

DISSERTATION

THREE ESSAYS ON THE ECONOMICS OF WATER RESOURCES AND CLIMATE
CHANGE

Submitted by

Joey Blumberg

Department of Agricultural and Resource Economics

In partial fulfillment of the requirements

For the Degree of Doctor of Philosophy

Colorado State University

Fort Collins, Colorado

Spring 2024

Doctoral Committee:

Advisor: Chris Goemans

Dale Manning
Jesse Burkhardt
Mazdak Arabi

Copyright by Joey Blumberg 2024

All Rights Reserved

ABSTRACT

THREE ESSAYS ON THE ECONOMICS OF WATER RESOURCES AND CLIMATE CHANGE

This dissertation investigates three broad topics in the economics of water resources and climate change. In the first chapter, I explore how changes in perceptions about water availability affect the adoption of conservation practices. I present a theoretical framework to examine how producer perceptions influence investment in irrigation efficiency, and a period of extreme drought and institutional change in Colorado is leveraged as a natural experiment to evaluate theoretical hypotheses empirically. The second chapter assesses the sensitivity of climate change impact estimates to the climate-economy functional form in agriculture. I accomplish this through the development of a long-run dataset of county-level weather and climate metrics, including hourly temperature measurements across all counties in the conterminous US, and demonstrate the consequences of multiple modelling approaches that are common in the literature. I also create a composite vulnerability index that integrates the magnitude and consistency of impacts across all defensible models to generate a comprehensive measure of climate risk to a county's agricultural sector. In the final chapter, I compare the economic efficiency of different water allocation mechanisms. A combination of optimization models and water supply simulations are employed to compare prior appropriation with and without water markets, and alternative, share-based mechanisms. I illustrate how the physical and institutional components of a river basin, such return flows, user seniority, heterogeneous value functions, and user locations (i.e., upstream or downstream), impact allocative performance.

ACKNOWLEDGEMENTS

The material in this dissertation is based upon work supported by the National Institute for Food and Agriculture under Award No. 2018-69011-28369 and the National Science Foundation under Grant No. 1828902.

TABLE OF CONTENTS

ABSTRACT	ii
ACKNOWLEDGEMENTS	iii
LIST OF TABLES	vi
LIST OF FIGURES	vii
 Chapter 1	
Perceived Water Scarcity and Irrigation Technology Adoption	1
1.1 Introduction	1
1.2 Theoretical Framework	5
1.2.1 Parameter Simulations	9
1.3 Data and Methods	10
1.3.1 Study Area	10
1.3.2 The Natural Experiment	12
1.4 Results	16
1.5 Conclusion and Policy Implications	20
1.6 Tables	25
1.7 Figures	29
1.8 Supplementary Analysis	36
1.8.1 Scrutinizing the Treatment Design	36
1.8.2 Addressing Threats to Identification	36
1.8.3 Nonlinear Impacts of Changing Perceptions	38
 Chapter 2	
Assessing the Sensitivity of Climate Change Impacts in Agriculture to the Climate-Economy Functional Form	48
2.1 Introduction	48
2.2 Estimating the Impacts of Weather and Climate on Economic Activity . . .	52
2.2.1 Climate versus Weather	53
2.2.2 The Panel Approach	53
2.3 Data	58
2.4 Empirical Methods	62
2.4.1 Base Model and Specifications of Weather Variables	62
2.4.2 Simulating Future Impacts	65
2.4.3 Model Selection	67
2.5 Results	68
2.5.1 The Marginal Effect of Temperature	68
2.5.2 Climate Change Impacts on Agricultural Productivity	70
2.5.3 Model Performance	72
2.5.4 Consequences of Modeling Choices for Targeted Policy and the Social Cost of Carbon	74
2.6 Conclusion	75
2.7 Tables	78
2.8 Figures	82

2.9	Supplementary Analysis	92
2.9.1	Model Selection Criteria	92
2.9.2	Crop Revenue and Agricultural GDP	92
2.9.3	Climate Model Sensitivity	92
2.9.4	Regional Models	93
2.9.5	Characterizing Vulnerability	96
Chapter 3	A Comparative Analysis of Water Allocation Institutions and Economic Ef- ficiency	106
3.1	Introduction	106
3.2	Theoretical Models of Water Allocation Institutions	112
3.2.1	Two Users	112
3.2.2	N Users	117
3.3	Results	123
3.3.1	Water Supply Simulations	123
3.4	Conclusion and Policy Discussion	128
3.4.1	The Colorado River	129
3.5	Tables	131
3.6	Figures	133
3.7	Supplementary Analysis	137
3.7.1	The Role of Return Flows in Water Allocation	137
3.7.2	Regression Exercises	137
3.7.3	Other Considerations	138

LIST OF TABLES

1.1	Summary of characteristics of treatment and control diversion structures at 50% curtailment increase cutoff, 2001	25
1.2	Difference-in-difference estimations of the impact of drought and institutional change on irrigation practices, 50% curtailment increase cutoff	26
1.3	Difference-in-difference estimations of the impact of drought and institutional change on crop-specific irrigation practices, 50% curtailment increase cutoff	27
1.4	Robustness of difference-in-difference estimations when controlling for groundwater use	28
2.1	Econometric specifications of $f(\mathbf{X}_{it}; \beta)$	78
2.2	Common functional forms of weather in the economics literature	79
2.3	The impact of growing season daily mean temperature and precipitation on agricultural productivity	80
2.4	Cross-sectional regression results assessing the distributional impacts of predicted climate change impacts on county net farm income per acre by model	81
3.1	Water allocation outcomes under different water supply scenarios; $R_i^s = 250$, $R_i^j = 200$, $\theta = 0.25$	131
3.2	Detailed parameters for water supply simulations	131
3.3	Parameters for a sample simulation with ten users and $\tilde{w} = 2, 252$	132
3.4	Cross-sectional regression results assessing the impacts of basin composition and water shortage on institutional efficiency relative to the Social Planner	139

LIST OF FIGURES

1.1	How perceptions of water availability impact the gross benefit of investment in improved irrigation technology	29
1.2	Monthly PDSI and frequency of calls by state engineer, Colorado Water Division 1 . .	30
1.3	Treatment and control diversion structure map at 50% curtailment increase cutoff, Colorado Water Division 1	31
1.4	Difference-in-difference estimations of the impact of drought and institutional change on irrigation practices, 50% curtailment increase cutoff	32
1.5	Nonlinearities in the impact of drought and institutional change on irrigation practices as average increase in curtailment increases from 0% to 150%	33
1.6	Difference-in-difference estimations of the impact of drought and institutional change on crop choice, 50% curtailment increase cutoff	34
1.7	Difference-in-difference estimations before and after propensity score matching, sprinkler %	35
1.8	Estimated perception parameters for control versus treatment diversion structures . . .	40
1.9	Pre-treatment placebo test as average increase in curtailment increases from 0% to 150%	41
1.10	Average Days Under Curtailment for all Treatment and Control Structures by Year . .	42
1.11	Difference-in-Difference Estimations, Reference Year 1976	43
1.12	Difference-in-Difference Estimations, Reference Year 1987	44
1.13	Difference-in-Difference Estimations, Reference Year 1997	45
1.14	Nonlinear Impacts of Perceptions on the Adoption of Sprinkler Technology, Total Sprinkler Acreage	46
1.15	Nonlinear Impacts of Perceptions on the Adoption of Sprinkler Technology, Sprinkler Acreage as a Percentage of Total Acreage	47
2.1	Differences in temperature measures between Orleans County, Louisiana and Pima County, Arizona. Each plot shows the mean of each weather variable across all years in the sample for each county.	82
2.2	Estimated upper temperature thresholds for harmful degree days for different crop types. The curves represent estimated idealized daily temperature curves for an average summer solstice for each county. The cross marks correspond to daily mean temperatures.	83
2.3	Changes in the distribution of daily mean temperature and precipitation during the growing season from 1980-2021 to 2060-2100 across all counties. Future climate data is from CMIP6 model projections under the SSP245 scenario.	84
2.4	Estimated marginal effects of temperature and precipitation on net farm income per acre. The coefficient point estimates with 95% confidence intervals correspond to the linear ($L = 1$) hourly bin specification with ERS FRR-by-year fixed effects.	85
2.5	How the marginal effect of an extra hour spent in the $[33, 36)^{\circ}\text{C}$ temperature bin changes as a county's climate becomes warmer by model specification.	86
2.6	How the marginal effect of a 1°C mean growing season temperature increase changes as a county's climate becomes warmer by model specification.	87

2.7	Comparison of predicted climate change impacts on US agricultural profits by model specification over 2060-2100 with bootstrapped 95% confidence intervals. Future climate data is from the CMIP6 ACCESS-CM2 model projection under the SSP245 scenario.	88
2.8	Comparison of predicted distributional impacts of climate change on county net farm income per acre by model specification with county and FRR-year fixed effects. Each density curve represents all point estimates of the average annual change in net farm income every county in the sample over 2060-2100. Future climate data is from the CMIP6 ACCESS-CM2 model projection under the SSP245 scenario.	88
2.9	Comparison of model performance by model fit and out-of-sample prediction. Tiles are colored by the best performing model by the following criteria: lowest RMSE for all out-of-sample prediction methods, lowest AIC and BIC, and highest within-adjusted R^2	89
2.10	Differences in projected impacts to net farm income per acre between the Hour Bins, Cubic and Growing Degree Days, 10-32 model specifications under six CMIP6 model projections. All future climate data assumes the SSP245 scenario.	90
2.11	Range of projected impacts across models to total net farm income by net farm income decile under six CMIP6 model projections. All future climate data assumes the SSP245 scenario. The range of impacts correspond only to models estimated with ERS Farm Resource Region-by-year fixed effects and omits the Seasonal Means, Linear specification.	91
2.12	Comparison of fit statistics by model specification for net farm income per acre.	92
2.13	Comparison of cross validation performance by by model specification for net farm income per acre.	93
2.14	Comparison of predicted climate change impacts on US gross crop revenue by model specification over 2060-2100 with bootstrapped 95% confidence intervals. Future climate data is from the CMIP6 ACCESS-CM2 model projection under the SSP245 scenario.	98
2.15	Comparison of median predicted climate change impacts on US agricultural GDP by model specification over 2060-2100 with bootstrapped 95% confidence intervals. Future climate data is from the CMIP6 ACCESS-CM2 model projection under the SSP245 scenario.	99
2.16	Comparison of median predicted climate change impacts on net farm income per acre by model specification and CMIP6 model over 2060-2100 with bootstrapped 95% confidence intervals. All future climate data assumes the SSP245 scenario.	100
2.17	Comparison of predicted climate change impacts on aggregate US agricultural profits by model specification and CMIP6 model over 2060-2100 with bootstrapped 95% confidence intervals. All future climate data assumes the SSP245 scenario.	101
2.18	Comparison of predicted distributional impacts of climate change on county net farm income per acre by model specification with county and FRR-year fixed effects and CMIP6 model. Each density curve represents all point estimates of the average annual change in net farm income every county in the sample over 2060-2100. All future climate data uses the SSP245 scenario.	102

2.19	Comparison of estimated marginal effects of temperature on net farm income per acre on counties east and west of the 100th meridian. On the left, coefficient point estimates with 95% confidence intervals correspond to the linear ($L = 1$) hourly bin specification. On the right, marginal effects curves correspond to the cubic ($L = 3$) growing season daily mean specification.	103
2.20	Comparison of predicted average annual impacts of climate change on aggregate US agricultural profits by model specification and regional aggregation over 2060-2100 with bootstrapped 95% confidence intervals. Future climate data is from the CMIP6 ACCESS-CM2 model projection under the SSP245 scenario. Cubic specifications for the binned models are omitted due to exceptionally large confidence intervals.	103
2.21	Comparison of predicted distributional impacts of climate change on county agricultural profits per acre by model specification and regional aggregation. Each density curve represents all point estimates of the average annual change in net farm income per acre for every county in the sample over 2060-2100. Future climate data is from the CMIP6 ACCESS-CM2 model projection under the SSP245 scenario.	104
2.22	Predicted impacts of climate change on county agricultural profits per acre by regional aggregation over 2060-2100 for the hour bins, linear model. Future climate data is from the CMIP6 ACCESS-CM2 model projection under the SSP245 scenario.	104
2.23	Vulnerability of county-level agricultural profits to climate change	105
3.1	Spatial heterogeneity in hydraulic connectivity of western HUC8 basins per Duan et al. (2018)	133
3.2	Visual representation of a hypothetical waterway with three users in which return flows from an upstream user are available to all downstream users.	134
3.3	Monte Carlo generated total value distributions under different water allocation mechanisms with two water users	134
3.4	Average total value generated and percentage of reallocated water by institution and number of water users	135
3.5	How priority location impacts the total amount of water available for trade	135
3.6	Distributions of total value losses relative to the Social Planner by institution, priority composition, and water shortage scenarios	136
3.7	Monte Carlo generated total value, average value, total water, and basin outflow distributions with 36 water users	140

Chapter 1

Perceived Water Scarcity and Irrigation Technology Adoption

1.1 Introduction

Scientists and policy makers are actively exploring efficient and sustainable resource planning under a changing climate (IPCC, 2021). Climate change describes a shift in the underlying distribution of weather patterns over a long period of time, and anticipated changes in temperature and precipitation are projected to escalate water scarcity, posing a risk to food production (Mancosu et al. 2015). This issue is particularly salient for arid regions, where water supplies are expected to be intensely affected by climate change (Lioubimtseva 2004) and agriculture is often dependent on irrigation. For areas that rely on irrigation water derived from snowpack, accelerated snowmelt will change the timing and quantity of water available during the growing season, increasing the risk of costly shortages (Adam et al. 2009). It is estimated that water shortages result in more annual crop loss than all pathogens combined, totaling \$30 billion in global production losses over the past decade (Gupta et al. 2020). As overall water availability changes, maintaining agricultural output will depend on how producers adapt. Adopting water conservation strategies is one possible mechanism.

