade (2005) confirms this observation. Compared to their use on other crops, agrochemicals are relatively ineffective at managing naranjilla pest and disease problems (Baez, Ochoa, and Ellis, 2003; Gallegos et al., 2003); these chemicals also cause environmental damage and pose long-term human health risks (Zilberman and Castillo, 1994).

Cost-effective integrated pest management (IPM) techniques have been developed to help naranjilla farmers (Shiki et al., 2003), yet Andrade (2005) finds IPM adoption has been slow. IPM messages tend to be complex and uptake usually requires extensive learning (Norton, Heinrichs, and Luther, 2005). In Ecuador, transmission of complex messages to naranjilla farmers is complicated due to their physical isolation and the absence of public agricultural extension since the early 1990s (Alwang et al., 2005). Ecuadorian policy-makers are now planning to re-invest in public agricultural extension services and are seeking efficient models for information transfer. While supplying naranjilla farmers with information exposes them to IPM practices, farmers may ignore it or put less effort into learning if they do not trust the information source. Farmers may wonder if financial, political or environmental objectives underlie the advice technicians provide. Farmers report they are skeptical of outside technicians who bring new agricultural technologies; they question if these technicians are promoting these technologies for personal financial gain independent of whether the product is appropriate. Other farmers suspect that the promises of successful new technologies come from technicians representing political agendas. The appearance of service delivery agents such as agricultural technicians, especially near elections, may be rewarded in the ballot box although technicians' advice may later be found to produce little value. Some farmers report that technicians have tried to convince them to reduce tree-cutting, which is viewed by farmers as a strategy to conserve natural resources at the expense of reducing harvests. Similarly, farmers may question whether IPM practices truly increase yields or if they are a strategy to prevent chemical run-off into surface waterways. From the farmer's perspective, technicians who are perceived to have environmental or other objectives, rather than farmer welfare as an objective, are self-regarding technicians. Accordingly, such technicians would prefer to devote resources towards consuming environmental services (or another competing objective) rather than on consumption that improves farmer welfare. Thus, naranjilla farmers may suspect the motives of technicians.

This problem is not unique to naranjilla farmers in Ecuador; substantial evidence shows that development agents often face incentives that are not aligned with the objectives of those they serve. Chaudhury et al. (2006) survey evidence in India showing that rural school teachers and healthcare providers are consistently absent from work; Easterly (2009) finds similar evidence in Africa. Labie, Méon, and Szafarz (2009) investigate how incentives may lead credit officers in microfinance institutions to discriminate against minority loan applicants despite a mission to serve minority populations. In addition to pecuniary benefits, our experience in the field suggests that program mission and other nonpecuniary motivations are present among agricultural technicians in Ecuador. In our framework we let service recipients recognize the different motivations service delivery agents may have. In particular we let farmers condition their effort to learn during an agricultural training on their trust in the motivations of information sources.

Our work also contributes to the growing literature on agricultural technology adoption. Recent work in the adoption literature has taken advantage of advances in behavioral economics and experimental methods to investigate psychological explanations of adoption behavior. Liu (2010) tests the role of prospect theory in the adoption of biotech rice in China by capturing experimental measures of risk preferences among farmers. Using a randomized intervention Duflo, Kremer, and Robinson (2009) test whether farmers in Kenya are present-biased as a reason for the slow adoption of fertilizer. Most closely related to our work is that of Cole et al. (2009) in India; they test whether trust in an insurance product representative affects adoption of rainfall insurance. In an ideal experiment, trust would be randomly assigned, but this is impossible. Absent this experiment, Cole et al. randomly vary whether the insurance product representative is endorsed by a local microfinance customer service agent. They also run an experiment in which insurance flyers are randomly assigned endorsements using references to Hindu or Muslim faith. While both endorsement treatments have positive impacts on insurance product uptake, the authors do not have a revealed measure of trust in the messenger. It is also not clear if these endorsements only confer trust or something additional such as authority or group membership. Our study complements their work in three ways. First, we use revealed measures of trust in information sources, the messenger. Second, we consider the role of trust via learning in adoption of an agricultural cultivation technology rather than a financial product. Finally, we consider how trust in the messenger relative to trust in other information sources affects adoption as opposed to only the importance of absolute trust in the messenger.

We hypothesize that learning from a technician, and adoption behavior more broadly, will depend on a farmer's trust in technicians and other community farmers. Using a stylized model as our guide, we consider measures of absolute and relative trust in technicians compared to community farmers to empirically test our thesis. We find evidence that farmers who attend a training and who trust technicians relatively more than community farmers score significantly higher on an exam based on the training than other farmers. We find no evidence that farmers' trust attitudes are related to factors that make them more or less knowledgeable ex-ante about the exam material; trust measures are not significantly correlated with exam scores for farmers who take the exam without the benefit of the training.

Based on our estimates the impact of trust on learning during training is relatively large; for example, farmers who trust technicians more than community farmers score almost one standard deviation higher on the exam than others. This result is timely given that Ecuador is moving forward with a renewed public agricultural extension system. Our finding suggests that the channel by which information is delivered is an important determinant of learning and acceptance. And while we say nothing about how learning translates into behavioral change, incorporating new information and updating beliefs is a necessary, if not sufficient, condition for such change.

The remainder of this chapter is organized as follows. We start with a brief discussion of how we define trust. Then we develop a simple model of farmer behavior and learning during training. Subsequently, we describe the data sources and the empirical framework. The article concludes with a discussion of results and conclusions.

4.1 Defining Trust

Following Gambetta (1988) and Schechter (2007) trust is defined as an individual's subjective probability that another person is motivated to perform an action beneficial to her. Said differently, trust in another person is the belief that an informal contract with him or her is incentive compatible. For clarity throughout the remainder of the paper, we reference someone who has motives to perform an action beneficial to a farmer as other-regarding behavior. We reference the absence of such motives as self-regarding behavior. In our framework both other-regarding and self-regarding individuals are pursuing their self-interests. However, for other-regarding individuals their interests include the pay-outs to farmers, while for self-regarding individuals their interests do not include the pay-outs of farmers, although they may include rewards from environmental, political or personal financial outcomes.

Next we distinguish between trust in motives and trust in competence. Trust in motives refers to an individual's beliefs that another individual is other-regarding or not. Trust in competence refers to an individual's beliefs that another person has credible expertise. By way of example, suppose Sam just bought a car from her grandparents and they told her they just replaced the brake pads. Now suppose Sam goes to change the oil and the mechanic says she needs new brake pads. It is doubtful that Sam will get new brake pads. It is not because Sam does not trust that the mechanic knows when brake pads are worn thin; it is because Sam does not trust the mechanic's motives, rather Sam trusts her grandparents' other-regarding motives.

Both farmer trust in the motives and in the competence of information sources may condition from whom she accepts agricultural advice, although we view the latter as less relevant to our study. We believe there to be insignificant variation among farmers with respect to their trust in the competence of technicians. Absent data on trust in competence we cannot empirically test this claim. Nonetheless, casual observation indicates that in our study area, and throughout Ecuador, higher education is uniformly held in high-esteem. Agricultural technicians have advanced training in agronomy, and so it is likely that farmers' respect the competence of technicians in agronomic matters. However, as we have already described, some naranjilla farmers may question the motives of technicians; hence, our interest is in how trust in motives impacts learning about technology.

4.2 The Trust Game

Measures of trust were obtained from a dataset originally collected in 2005 to study the correlates of trust attitudes. Trust measures are based on the investment game, also commonly called the trust game (Berg, Dickhaut, and McCabe, 1995), in which there are two players, an investor and trustee. The investor starts with a given endowment. The investor decides to send part, all or none of her endowment to an anonymous person, the trustee. Any money the investor passes to the trustee is tripled. The trustee has the opportunity to return part, none or all of the money to the investor. None of these plays are observed by participants; in addition, the investors do not know how much money is returned to them from the trustees until after all plays have been made.

Since the trustee's play is anonymous and the game is not repeated, we expect the trustee to return nothing to the investor under the assumption that individuals are self-interested. From backward induction the investor identifies the trustee's 'return nothing' strategy. Hence, the only sub-game perfect Nash Equilibrium is that the investor sends nothing because she expects the trustee to return nothing. In laboratory experiments, participants rarely play according to this strategy (Barr, 2003; Glaeser et al., 2000; Karlan, 2005; Schechter, 2007). Investors typically send money. Social scientists often attribute this deviation from expected play to trust (Camerer, 2003). Farmers play the trust game as the investor against community farmers and against agricultural technicians. Since the farmers know which class of trustee (community farmer or technician), but not the specific individual with whom they are playing, measures of trust in community farmers and trust in technicians are obtained[1]. We employ the trust game with anonymous partners because it mimics the interactions that often typify relationships between naranjilla farmers and agricultural technicians in Ecuador. Under the current extension system a technician might visit a community for a one time training and never return.

4.3 Modeling Farmer Behavior in the Agricultural Training

To motivate and direct our empirical tests we develop simple models of farmer behavior during the training session. Like all models, the ones we present are incomplete, but we find ours informative

[1]Recently social scientists have questioned whether the action of the sender in the investment game is a measure of trust or just a measure of the propensity to gamble. In a laboratory setting, Eckel and Wilson (2004) find that experimental measures of trust are not correlated with experimental or survey measures of risk; however, in lab experiments taken to the field, both Karlan (2005) and Schechter (2007) present evidence that experimental trust measures are partially determined by, or at least associated with, risk behavior. To address this we control for risk in the regression analysis.

as a guide to the economic logic that inspired the empirical analysis.