Understanding what motivates producers to conserve water resources is important for future planning. While there exists a rich literature on potential conservation strategies (e.g., Howden et al., 2007), little emphasis has been placed on the role of perceptions that influence the adoption of these strategies. One reason for the lack of research in this area is that identifying events in the natural world that change perceptions, and subsequently change behavior, can be challenging. Some literature suggests that personally experiencing an extreme weather event can change perceptions about climate change and increase the inclination to adopt conservation strategies (Spence et al. 2011, Wang 2017, Wang & Lin 2018). Maddison (2007) finds that many farmers in Africa perceive climate change to be real, yet some still do not respond in their practices. While these studies provide important insights into attitudes on climate change, they rely on cross-sectional survey data and do not track systematic behavioral changes over time. Overall, literature investigating how perceptions impact behavior using observational, non-survey data is scant. In this article, we explore how changing perceptions about water scarcity affects conservation investment decisions for agricultural producers. A simple theoretical framework is developed to demonstrate the conditions under which a producer's perception of water availability would incentivize investment in improving irrigation efficiency. Then, a unique period of extreme drought and institutional reform in Colorado is leveraged as a natural experiment to compare empirical results to simulations from the theoretical model.

Perceptions about water availability play a critical role in decision-making for producers that are dependent on irrigation. In the United States, western states (the American West) account for 81% of total irrigation withdrawals (Dieter et al. 2018), and productivity in many areas is dependent on surface water from snow runoff. In this area, rising temperatures cause more precipitation to fall as rain instead of snow, which reduces snowpack depth and changes the seasonality of runoff (EPA 2016). Several studies have examined recent hydrological changes in the American West, documenting trends in earlier snowmelt-driven stream flows and declines in April snowpack (Mote et al. 2005, Hamlet et al. 2005, Mote 2006, EPA 2016). In addition to increasing temperature and evaporation trends, monthly projections of the Palmer Drought Severity Index (PDSI) suggest that

climate change will amplify the length and severity of droughts while also hindering the recovery of macro-scale water supplies (Gutzler & Robbins 2011). One mechanism for agricultural producers to adapt to these trends is to adopt more water-efficient, pressurized irrigation systems like sprinkler or drip (Howden et al. 2007, Frisvold & Bai 2016), but gravity systems are still prevalent throughout the American West partially due to high costs of sprinkler systems and relatively low water prices. Carey & Zilberman (2002) use a stochastic dynamic model to demonstrate how uncertainty in water supplies creates an option value and deters irrigation technology adoption unless the expected present value of the investment exceeds the cost by a large margin. However, some empirical evidence has shown that farmers adopt new technologies to hedge against production risk (e.g., Koundouri et al., 2006). The present article provides insights into the disconnect between some theoretical predictions and empirical evidence surrounding investment in irrigation technology.

Presumably, producers hold some belief about the probability distribution of input shocks. When considering multi-year investments in water conservation technologies, a farmer likely assesses of the probability of a water shortage. Ji & Cobourn (2021) provide an intuitive framework of expectation formation, proposing that perceptions about weather—or water supply conditions—develop with an increased bias toward recent events. Therefore, experiencing a disproportionately extreme event triggers a larger revision to expectations, and a subsequent series of events closer to long-run averages would be necessary to decrease the perceived likelihood of another extreme event. They corroborate their theoretical hypotheses empirically, finding that weather shocks significantly impact short-run planting decisions for farmers. Similarly, Cobourn et al. (2021) demonstrate that irrigators anticipating water shortages are more likely to fallow land and plant drought-resilient crops. Complementing these recent studies that focus on short-run responses (i.e., yearly planting decisions), our attention lies on long-run responses (i.e., investment in infrastructure). Our novel dataset of 70 years of water right curtailment recordings alleviates our reliance on weather data in estimating producer expectations of water availability. We are able to pinpoint producers that directly experienced shortages, allowing us to identify changes in per-

ceptions and subsequent long-run improvements in water-use efficiency via irrigation technology adoption.

The present article contributes to the relevant literature in two aspects. First, we develop a theoretical model to analyze the conditions under which an agricultural producer's perception of a possible water shortage would incentivize investing in a more water-efficient irrigation system. The model framework captures how risk is perceived for farmers operating under a priority-based water allocation institution through two parameters: (i.) the probability that water supply will be curtailed in a given year and (ii.) if curtailed, the intensity of the water loss. We then consider a probable range of model parameters to identify when the benefit of investing in more efficient irrigation infrastructure is highest. Since our framework captures the nuances of a priority-based water rights regime, insights on how investment decisions are influenced by perceptions are particularly applicable to the American West, though similar regimes exist throughout the world. Our theoretical model shows that the net benefit of adopting more efficient irrigation technology increases in the probability of curtailment, holding all else constant. However, changes in the expected amount of water received (when curtailed) impacts net benefits non-monotonically.

Second, we capture changing perceptions of water shortage risk for farmers empirically in northeast Colorado using a comprehensive panel dataset of irrigated cropland, agricultural water rights, and curtailment recordings. For many producers, increased water scarcity will change the perceived reliability of their water right portfolios. In Colorado, producers with historically secure water rights are facing increases in curtailment due to institutional changes resulting from litigation and sustained drought in the early 2000s (Waskom 2013). Our empirical context provides us with a unique opportunity to leverage a natural experiment to identify a change in perceptions about the reliability of a water supply, which allows us to measure how those changes affect decisions to adapt to increasing scarcity. Treatment and control groups are determined by water right curtailments during the early 2000s shock relative to historical droughts. Results of a difference-in-difference analysis indicate that the treatment group, i.e., those who experienced an unprecedented increase in curtailment, adopted more water-efficient irrigation systems at signifi-

cantly higher rates than the control group. Additionally, cropland with corn experienced the largest increases in irrigation efficiency improvements in years following the shock relative to other crops, although total corn acreage declined. Corn is considered more sensitive to water stress than other popular crops grown in the region, such as alfalfa or wheat, further indicating that the shock incentivized a change in practices to hedge against production risks. Some producers in our study area supplement their surface water irrigation practices with groundwater, yet we find that changes in groundwater use did not differ substantially between treatment and control groups beyond years immediately following the shock. This is, in part, due to the conjunctive governance of surface water and groundwater in Colorado. These empirical findings provide fresh evidence of the link between updating perceptions and conservation investment behavior.

The remainder of the article is organized as follows. First, we describe a theoretical framework used to analyze the impact of perceptions on irrigation technology adoption in the context of prior appropriation, which is the dominant water allocation system in the American West. We then discuss our study area and the period of extreme drought and institutional reform that we leverage as a natural experiment, followed by a description of the data and modeling approach. In the final sections we present the estimation results, analyze their robustness, and conclude by discussing policy implications.

1.2 Theoretical Framework

When adapting to increasing water scarcity, farmers face a menu of potential strategies to reduce their overall water dependency. Drought-tolerant crop varieties and species can be planted in lieu of water-intensive ones. Deficit irrigation on large quantities of land or increased irrigation on a reduced quantity of land offer opportunities for water savings. Technologies that harvest and store water or reduce conveyance losses can increase average supplies. Improving the application efficiency of an irrigation system can reduce the amount of diverted water necessary to achieve full evapotranspiration. The advent of water markets and water-sharing agreements allow farmers to

diversify their income through selling water or to hedge against drought risk by buying water. In general, the actual costs and benefits of different adaptation strategies from this suite of options depends on a farmer's characteristics, such as geographic location. However, farmers' *expected* costs and benefits, which ultimately drive adoption decisions, can vary according to perceptions about water scarcity that may not systematically correlate with observed demographics. Otherwise similar farmers with different perceptions may exhibit a vastly different willingness to invest in practices that reduce water use. In the following theoretical framework, we focus solely on how perceptions drive the decision to improve irrigation efficiency. Although our theoretical model is presented in the context of a producer evaluating the expected benefits from adopting an irrigation technology, the findings apply more generally with respect to how perceptions of scarcity influence producer decisions to invest in technologies and/or practices that improve water use efficiency.

To examine the impact of perceptions on conservation investment, we develop a theoretical model describing a producer's decision to improve the efficiency of his irrigation infrastructure. We adopt the conceptual framing of a producer's irrigation water supply under a prior appropriation system from Li et al. (2019), with some simplification. After summarizing prior appropriation, we show a general condition characterizing the net benefit of investment in a conservation technology. We then impose assumptions on the parameters to estimate the impact of perceptions about input availability on the investment decision.

Water allocation in most of the American West is governed by a system of prior appropriation, a legal framework that rules over all water use. To divert water under prior appropriation, one must obtain a water right from a court or purchase an existing right. Water rights are usufructuary, meaning that the rightsholder does not own the water itself but the right to divert and use it. Rights are ranked in a hierarchy of priority determined by the date on which a user first appropriated and diverted water for beneficial use, colloquially phrased as "first in time, first in right." Owners of agricultural water rights cannot use more water for irrigation than what is decreed by their right, and when basin water supplies are insufficient to fulfill all decreed water rights, rightsholders with older water rights have priority over users with newer rights. In the state of Colorado, water rights

are curtailed through a system of administrative “calls.” When inflows are insufficient to satisfy all decreed water rights, the State Engineer places a “call” on a stream which curtails the ability for junior water rights holders to divert. The administrative call communicates a priority level required to continue diverting water. In essence, when senior water rights are unable to achieve their decreed allotment due to insufficient supplies, all junior upstream users must temporarily stop diverting to make more water available (Getches 2009).

Consider a producer operating under a system of prior appropriation who uses water to grow crops. The producer owns a water right with a fixed priority level and a maximum amount of $\bar{w}$ units of water that may be diverted from a specified stream. Irrigation water w available to the producer to grow crops over a growing season is a random variable that takes the form

$$w = \begin{cases} \bar{w}, & S \geq V \\ \delta \bar{w}, & S < V \end{cases} \quad (1.1)$$

where S is a stochastic stream supply term, corresponding to the total quantity of water available for diversion by all water rights holders, and V is the total supply necessary within the stream system for the producer to divert the maximum quantity of water associated with the water right. If $S < V$, the producer’s water right is called, and he receives a proportion $\delta \in [0, 1)$ of the total allotment. We further assume a relationship between irrigation water and crop yield equal to

$$y(w, \lambda, \alpha) = \begin{cases} (\lambda w)^\alpha, & 0 \leq \lambda w < w_m \\ y_m, & \lambda w \geq w_m \end{cases} \quad (1.2)$$

where $y(w, \lambda, \alpha)$ is the total quantity of output, $\lambda \in (0, 1)$ is an irrigation efficiency coefficient, y_m is maximum yield, w_m is the net irrigation requirement for maximum yield, and $\alpha \in (0, 1)$ is a shape parameter.¹

¹Our choice of functional form attempts to exhibit the typical relationship between total seasonal irrigation and crop yield as represented on page 4 of Foster & Brozović (2018). We assume no yield when no water is applied, i.e., $y(w = 0, \lambda, \alpha) = 0$, since irrigated crop varieties in Colorado are often not drought tolerant.

Now consider the case where the producer has an existing low-efficiency, flood irrigation system and can invest in a high-efficiency, sprinkler system. The producer can pay an annualized cost of the upfront capital investment, c_s , that would increase irrigation efficiency from λ_f to λ_s . Assume this producer's objective is to maximize expected profit by first choosing whether to invest in the new irrigation system, taking prices as given, and then applying water to his fields after w is realized. The profit function after realization is composed of the per unit price of output p , output $y(w, \lambda, \alpha)$, and some fixed cost of production k ,

$$\pi = py(w, \lambda, \alpha) - k. \quad (1.3)$$

The producer assumes a probability that his water right will be called in a given year, $P(S < V) = \theta \in [0, 1]$, and the magnitude of water loss, $\delta \in [0, 1)$, should the call occur. Given his perception of parameters θ and δ , and conditional on efficiency, the producer's expected profit prior to the realization of w is

$$\mathbb{E}[\pi] = p[(1 - \theta)y(\bar{w}, \lambda, \alpha) + \theta y(\delta\bar{w}, \lambda, \alpha)] - k. \quad (1.4)$$

The decision to invest in the new irrigation system is modeled as binary, so the producer chooses between only two profit functions. For simplicity, we examine the payoff of investing for a single period case. The annualized expected net benefit of investment is the difference between the expected profit under a sprinkler system (λ_s) and a flood system (λ_f):

$$\mathbb{E}[\pi_s] - \mathbb{E}[\pi_f] = p\{(1 - \theta)[y(\bar{w}, \lambda_s, \alpha) - y(\bar{w}, \lambda_f, \alpha)] + \theta[y(\delta\bar{w}, \lambda_s, \alpha) - y(\delta\bar{w}, \lambda_f, \alpha)]\} - c_s, \quad (1.5)$$

and assuming $\bar{w}$ is the amount of water necessary for maximum yield with flood irrigation,² i.e., the marginal productivity of water is zero beyond $\bar{w} = y_m^{1/a}/\lambda_f$, equation (5) is reduced to

$$\mathbb{E}[\pi_s] - \mathbb{E}[\pi_f] = p\theta[y(\delta\bar{w}, \lambda_s, \alpha) - y(\delta\bar{w}, \lambda_f, \alpha)] - c_s. \quad (1.6)$$

²Water allotments under prior appropriation are determined by the historical consumptive use of the activity allowed by the water right, so this assumption is appropriate in this context.