We consider two sources of information: technicians and community farmers. Also, we let agricultural information sources be of two types: those who care about farmers (other-regarding) and those who care about a personal interest (self-regarding). An other-regarding farmer may be willing to trade-off a personal pay-out to increase the pay-out of another farmer. For example, one can imagine an other-regarding farmer foregoing a day of work in his or her own field to help her neighbor learn how to use a new pesticide. Self-regarding farmers would not be willing to make such as sacrifice; in fact, a self-regarding farmer might withhold information about a new technology from others in order to reduce competition. We argue that such people would return less money to the investor while playing as the trustee in the trust game. Similarly, compared to the self-regarding technician, an other-regarding technician naturally has objectives that are relatively well-aligned with farmers. He or she learns and teaches practices that are of high-value to farmers. However, a self-regarding technician who cares about other objectives, for example, financial, political or environmental objectives, has incentives that are not as well-aligned with the farmer. He or she learns and teaches practices that are of relatively low-value to farmers. Similarly, the other-regarding farmer offers advice that is of relatively high-value, while the self-regarding farmer may offer low-value advice that lets him free-ride on other farmers' investments in pest and disease reduction.

We let each farmer i form a subjective probability $\theta_i^s \in [0, 1]$ that a given information source is other-regarding, and $1-\theta_i^s$ that a given information source is self-regarding; s denotes if the information souce is a technician (t) or farmer (f).

4.3.1 Model of Learning from Community Farmers and Technicians

The value of advice, v, varies across information sources and assumes values v_H^t, v_L^t v_H^f, or v_L^f. The superscript on v indicates whether the source is a technician (t) or farmer (f), while the subscript indicates whether the source is other-regarding (H) or self-regarding (L). As already suggested other-regarding information sources offer relatively high-value advice and self-regarding information sources offer relatively low-value advice; hence, we assume $v_H^t > v_L^t$ and $v_H^f > v_L^f$.

The return to effort when learning from a given information source s can be written as $v^s \cdot g(e)$ where e is cumulative effort. We assume there is overlap between advice offered by community farmers and in the training; the return to effort in the training is likely diminished as a consequence of past effort. Therefore, e is the sum of effort exerted learning from community farmers prior to the training and from technicians during the training, $e = e^f + e^t$. We suppose diminishing marginal returns to learning in effort ($g'(e_i) > 0$, $g''(e_i) < 0$) for farmer i. Also, we let effort be costly and the marginal cost of effort during the training to be increasing in effort ($c''(e_i^t) > 0$) where the superscript t indicates effort exerted during the training. We assume farmers are utility maximizing with respect to effort. Absent technicians, farmer i chooses the optimal level of effort to exert learning from community farmers by solving:

$$\max_{e_i^f} \, (\theta_i^f v_H^f + (1-\theta_i^f) v_L^f) g(e_i^f) - c(e_i^f). \tag{4.1}$$

Note that the first term captures the expected benefits of information offered by community farmers which is scaled by effort, and the second term captures the cost of exerting effort to learn from community farmers.

Now suppose an agricultural technician makes an unexpected visit to the community and provides an agricultural training. Farmer i has already solved the problem described in expression (4.1) and exerted effort e_i^{f*} learning about naranjilla cultivation from community farmers; hence, due to diminishing marginal returns to effort it is less worthwhile to exert additional effort during the training. Formally, farmer i's maximization problem in the training is:

$$\max_{e_i^t} \; (\theta_i^t v_H^t + (1-\theta_i^t)v_L^t)g(e_i^t + e_i^{f*}) - c(e_i^t). \tag{4.2}$$

Again, note that $g(e_i^t + e_i^{f*})$ captures the fact that after already exerting effort to learn from community farmers the return to effort during the training is diminished (i.e., $g'(e_i^{f*}) > g'(e_i^t + e_i^{f*})$). Taking the derivative of expression (4.2) with respect to e_i^t we obtain the following first order condition:

$$(\theta_i^t v_H^t + (1-\theta_i^t)v_L^t)g'(e_i^t + e_i^{f*}) - c'(e_i^t) = 0. \tag{4.3}$$

Totally differentiating equation (4.3) with respect to e_i^t and θ_i^t we find:

$$[(\theta_i^t v_H^t + (1-\theta_i^t)v_L^t)g''(e_i^t + e_i^{f*}) - c''(e_i^t)]de_i^t + [(v_H^t - v_L^t)g'(e_i^t + e_i^{f*})]d\theta_i^t = 0. \tag{4.4}$$

Similar to before, re-arranging the expression we find:

$$\frac{de_i^t}{d\theta_i^t} = -\frac{(v_H^t - v_L^t)g'(e_i^t + e_i^{f*})}{(\theta_i^t v_H^t + (1-\theta_i^t)v_L^t)g''(e_i^t + e_i^{f*}) + (-c''(e_i^t))} > 0. \tag{4.5}$$

It is easy to verify $\frac{de_i^t}{d\theta_i^t} > 0$; both factors in the numerator are positive; and both terms in the denominator are negative. Applying the negative sign to their ratio yields the observed inequality.

An additional interesting feature of this model is that because e_i^{f*} is a function of θ_i^f we can solve for $\frac{de_i^t}{d\theta_i^f}$. We totally differentiate our first order condition with respect to e_i^t and θ_i^f to obtain:

$$[(\theta_i^t v_H^t + (1-\theta_i^t)v_L^t)g''(e_i^t + e_i^{f*}) - c''(e_i^t)]de_i^t + [(\theta_i^t v_H^t + (1-\theta_i^t)v_L^t)g''(e_i^t + e_i^{f*})(\frac{de_i^f}{d\theta_i^f})]d\theta_i^f = 0. \tag{4.6}$$

Re-arranging this expression we find:

$$\frac{de_i^t}{d\theta_i^f} = -\frac{(\theta_i^t v_H^t + (1-\theta_i^t)v_L^t)g''(e_i^t + e_i^{f*})(\frac{de_i^f}{d\theta_i^f})}{(\theta_i^t v_H^t + (1-\theta_i^t)v_L^t)g''(e_i^t + e_i^{f*}) + (-c''(e_i^t))} < 0. \tag{4.7}$$

It is easy to verify $\frac{de_i^t}{d\theta_i^f} < 0$; the first and last factors in the numerator are positive while the middle factor ($g''(e_i^t + e_i^{f*})$) is negative so that the numerator is negative; in the denominator the first term is negative while the second term is negative; hence, the denominator is also negative. Applying a negative sign yields the observed inequality.

If effort to learn from farmers enters into the farmer's decision problem as described then effort in the training increases with trust in technicians and decreases with trust in community farmers. Combining these results also delivers the result that effort in the training increases as the difference between trust in technicians and community farmers widens. This occurs because as trust in farmers (θ_i^f) increases then effort to learn from farmers (e_i^{f*}) also increases, which decreases the value of effort to learn during the training (e_i^t); at the same time as trust in technicians (e_i^t) increases then so does effort during training (e_i^t). Together these results imply: $\frac{de_i^t}{d(\theta_i^t - \theta_i^f)} > 0$.

4.3.2 Relative Trust & Selection into Learning Regimes

Farmers in the Pastaza province of the Ecuadorian Upper Amazon Basin are not unfamiliar with naranjilla cultivation–there exist profit-maximizing agricultural practices that account for the variety of plant pests and diseases a naranjilla farmer may face. When alternative practices are introduced, these practices may be at odds with this optimal mix; otherwise, there would be no reason to introduce outside practices. Our second model of farmer learning during training highlights this point. Trust in agricultural technicians relative to community farmers may affect decisions about how much effort to exert at learning during training, this is because agricultural advice from different sources of information are some times at odds.

In the empirical analysis we consider advice on the uses of agrochemicals as an illustration of this tension: neighboring farmers may advise that a particular agrochemical is useful for controlling nematodes while the technician may argue that the agrochemical is ineffective at controlling nematodes (or vice versa). In this extreme scenario, the farmer can only accept the advice of one information source, although the farmer does not know which is the more valuable advice. We assume farmers who are faced with conflicting advice in the training rely on the advice of the more trusted information source and forgo the advice of the other information source.

To address this possibility, we endogenize the choice of selecting into learning from community farmers versus learning from technicians. Farmer i's choice of learning regime depends on the expected value of advice from community farmers compared to the expected value of advice from technicians. Formally, farmer i compares the following expected values:

$$\mathbb{E}(\text{Advice from farmers}) = \theta_i^f v_H^f + (1 - \theta_i^f) v_L^f \tag{4.8}$$

$$\mathbb{E}(\text{Advice from technicians}) = \theta_i^t v_H^t + (1 - \theta_i^t) v_L^t. \tag{4.9}$$

In this framework farmer i chooses to learn from the information source offering the greater of the values expressed in equations (4.8) and (4.9)[2]. We derive a threshold, $\bar{\theta}_i(v_H^t, v_L^t, v_H^f, v_L^f, \theta_i^f)$, $\bar{\theta}_i$ henceforth, at which farmer i switches learning regimes if $\theta_i^t > \bar{\theta}_i$. That is, $\bar{\theta}_i$ is the point when farmer i decides to opt for the advice from technicians and forgo the advice of his or her fellow farmers. For example, farmers who came to the field day for the promised financial compensation and are untrusting of technicians are likely to value the advice of community farmers more than technicians; farmers attracted to the field day for the technician's training rather than the financial compensation are likely to value the advice of technicians more than community farmers. We write

[2]Farmer i may anticipate the cost of effort as well as how the level of effort during training will scale learning; if so, then these should appear in equations (4.8) and (4.9). However, we choose to suppress these effects of effort under the assumption that farmers employ a non-standard decision-making process in which they first choose who to trust, and then choose the optimal level of effort.

down the necessary inequality using equations (4.8) and (4.9):

$$\theta_i^t v_H^t + (1-\theta_i^t) v_L^t > \theta_i^f v_H^f + (1-\theta_i^f) v_L^f. \tag{4.10}$$

Collecting like terms on both sides of the inequality we obtain:

$$\theta_i^t (v_H^t - v_L^t) + v_L^t > \theta_i^f (v_H^f - v_L^f) + v_L^f. \tag{4.11}$$

Subtracting v_L^t from both sides of the inequality and then dividing both sides by $(v_H^t - v_L^t)$ we find:

$$\theta_i^t > \frac{\theta_i^f (v_H^f - v_L^f) + v_L^f . - v_L^t}{(v_H^t - v_L^f)} \tag{4.12}$$

Re-arranging the right-hand side of the inequality we identify our threshold as:

$$\bar{\theta}_i = \frac{v_L^f - v_L^t}{v_H^t - v_L^t} + \frac{v_H^f - v_L^f}{v_H^t - v_L^t} \cdot \theta_i^f = \beta_0 + \beta_1 \cdot \theta_i^f \tag{4.13}$$

To avoid unnecessary quantitative complication, our approach is to assume different qualitative cases for $(v_H^t, v_L^t, v_H^f, v_L^f)$, always preserving $v_H > v_L$ independent of the information source. As a baseline we assume the value of advice from farmers and technicians to be identical ($v_H^t = v_H^f$ and $v_L^t = v_L^f$). In this case we observe $\bar{\theta}_i = \theta_i^f$ so that the comparison between (4.8) and (4.9) simplifies to comparing θ_i^f and θ_i^t; this is the previously described scenario where farmers rely on the advice from the information source they trust the most. In the jackknife case when $\theta_i^f = \theta_i^t$ we assume farmers accept the advice of farmers and ignore the advice of technicians (i.e., exert no effort during training) since they have already sunk costs into learning from farmers. Clearly, $\bar{\theta}_i$ will vary according to the values of $(v_H^t, v_L^t, v_H^f, v_L^f, \theta_i^f)$. Our insight is that farmers with $\theta_i^t > \bar{\theta}_i$ exert effort and behave according to the model described in section 4.3.1.

To summarize, relative trust attitudes govern a selection process whereby farmers decide whether to listen to community farmers or technicians. As in the previous model, the inequalities (4.5) and (4.7) indicate farmers exert more effort as trust in technicians increases and less effort as trust in community farmers increases. Combining these results as in the end of section 4.3.1, we predict that if $\theta_i^t > \bar{\theta}_i$, then as the absolute difference between trust in technicians and community farmers increases so does effort in the training and thus learning; if $\bar{\theta}_i \geq \theta_i^t$, the farmer exerts no effort.

4.4 Data

4.4.1 Background

Data was orginally collected in 2005 as part of study examining the correlates of trust, and the dataset contains no personal identifiers. The sample of participants came from five communities in two adjacent counties in the Pastaza Province of the Ecuadorian Upper Amazon Basin. We highlight features that are unique to this field data. In July 2005 local leaders informed farmers that there would be a training to help combat naranjilla pests and diseases as well as an activity in which farmers would earn roughly 5-15 dollars; flyers were distributed with the same information. So many farmers participated in each field day that it was necessary to divide farmers into two groups. One group attended a training session given by a technician, participated in the trust games and then completed a short survey; we reference this group as the treatment group. The other group participated in the trust games, completed the short survey and then attended the training; we reference this group as the control group. A short exam was administered during the survey based on material presented in the training. The farmers in the treatment group completed the exam after the training, while the farmers in the control group completed the exam without the benefit of the training. The final sample used in our analysis is composed of 140 naranjilla farmers.

We also point out that farmers participated in multiple trust games. Each farmer participated in three separate trust games, each with a different anonymous partner, and each with the following order. Farmers played as the investor partnered with an anonymous farmer; as the investor partnered with an anonymous technician; and as the trustee partnered with an anonymous farmer. It is important to note that farmers made each decision without any knowledge of the outcomes from the other games. That is, there was no feedback to farmers between plays so that no income effects or learning were possible from observed outcomes.

An obvious concern with the ordering of the field day events is that the agricultural training was given by an agricultural technician; this may have affected trust attitudes towards technicians which were subsequently measured. We address this concern in the empirical section.

4.4.2 Sample Summary

The sample of farmers is diverse in age, education, wealth, family size, and landholdings and a significant portion were female (26 percent), indigenous (31 percent), Catholic (28 percent), or immigrants into the community (9 percent). About 18 percent of our sample reported having adopted agrochemicals based on advice received from a technician in the past. Current use of methamidophos pesticide was reported by 38 percent of participants; 26 percent reported current use of carbofuran pesticide; 11 percent reported use of chlorothalonil fungicide; and 8 percent reported use of 2, 4-D butyl ester herbicide. If assignment into the 'treatment' and 'control' groups was as good as randome then the two groups should not differ systematically. Table 4.1 presents the means and standard deviations of the covariates for the two groups. We observe that randomization was reasonably successful, although with a small sample size it is difficult to detect differences in means. Significantly more immigrant farmers are found in the control group, but the other covariates are relatively similar. Noteably, in both groups, just under 20 percent reported that they adopted use of a pesticide from a technician in the past. In addition, farmers in both groups appear to have similar risk attitudes as was measured in a simple betting game. We also see that farmers in both groups returned just under 40 percent of what they received as trustee in the trust game; close to 70 percent of farmers in both groups reported volunteering in the community.

Table 4.1: Summary statistics for control and treatment groups

	Control		Treatment		
Variable	Mean	S.D.	Mean	S.D.	P-Val.
Demographic and Household Controls					
Education (highest grade completed)	6.76	(3.39)	6.36	(2.91)	0.455
Completed primary school (d)	0.8	(0.048)	0.8	(0.048)	1.00
Age	37.71	(14.36)	40.23	(13.15)	0.282
Male (d)	0.79	(0.41)	0.70	(0.46)	0.249
Indigenous (d)	0.37	(0.49)	0.26	(0.44)	0.147
Catholic (d)	0.26	(0.44)	0.30	(0.46)	0.575
Immigrant (d)	0.14	(0.35)	0.04	(0.20)	0.045
Total household members	6.46	(3.05)	6.07	(2.86)	0.442
Wealth index (integers 0-5)	2.66	(1.79)	2.90	(1.73)	0.416
Farm Controls					
Reported owning land (d)	0.53	(0.50)	0.66	(0.48)	0.123
Farm size (hectares)	45.21	(29.40)	40.62	(29.69)	0.360
Previous adoption from technician (d)	0.16	(0.37)	0.20	(0.40)	0.511
Reported methamidophos use (d)	0.41	(0.50)	0.34	(0.48)	0.387
Reported carbofuran use (d)	0.31	(0.47)	0.20	(0.40)	0.124
Reported chlorothalonil use (d)	0.11	(0.32)	0.11	(0.32)	1.000
Reported 2, 4-D butyl ester use (d)	0.04	(0.20)	0.11	(0.32)	0.118
Risk, Altruism and Experiment Controls					
Bet in risk game (integers 0-5)	2.27	(1.46)	2.56	(1.44)	0.247
Average share returned in trust game	0.38	(0.13)	0.38	(0.15)	0.979
Reported volunteering in community (d)	0.65	(0.06)	0.72	(0.05)	0.362
Experiment session size	18.21	(5.19)	19.52	(5.18)	(0.136)
Technician to farmer ratio	0.48	(0.16)	0.44	(.14)	(0.121)
Observations	70		70		

Note: (d) indicates a dummy variable. The first and second columns present the mean and standard deviation of the control sample; the third and fourth columns present the mean and standard deviation of the treatment sample. The fifth column reports the p-value from a simple comparison of means t-test.

4.4.3 Summary Statistics of Trust Measures

The trust game results from the data used in this paper are similar to those of Schechter (2007), who played the same trust and risk games in Paraguay. An important difference between these two studies is that the Ecuador study has farmers play the trust game with technicians in addition to playing the trust game with members of their own community. In the Paraguay study, investors sent about 46 percent of their initial endowment, while in the Ecuador study the investors sent about 45 percent of their initial endowment to other community farmers and about 46 percent to technicians. Schechter's study, like in Ecuador, found that trust did pay; trustees returned about 38 percent of what they received. A Wilcoxon-Signed Rank Test fails to reject the null hypothesis that the median value sent to community members in this experiment was identical to the median value sent to technicians; however, we find that only 32 percent of the farmers played the trust game the same with a technician as they did with a farmer from their community. Moreover, an Epps-Singleton (Epps and Singleton, 1986) test indicates that the difference in distribution of money sent is significant at the 5 percent level (p-value 0.023). This suggests that farmers do not view the trust game against community farmers as the same as the trust game against technicians.