Lastly, we assume that a producer adopts the technology if the net benefit of investment is greater than zero:

$$p\theta[y(\delta\bar{w}, \lambda_s, \alpha) - y(\delta\bar{w}, \lambda_f, \alpha)] - c_s > 0 \quad (1.7)$$

or after rearranging,

$$\theta[y(\delta\bar{w}, \lambda_s, \alpha) - y(\delta\bar{w}, \lambda_f, \alpha)] > \frac{c_s}{p}. \quad (1.8)$$

The left-hand side of equation (8) is the difference in yields when water is called multiplied by the probability that water is called. It represents the expected gross benefit of technology adoption. As the difference increases, producers become more likely to adopt the technology. The right-hand side describes the ratio of cost to output price. As the cost of the investment increases, producers are less likely to invest. As price increases, the net benefit of adoption becomes higher, and producers become more likely to invest. A key feature of this model is that the benefit of technology adoption comes only from reducing down-side risk. The adoption of a more efficient irrigation technology allows the producer to achieve a higher yield per unit of water when he does not receive the entirety of his water right. The highest priority farmer has $\theta = 0$ and an expected gross benefit equal to zero.

1.2.1 Parameter Simulations

We now examine how perceptions of θ and δ incentivize adoption of the efficient technology by parameterizing the left-hand side of (8). We are only concerned with identifying where a producer would have the highest likelihood of adopting, so we focus on the range of parameters in which the gross benefits of adoption are highest. When the gross benefits of adoption are highest, we would expect adoption to be more likely.

First, we assume the following parameter values: $\lambda_f = 0.5$, $\lambda_s = 0.9$, $y_m = 6$, and $\alpha = 0.5$. Irrigation application efficiencies are used from Bauder et al. (2014) and maximum yield can be interpreted as tons of corn per acre. We then calculate the left-hand side of (8) over the range of plausible values of θ and δ to generate a heat map displaying the areas in which the gross

benefit of adoption is greatest given the producer's perceptions (Figure 1.1). Each point on the heat map corresponds to a possible combination of θ (probability of a call) and δ (magnitude of shortage) for an individual producer. The background shading at each point corresponds to the gross benefit, or increase in expected yield, associated with that combination of θ and δ . Darker, purple (lighter, yellow) areas are associated with lower (higher) gross benefits, so a change that results in movement from a dark area to a lighter area would result in an increase in expected benefit. The impact of θ is straightforward and monotonic. Pick one point along the horizontal axis and hold δ constant, and each point directly above (increasing θ) lies on a lighter area on the map. In other words, as a possible call becomes more likely, the benefit of improving water application efficiency increases monotonically. If the perceived probability of a call is 0, there is no incentive to invest.

The impact of δ is less straightforward. When holding θ constant and moving along points from left to right, the benefit is greatest around $\delta = 0.55$, after which the benefit decreases. When the producer expects to receive nearly all or nearly none of his water during a call, the benefit of improving water application efficiency approaches 0. If a producer were to experience a change in perception that moved him from a dark to light area on Figure 1.1, we would expect an increased likelihood of investing in the high-efficiency, sprinkler system. The region in which producers are most likely to adopt the new technology occurs when the probability of a call is perceived as high, and the volume of water lost during a shortage is about half of the full right. In the empirical section of this paper, we investigate these theoretical predictions.

1.3 Data and Methods

1.3.1 Study Area

In Colorado, The Water Right Determination and Administration Act of 1969, C.R.S. 37-92 et seq. (1969), designated seven water divisions based on drainage characteristics, each staffed with its own division engineer and water judge. Water Division 1 (WD1), the study area for this analysis,

is highly dependent on surface water and contains the South Platte River basin (SPRB), Republican River basin, and Laramie River basin. The Colorado Water Plan (CWCB, 2015) provides extensive detail on all basins and water divisions, and here we summarize the details relevant to our analysis. The SPRB alone is home to approximately 80% of Colorado's population while also having the largest proportion of irrigated agriculture. Irrigated agriculture accounts for approximately 85% of total water diversions within the basin, with water supplies originating in mountain snowpack along the Continental Divide. Farmland in WD1 typically receives less than 8 inches of precipitation during the growing season (Schneekloth & Andales 2017). In addition to 1.4 million acre-feet of average-annual native flow volume, the basin receives an additional 500,000 acre-feet in transmountain diversions. Overall, the basins in WD1 are over-appropriated, meaning the total allotted volume of water rights exceeds the current average supply, and many irrigation season water rights are continuously out of priority.

WD1 provides a relevant case study of many arid regions that are experiencing water scarcity concerns coupled with irrigation-dependent agriculture and fast growing populations. The 17 states wholly or partially west of the 100th meridian in the conterminous United States all utilize a strict or hybrid prior appropriation water rights regime (Leonard & Libecap 2019) and depend on irrigation water for agricultural production. Of these 17 states, seven were among the top ten fastest growing states in percent growth from 2020 to 2021 (US Census Bureau 2021). Increased water scarcity due to climate change combined with increasing demands for urban uses place significant pressures on agricultural production in these areas. While there is considerable heterogeneity in producers across the American West, many face similar problems to those represented in this study.

Agricultural producers in WD1 face uncertainty in water availability from two predominant sources. The first source is the variability in water supplies under a changing climate. The second source is institutional, as water administration is complex and constantly evolving. Colorado is experiencing rapid population growth, with increasing water demands for municipal, industrial, recreational, and environmental uses, and the administration of water law frequently undergoes changes from new legislation and court rulings as new problems emerge (Jones & Cech 2009).

1.3.2 The Natural Experiment

In addition to designating water divisions, the 1969 Act determined that groundwater was to be regulated in conjunction with surface water under prior appropriation. The act introduced “augmentation plans” that allow for out-of-priority diversions so long as sufficient replacement water is supplied to prevent injury to senior users. Such plans are required to be approved through a decree of a district water court,³ but the State Engineer was granted the ability to temporarily approve substitute water supply plans (SWSPs). SWSPs were essentially augmentation plans that could be renewed on an annual basis without official approval from the courts. Consequently, many junior users neglected to formally seek court adjudication and relied on the State Engineer for continued water use under SWSPs (Waskom 2013). SWSPs were predominantly utilized by groundwater users who would collectively provide replacement water through recharge ponds or reservoirs. Throughout the 1980s and 1990s, groundwater users in particular were accused of providing inadequate replacement water (Waskom 2013), however exceptional precipitation and snowpack (McKee et al. 2000) veiled potential water shortages. Nearly two decades of abundant water supply meant there was little incentive to impose change within the system.

Then, in 1999-2000, Colorado experienced an unexpected combination of low winter snow accumulation and above average spring and summer temperatures that led to drought conditions across the state (Pielke et al. 2005). This revealed that existing replacement efforts under SWSPs did not adequately cover shortfalls in water availability, and as a result, litigation was launched between two water users over misuse of SWSPs. The result of *Empire Lodge Homeowner’s Association v. Moyer*, 39 P.3d 1139 (Colo. 2001) declared that the State Engineer did not have legal authority to approve SWSPs on an annual basis and shifted more oversight of water replace-

³Colorado water courts are specialized state courts with water judges appointed by the state Supreme Court. Water judges have jurisdiction over all water use and administration within their water division. See <https://www.courts.state.co.us/Forms/PDF/JDF%20301W.pdf> for the application and detailed requirements for approval of an augmentation plan.

ment plans to the water courts.⁴ Although this ultimately led to the permanent curtailment of many groundwater rights, it had a direct impact on surface water. First, producers faced increased dependence on uncertain surface water supplies during the summer months. Additionally, the number of formally decreed augmentation plans that require records of actual diversions increased dramatically in subsequent years (Waskom 2013). Since the basins in WD1 are over-appropriated, net surface water diversions could not increase in practice. As more water rights recorded daily diversions, the State Engineer had a better understanding of actual surface water supplies, and the likelihood of calls along mainstream rivers increased.

After the institutional change, drought conditions persisted through 2009 with the most intense period occurring in 2002. In 2002, all of Colorado was in extreme drought conditions, and April snowpack was estimated at 52% of the previous 30-year average (Pielke et al. 2005). PDSI levels for WD1 reached -6 (Figure 1.2, top panel), a classification of drought categorized by widespread crop losses and severe water shortages that result in water emergencies. The newly increased reliance on surface water, better records for actual diversions, and unprecedented drought conditions resulted in a permanent change to the call regime (Figure 1.2, bottom panel). The average number of days under call from 2002-2012 was two to four times that of 1982-2001 for districts within WD1 (Waskom 2013, pp. 149-152). The change in oversight for out-of-priority diversions, combined with an unprecedented decrease in surface water supply, created an exogenous shock to the distribution of surface water available to relatively junior water rights.

To exploit the exogenous change in surface water availability for some users, we compile an extensive dataset for WD1 on irrigated cropland, irrigation technology, agricultural surface water rights, call recordings, and population across eight observation years (1976, 1987, 1997, 2001, 2005, 2010, 2015, 2020). County-level population data is available through the Colorado State Demography Office and the remainder of the data from Colorado's Division of Water Resources

⁴For more information on SWSPs and *Empire Lodge Homeowners v. Moyer*, see the "Guidance Documents" available at <https://dwr.colorado.gov/services/water-administration/water-supply-plans-and-administrative-approvals>.

HydroBase software.⁵ Information on individual water rights includes water source, point of diversion, water use type, maximum flow volume, appropriation date, and priority number. The priority number ranks all water rights in terms of seniority, determined by rights' appropriation and court adjudication dates. Information on irrigated cropland includes acreage, point of diversion, and crop type. Irrigation technology at the field level describes if a field irrigates using flood or sprinklers. Water rights, irrigation technology, and irrigated acres can be matched to a diversion structure, such as a ditch or canal, however we cannot identify the individual parcels owned by a specific water right holder. Therefore, we aggregate information to the diversion structure as the unit for analysis. Altogether we construct a balanced panel of 434 diversion structures.

Since 1950, all administrative calls by the State Engineer have been recorded, which we use for our treatment design. Annual information on the length of curtailment for each water right allows us to define treatment and control groups by losses during the 2000s drought relative to historic droughts in the 1950s and 1970s (McKee et al. 2000). The average number of curtailed days per year in drought period d , C_d , during the growing season (April-October) is calculated over the "historic" drought years (1950-1956 and 1974-1978) and the "recent" drought years (2000-2009) for all water rights sharing a diversion structure. Our theoretical model demonstrates that a change in perceptions about water availability can impact the expected gross benefit of investing in more efficient irrigation technology, and we assume that producers developed a perception about the reliability of their water rights from the intensity of their curtailment during historic droughts. Diversion structures that experienced an increase in average curtailment C_d during the recent drought period are placed into the treatment group:

$$\text{Treatment} = \begin{cases} 1, & \Delta C_d > X\% \\ 0, & \Delta C_d \leq 0\% \end{cases} \quad (1.9)$$

⁵See <https://cdss.colorado.gov/software/hydrobase>.

where $\Delta C_d = \frac{C_{recent} - C_{historic}}{C_{historic}} * 100$. We remain agnostic to the magnitude of increase that would trigger a change in perceptions and allow X to vary from 0% to 150%. This approach allows us to detect potential heterogeneity in the response of the treatment group as the average increase in curtailment becomes larger. It is important to note that as X increases, diversion structures that experienced changes between 0% and $X\%$ are dropped from the analysis.

To examine the impact of the shock on the number of irrigated acres at diversion structure i in year t with technology j , y_{it}^j , we estimate the following difference-in-difference models:

$$y_{it}^j = \sum_t^{\tau} \beta_t^j D_i T_t + \omega^j x_{it} + \alpha_t^j + \gamma_i^j + \varepsilon_{it} \quad (1.10)$$

where j denotes the technology-specific model (i.e., sprinkler or flood), $D_i = 1$ if structure i is in the treatment group and 0 otherwise, and T_t is an indicator equal to 1 if $t = \text{year } T$ and 0 otherwise. The term x_{it} is county population, and α_t and γ_i are year and diversion structure fixed effects to control for time trends and omitted variables. Lastly, ε_{it} is the error term clustered at the diversion structure. As a placebo test, $D_i T_t$ includes all panel years, excluding the reference year of 2001, to investigate differences prior to and after the natural experiment. Hereinafter we will refer to years 1976, 1987, and 1997 as “pre-treatment” and years 2005, 2010, 2015 and 2020 as “post-treatment.”