Table 4.2: Description of trust measures

Measures of trust	Description of variable	Control	Treatment	P-Val.
Trust in community farmers (0-5)	Amount sent by the investor to an anonymous community farmer.	2.13 (1.10)	2.39 (1.32)	0.21
Trust in technicians (0-5)	Amount sent by the investor to an anonymous technician.	2.19 (1.50)	2.41 (1.73)	0.36
Conditional trust in community farmers (0-5)	Trust in community farmers if trust in technician is more than in community farmers; otherwise, zero.	0.60 (1.03)	0.47 (1.07)	0.47
Conditional trust in technicians (0-5)	Trust in technicians if trust in technician is more than in community farmers; otherwise, zero.	1.23 (1.83)	1.10 (1.83)	0.68
Difference in trust (-5-5)	Trust in technicians minus trust in community members.	0.057 (1.74)	0.028 (1.79)	0.92
Conditional difference in trust (-5-5)	Trust in technicians minus trust in community members if trust in technician is more than in community farmers; otherwise, zero.	0.63 (1.18)	0.63 (1.13)	0.99
Trust technicians more than farmers (0-1)	A dummy variable equal to 1 if the investor sent more to technician than to the community farmer; 0 otherwise.	0.34 (0.48)	0.29 (0.46)	0.47

Notes: We report the range of values the given trust measure may assume in parantheses next to the variable name. We report the mean value for each trust measure according to control and treatment groups. Sample standard deviations are reported in parantheses beneath the mean values. The final column reports the p-value from a comparison of means t-test between the control and treatment groups.

In table 4.2 we see farmers sent slightly more in the trust game in the treatment group than in the control group, that is, after attending the training with the agricultural technician and farmers.

This difference in play between treatment and control observations is not statistically significant. Furthermore, this difference in play between treatment and control groups is observed both when farmers played with technicians and with other farmers. If there is an undetectable effect of treatment on play as investor in the trust game, then the effect does not appear different for investor play with technicians compared to play with other farmers. We also explore the relationship between trust attitudes and treatment status using regression, none of the trust attitudes show a statistically significant relationship with treatment status. This is evidence that trust attitudes were not affected by the treatment, and the treatment can thus be used to identify the effect of trust on learning during training.

However, only comparing means across treatment and control groups can be misleading since there could be a heterogeneous response to training. For example, some farmers may find the training very useful, which may increase their trust in technicians. Other farmers may be dissatisfied with the training and trust technicians less. As a result we may observe no differences in trust levels between treatment and control groups. One test for a mean preserving heterogeneous response to treatment is to compare variances in trust levels across treatment groups. If we observe significantly different variances this may suggest heterogeneous responses. Figure 4.1 presents histograms of trust measures we use in our subsequent analysis. The bars in these histograms represent the distributions of trust attitudes for the whole sample (all). The solid lines represent the distributions of trust attitudes for the treatment group and the dashed lines represent the distributions for the control group. Figure 4.1: 1(a) depicts the distribution of trust in technicians, and indicates that 19 percent of investors sent nothing to technicians in the trust game while 10 percent sent their entire endowment to technicians. Figure 4.1: 1(b) shows the distribution of the farmers' difference in trust between technicians and community farmers, and provides evidence that farmers did not treat the two trust games identically. As mentioned before, only 32 percent of investors played identically in the two games and just over 30 percent had a difference in play of two dollars or more. Figure 4.1: 1(c) indicates that 70 percent of investors sent at least as much to community farmers as they sent to technicians; 30 percent trusted technicians more than community farmers. In figure 4.1: 1(d) we see the distribution of what we call the conditional difference in trust; consistent with 1(c), 70 percent of investors sent as much to community farmers in the trust game as they sent to technicians. About 15 percent of investors sent one more dollar to technicians than to community farmers; just under 10 percent of investors sent two more dollars to technicians in the trust game than to community farmers.

In summary, figures 4.1: 1(a)-1(d) show that the distributions of different trust measures are similar across treatment and control groups; and in table 4.1 the means are not statistically different in the treatment and control groups, and the variances are nearly identical. Taken together, we have evidence that this one training interaction with an agricultural technician did not have an effect on the farmers' trust attitudes.

Figure 4.1: Histograms of trust measures

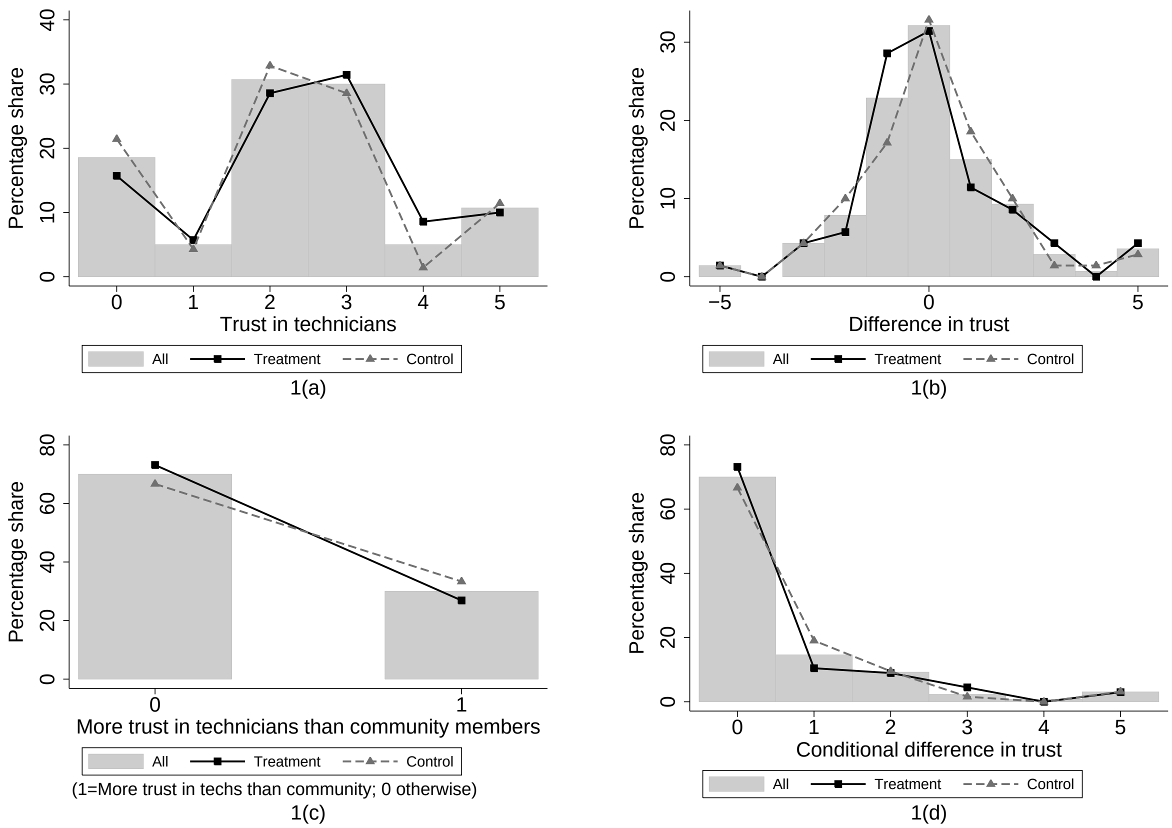

4.4.4 Summary of Exam Scores

Methamidophos and carbofuran (pesticides), chlorothalonil (fungicide) and 2, 4-D butyl ester (herbicide) are the four most common chemicals used in the area. In the 2005 survey, farmers were asked twenty-four yes or no questions about these four agrochemicals. We computed an exam score for each farmer based on the number of answers consistent with the information presented in the training. The average exam score was 48.2 percent (s.d. 16.9 percentage points); the average exam score for those in the treatment group was 49.7 percent (s.d. 18.4 percentage points) and the average exam score for those in the control group was 46.5 percent (s.d. 15.3 percentage points). A simple comparison of means t-test shows that the treatment was not successful on its own (with a p-value of 0.263), although this does not preclude treatment effects for specific sub-groups such as farmers with more trust in technicians.

4.5 Empirical Framework

4.5.1 Estimation Strategy

Our first empirical prediction is that farmers choose their optimal level of effort to learn during training depending on their trust in technicians. Due to diminishing marginal returns to effort, learning during training may also depend on how the farmer trusts other farmers. Accordingly, the first testable hypotheses are that conditional on being in treatment, a farmer's absolute trust in technicians has a positive effect on exam scores; a farmer's absolute trust in community farmers has a negative effect on exam scores[3]. Furthermore, in a placebo-type exercise we test that absolute trust in either technicians or community farmers has no effect on exam scores for those in the control group. Absolute trust in technicians ("Trust in technicians") is measured by the amount of money a farmer sends to the technician in the trust game; absolute trust in community farmers is measured in a similar way.

However, since community farmers and technicians may offer advice that is at odds, learning during training may depend on whether a farmer i's trust in technicians, represented by θ_i^t, is greater than some threshold $\bar{\theta}_i$. With $\bar{\theta}_i = \theta_i^f$, trust attitudes affect learning during training if farmer i trusts technicians more than community farmers. To test this empirically we interact each of our trust measures with an indicator variable for whether the farmer trusts technicians more than community farmers; we reference these as the conditional trust measures. Conditioned on trusting technicians more than community members, trust in technicians should increase learning and trust in farmers decrease learnining. As a sensitivity analysis we re-do the analysis letting $\bar{\theta}_i$ assume different values.

We use Ordinary Least Squares (OLS) estimation to test our hypotheses. To test each of our hypotheses we take advantage of the control group to form the counterfactual of how, ceteris paribus, farmers with similar trust attitudes would have scored had they not been treated (i.e., attended the training provided by the technician prior to responding to the questions in our survey).

Our main estimating equation is:

$$\text{Exam}_{ik} = \alpha_0 + \alpha_1 \text{T}_i + \alpha_2 \widetilde{\text{Trust}}_i^{tech} + \alpha_3 \text{T}_i \cdot \widetilde{\text{Trust}}_i^{tech} + \alpha_4 \widetilde{\text{Trust}}_i^{farm} + \alpha_5 \text{T}_i \cdot \widetilde{\text{Trust}}_i^{farm} + \eta_k + \varepsilon_{ik}. \tag{4.14}$$

where T_i takes on 1 if the farmer i is in treatment and zero otherwise. $\widetilde{\text{Trust}}_i^s$ represents the trust measures of interest with i indicating the farmer and s indicating the information source, and η_k is a community-level fixed effect.

The measures we obtain from the trust game are likely imperfect since they probably measure idiosyncrasies specific to each farmer. In this case our trust measures can be written as the sum of farmer i's true underlying trust in trustee s and an idiosyncratic component specific to farmer i, μ_i, yet common across play with technician and farmer trustees. Our measure of farmer i's trust in a trustee can be written as $\widetilde{\text{Trust}}_i^s = \text{Trust}_i^s + \mu_i$. Also, there may be a component common to both trust in technicians and community farmers, ν_l, due to the ordering of events for treatment and control groups. For instance, farmers may have behaved in a slightly more trusting fashion later in the day as is suggested by the data so that $\widetilde{\text{Trust}}_i^s = \text{Trust}_i^s + \mu_i + \nu_l$. If μ_i and/or ν_l are uncorrelated

[3]One can argue that total knowledge will be increasing in trust in peers because farmers who have learned from their peers are likely to have higher levels of knowledge. However, we assume an empirical measure of knowledge that is specific to the intervention so that past learning is not tested.

with the outcome variable as is likely the case, then our estimates from equation (4.14) will be attenuated and imprecise. While measurement error in an independent variable attenuates the point estimate in the bivariate case, this may not be borne out if there are multiple independent variables; however, the standard errors will necessarily increase when measurement error is present. Fortunately, our theoretical model also delivers the result that as difference in trust increases then so should learning during training, and we use this fact to sharpen our analysis. Taking the difference in trust between technicians and community farmers eliminates μ_i and ν_l so that we have: $\widetilde{\text{Trust}}_i^{tech} - \widetilde{\text{Trust}}_i^{farm} = \text{Trust}_i^{tech} + \mu_i + \nu_l - \text{Trust}_i^{farm} - \mu_i - \nu_l = \text{Trust}_i^{tech} - \text{Trust}_i^{farm}$. Based on this we estimate equation (4.14) using difference in trust, $\text{Trust}_i^{tech} - \text{Trust}_i^{farm}$, rather than estimating trust in technicians and community farmers separately; this estimation should address the described measurement error. The drawback to this approach is that estimating the relationship between difference in trust and learning during training is more restrictive since it implies a linear relationship in trust in technicians that is equal and opposite in sign to that of trust in community farmers. We will see the data indicate such a restriction is plausible.

4.5.2 Sensitivity analysis and sources of potential bias

To probe the sensitivity of our results, we run numerous robustness checks. Of particular concern is that trust attitudes may be correlated with an omitted variable such as ability to learn during training. However, our results are not sensitive to the inclusion of educational attainment in our regression analysis, which is arguably a good proxy for ability to learn in a training setting[4]. We also estimate alternative specifications of equation (4.14) that include demographic and household controls: farmer's age, age squared, gender, membership to an indigenous group, membership to the Catholic church, immigrant status, household size, and a wealth index. In addition,we control for farm characteristics such as whether or not the farmer reports owning land, the amount of land in hectares, pesticide-use controls including whether the farmer reported adopting a pesticide from a technician in the past, and whether the farmer reported currently using the agrochemicals mentioned on the exam. Next, we add measures of risk (amount bet in the risk game) and altruism (average share returned as the second mover in the trust game and volunteerism in the community). Finally, we control for characteristics of each experimental session, such as session size and the ratio of technicians to farmers.

Our list of control variables is, of course, not exhaustive. For instance, trust in the competence of an information source, as discussed earlier, is also likely to affect a farmer's decision on who they should rely for advice. Our argument is that trust in competence is highly correlated with education level. Since there is little variation in education level among technicians and because each community only has one aggregate educational endowment we we assume farmers have homogeneous trust in competence for each information source. However, if trust in competence is heterogeneous and is both positively correlated with learning during training and trust in motives (as measured by the trust game), then our point estimates on trust will likely be upward biased. We know of no evidence suggesting that investor play in the trust game measures trust in competence

[4]We consider two measures of educational attainment in the regression analysis: 1) An indicator variable denoting whether the farmer completed primary schooling (6th grade), and 2) a finer measure, the highest grade level completed. Because of measurement error in the second measure we only present results using the first measure; however, the main results are invariant to the education variable used.

rather than or in addition to trust in motives. Furthermore, anecdotal observation suggests there are all types of people in the world: competent people steal and so do incompetent ones; similarly, incompetent people return lost wallets and so do competent ones. Nonetheless, if we do not believe homogeneity of trust in competence for our sample or orthogonality of trust in motives and trust in competence then our estimation strategy is still vulnerable to omitted variable bias.

Finally, two concerns with the research design which we have yet to address relate to the ordering of trust games and the fact that the technicians who participated in the experiment were not physically present. Studies in experimental economics have shown that the ordering of games can affect decision-making as can other details of the research design. If either the ordering of the games or the physical absence of participating technicians affects play then there is measurment error in our measure of trust in technicians. However, it is reasonable to assume that any measurement error induced by these design details are unrelated to underlying trust attitudes or learning during training. If the described measurement error exists then it would attenuate our point estimates on trust and increase the standard errors–these would work against finding a positive result.

4.5.3 Standard Error and Pivotal Statistic Adjustments

Statistical inference is complicated due to the clustered structure of the data. Observations with-in a training session may be correlated so that it is important to account for the possibility of intra-cluster correlation. Per Bertrand, Duflo, and Mullainathan (2004) clustering the standard errors is the common way to address this concern; however, they show that when there are a small number of clusters (ten or less), the performance of statistical inference using cluster-robust standard errors is unreliable. Cameron, Gelbach, and Miller (2008) find the same result and try a variety of adjustments to standard errors in a monte carlo simulation, including several bootstrapping methods, to evaluate which adjustments offer reliable performance when making statistical inference. They suggest using the wild cluster-bootstrap method to obtain pivotal t-statistics as this offers asymptotic refinement. We follow their advice, and compute a wild cluster-bootstrap pivotal t-statistic for each regression coefficient; then the analytically computed cluster-robust standard error is used with the adjusted pivotal t-statistic to perform hypothesis testing.

4.6 Results and Discussion

Our results are presented in tables 4.3 and 4.4; the dependent variable for all regressions is "Exam score". Point estimates for the covariates of interest are presented, while the results of the numerous control variables are suppressed. The cluster-robust standard errors are presented in parantheses beneath the point estimates and the wild cluster-bootstrap pivotal statistics used for hypothesis testing at the 95% confidence level are presented in square brackets beneath these standard errors. If we replace these with heteroskedastic-consistent robust standard errors and our usual t-statistics the results reported in tables 4.3 and 4.4 are similar. In table 4.3 we present the results from regressing the exam score on trust in technicians and trust in community farmers, while in table 4.4 we look at the difference in trust or a dummy variable indicating if the farmer trusted technicians more than community farmers in the trust game.

When interpreting the regression coefficients the total effect of trust on learning is given by the sum of the point estimate on trust and on trust interacted with treatment. However, from the policy perspective we are interested in the coefficient on trust interacted with treatment alone to assess potential policy impacts. The coefficient on trust by itself is only a check that those with high levels of trust did not have greater knowledge on material tested in the exam before attending the field day. Furthermore, we do not believe an exogenous shock to farmers' trust attitudes would impact knowledge of agricultural practices without the opportunity to excercise trust. For example, if an extension office were to randomly assign a technician to attend community festivals in community X then this may likely increase trust in the attending technician. However, the technician's mere trust-inducing presence at the festival would have no effect on knowledge. On the other hand, if an opportunity to discuss agrochemicals were to arise between a technician and a community farmer at the community festival then increased trust may lead to more knowledge transfer. With this in mind we turn our attention to a more detailed discussion of tables 4.3 and 4.4.