For simplicity, we summarize the sample characteristics of treatment and control groups when $X = 50\%$ in Table 1.1 and explore the affects of different levels of X in the next section. Statistics for 2001 are reported to provide a snapshot of the sample just before the natural experiment, and it is also used as the reference year for our regression analysis. From the data presented it is apparent that larger diversion structures with slightly more junior water rights were disproportionately impacted by the shock. To investigate if treatment structures are correlated spatially, we present a map of the same treatment and control structures in Figure 1.3. The location of treatment structures provides evidence that the shock was not localized to a specific area. We find treatment structures in both urban and rural areas and along a variety of different streams.

1.4 Results

Coefficient estimates from (10) with corresponding cluster-robust standard errors when $X = 50\%$ are reported in Table 1.2. We estimate four iterations of the model with different dependent variables: the number of irrigated acres with flood technology, the number of irrigated acres with sprinkler technology, sprinkler acres as a percentage of total irrigated acres, and total irrigated acres. We include the percentage of sprinkler acres to ensure estimates in the first two columns are not deceptively inflated by the behavior of larger diversion structures in our sample. Insignificant estimates for the pre-treatment variables in the first three columns indicate that differences in the outcome variables between treatment and control groups are not statistically distinguishable from zero prior to the shock. This provides suggestive evidence that the treatment and control groups satisfy the parallel trends assumption. We present coefficients for the treatment variables graphically in Figure 1.4, with 95% confidence intervals, to check for the existence of differential pre-trends visually. Dashed confidence intervals indicate overlap with zero. In years after the shock, estimates become significant and increase in magnitude, suggesting that a change in behavior persisted for over a decade. By 2020, the average treatment structure adopted sprinkler technology on 889 more acres than the average control structure. This amounts to an 18.7% point increase in irrigated land with sprinkler technology. Applying this estimate to the entire treatment group, the shock incentivized an increase of over 63,000 sprinkler-irrigated acres in our study area as of 2020.

Surprisingly, there is no statistically significant impact on total irrigated acreage. Although WD1 is experiencing an overall decline in irrigated acreage (CWCB 2015), the rate at which land is leaving production is comparable between the treatment and control groups. This suggests that the treatment group responded to the shock to water availability through more efficient use of the input on the intensive margin. The overall decline in irrigated acres across the basin is perhaps better explained by the population variable. The negative and significant coefficient for population suggests that an increase in population places downward pressure irrigated acres within a county.

This finding is consistent with large cities in Colorado buying agricultural water rights to meet increasing municipal demands (Pritchett et al. 2008).

Figure 1.5 displays coefficient estimates with 95% confidence intervals for the post-treatment years when X is allowed to vary from 0% to 150%. For the Sprinkler Acres and Sprinkler % models, treatment structures display increasing rates of sprinkler technology adoption as the curtailment cutoff increases from 0 to 90% and 75% respectively. After this point, the treatment group substantially diminishes in size and coefficient estimates become less stable, although they remain significant. These results suggest heterogeneity in treatment effects, as producers responded somewhat proportionately to the magnitude of the shock that they experienced. Estimates for the pre-treatment years are available in the Supplementary Analysis and show consistent insignificance.

In addition to irrigation technology, agricultural producers can respond to water scarcity by planting less water-intensive crops. We estimate crop-specific models using the same specification as (10) while limiting the dependent variable to total and sprinkler irrigated acres with corn, alfalfa, and wheat. We exclude results for grass pasture as there is very little sprinkler irrigated pasture in our sample. Results from the crop-specific models when $X = 50\%$ are presented graphically in Figure 1.6, again with 95% confidence intervals, and cluster-robust standard errors are available in Table 1.3. Regression results indicate that corn was the crop that experienced the biggest increase in sprinkler-irrigated land as a result of the shock. On average, corn acreage accounted for 60-65% of the increase in sprinkler acreage for all post-treatment years.⁶ This result holds in 2005 despite the significant average decrease of 156 total corn acres, which was a potential short-run response to the shock. Between the three crops, corn is generally more sensitive to drought than alfalfa or wheat (Lobell et al. 2014), making this result consistent with risk-mitigating behavior. By 2015, we find significant and positive differences for alfalfa and wheat in addition to corn for the Sprinkler Acres specification. With the exception of corn in 2005, we find no significant differences in total

⁶This was estimated by dividing the coefficient estimates from the corn-specific Sprinkler Acres model (Table 1.3, column 2) by the coefficient estimates from the total Sprinkler Acres model (Table 1.2, column 2).

acres post-treatment for each crop. This suggests that adjustments to the change in relative scarcity were made on the intensive margin (adjusting water application per acre) rather than the extensive margin (retiring cropland).

We also investigate the potential impacts on irrigated acreage that is supplemented with groundwater. The average number of estimated acres supplemented with groundwater in Table 1.1 indicates that the treatment group utilizes more groundwater to augment their irrigation practices. Since the institutional change in the early 2000s resulted in the curtailment of many groundwater rights, it is important to scrutinize what changes in sprinkler adoption can be attributed to changes in surface water versus groundwater availability. We first control for groundwater supplemented acreage in the Sprinkler Acres and Sprinkler % models to check for loss of significance and magnitude of the treatment effects, and then estimate one additional model with groundwater supplemented acres as the dependent variable. Estimates of groundwater supplemented acreage were omitted from the primary regressions due to endogeneity concerns and potential measurement error, particularly because attenuation bias due to measurement error is amplified in fixed effects estimations (Johnston & DiNardo p. 404, 2009). Regression results for the groundwater models are presented in Table 1.4, where columns (2) and (4) correspond to the models with the added groundwater control variable and column (5) to the model with groundwater supplemented acres (GW Acres) as the dependent variable. The only qualitative change to the results from the primary regressions is the loss of significance of the Treatment*2005 variable for the Sprinkler Acres model. Otherwise, the longer-term trends and Sprinkler % results remain qualitatively unaffected. For the GW Acres model, we find a significant decrease in groundwater supplemented acreage for the treatment group in 2005, which reflects the immediate curtailment of groundwater rights after the shock. However, estimates for Treatment*2010 and Treatment*2015 are not statistically distinguishable from zero, indicating that longer-term changes in groundwater use did not consistently differ between the treatment and control groups. It is therefore likely that the significant increases in sprinkler adoption in the treatment group was a mechanism to adapt to long-run changes in surface water availability due to the shift in the call regime.

In summary, we observe a short-run response to the shock in the reduction of total corn acreage and a long-run response in the increased and consistent adoption of sprinkler technology. To examine what this implies for potential water use, we use seasonal crop-water demands for corn and alfalfa to make a back-of-the-envelope calculation of the reduction in water required for full crop yields for treatment structures. First, we multiply the 2020 Sprinkler Acres coefficient estimates in Table 1.3 for corn and alfalfa by the number of treatment structures to estimate the total increase in crop-specific sprinkler-irrigated acreage. Next, we calculate the difference in the seasonal net-irrigation requirement, accounting for precipitation and soil moisture typical to northeastern Colorado, for an acre of each crop with flood irrigation versus sprinkler irrigation.⁷ The difference for each crop is then multiplied by the estimates from the first step. In total, we estimate a potential reduced seasonal irrigation demand for water diversions of 50,000 acre-feet or 16 billion gallons of water across WD1 by 2020 attributable to the change in expectations of water availability. The average Colorado household uses about 0.5 acre-feet of water per year (Waskom & Neibauer 2014), so the demand reduction is roughly equivalent to the yearly water demands of 100,000 households.⁸

In Figure 1.4, there is some evidence of trends prior to the shock given the direction of coefficient estimates across time, particularly for the Sprinkler % model. One might attribute these trends to differences in the distribution of seniority (Table 1.1) between the treatment and control water rights. Given our theoretical results, it is reasonable to assume that junior water right holders would invest more in water efficient technologies than senior water rights holders, regardless of the shock to surface water availability in the 2000s. If that is the case, then our coefficient estimates could be biased. We test this supposition by limiting our sample to similar treatment and control structures and re-estimating the Sprinkler % model. We use a propensity score matching algorithm using the minimum, median, and maximum appropriation year for the water rights associated with a structure to make the distribution of all water rights between treatment and control groups as

⁷Net crop water requirements are calculated from data presented in Schneekloth & Andales (2017).

⁸This comparison is made only to provide perspective on the volume of water. According to Colorado water law, water "saved" via irrigation efficiency gains cannot be reused or sold.

similar as possible.⁹ Results from the matching exercise are presented in Figure 1.7. The treatment group is smaller than the control group, so we first match every treatment structure with two similar control structures (second column) and then one-to-one (third column). The first row of Figure 1.7 displays a smoothed density curve of the seniority distribution for the water rights in the total sample and the two matched samples, and the second row displays coefficient estimates corresponding to the sample directly above. In each of the regressions using matched samples, pre-treatment coefficients move closer to zero and post-treatment coefficients remain positive and trend upward. In the one-to-one matched regression, the distribution of water rights are nearly identical, however the post-treatment coefficient for 2005 loses significance likely due to the reduced sample size. This exercise provides evidence that our main results are not driven by differences in seniority among the water rights at treatment and control diversion structures. Additional robustness checks and analysis relating to treatment design, model specification, and nonlinear impacts from the shock can be found in the first three sections of the Supplementary Analysis, respectively.

1.5 Conclusion and Policy Implications

In this article, we explore the impact of perceived input scarcity on conservation investment decisions. We develop a theoretical model to examine the conditions under which an agricultural producer’s perception of water shortages would incentivize investment in a more efficient irrigation technology. A numerical exercise is used to demonstrate a range of perceptions that maximize the gross benefit of investing in irrigation efficiency, and we test our theoretical predictions empirically. A period of severe drought and institutional change in Colorado that led to a change in expectations about the availability of irrigation water is leveraged as a natural experiment. Results suggest that agricultural producers who experienced an unprecedented shock to their irrigation water supply transitioned more land from low- to high-efficiency irrigation systems in the following decade. Our

⁹We use the MatchIt package in R to perform a greedy nearest neighbor matching algorithm. Details can be viewed at <https://cran.r-project.org/web/packages/MatchIt/MatchIt.pdf>.

analysis provides evidence that input shocks can trigger investment in efficiency due to changes in perceptions.

This research has limitations that must be acknowledged. Subsidy programs such as the Environmental Quality Incentives Program (EQIP) that can significantly reduce the costs of investment may affect conservation decisions. Although we can observe general rates of adoption through land use changes, we do not know producer-specific costs of a sprinkler system. We also cannot observe conservation practices beyond irrigation technology and crop choice in our data. For example, when evaluating EQIP enrollment, Wallander et al. (2013) found that many drought-facing producers adopted tillage practices that conserve soil moisture. Lining or replacing irrigation ditches to reduce seepage is another practice identified as water saving by EQIP that we are unable to detect.

In Colorado, water rights can be bought and sold, and a distinct feature of our study area is the presence of active water markets. Most market activity consists of municipal and industrial buyers and agricultural sellers. Some rights were undoubtedly traded during our study period, which we are unable to track. We can only observe the decreed uses of a particular water right as they exist today, and although we limited our analysis to water rights that have a decreed agricultural use, some have gained additional uses through previous transactions. It is possible that not all water right owners with an agricultural water right are using their water for agricultural production in a given year. Although we cannot identify which water rights were sold, we find that changes in total irrigated acreage did not differ substantially between treatment and control groups. This provides some evidence that agricultural water rights are being sold at similar rates across all diversion structures, regardless of the heterogeneous impacts of the shock.

Concerning water right transactions, it is important to note that improving irrigation efficiency does not generally reduce the value of a water right. One aspect of prior appropriation is that water rights may be forfeited if the owner consistently fails to apply the water to a beneficial use, otherwise known as "use it or lose it." This component however only applies to the *consumptive use* determined by the water right. In the case of a farmer, the consumptive use of his water

right is determined by the annual documented evapotranspiration of his crops, not the total amount of water diverted. Since improving irrigation efficiency only reduces the amount necessary for diversion and not the beneficial, consumptive use, the water right's value should not be affected. In the case of a water right transfer, the transferee buys only the right to the consumptive use regardless of the transferor's former diversion amounts. In general, the "use it or lose it" rule is not a true barrier to improving irrigation efficiency, although it is potentially perceived that way by some (Waskom et al. 2016).

Another important characteristic of our study area is that all surface water and most groundwater resources are administered similarly under prior appropriation, which is not uniformly the case across the American West. For example, some areas in Arizona, California, and Texas still abide a "rule of capture" for groundwater, so land owners who can drill into an aquifer can pump uninhibited. In Colorado, the conjunctive governance of surface and groundwater effectively limits their substitutability, so agricultural producers cannot rely on increased groundwater pumping when surface water supplies are low during drought. This lack of substitutability certainly affected producers' willingness to invest in technology to use surface water more efficiently. Groundwater aquifers are often exhaustible in practice, since they can take long periods of time to replenish naturally. Inhibiting the ability to excessively pump groundwater during drought may prompt an earlier adoption of water conserving technologies. Improving the use efficiency of renewable surface water supplies before exhausting limited groundwater resources may increase the longevity of agricultural production under climate change.

Lastly, it is worth noting that hydrological systems are exceptionally complex, and any changes to how and when water is diverted has common property resource implications. Water is considered a public good under prior appropriation, and water rights are usufructuary. If downstream users in a basin are reliant on return flows, i.e., the water that returns to the system after human use, reducing upstream diversions by improving irrigation efficiency can impact their water availability. For example, a downstream user may rely on the return flows from seepage due to upstream flood irrigation. In some cases, it may not be clear if the adoption of efficient application technologies

improves system-level performance. An area of future research that warrants attention is evaluating how uncertainty in return flows impact the overall efficiency of a basin. Return flows can be difficult to track and can vary in their amounts depending on the crop being grown, soil type, weather conditions, and when the water is applied. This added uncertainty can make a system more difficult to manage, all else equal. High efficiency irrigation technologies however increase the control that a producer has to target water to a crop, which reduces the uncertainty in the value added from a unit of water that could have otherwise been applied with a low efficiency technology. Scrutinizing these uncertainties and understanding how incentives for efficient water use are aligned across producers within a basin is crucial for agricultural sustainability.

Despite some limitations, our results are generally informative and have important policy and water management implications. First, drought in arid regions is expected to worsen under a changing climate, and perceptions will play a critical role in the future adoption of conservation practices in agriculture. Neglecting how costs and benefits are perceived when assessing the effectiveness of programs designed to encourage conservation efforts could provide policy makers with misleading information. For example, if a policy maker is considering the implementation of a subsidy program to promote the adoption of water-conserving technologies, it is critical to understand whether non-adoption is driven by conventional cost hurdles or by perceptions about necessity. If the latter is the driving factor, efforts to accelerate revisions to perceptions to align with actual shortage distributions, *before* the realization of costly weather disruptions, could bolster a more efficient path to adoption. This may be an opportunity for agricultural extension programs to address and build perceptions about water scarcity in arid climates. Surveys and qualitative interviews can be administered to local farmers to gauge perceptions about climate change, drought risk, and the efficacy and necessity of adaptation strategies. If climate change risk is perceived as negligible, awareness campaigns tailored to communicating water scarcity concerns in localized areas may be effective at accelerating changes. If climate change is perceived as a real risk, communicating the benefits of increasing water use efficiency and providing better information on the possibility of future water shortages can enable producers to minimize their down-side risk. Highlighting

the conservation practices of local farming operations may also facilitate changes in perceptions, as the behavior of neighbors has been found to be influential in adoption behavior (Case 1992). Once climate change perceptions align with a need to improve water use efficiency, disseminating opportunities that reduce costs of implementation, such as EQIP participation, can hasten the path to adoption.

1.6 Tables

Table 1.1: Summary of characteristics of treatment and control diversion structures at 50% curtailment increase cutoff, 2001

Variables:	Control		Treatment	
	Mean	Std Dev	Mean	Std Dev
Irrigated Land (Acres)				
Flood Technology	1,043.41	3,695.46	2,777.73	5,759.35
Sprinkler Technology	272.60	1,453.47	907.31	2,848.66
Total	1,316.01	4,979.66	3,685.04	8,080.56
Groundwater Supplemented ^a	351.49	1,718.50	2,088.92	4,741.83
Crop Varieties (Acres)				
Corn	405.01	1,808.84	1,651.35	3,843.15
Alfalfa	452.90	1,990.10	1,166.14	2,586.79
Grass Pasture	239.37	501.71	339.72	701.21
Wheat	100.27	436.50	194.89	575.85
Other ^b	39.47	221.30	110.98	408.36
Water Rights Data ^c				
Appropriation Year	1880	13.18	1892	24.18
Number of Rights	5.92	11.24	2.94	3.63
County Population	211,001.8	136,786.1	147,749.5	122,041.8
Number of Structures	338		71	

^aHydroBase provides estimates of surface water irrigated acreage that is supplemented with groundwater.

^bOther crops include sugar beets, dry beans, and assorted vegetables.

^cRefers only to water rights with decreed agricultural uses.

Table 1.2: Difference-in-difference estimations of the impact of drought and institutional change on irrigation practices, 50% curtailment increase cutoff

Variables:	(1) Flood Acres	(2) Sprinkler Acres	(3) % Sprinkler	(4) Total Acres
Treatment*1976	445.0 (250.9)	-303.9 (191.0)	-0.027 (0.022)	141.1 (89.4)
Treatment*1987	202.5 (146.4)	-207.0 (115.0)	-0.012 (0.013)	-4.57 (69.1)
Treatment*1997	129.8 (74.5)	-134.9 (76.2)	0.007 (0.009)	-5.10 (35.8)
Treatment*2005	-420.6** (152.8)	262.6** (87.9)	0.048*** (0.012)	-158.0 (82.5)
Treatment*2010	-593.2** (199.6)	526.4** (166.3)	0.076*** (0.020)	-66.8 (79.2)
Treatment*2015	-892.7** (308.1)	743.8** (241.0)	0.115*** (0.026)	-148.9 (105.1)
Treatment*2020	-1,365.4** (516.2)	889.0** (293.0)	0.187*** (0.033)	-476.4 (273.0)
County Population	-0.001 (0.001)	-0.0007 (0.001)	-2.11×10^{-7} * (1.07×10^{-7})	-0.002*** (0.0004)
Fixed effects:				
Diversion Structure	✓	✓	✓	✓
Year	✓	✓	✓	✓
Observations	3,272	3,272	3,272	3,272
Adjusted R ²	0.879	0.797	0.721	0.987
Within Adjusted R ²	0.023	0.026	0.057	0.008

Diversion Structures: 409, Time Periods: 8, Reference Year: 2001

Standard errors (in parentheses) are clustered at the diversion structure level.

Signif. Codes: ***: 0.001, **: 0.01, *: 0.05

Table 1.3: Difference-in-difference estimations of the impact of drought and institutional change on crop-specific irrigation practices, 50% curtailment increase cutoff

Variables:	Corn		Alfalfa		Wheat	
	(1) Total Acres	(2) Sprinkler Acres	(3) Total Acres	(4) Sprinkler Acres	(5) Total Acres	(6) Sprinkler Acres
Treatment*1976	303.9 (190.8)	-85.0 (58.4)	-301.6 (167.2)	-189.5 (111.7)	-99.9 (55.5)	-33.2 (21.8)
Treatment*1987	95.2 (123.5)	-75.2 (49.0)	-369.3* (159.9)	-163.4 (92.6)	47.3 (36.8)	-24.2 (14.4)
Treatment*1997	0.686 (123.2)	-62.5 (60.8)	-236.4 (122.0)	-117.6 (66.8)	22.2 (31.9)	-9.03 (13.0)
Treatment*2005	-155.7* (69.5)	160.3** (51.7)	-152.3 (120.0)	20.8 (37.3)	-30.8 (34.0)	-5.11 (7.29)
Treatment*2010	-36.3 (54.6)	346.0** (104.5)	-229.3 (135.9)	53.4 (40.8)	88.5 (53.7)	73.9 (38.0)
Treatment*2015	-114.6 (81.2)	465.6** (147.2)	-164.3 (109.6)	147.2* (57.7)	46.0 (42.9)	76.8* (35.6)
Treatment*2020	-31.9 (65.6)	650.4** (216.7)	-92.8 (92.9)	237.1** (91.1)	-73.5 (65.8)	-6.61 (17.8)
County Population	-0.001** (0.0004)	-0.0005 (0.0005)	-0.0009** (0.0003)	-0.0002 (0.0005)	-0.0001 (0.0001)	1.5×10^{-5} (9.58×10^{-5})
Fixed effects:						
Diversion Structure	✓	✓	✓	✓	✓	✓
Year	✓	✓	✓	✓	✓	✓
Observations	3,272	3,272	3,272	3,272	3,272	3,272
Adjusted R ²	0.958	0.767	0.922	0.738	0.771	0.524
Within Adjusted R ²	0.008	0.033	0.007	0.015	0.010	0.022

Diversion Structures: 409, Time Periods: 8, Reference Year: 2001

Standard errors (in parentheses) are clustered at the diversion structure level.

Signif. Codes: ***: 0.001, **: 0.01, *: 0.05

Table 1.4: Robustness of difference-in-difference estimations when controlling for groundwater use

Variables:	(1) Sprinkler Acres	(2) Sprinkler Acres	(3) Sprinkler %	(4) Sprinkler %	(5) GW Acres
Treatment*1976	-303.9 (191.0)	-279.1 (170.8)	-0.027 (0.022)	-0.026 (0.022)	17.5 (33.2)
Treatment*1987	-207.0 (115.0)	-269.5 (136.7)	-0.012 (0.013)	-0.015 (0.013)	-44.0 (35.8)
Treatment*1997	-134.9 (76.2)	-169.0 (95.8)	0.007 (0.009)	0.005 (0.009)	-24.0 (19.5)
Treatment*2005	262.6** (87.9)	-125.4 (146.2)	0.048*** (0.012)	0.032* (0.013)	-273.5** (100.3)
Treatment*2010	526.4** (166.3)	330.5* (172.4)	0.076*** (0.020)	0.068*** (0.020)	-138.0 (85.7)
Treatment*2015	743.8** (241.0)	619.2** (224.2)	0.115*** (0.026)	0.109*** (0.026)	-87.8 (58.6)
Treatment*2020	889.0** (293.0)	576.7* (251.0)	0.187*** (0.033)	0.174*** (0.033)	-220.1* (102.6)
County Population	-0.0007 (0.001)	-0.002 (0.001)	-2.11×10^{-7} * (1.07×10^{-7})	-2.46×10^{-7} * (1.07×10^{-7})	-.0007** (.0002)
GW Acres	- -	-1.42*** (0.303)	- -	-5.9×10^{-5} *** (8.73×10^{-6})	- -
Fixed effects:					
Diversion Structure	✓	✓	✓	✓	✓
Year	✓	✓	✓	✓	✓
Observations	3,272	3,272	3,272	3,272	2,877
Adjusted R ²	0.797	0.828	0.720	0.727	0.987
Within Adjusted R ²	0.026	0.174	0.057	0.078	0.018

Diversion Structures: 409, Time Periods: 8, Reference Year: 2001

Standard errors (in parentheses) are clustered at the diversion structure level.

Signif. Codes: ***: 0.001, **: 0.01, *: 0.05

1.7 Figures




Figure 1.1: How perceptions of water availability impact the gross benefit of investment in improved irrigation technology




Figure 1.2: Monthly PDSI and frequency of calls by state engineer, Colorado Water Division 1




Figure 1.3: Treatment and control diversion structure map at 50% curtailment increase cutoff, Colorado Water Division 1




Figure 1.4: Difference-in-difference estimations of the impact of drought and institutional change on irrigation practices, 50% curtailment increase cutoff




Figure 1.5: Nonlinearities in the impact of drought and institutional change on irrigation practices as average increase in curtailment increases from 0% to 150%




Figure 1.6: Difference-in-difference estimations of the impact of drought and institutional change on crop choice, 50% curtailment increase cutoff




Figure 1.7: Difference-in-difference estimations before and after propensity score matching, sprinkler %

1.8 Supplementary Analysis

1.8.1 Scrutinizing the Treatment Design

To further investigate our treatment design at $X = 50\%$, we approximate the parameters defined in the theoretical model and impose the values on the heat map (Figure 1.8). Although treatment and control groups were determined by impacts during only drought years, we estimate perceptions using three 16-year periods that include both dry and wet years. Drought shocks are random, and a producer would develop perceptions about the probability of a call conditional on a variety of weather realizations. The estimate for the probability of a call (θ) is the average number of years a water right at a given structure was called during the period, divided by the length of the period. The estimate for the proportion of water received when called (δ) is the average of $(1 - \frac{\text{days under curtailment}}{\text{growing season days}})$ for the years in which a water right was called. From Figure 1.8 it is evident that most treatment structures shifted from darker to lighter areas in the period containing the shock (2000-2015), indicating a movement from low to high gross benefits from adopting a more water-efficient irrigation technology. The control structures do not exhibit the same movement. Although many control structures lie in an area that predicts a high benefit of adoption, our treatment designs aims to capture a change in perceived water availability. The relatively stable parameter estimates for the control group indicate they did not experience the drought shock to the same degree, as compared to prior droughts, as the treatment group. Next, as a placebo test, we examine the coefficient estimates for pre-treatment years as the treatment cutoff varies in Figure 1.9. Confidence intervals at most cutoffs overlap with zero, providing some evidence of adherence to the parallel trends assumption.

1.8.2 Addressing Threats to Identification

Here we address the possibility of multiple or staggered treatments. In our main econometric specification, we leverage the shift in the call regime as a singular treatment event. However, drought varies in intensity from year to year, and it is possible that some treatment structures were

impacted differently in years post-2002. If treatments are heterogeneous and staggered across time, then our model would be misspecified, and we would instead need to employ a difference-in-difference design suitable for estimating average treatment effects with two-way fixed effects and heterogeneous treatments (Callaway & Sant’Anna 2021). We explore this possibility graphically and further clarify the aim of our current treatment design. We first draw attention to the bottom right panel of Figure 1.8, which displays parameter estimates from the theoretical model for all treatment structures for the 2000-2015 time period. We find that all treatment structures move similarly on the heat map, in aggregate. Disaggregating to yearly impacts, Figure 1.10 shows the average days under curtailment for every diversion structure in our sample by year. Each point represents the average number of days under curtailment across all water rights associated with a given structure in that year. Grey points correspond to dry years ($PDSI < 0$) and black points correspond to wet years ($PDSI > 0$). From the bottom panel of Figure 1.10, it appears that treatment structures are consistently curtailed more in dry years post-2002 than any dry year previous. Although calls at different treatment structures may have varied across years, we argue that our treatment design aims to capture the singular and systematic shift in the way calls are administered post-2002. In other words, the treatment captures a change in perceptions about water supply certainty due to a distributional change in the call regime rather than year to year drought severity, which makes our model specification appropriate.