Moving across the columns in table 4.3, we see models with different sets of covariates and different measures of trust. We first consider the point estimates on "Treatment". Consistent with our comparison of means t-test for exam scores between treatment and control groups, the point estimates on "Treatment" are largely insignificant. Some of the point estimates on "Treatment" appear negative and large in size, although this is artificial since increases in the treatment point estimate are almost perfectly off-set by the point estimate of education interacted with treatment. We now draw attention to the point estimates on "Trust in technicians" and "Trust in farmers". We see that these trust attitudes by themselves are not significantly related to performance on the exam. This result indicates that trust attitudes are not proxying for prior knowledge of exam material; more trusting farmers do not come to the field day with significantly higher or lower levels of knowledge. Therefore, we have evidence against reverse causality; it is an unlikely story that past experiences with technicians presenting the same exam material affected trust attitudes. Consistent with this finding, farmers reported that they had not been to a similar naranjilla training in the past.

Next we direct our attention to the point estimates associated with the variable label "Trust in tech*Treatment" in table 4.4. The first three columns (A.1)-(A.3) of table 4.3 test the model described in section 4.3.1 based on unconditional trust attitudes. We see in column (A.1) that the coefficient estimates are of the expected sign; positive for "Trust in tech*Treatment" and negative for "Trust in farm*Treatment". The point estimate on "Trust in tech*Treatment" is significant at the 95% level. Moving across the columns we check if our results are sensitive to adding covariates. In column (A.2) the coefficient on "Trust in tech*Treatment" is significant at the 95% level; however, in our preferred specification shown in column (A.3), which controls for relevant observables, this significance is lost. In column (A.3) we see a point estimate of 3.339 on "Trust in tech*Treatment" and -2.942 on "Trust in farm*Treatment", which are opposite and approximately equal in sign. Hence, assuming an empirical model replacing these two trust measures with their difference imposes a plausible linear restriction. If there is measurement error in our trust measures as described in section 4.5.1 then re-estimating columns (A.1)-(A.3) may improve the precision of our estimates.

Table 4.3: Regress exam score on trust measures

Dependent variable: Exam Score (Mean: 48.15, S.D.: 16.93)						
Trust Measure	A: Unconditional Trust			B: Conditional Trust		
	(A.1)	(A.2)	(A.3)	(B.1)	(B.2)	(B.3)
Treatment	.336	-9.522	-5.773	-1.098	-11.851	-9.976
	(8.405)	(10.615)	(12.015)	(1.033)	(7.195)	(7.781)
	[2.384]	[2.432]	[2.599]	[2.128]	[3.816]	[2.242]
Trust in technicians	-1.338	-1.386	-.824	-.740	-.886	-.266
	(.629)	(.566)**	(.822)	(1.314)	(1.163)	(1.083)
	[3.616]	[2.334]	[2.079]	[1.664]	[1.526]	[1.736]
Trust in tech*Treatment	3.177	3.479	3.339	2.876	4.050	3.993
	(1.777)*	(1.795)*	(2.414)	(2.004)	(1.833)**	(1.546)**
	[1.879]	[2.024]	[2.263]	[2.355]	[1.875]	[2.051]
Trust in farmers	-.471	-.483	-.409	-.835	-.766	-1.519
	(2.117)	(2.193)	(1.954)	(2.530)	(2.425)	(2.056)
	[3.858]	[3.562]	[3.244]	[1.919]	[1.802]	[1.893]
Trust in farm*Treatment	-1.922	-2.041	-2.942	1.586	-.453	-.562
	(2.571)	(2.525)	(1.871)	(3.509)	(3.277)	(3.642)
	[2.619]	[2.574]	[2.327]	[2.550]	[2.188]	[2.500]
Completed primary school		-1.351	-4.597		-2.012	-5.199
		(3.159)	(6.082)		(3.291)	(5.987)
		[1.326]	[2.031]		[2.881]	[2.551]
Primary edu*Treatment		11.845	13.411		13.078	14.387
		(8.182)	(9.261)		(8.843)	(9.969)
		[3.852]	[2.782]		[5.051]	[2.280]
P-val.–tot. effect of trust in tech	0.319	0.271	0.237	0.210	0.061	0.036
Obs.	140	140	140	140	140	140
R^2	.157	.188	.262	.167	.199	.265
Community FE	Yes	Yes	Yes	Yes	Yes	Yes
Demo. and HH controls	No	No	Yes	No	No	Yes
Farm controls	No	No	Yes	No	No	Yes
Risk and altruism	No	No	Yes	No	No	Yes
Session controls	No	No	Yes	No	No	Yes

Notes: The dependent variable in each regression is exam score. Moving across the columns, we present models with different sets of covariates and different measures of trust. Columns (A.1)-(A.3) test the model described in section 4.3.1 based on unconditional absolute trust attitudes. Subsequent columns test the model described in section 4.3.2 based on conditional absolute trust attitudes. The "Education" measure used in these regressions is a indicator variable denoting completion of 6th grade (primary schooling). Robust standard errors clustered at the experiment session level are reported in parantheses beneath the reported coefficient estimates; wild cluster-bootstap pivotal t-statistics are reported in square brackets beneath the reported standard errors. *-significance at the 90% confidence level, **-significance at the 95% confidence level and ***-significance at the 99% confidence level.

Subsequent columns test the model described in section 4.3.2 based on conditional trust attitudes. The magnitude of the coefficients on "Trust in tech*Treatment" are slightly larger and are more precisely estimated than the point estimates from the first three columns; the point estimate in our preferred specification (B.3) is significant at the 95% level. The point estimate on "Trust in farm*Treatment" in our preferred specification is negative but close to zero. This result could be an artifact of the selection process: trust in farmers is on average less for farmers who trust technicians more than community farmers. According to our model, farmers who trust community farmers less would exert less effort to learn from them. Depending on the shape of the relationship between trust and effort, any negative effect of trust in community farmers on effort to learn during training that is observed in columns (A.1)-(A.3) table 4.3 may be diminished for the subsample of farmers who trust technicians more than community farmers. We would not necessarily expect the point estimates for "Trust in tech*Treatment" and "Trust in farm*Treatment" to be equal in magnitude; this is inconsistent with using the difference in trust specification. However, the point estimates on "Trust in farm*Treatment" are unstable and imprecisely estimated, which may be due to the sparsity of data in the subsample of farmers who trust technicians more than community farmers, rather than representing true differences in the magnitudes of the effects between "Trust in tech*Treatment" and "Trust in farm*Treatment". If this is the case, then re-estimating columns (B.1)-(B.3) using the conditional difference in trust measure may yield more precise estimates; we will check this in the next table of results.

Table 4.4 reports the results using the relative trust measures. In columns (A.1)-(A.3) we see the point estimates from regressing the exam score on the unconditional difference in trust. The point estimate on "Relative trust*Treatment" is more precisely estimated than the corresponding two estimates on the unconditional trust attitudes from table 4.3. Indeed, the point estimate is significant at the 95% level in all specifications, and the point estimates increase in magnitude as we add covariates. Columns (B.1)-(B.3) of table 4.4 show the point estimates from regressing the exam score on the conditional difference in trust. These estimates are slightly larger and more precisely estimated than those for the unconditional difference in trust. Together, columns (A.1)-(B.3) are consistent with our thoughts on measurement error in the trust variables.

In the final three columns we replace our relative trust measure with an indicator variable denoting whether a farmer trusted technicians more than farmers in the trust game. If we believe measurement error remains in our trust measures then this is arguably the simplest measure we can construct to consider relative trust attitudes, and so may remove noise from our explanatory variable. This measure of relative trust is also appealing because of its comparability to our indicator variable denoting whether a farmer completed primary education. We observe that farmers who trust technicians more than community farmers score approximately 15 percentage points higher on the exam. These point estimates are precisely estimated and significant at the 95% confidence level.