Next, we examine the robustness of our econometric results to adjustment of the reference year. The reference year of 2001 was chosen as it is the most recent year prior to treatment, but its omission inhibits the ability to fully investigate possible anticipatory behavior. If water right owners were anticipating the shift in the call regime and took action prior to its realization, the parallel trends assumption would not hold and our results could be biased. Figure 1.11, Figure 1.12, and Figure 1.13 present coefficient estimates with 95% confidence intervals for the main econometric models (Table 1.2) with reference years 1976, 1987, and 1997, respectively. For the Flood Acres, Sprinkler Acres, and Sprinkler % models, results do not change qualitatively across reference years and pre-treatment coefficient estimates are consistently insignificant. Most importantly, the

insignificant coefficient estimates for 2001 indicate no substantial behavioral differences between treatment and control diversion structures in the year immediately preceding the shock.

Lastly, we discuss the model specification, in particular the choice of year fixed effects versus a drought severity control variable (e.g., PDSI) across our study area. Localized weather conditions influence crop water demands as well as the amount of effective precipitation available to plants. As with much of the American West, the timing and quantity of surface water supplies available to producers for irrigation is based largely on winter (quantity of snow) and spring/early summer (timing of snowmelt) weather conditions in the mountains outside of our study area (as much as 200 miles away from the planting location and/or in another river basin). While PDSI is an effective measure of long-term (18 month) drought, it reflects local conditions. Within a time period, there is little variation across our study area, and the variation that does exist would not explain differences in surface water availability across producers (or the exogenous shock to perceived water right reliability that we observe). Our goal in this paper is to identify the effect of changing perceptions about surface water availability on producer behavior. Including year fixed effects allows us to estimate the effect of the unexpected shortages beginning around 2002 while controlling for differences in weather conditions (along with market conditions and any other factors constant across space) that may exist across the study period.

1.8.3 Nonlinear Impacts of Changing Perceptions

From the theoretical model, we determined that changing perceptions can impact adoption nonlinearly depending on the movement of θ and δ . The increase in expected gross benefit from an increase in θ could potentially be nullified by either an increase or decrease in δ , depending on the starting combination. Looking again at the right-hand column of figure Figure 1.8, we would expect treated units that moved to the lightest areas to have higher rates of adoption. We investigate this hypothesis informally by imposing total changes in sprinkler acreage (Figure 1.14) and sprinkler acreage as a percentage of total acreage (Figure 1.15) for each structure on the bottom

right panel of Figure 1.8.¹⁰ A larger point indicates a bigger increase in sprinkler technology adoption. In both figures, there appears to be a greater concentration of high adoption rates near the lighter areas, where the gross benefits of adoption are predicted to be highest.

¹⁰Change in sprinkler acreage as a percentage of total acre is calculated as a difference in percentage points.




Figure 1.8: Estimated perception parameters for control versus treatment diversion structures




Figure 1.9: Pre-treatment placebo test as average increase in curtailment increases from 0% to 150%




Figure 1.10: Average Days Under Curtailment for all Treatment and Control Structures by Year




Figure 1.11: Difference-in-Difference Estimations, Reference Year 1976




Figure 1.12: Difference-in-Difference Estimations, Reference Year 1987




Figure 1.13: Difference-in-Difference Estimations, Reference Year 1997




Figure 1.14: Nonlinear Impacts of Perceptions on the Adoption of Sprinkler Technology, Total Sprinkler Acreage




Figure 1.15: Nonlinear Impacts of Perceptions on the Adoption of Sprinkler Technology, Sprinkler Acreage as a Percentage of Total Acreage

Chapter 2

Assessing the Sensitivity of Climate Change Impacts in Agriculture to the Climate-Economy Functional Form

2.1 Introduction

Anthropogenic global warming has manifested in the increased frequency and severity of heat-waves and altered precipitation patterns (Stott 2016), both of which pose significant challenges to humans and the environment (Butler 2018). As such, the relationship between climate and the economy has garnered considerable attention over the past two decades in the applied economics literature. Recent advances in panel data econometrics and the advent of new geospatial weather datasets have allowed economists to exploit fluctuations in short-term weather, and to a lesser extent longer-term changes in climate, to discern an underlying climate-productivity relationship. Econometric estimates of such relationships, and the subsequent projections of future climate impacts, are used to inform the social cost of carbon as well as damage functions for integrated assessment models (National Academies of Sciences et al. 2017). Many studies focus on relationships between weather variables and broad economic indicators like gross domestic product

(Burke et al. 2015, Newell et al. 2021) or per capita income (Deryugina & Hsiang 2017), however there are limitations. For one, the choice of models and assumptions made by researchers can drastically influence predictions about future economic damages from climate change (Newell et al. 2021). Global studies may also suffer from inaccuracies in weather data, especially in developing countries (Auffhammer & Schlenker 2014, Michler et al. 2020), or oversimplification from averaging temperature data across diverse regions within a country. These challenges could result in misestimations of climate change's impact on economies, both in terms of the magnitude of the impact as well as the distribution of impacts across space.

Agriculture, fundamentally tied to climatic variables like temperature and precipitation, is particularly vulnerable to climatic change. There exists a robust literature on the impacts of climate change on agricultural productivity (e.g., Deschênes & Greenstone, 2007; Schlenker & Roberts, 2009; Walthall et al., 2013; Schauburger et al., 2017; Mérel & Gammans, 2021), and practitioners implement a diverse range of methodologies, sometimes yielding conflicting findings. While this breadth of research aims to inform mitigation and adaptation policies, the sizeable range in estimated marginal effects and future damages may not be helpful to policymakers. Additionally, research addressing the *distributional* effects of climate change on agriculture—i.e., heterogeneity in impacts across regions—remains limited. Identifying these distributional impacts is crucial for developing equitable policies and interventions that address the challenges faced by vulnerable farmers and communities.

Using agriculture in the conterminous United States (US) as an illustrative example, the primary purpose of the present research is to estimate the impacts of climate change on economic activity and assess the sensitivity of estimates to the functional form of weather. The agricultural and food sectors accounted for 5.4% of gross domestic product and 10.5% of employment in 2021 in the US (Kassel & Martin 2023). Given the significant contribution of the agricultural sector to both economic production and livelihoods, understanding the implications of climate change is paramount. The secondary purpose of this study is to evaluate the distributional impacts of climate change across US counties, provide evidence for the most appropriate model specifications,

and investigate the potential consequences of unsupported modeling choices for targeted policy intervention.

The availability of rich data sources for both weather and economic outcomes at a relatively fine spatial scale, combined with substantial heterogeneity in climates and economies in the lower 48 states, makes the US an ideal context for evaluating the climate-economy relationship. Furthermore, given the wide range of climates, the US provides a unique opportunity to evaluate heterogeneity in weather responses across diverse agricultural systems. The extent to which global climate change affects individual regions hinges on geographic location and the adaptive capacity of human and environmental systems. Certain regions may be more susceptible to the adverse effects of climate change due to their physical, social, or economic characteristics, while others may possess a greater ability to manage these changes. For example, regions with higher average temperatures or arid climates might experience intensified heatwaves and droughts, whereas coastal areas could be more vulnerable to flooding and sea level rise. As weather disruptions and their consequences are not evenly distributed, the effectiveness of policies aimed at enhancing climate resilience depends on local leaders, resource managers, and individual citizens understanding the climate risks in their own communities. In general, recognizing the distributional impacts of climate change across and within regions is critical for devising targeted adaptation and mitigation strategies.

To date, many studies assessing the economic damages resulting from climate change tend to estimate average impacts on a macro-scale, such as at the sector-, country-, or global-level, often overlooking the spatial heterogeneity of impacts within countries or regions. Such studies may offer general policy recommendations, such as seen in the Intergovernmental Panel on Climate Change "Summary for Policymakers" (IPCC, 2021), that are not be adequately tailored to the needs of decision-makers at local levels. For instance, in the US, risk management strategies in agriculture-dependent economics in the drought-affected, arid west will differ significantly from those addressing rain-fed regions. Consequently, the information needs of local policymakers in these distinct areas will also diverge considerably. By recognizing and accounting for such differ-

ences, tailored policy interventions can be designed to effectively address the unique challenges faced by individual regions.

The methods and findings presented in this paper advance the existing literature on climate change and agricultural productivity in several ways. First, the panel datasets employed in this analysis are among the most comprehensive to date, combining over 40 years of weather data with county-level economic outcomes from the Bureau of Economic Analysis (BEA). New agricultural data available from the BEA Regional Economic Accounts alleviates a previous dependence on the quinquennial Census of Agriculture, allowing variability in five times as much weather data to be exploited for identification, providing fresh and robust insights into the climate-agriculture relationship. I also draw from the climatology literature to generate idealized daily temperature curves, incorporating latitude and time of year, to interpolate hourly temperatures over the 42-year horizon and create more nuanced measures of exposure. Second, I estimate a series of two-way fixed effects regressions incorporating different functional forms of weather prevalent in the economics literature to assess the sensitivity of future impacts to model specification. Historical data limitations have precluded the estimation of complex and flexible econometric models due to lack of statistical power, which I address in this study. Future damages are then estimated using the latest projections from global climate models within the Coupled Model Intercomparison Project 6 (CMIP6). When projecting average annual impacts to the entire conterminous US, preferred models consistently estimate losses of \$10-20 billion in agricultural profits (15-30% of current levels) by the end of the century. However, when evaluating the *distributional* impacts of climate change, different model specifications produce inconsistent estimates both within and across counties. Model specifications including flexible, nonlinear terms that allow the marginal effects of temperature to vary across space reveal more diverse and intuitive impacts. Additionally, these flexible models tend to perform the best both in terms of model fit and out-of-sample prediction. Overall, this study offers invaluable insights for practitioners and policymakers seeking to analyze challenges posed by climate change within the agricultural sector.

The remainder of the paper is organized as follows. I first provide an overview of how climate and weather are defined in the context of this paper and the relevant literature as a whole¹¹, and provide a synthesis of the common modelling approaches for identifying weather and climate impacts in agriculture. I then outline the development of a comprehensive 42-year panel of county-level US weather data to replicate these models. Third, I compare specifications theorized to estimate both short-run responses to weather perturbations or long-run responses to climate change net of adaptation, depending on the literature viewed most convincing by the reader. Models are judged pragmatically in terms of interpretability of marginal effects, according to fit statistics and model selection criteria, and by out-of-sample prediction. I also estimate non-marginal changes in agricultural outcomes using climate projections to evaluate how predictions across model specifications vary in terms of total average and distributional impacts. Finally, I discuss implications for practitioners and policymakers.

2.2 Estimating the Impacts of Weather and Climate on Economic Activity

This section reviews the existing literature seeking to explore how weather and climate affect economic activity, with a particular focus on agricultural outcomes. It begins by distinguishing between climate and weather, clarifying the distinction between the long-term patterns of climate and the short-term fluctuations of weather, and their respective impacts on economic activity. Following this, I provide an overview of the methodologies used in current literature to isolate the effects of climate change from random weather variations. This includes a discussion on the challenges of differentiating between short-term weather shocks and long-term climate shifts, and how these factors are incorporated into econometric models to assess their impact on economic outcomes.

¹¹Definitions of climate and weather are largely synthesized from Hsiang (2016).

2.2.1 Climate versus Weather

Climate refers to the long-term patterns and averages of atmospheric conditions at a specific location. For any point in space, climate can be represented by a joint probability distribution which characterizes potential atmospheric conditions relevant to economic circumstances at that location, over a certain period. This distribution encompasses key parameters that detail potential conditions, like temperature, precipitation, and humidity. These parameters form a vector which signifies the *expected* conditions, defining the climate during an interval of time. In contrast, weather represents the conditions that are *realized* during a certain interval. Using the same time and space variables, weather is the empirical distribution that details the actual atmospheric states. For instance, weather might capture the actual average temperature over a month, while the climate would represent the theoretical average temperature range for that month.

Clarifying further, while weather provides a snapshot of atmospheric conditions experienced over short periods (like a single day or month), climate gives a comprehensive view of expected conditions over longer periods, often decades. Although historically, the average climate was sometimes defined over 30 years, this does not set a strict boundary. These expected conditions may even change yearly due to phenomena like ocean currents. This implies that while weather is the condition we experience daily, climate is the broader set of expectations for those conditions.

2.2.2 The Panel Approach

Climate change, using the definition above, describes a shift in the underlying distribution of weather patterns over an extended period of time. Periods of shifting weather patterns have occurred in the past, however current rates of change are more extreme than any other period in the last 65 million years (Diffenbaugh & Field 2013). While the economic impacts of climate change has been of interest since the early 1990s (e.g., Cline et al. 1992, Nordhaus 1992, 1994, Mendelsohn et al. 1994), the last decade has seen substantial growth in studies linking weather to economic outcomes. This modern literature chiefly uses panel data methods to identify the impacts of “climate change” from random variations in weather (Deschênes & Greenstone 2007,

Schlenker & Roberts 2009, Deryugina & Hsiang 2017, Newell et al. 2021), employing variations of the following baseline econometric model:

$$y_{it} = \mathbf{f}(\mathbf{X}_{it}; \beta) + \theta_i + \gamma_t + h_i(t) + \varepsilon_{it}, \quad (2.1)$$

where y_{it} is an economic outcome for spatial unit i in time period t , $\mathbf{f}(\mathbf{X}_{it}; \beta)$ is the functional form of realized weather outcomes X_{it} with β being the vector of parameters to be estimated, θ_i is a spatial unit fixed effect, γ_t is a time period or region-by-time period fixed effect, $h_i(t)$ is a spatial unit specific time trend, and ε_{it} is the error term. The primary goal of this approach is to saturate the model with fixed effects and time trends to control for omitted variables and isolate the average affect of weather on the outcome of interest.

In a model where all weather variables enter linearly, the two-way fixed effects estimator effectively takes the following form:

$$y_{it} - \bar{y}_i - \bar{y}_t + \bar{y} = \beta(X_{it} - \bar{X}_i - \bar{X}_t + \bar{X}) + \varepsilon_{it} \quad (2.2)$$

where β coefficients are identified from within-spatial unit deviations in weather from unit averages after adjusting for shocks common to all units in a given time period. A key contention in this approach is the tradeoff between credible causal inference through the use of fixed effects and estimating the long-run impacts of climate change (Lemoine 2018, Kolstad & Moore 2020, Mérel & Gammans 2021). A common criticism of Equation 2.1 is that the fixed effects absorb a potentially time invariant climate (seen in Equation 2.2 as $\bar{X}$) and therefore does not account for adaptation. In other words, panel models with fixed effects only capture responses to weather shocks, which may not be necessarily identical to permanent shifts in climate (Schlenker & Lobell 2010). To accurately assess the economic implications of *climate change*, as opposed to weather, it is important to estimate the impacts net of adaptation efforts undertaken by economic agents to counteract negative climate effects. Failing to do so may lead to overestimation of impacts, poten-

tially misinforming policymakers and resulting in an inefficient allocation of resources for further adaptation measures (Newell et al. 2021).

Difficulties in isolating the impact of climate on economic outcomes arise from both direct and indirect channels through which climate can affect economic activity. Weather realizations, such as heatwaves, storms, or droughts, can have immediate and direct impacts on economic outcomes. For example, a heatwave might damage crops, leading to a reduction in agricultural production. People can respond to these short-term weather events by taking adaptive measures, like applying more water to crops or otherwise adjusting their agricultural practices. However, the direct impacts of weather perturbations can still disrupt economic activity, despite short-term responses. Secondly, individuals can interpret weather events as informative signals, which can influence their beliefs about future climate conditions and risks. These beliefs, in turn, can shape long-run investments and adaptation measures. For instance, if individuals expect more frequent heatwaves or flooding in the future, they may invest in more resilient infrastructure, adopt new technologies, or change their land-use patterns. In the case of expected drought, Blumberg et al. (2024) find that perceptions of future water availability can drive investment in more efficient irrigation systems to mitigate risk. These long-term adaptation measures can impact economic outcomes indirectly by influencing productivity, labor markets, and resource use. Thus, given the direct and indirect pathways through which climate can influence society, the causal relationship between climate and economic outcomes is not easily identifiable.

Of studies that have employed Equation 2.1, it is sometimes argued that a nonlinear specification of weather variables can either completely or partially capture all effects of long-run adaptation. Deryugina & Hsiang (2017) present a theoretical framework that illustrates how the marginal effect of weather and the marginal effect of climate are equal when an economy is in competitive equilibrium. They therefore argue that within an appropriate context (and a sufficiently “large” panel of data) the marginal effect of climate on economic output can be exactly empirically identified using short-run fluctuations in weather. The authors demonstrate the equivalency of the marginal effects of weather and climate by use of the Envelope Theorem and Gradient Theorem,

and their finding is only appropriate when estimating the effects of climate on an outcome that constitutes an arguably *maximized* quantity. In the context of the present paper, their findings may apply to some of the explored outcomes.

Others acknowledge that incorporating nonlinearities in explanatory weather variables within a fixed effects panel regression may produce estimates that partially reflect long-run adaptation (Blanc & Schlenker 2017, Kolstad & Moore 2020).¹² The extent to which current panel methods can account for adaptation is formally reviewed by Mérel & Gammans (2021), who provide conditions under which long-run responses to climate can be captured. First, they argue that the common conception that panel fixed effects estimates are upwardly biased due to a failure to account for long-term adaptation is conceptually incorrect in a quadratic weather specification. They demonstrate that the bias is dependent on the distribution of weather variables' skewness and can be either positive or negative. Secondly, they show that the estimated coefficients from a panel fixed-effects model with a quadratic weather variable specification are a weighted average of long-term and short-term response functions. These estimates will be closer to the true long-term (short-term) parameters when the time-series variation in weather variables is small (large) relative to cross-unit climate variation. Thus, substantial differences in climate across units relative to weather shocks within units are required to capture the long-run effect of climate.¹³

In addition to the short-run weather versus long-run climate impacts debate, the functional form of Equation 2.1 takes many different flavors in the literature. This is largely due to the lack of consolidated theory on a "true" climate-economy relationship, so the empirical specifications of econometric models vary substantially across contexts. One common variation is the difference in the specification of the dependent variable, expressed in either levels (Schlenker & Roberts

¹²In fixed effects regressions, identification hinges on within-group variation. However, fixed effects models with quadratic specifications include a component of between-group variation that enters identification (McIntosh & Schlenker 2006).

¹³This is similar to the findings in Deryugina & Hsiang (2017), whose data requirements for identification are a "large panel of similar economies that (i) span the space of possible climates... (ii) are sufficiently densely packed in this space such that integration between positions is reasonable, and (iii) experience short-run weather disturbances that are not 'too large' in the sense that they do not perturb the economy so far from its equilibrium that the Envelope Theorem no longer applies."

2009, Deryugina & Hsiang 2017) or growth rates (Dell et al. 2012, Burke et al. 2015). An impact on levels might suggest a one-time shift or immediate impact due to a weather event or longer-term climate change, whereas an impact on growth rates suggests a compounding effect over time, which may be more or less concerning depending on the context. Newell et al. (2021) find that using levels versus growth rates can have a sizeable impact on the estimate of climate change impacts: changing size, sign, and significance. Also, acknowledging that these models are primarily used to predict future outcomes under different climate projections, the authors conduct a thorough analysis using different cross validation methods for model selection. Similarly, in the context of Italy, Olper et al. (2021) find that variation in model specification induces considerable uncertainty, however results tend to be more robust when an outcome variable is expressed in levels.

The questions of lagged effects of weather and how to properly account for unobserved variables have proven to be additional points of contention in the literature. While some scholars argue that lagged temperature and precipitation events impact growth and level outcomes (Dell et al. 2012), others contend only a contemporaneous influence (Burke et al. 2015). This divergence highlights the challenge in discerning the persistence and magnitude of weather's influence on economic activity. Regarding unobserved and potentially confounding factors, striking the right balance in controls can be difficult due to varying growth and productivity potential inherent to different countries and regions. Researchers typically implement some combination of parametric time trends and fixed effects. Region-specific time trends can account for some heterogeneity in endowments and institutions that may influence things like differences in technological growth, but may induce over-fitting (Newell et al. 2021). Region and time (and in some cases, region-by-time) fixed effects on the other hand flexibly account for common trends and time-specific shocks, but do not account for unit-specific trends. Moreover, although an abundance of fixed effects may aid in the identification of a plausibly causal effect, they may absorb too much variation across time and space to discern any relationship (Fisher et al. 2012). Notwithstanding the associated trade-offs, properly accounting for unobserved factors is crucial, as their omission can obscure the true relationship between climate variables and economic outcomes.

The final dimension in which Equation 2.1 can vary, and the central focus of the present paper, is the aggregation and transformation of weather data to generate the functional form of $f(\mathbf{X})$. Narrowing the focus to an agricultural context, Table 2.2 summarizes how the specification of weather variables varies across agricultural outcome variables in recent and prominent papers. Predominantly focused on temperature and precipitation, economists have long been interested in econometrically identifying the impact of weather on crop yields, land values, and farm profit. Higher temperatures can impact agriculture either by directly impeding plant growth or by diminishing labor supplies or labor productivity.¹⁴

2.3 Data

The advent of rich and spatially disaggregated data sources for both economic and climatic information has brought forth new opportunities for micro-scale analysis in the US. Government agencies and programs such as the BEA and Quarterly Census of Employment and Wages provide monthly, quarterly, and annual data on county-level employment, wages, income, and gross domestic product. Additionally, a suite of climate models offer daily meteorological data at a high-spatial resolution, both historical and projected.¹⁵ Combining these sets of data can provide useful information spanning decades. For this research, all weather and climate metrics are generated using gridMET (Abatzoglou 2013), a daily, 4-km x 4-km gridded weather dataset covering the conterminous US from 1979-now. The dataset merges the spatial characteristics of PRISM's gridded climate data with the temporal qualities and additional variables from the NLDAS-2 reanalysis. This integration yields a high-resolution, temporally consistent gridded dataset of surface mete-

¹⁴Graff Zivin & Neidell (2014) find that outdoor industry workers reduce labor time allocation at daily maximum temperatures above 85°F, effectively decreasing overall labor supply. Regarding labor productivity, Stevens (2019) show that higher piece-rate wages do not incentivize increased productivity at elevated temperatures for blueberry pickers in California.

¹⁵Currently there are many options for historical data, such as NASA's Daymet, the PRISM Climate Group out of Oregon State University, or the Climatology Lab.

orological variables. Validation of gridMET was performed against a broad network of weather stations, including RAWS, AgriMet, AgWeatherNet, and USHCN-2.¹⁶

The weather variables pertinent to this analysis include minimum and maximum ambient air temperature and precipitation accumulation. First, as gridMET provides gridded daily weather data, I compute an area-weighted mean of all grid cells that fall within a county's borders to generate county-day observations from which the weather metrics in Table 2.2 are created.¹⁷ From the county-day observations, I interpolate hourly temperatures by fitting an idealized temperature curve to minimum and maximum temperatures. The idealized daily temperature curve uses a sine curve for daytime warming and a logarithmic decay function for nighttime cooling. Daylight hours, a parameter in the interpolation process, are estimated from the time of year and a county's latitude.¹⁸ This process allows for the generation of both growing degree days and temperature bins.

Degree days (DD) are a measure of heat accumulation used by agronomists to estimate plant development rates. They are based on the idea that beneficial growth only happens within a certain temperature range: a minimum threshold temperature, below which growth is negligible, and a maximum threshold, above which the rate of growth does not increase and may even decrease. The focus in the economics literature centers around identifying the impact of growing degree days (ideal growing conditions, GDD) and harmful degree days (detrimental growing conditions, HDD). Each DD measure represents the cumulative number of temperature degrees per hour between two thresholds over a 24-hour period. Using the daily hourly temperature data, these measurements of

¹⁶For detailed information on the validation process (Abatzoglou 2013).

¹⁷This process may introduce more measurement error for some counties than others. Compared to eastern counties, western counties are often larger and contain a more diverse topography. Consequently, aggregating grid cells within the borders of a large county may not fully capture distinct climatic conditions, potentially resulting in less precise weather pattern representations.

¹⁸Hourly temperatures are interpolated using the chillR package in R (Luedeling & Fernandex 2022), which implements the methods described in Linvill (1990).

thermal-time can be expressed mathematically as:

$$CDD = DD_{(-\infty, H^L]} = \sum_{h=1}^{24} \max(0, H^L - H_h), \quad (2.3)$$

$$GDD = DD_{(H^L, H^U]} = \sum_{h=1}^{24} [\max(0, H_h - H^L) - \max(0, H_h - H^U)], \quad (2.4)$$

$$HDD = DD_{(H^U, \infty]} = \sum_{h=1}^{24} \max(0, H_h - H^U), \quad (2.5)$$

where CDD represents “cooling degree days,” H^L and H^U correspond to the lower and upper temperature thresholds outside of which plant growth may be harmed, and H_h denotes the hourly temperature at the h^{th} hour. The max function in these formulas ensures that only positive differences are considered. The choice of H^L and H^U are often data-driven and determined separately by crop type (e.g., Schlenker & Roberts, 2009).

From the hourly temperatures I also sum over all days in the growing season to calculate the total amount of time spent in 1°C temperature bins. The result is z_{it} , a 1-by- J vector of temperature exposures for county i in time period t , where J is the number of bins. Temperatures are binned from -55°C to 55°C (-67°F to 131°F), so

$$z_{it} = (z_{it,-55}, \dots, z_{it,55}), \quad (2.6)$$

where $z_{it,j}$ is the number of hours spent within $[j^\circ\text{C}, (j+1)^\circ\text{C})$. I then use z_{it} to compute wider temperature bins, denoted $HT_{it,k}$, that capture the number of hours spent within a sequence of bins $[\underline{j}, \bar{j})$. The expression for $HT_{it,k}$ is

$$HT_{it, [\underline{j}, \bar{j})} = \sum_{j=\underline{j}}^{\bar{j}} z_{it,j}. \quad (2.7)$$

In total, HT is calculated for $K = 14$ °C ranges: 1. $(-\infty, 0)$, 2. $[0, 3)$, 3. $[3, 6)$, ..., 12. $[30, 33)$, 13. $[33, 36)$, 14. $[36, \infty)$.