We note that the point estimates between the difference in trust measures and the indicator variable denoting more trust in technicians than community farmers are different in size; this makes sense because these measures take on different ranges of values. After normalizing the results, the point estimates are all of a similar magnitude. Furthermore, the effect of these trust attitudes on

Table 4.4: Regress exam score on relative trust measures

	Dependent variable: Exam Score (Mean: 48.15, S.D.: 16.93)								
Trust Measure	A: Unconditional Difference in Trust			B: Conditional Difference in Trust			C: Trust tech more than farmers		
	(A.1)	(A.2)	(A.3)	(B.1)	(B.2)	(B.3)	(C.1)	(C.2)	(C.3)
Treatment	3.015	-5.996	-4.853	.539	-11.190	-9.382	-1.073	-10.848	-8.865
	(1.352)	(6.650)	(7.308)	(1.439)	(6.917)	(7.598)	(.930)	(6.963)	(7.553)
	[2.747]	[3.663]	[2.234]	[2.761]	[4.284]	[2.560]	[1.779]	[3.609]	[2.353]
Relative Trust	-.797	-.826	-.487	-1.380	-1.526	-.897	-3.030	-3.300	-2.674
	(.846)	(.827)	(.895)	(.802)	(.738)**	(.850)	(1.706)	(1.759)	(1.833)
	[5.819]	[4.325]	[2.402]	[2.164]	[1.857]	[1.682]	[2.591]	[2.625]	[2.002]
Relative Trust*Treatment	2.885	3.116	3.369	4.058	5.154	4.897	13.860	14.767	15.219
	(1.105)**	(1.223)***	(1.306)***	(1.768)*	(1.648)***	(1.459)***	(4.554)***	(5.185)***	(5.667)***
	[2.040]	[1.842]	[1.916]	[2.705]	[2.042]	[1.927]	[2.452]	[2.133]	[1.958]
Completed primary school		-.760	-4.443		-1.602	-4.534		-1.364	-4.640
		(2.847)	(6.074)		(2.861)	(6.005)		(2.955)	(5.660)
		[1.493]	[2.115]		[2.695]	[2.440]		[2.614]	[2.421]
Primary edu*Treatment		11.318	13.079		13.844	14.991		11.950	12.982
		(8.085)	(9.193)		(8.688)	(9.791)		(8.295)	(9.290)
		[4.519]	[2.281]		[5.933]	[2.913]		[4.858]	[2.384]
P-val.–tot. effect of rel. trust	0.040	0.053	0.023	0.134	0.041	0.021	0.036	0.048	0.039
Obs.	140	140	140	140	140	140	140	140	140
R^2	.152	.183	.259	.148	.188	.257	.169	.2	.269
Community FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Demo. and HH controls	No	No	Yes	No	No	Yes	No	No	Yes
Farm controls	No	No	Yes	No	No	Yes	No	No	Yes
Risk and altruism	No	No	Yes	No	No	Yes	No	No	Yes
Session controls	No	No	Yes	No	No	Yes	No	No	Yes

Notes: Moving across the columns, we present models with different sets of covariates and different measures of trust. Robust standard errors clustered at the experiment session level are reported in parantheses beneath the reported coefficient estimates; wild cluster-bootstap pivotal t-statistics are reported in square brackets beneath the reported standard errors. *-significance at the 90% confidence level, **-significance at the 95% confidence level and ***-significance at the 99% confidence level.

learning during training is large in magnitude; for example, farmers who trusted technicians more than community members scored almost one standard deviation higher on the exam after attending the training than other farmers.

One may wonder how the choice of threshold affects our results related to the conditional trust measures. Recalling the threshold equation (4.13), we assume different values of β_0 and β_1 and re-estimate our regressions using the conditional trust measures. We obtain results that are qualitatively consistent with those presented in tables 3 and 4, although there are differences in magnitude and standard errors. For example, when we use a threshold based on $\beta_0 = -1$ and $\beta_1 = 1$, using the specification with fixed effects and controls we find a coefficient on conditional trust in technicians of 4.68 with a t-statistic of 2.46 and a coefficient on trust in farmers of -4.23 with a t-statistic of 2.18. These estimates are consistent with the validity of using the conditional difference in trust measure which imposes a linear restriction. Replacing the conditional trust measures in technicians and farmers with conditional difference in trust we find a point estimate of 5.37 with a t-statistic of 5.80; if we include the direct effect of trust then the total effect is still significant at the 95% level. As we would expect, the magnitudes of the point estimates on our conditional trust measures change depending on the threshold; however, the qualitative relationship between trust attitudes and learning during training remains robust during this excercise.

Our focus has been on the point estimates of our trust attitudes interacted with treatment. This reflects the fact that farmers were tested on material presented in the training only, rather than knowledge previously obtained from technicians. Once again, a random shock to trust in technicians would not alone change knowledge about agrochemicals–there must be an opportunity for farmers to [not] place trust in a technician's advice. Nonetheless, one might argue that the point estimate of trust attitudes without the interaction should be included when calculating the impacts of trust, especially since the point estimates on the different trust measures without the interaction are generally negative. If the correlation between trust attitudes alone and exam scores is considered a direct effect then it may negate the benefits of trust during training. To investigate this we test whether the sum of the direct effect of trust plus the effect of trust interacted with treatment is significantly different from zero. The p-values from these tests are presented at the bottom of tables 4.3 and 4.4; we observe that the results remain significant even when taking into acount the possibility of negative direct effects.

Overall, the result is robust, although our finding may still be vulnerable to the criticism of omitted variable bias. Also, given the timing of the trust games relative to treatment there exists the possibility of reverse causality; however, as shown in section 4.4.3 the evidence indicates treatment did not induce differences in play during the trust games.

Finally, there is evidence that those who completed their primary schooling scored higher on the exam after having attended the training. The point estimates on "Edu.*Treatment" in table 4.3 and table 4.4 range from 11.31 to 14.99 depending on the specification, but these point estimates are all imprecisely estimated with none being significant at the 90% confidence level. The effect of completing primary schooling on learning during training is not identified since it is impossible to distinguish between this measure of educational attainment and ability. Nonetheless, the sign of this measure (whether it is capturing education or ability) interacted with treatment is as we would expect; and interestingly, is of a similar (though slightly smaller) magnitude relative to the point estimates on our relative trust measures interacted with treatment. This suggests that while human capital may help a farmer learn, its effect may be more than off-set when the farmer does not trust the institutional environment and those affiliated with it.

4.7 Conclusions

Naranjilla farmers in Ecuador may question the motives of agricultural technicians and community members. Trust in the provider of information may play an important role in knowledge spread particularly in cases where the message is complex. To generate testable predictions, we develop a simple model of farmer learning during agricultural training. Our results are consistent with the theory that trust impacts farmer learning. If valuable information is delivered via institutions or persons in which potential beneficiaries have low trust, then the information dissemination effort may be ineffective.

This finding has two policy implications for a renewed public agricultural extension system in Ecuador. First, effort can be made to identify trusted sources of agricultural information. For example, in the initial phase of renewed extension services, program designers can undertake case studies to gauge the relative trust potential beneficiaries place in various channels by which information could be delivered. Our results indicate that dual information dissemination channels may be preferred to single channels. Using both technicians and community farmers to spread new information may be a good strategy as some farmers will want to learn from technicians while others will prefer to learn from community members. Second, effort can be made to build trust in those channels by which information or services will be delivered. Information delivery programs might be structured to build trust gradually over time. For example, messages that have large and relatively certain impacts on incomes should be delivered first; more risky and complex messages could be delayed until trust grows. The presence of the technician in an area over time should also build trust.

A problem with stationing technicians over many years in remote villages throughout Ecuador is the prohibitive cost. However, new technology may allow agricultural extension services to overcome both trust and cost concerns. For example, in India, Gandhi, Veeraraghavan, and Toyama (2009) report that the non-governmental organization Digital Green is using low-cost digital video recording devices and community-mediated instruction for technology dissemination and training. Digital Green's method is to use a participatory approach for content production and to create locally generated digital video databases featuring trained agricultural technicians and community farmers. Furthermore, in each community a trusted farmer is employed to mediate and encourage discussion of the locally generated training videos. Finally, when Digital Green enters a community they use a carefully devised training schedule to initiate the process; they usually start with relatively simple technologies that provide immedate pay-offs and then move to more complex technologies that have long-term returns (Gandhi, Veeraraghavan, and Toyama, 2009). The advantage is that the local participatory approach fosters trust, while the use of low cost digital video recording devices keeps costs reasonable. Of course, this is a single example of how trust might be addressed, and evidence is required to evaluate effectiveness before a program's method is widely used. Agricultural extension services in Ecuador would be wise to consider piloting training programs before scaling-up operations.

At the present, indications are that local governments in Ecuador will take an increased role in providing agricultural extension services (Barrera et al., 2008), and this may help strengthen trust in the messenger. For the success and longevity of a public agricultural extension system in Ecuador and elsewhere, the importance of installing faith in the renewed system should not be underestimated–it is essential that those who deliver agricultural information are trusted sources.

There are useful ways to build on this current work. Our main dependent variable, knowledge about agrochemical use, is an intermediate input to the production process. It would be interesting to see if trust attitudes and attending training actually changed use of physical inputs, that is, whether knowledge led to action. In addition, it may be worthwhile to collect data that helps trace out the relationship between trust and effort, and then effort to learning and adoption.

Our results are of interest to policy-makers in a wide variety of poverty-reduction programs where trust in program messengers is essential for success; these include projects related to agricultural extension efforts to public health interventions to microfinance programs. Although the ability to make specific policy recommendations from this study is limited, the evidence is provoking and warrants follow-up in the field. Thus, another way to build upon this work is through a randomized field experiment that rolls out agricultural extension via different information sources. This type of research design would have more exacting implications for who should disseminate information (technician versus community farmer) and would help answer other policy relevant questions such as the optimal timing and frequency of visits.

Bibliography

Adegbola, P., and C. Gardebroek. 2007. "The effect of information sources on technology adoption and modifcation decisions." *Agricultural Economics* 37:55–65.

Alwang, J., S. Weller, G. Sanchez, L. Calderon, R. Edwards, S. Hamilton, R. Williams, M. Ellis, C. Suarez, V. Barrera, C. Crissman, and G. Norton. 2005. *Globalizing Integrated Pest Management: A Participatory Research Process, chapter Developing IPM Packages in Latin America.*. Blackwell Press.

Andrade, R. 2005. "Analisis de la Situacion Agro Socio Economica de los Productores de Narajilla." MS thesis, Universidad Catolica de Quito.

Baez, J., J. Ochoa, and M. Ellis. 2003. "Reaction of members of the lasiocarpa section to infection by Fusarium oxysporum causing naranjilla vascular wilt in Ecuador." In J. Alwang, C. Suarez, and V. Barrera, eds. *Overview of the South America Site in Ecuador*. Office of International Research, Education, and Development, Virginia Tech.

Barr, A. 2003. "Trust and Expected Trustworthiness: Experimental Evidence from Zimbabwean Villages." *The Economic Journal* 113:614–630.

Berg, J., J. Dickhaut, and K. McCabe. 1995. "Trust, Reciprocity, and Social History." *Games and Economic Behavior* X:122–142.

Bertrand, M., E. Duflo, and S. Mullainathan. 2004. "How much should we trust differences-in-differences estimates?" *Quarterly Journal of Economics* 119:249–275.

Bouma, J., E. Bulte, and D. Van Soest. 2008. "Trust and cooperation: Social capital and community resource management." *Journal of Environmental Economics and Management* 56(2):155–166.

Brozovic, N., D. Sunding, and D. Zilberman. 2007. "Estimating business and residential water supply interruption losses from catastrophic events." *Water Resources Research* 43.

Camerer, C. 2003. *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton University Press.

Cameron, C., J. Gelbach, and D. Miller. 2008. "Bootstrap-based improvements for inference with clustered errors." *The Review of Economics and Statistics* 90:414–427.

Cayan, D., E. Maurer, M. Dettinger, M. Tyree, and K. Hayhoe. 2008. "Climate change scenarios for the California region." *Climatic Change* 87:21–42.

Chaudhury, N., J. Hammer, M. Kremer, K. Muralidharan, and H. Rogers. 2006. "Missing in Action: Teacher and Health Worker Absence in Developing Countries." *Journal of Economic Perspectives* 20(1):91–116.

Chay, K., and M. Greenstone. 2005. "Does Air Quality Matter? Evidence from the Housing Market." *Journal of Political Economy* 113:49.

Chong, W., T. Phipps, and L. Anselin. 2003. "Measuring the benefits of air quality improvements: a spatial hedonic approach." *Journal of Environmental Economics and Management* 45:24–39.

Chung, F.e.a. 2005. "Progress on Incorporating Climate Chnage into Management of California's Water Resources." Working paper, (www.climatechange.ca.gov).

Cole, S., X. Gine, J. Tobacman, P. Topalova, R. Townsend, and J. Vickery. 2009. "Barriers to household risk management: Evidence from India." *Federal Reserve Bank of New York: Staff Report no 373*, pp. .

Crouter, J. 1987. "Hedonic estimation applied to a water rights market." *Land Economics* 63:259–271.

Duflo, E., M. Kremer, and J. Robinson. 2009. "Nudging farmers to use fertilizer: theory and experimental evidence from Kenya.", pp. .

Easterly, W. 2009. "Can the West Save Africa?" *Journal of Economic Literature*, pp. , Forthcoming.

Eckel, C., and R. Wilson. 2004. "Is Trust a Risky Decision?" *Journal of Economic Behavior and Organization* 55:447–465.

Epps, T., and K. Singleton. 1986. "An Omnibus Test for the Two Sample Problem Using the Empirical Characteristic Function." *Journal of Statistical Computation and Simulation* 26:177–203.

Faux, J., and G. Perry. 1999. "Estimating irrigation water value using hedonic price analysis: A case study in Malhuer County, Oregon." *Land Economics* 75:440–452.

Foley, J.e.a. 2005. "Global consequences of land use." *Science* 309:570–574.

Gallegos, P., C. Asaquibay, L. Shiki, and R. Williams. 2003. "Integrated Pest Management Methodologies to Control Barrenador del Fruto in Fruits of Naranjilla." In J. Alwang, C. Suarez, and V. Barrera, eds. *Overview of the South America Site in Ecuador*. Office of International Research, Education, and Development, Virginia Tech.

Gambetta, D. 1988. *Can We Trust? In: Trust: Making and Breaking Cooperative Relations*, D. Gambetta, ed. Blackwell.

Gandhi, R., R. Veeraraghavan, and K. Toyama. 2009. "Digital Green: Participatory Video and Mediated Instruction for Agricultural Extension." *Information Technologies & International Development* 5:1–15.

Glaeser, E., D. Laibson, J. Scheinkman, and C. Soutter. 2000. "Measuring Trust*." *Quarterly Journal of Economics* 115:811–846.

Godtland, E., E. Sadoulet, A. De Janvry, R. Murgai, and O. Ortiz. 2004. "The Impact of Farmer-Field-Schools on Knowledge and Productivity: A Study of Potato Farmers in the Peruvian Andes." *Economic Development and Cultural Change* 52:129–58.

Greenstone, M., and J. Gallagher. 2008. "Does hazardous waste matter? Evidence from the housing market and the Superfund Program." *Quarterly Journal of Economics* 123:951–1003.

Hanemann, M., L. Dale, S. Vicuna, D. Bickett, and Dyckman. 2006. "The economic cost of climate change impact on California water: A scenario analysis." Working paper, California Energy Commission CEC-500-2006-003.

Hartman, L., and R. Anderson. 1962. "Estimating the value of irrigation water from farm sales data in Northeastern Colorado." *Journal of Farm Economics* 44:207–213.

Hayhoe, K.e.a. 2004. "Emission pathways, climate change, and impacts on California." *Proceedings of the National Academy of Sciences* 101:12322–12427.

Karlan, D. 2005. "Using Experimental Economics to Measure Social Capital and Predict Financial Decisions." *American Economic Review* 95:1688–1699.

Kuminoff, N.V., C.F. Parmeter, and J.C. Pope. 2010. "Which hedonic models can we trust to recover the marginal willingness to pay for environmental amenities?" *Journal of Environmental Economics and Management* 60:145–160.

Labie, M., P. Méon, and A. Szafarz. 2009. "Discrimination in Microfinance: The Role of Credit Officers." *Working Papers CEB*, pp. .

Leggett, C., and N. Bockstael. 2000. "Evidence of the Effects of Water Quality on Residential Land Prices." *Journal of Environmental Economics and Management* 39:121–144.

Liu, E. 2010. "Time to change what to sow: risk preferences and technology adoption decisions of cotton farmers in China." *PhD diss. Princeton University*, pp. .

McDonald, R.e.a. 2011. "Urban growth, climate change, and freshwater availability." *Proceedings of the National Academy of Sciences* 108:6312–6317.

Mendelsohn, R., W.D. Nordhaus, and D. Shaw. 1994. "The impact of global warming on agriculture: A Ricardian analysis." *American Economic Review* 84:753–771.

Moore, M., A. Mulville, and M. Weinberg. 1996. "Water allocation in the American West: Endangered fish versus irrigated agriculture." *Nat. Resources J.* 36:319.

Moser, C., and C. Barrett. 2006. "The complex dynamics of smallholder technology adoption: The case of SRI in Madagascar." *Agricultural Economics* 35:373–388.

Norton, G., E. Heinrichs, and G. Luther. 2005. *Globalizing integrated pest management: a participatory research process*. Wiley-Blackwell.

Ochoa, J., B. Yangari, V. Galarza, J. Fiallos, and M. Ellis. 2001. "Vascular wilt of common naranjilla (Solanum quitoense) caused by Fusarium oxysporum in Ecuador." Working paper, Plant Health Progress.

Petrie, R., and L. Taylor. 2007. "Estimating the value of water use permits: A hedonic approach applied to farmland in the Southeastern United States." *Land Economics* 83:302–318.

Rosen, S. 1974. "Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition." *Land Economics* 82:34–55.

Schechter, L. 2007. "Traditional Trust Measurement and the Risk Confound: An Experiment in Rural Paraguay." *Journal of Economic Behavior and Organization* 62:272–292.

Schlenker, W., W.M. Hanemann, and A.C. Fisher. 2006. "The impact of global warming on U.S. agriculture: An econometric analysis of optimal growing conditions." *Review of Economics and Statistics* 88:113–125.

—. 2007. "Water availability, degree days, and the potential impact of climate change on irrigated agriculture in California." *Climatic Change* 81:19–38.

Selby, H. 1945. "Factors affecting value of land and water in irrigated land." *Journal of Land and Public Utility Economics* 21:250–258.

Shiki, L., J. Viteri, J. Ochoa, R. Williams, and M. Ellis. 2003. "Developing IPM programs for naranjilla in Ecuador." In J. Alwang, C. Suarez, and V. Barrera, eds. *Overview of the South America Site in Ecuador*. Office of International Research, Education, and Development, Virginia Tech.

Tu, Q., and E. Bulte. 2010. "Trust, Market Participation and Economic Outcomes: Evidence from Rural China." *World Development*, pp. .

Vicuna, S. 2006. "Predictions of Climate Change Impact on California Water Resources Using CALSIM-II: A Technical Note." Working paper, California Climate Change Center CEC-500-2005-200-SF.

Zilberman, D., and F. Castillo. 1994. "Economic and Health Consequences of Pesticide Use in Developing Country Agriculture: Discussion." *American Journal of Agricultural Economics* 76:603–604.