In a similar spirit to calculating time spent in temperature bins using the hourly data, bins are also generated for the number of days in which mean temperature or precipitation accumulation falls within a certain range. Daily mean temperature ranges, denoted $DT_{it,k}$, capture the number of days in which the daily mean temperature falls within bin k . In total, $DT_{it,k}$ is calculated for $K = 12$ °C ranges: 1. $(-\infty, 0)$, 2. $[0, 3)$, 3. $[3, 6)$, ..., 11. $[27, 30)$, 12. $[30, \infty)$. Precipitation ranges, denoted $P_{it,m}$, capture the number of days in which total precipitation falls within bin m . In total, $P_{it,m}$ is calculated for $M = 6$ mm ranges: 1. 0, 2. $(0, 1]$, ..., 6. $(4, 5]$, 7. $(5, \infty)$. Weather data is aggregated over different time intervals depending on the availability of other economic data.

Data on county-level farm income, crop revenue, and agricultural GDP from the BEA is available annually. These measures are based on Agricultural Census data and sample surveys administered by the US Department of Agriculture. The primary outcome variable of interest for this analysis is *net farm income*, which is measured as total cash receipts and other income less total production expenses. This measure includes additions to inventories and transfers such as subsidies and crop insurance payments and excludes sales *out* of inventory to create a measure that reflects only the value of a single year of production.¹⁹ Net farm income can be negative, so I normalize it by dividing by the number of agricultural acres within a county to generate *net farm income per acre* using data from NASS QuickStats. Crop revenue is similarly normalized by acreage and GDP via log transformation. Overall, I compile several panel datasets that vary both in time step and the number of counties for which data is available from 1980-2021.

Figure 2.1 exemplifies how different specifications of temperature capture different levels of variability both within and across counties by comparing Pima County, Arizona (inland, arid desert) with Orleans County, Louisiana (coastal, humid). Unsurprisingly, a simple measure like mean growing season temperature (top left panel) washes out a substantial amount of variability compared to the more comprehensive distribution generated from temperature binning (bottom

¹⁹This measure of farm profit partially circumvents the valid criticisms of using profit as an outcome variable, as discussed in Fisher et al. (2012). Because this measure intends to capture only the value of a single year of production, profit “smoothing” through storage that may attenuate estimates toward zero is not an issue here.

panels). The choice of temperature specification may be dependent on the outcome variable of interest. The traditional approach in assessing the weather's impact on agriculture has predominantly used degree days (top right panel), although this literature typically focuses on yields. When examining a more aggregate measure of productivity like income, degree days may not be as appropriate. One drawback of degree days is the dependence on single threshold values to define heat accumulation, a simplification that may not adequately reflect the thermal requirements and tolerances of different crops. Figure 2.2 displays idealized daily temperature curves again for Pima and Orleans county along with estimated temperature thresholds above which yields are harmed for certain crops. This figure illustrates that the ideal temperature for one crop may overlap with stress-inducing thresholds for another. Total farm income within a county depends on a portfolio of crops and production practices across operations that may be impacted by temperature in a more nuanced manner, and binned temperature measures allow impacts to vary across different temperature ranges. Furthermore, net farm income is not only driven by crop yields; it is also influenced by labor and capital productivity, both of which can be temperature-sensitive. High temperatures, for instance, may reduce worker productivity or increase the costs associated with cooling equipment and storage areas to maintain product integrity. These impacts may not be adequately captured by the static thresholds of classical degree days.

2.4 Empirical Methods

2.4.1 Base Model and Specifications of Weather Variables

To assess how different functional forms of weather influence estimates of climate impacts, I begin with the following base econometric specification:

$$y_{it} = \mathbf{f}(\mathbf{X}_{it}; \beta, \alpha) + \theta_i + \gamma_{rt} + \varepsilon_{it} \quad (2.8)$$

where y_{it} is the economic outcome of interest for county i in year t , and $\mathbf{f}(\mathbf{X}_{it}; \beta, \alpha)$ represents the functional form of the relationship between weather variables collected in matrix $\mathbf{X}_{it}$ and y_{it} , with β (temperature) and α (precipitation) being the vectors of parameters to be estimated. The term θ_i denotes the county-specific fixed effects that capture unobserved heterogeneity that is constant over time but varies across counties. This could include geographic, economic, or social characteristics that are unique to each county but do not change over the period of study. The term γ_{rt} denotes year fixed effects or region-by-year fixed effects, controlling for common shocks that affect all counties within a region during year t , such as abrupt changes in input and output prices. Lastly, ε_{it} is the idiosyncratic error term. Similar to the county-level analysis in (Deryugina & Hsiang 2017), standard errors are clustered in two dimensions: within counties and within state-by-year. Clustering within counties addresses the issue of autocorrelation within each county over time. This is critical because economic outcomes in a county are not isolated incidents and are often influenced by past events and conditions, leading to a natural correlation in observations across different years within the same county. Clustering within state-by-years aims to account for the spatial correlation that may exist across different counties within the same state. Counties in a particular state are likely to experience similar policy environments, economic conditions, and climatic influences, leading to potential correlations in their economic outcomes. By simultaneously clustering standard errors in these two dimensions, this approach accounts for both the spatial dependencies across counties within a state and the temporal dependencies within each county.

By saturating Equation 2.8 with fixed effects, this base specification aims to isolate the impact of weather variations on economic outcomes by controlling for unobserved confounders that could bias the estimates. It allows for a causal interpretation of the estimated parameters under the assumption that all other factors that could influence the economic outcome are either controlled for by the fixed effects or captured by the error term. Table 2.1 displays the different econometric specifications of $\mathbf{f}(\mathbf{X}_{it}; \beta, \alpha)$ compared in this analysis. The primary coefficients of interest across models are the β 's, which convey the effects of temperature on net farm income per acre. The omitted ranges for the daily and hourly binned models capture outdoor temperatures that lie within the

thermal comfort zone, or “room temperature” (Heal & Park 2020). To recover possible nonlinear relationships between temperature and a given economic outcome, models allow for nonlinearities across variable dimensions up to order L .

Following Deschênes & Greenstone (2007, 2012), different regional variations of the region-by-year fixed effect γ_{rt} are included across model specifications. Each regional aggregation captures unique geographical and climatic characteristics that may experience region-specific shocks:

1. **East versus West of the 100th Meridian:** This division represents the historical and climatological boundary between irrigated and rainfed agriculture, marking the transition from the humid eastern United States to the more arid West.
2. **Census Regions:** Defined by the US Census Bureau, these four regions are Northeast, Midwest, South, and West. The classification is based on geographic proximity and shared economic activities, reflecting broad climatic and agricultural practice zones.
3. **ERS Farm Resource Regions (FRR):** Nine regions established by the Economic Research Service. These regions are delineated based on county clusters that share similar agricultural practices, farm characteristics, topography, climate, and other factors relevant to farming and agricultural productivity.
4. **USDA Land Resource Regions (LRR):** Twenty-eight regions delineated based on detailed assessments of soil, climate, topography, and land use, intended to capture the most localized variations relevant to agricultural outcomes.
5. **Lower 48 States:** This delineation allows for capturing state-specific policies, economic structures, and environmental conditions that could uniquely influence agricultural outcomes and responses to climate change.

In employing region-by-year fixed effects, there are notable trade-offs to consider, particularly as the regions become increasingly disaggregated. On one hand, finer disaggregation allows the fixed effects to more precisely control for localized shocks, which is especially critical in contexts where climate and agricultural systems exhibit significant local variation. This increased specificity can enhance the accuracy of the model in capturing unique environmental, economic, or

policy-related factors that might otherwise confound the relationship between weather variations and economic outcomes. On the other hand, as the regions become more narrowly defined, the scope for exploiting inter-regional variation in weather diminishes. This reduction in variation can potentially limit the ability to identify causal relationships, as the model becomes more sensitive to idiosyncratic local factors and may overlook broader trends or influences. Consequently, while more granular region-by-year fixed effects enhance control over localized factors, they simultaneously reduce the degrees of freedom. Thus, the choice of regional aggregation level in the model should be balanced to optimize both the control over local confounders and the capacity to discern wider causal patterns. As such, the preferred models by intuition are those with FRR-year and LRR-year fixed effects, which attempt to strike a balance between capturing regional agricultural practices and environmental conditions that may experience similar shocks while maintaining enough variability across counties within a region for identification.

2.4.2 Simulating Future Impacts

The estimates obtained from various specifications of Equation 2.8 provide vectors of coefficients $\hat{\beta}$ and $\hat{\alpha}$ that encapsulate the average effects of weather on economic outcomes. These coefficients are used in conjunction with climate projections to simulate potential future changes in economic measures, such as aggregate farm profits. To quantify the average annual impact of projected future climate conditions on a county-by-county basis for net farm income, I employ the following expression:

$$\Delta y_i = \mathbf{f}(\Delta \bar{\mathbf{X}}_i; \hat{\beta}, \hat{\alpha}) \quad (2.9)$$

where $\Delta \bar{\mathbf{X}}_i$ is defined as the difference between the projected weather outcomes $\bar{\mathbf{X}}_{i,projected}$ and the historical averages $\bar{\mathbf{X}}_{i,historical}$. This differential, $\mathbf{f}(\Delta \bar{\mathbf{X}}_i; \hat{\beta}, \hat{\alpha})$, effectively estimates the county-level economic impact per acre under future climate scenarios (spanning 2060-2100) relative to historical conditions (1980-2020), under the assumption that no adaptation or mitigation efforts

are undertaken. The total average annual impact across all counties is then calculated as:

$$\Delta Y = \sum_{i=1}^N (\Delta y_i \times acres_i). \quad (2.10)$$

This aggregate measure, ΔY , provides an overall assessment of the climatic impact on the economic variable of interest, projected into the future. Future impacts are estimated using the SSP245 scenario, part of the Shared Socioeconomic Pathways (SSPs) used in CMIP6 climate models, which represents a middle-of-the-road pathway where the world follows a path where social, economic, and technological trends do not significantly deviate from historical patterns. Figure 2.3 displays the projected changes in daily mean precipitation and temperature during the growing season from 1980-2021 to 2060-2100 under seven different CMIP6 climate models. Averaged across all county-year observations, mean daily temperature during the growing season is projected to increase from 20.1 to 23.8°C using the ACCESS-CM2 model, which characterizes a middle-of-the-road warming projection. Changes in precipitation are far less pronounced.

To gauge how models may predict differently across space, I run a series of cross-sectional regressions:

$$\Delta y_i = \alpha + \beta G_i + \varepsilon_i \quad (2.11)$$

where Δy_i is the predicted impact to net farm income per acre in county i and G_i is a vector of county characteristics. The goal of this exercise is not to explain variation in Δy_i , but to assess the implications of targeted policy using predictions from different models. One can envision a system where payments to support climate adaptation measures are allocated to counties proportional to their predicted damages. If the models outlined in Equation 2 provide systematically different estimates within and across counties, then choosing the “wrong” model may inadvertently disfavor certain communities. In the worst case, the “wrong” model may underestimate impacts in counties that depend on agriculture for their total economic production. As such, G_i includes census region dummy variables as well as indicators of agricultural dependency from the most recent Census of

Agriculture: agricultural GDP as a percentage of total GDP, farm income as a percentage of total income, and farm employment as a percentage of total employment.

2.4.3 Model Selection

Due to uncertainties about which variation of Equation 2.8 best captures the data generating process, model performance is an important criterion to consider for model selection. The performance of these econometric models is critical as they are often used to inform policy by projecting future outcomes under varying climate scenarios. A traditional choice of judging model performance is through in-sample fit statistics such as adjusted R^2 , Akaike Information Criterion (AIC), and Bayesian Information Criterion (BIC), which impose parametric penalties for weakly correlated covariates. However, in-sample fit statistics can favor over-fitted models, especially when these models include high-ordered polynomials (as is the case in this analysis).

An alternative to *in-sample* fit statistics is the use of cross validation (CV), which involves dividing data into training and test sets to assess *out-of-sample* prediction. CV is non-parametric and avoids the *ad hoc* penalties of in-sample measures, arguably providing a more accurate reflection of a model's predictive accuracy. As such, I employ forecast, backcast, and K-fold CV approaches as outlined in Newell et al. (2021). The forecast approach involves estimating each model on the first 21 years of data and predicting into the latter 21 years of data, and vice versa for the backcast approach. One challenge in panel data settings is that CV methods cannot generate parameter estimates for fixed effects not present in training data, however Newell et al. (2021) propose a technique to handle this by removing them through demeaning both dependent and explanatory variables prior to estimation. The K-fold approach can account for fixed effects by splitting the data into training and test sets via random sampling, however this method neglects the time-series nature of the data. Compared to the forecast and backcast approach, K-fold may yield overly optimistic predictive power.

2.5 Results

This section analyzes how the marginal effect of temperature varies across different models and presents findings from model selection exercises. Primarily focusing on *net farm income per acre* as the dependent variable, I explore a total of 66 model specifications derived from 11 functional forms of weather and 6 levels of fixed effects. Additionally, this section discusses the projected impacts of climate change on agricultural productivity, highlighting how these impacts differ across regions and models.

2.5.1 The Marginal Effect of Temperature

I first present the results for the “simple” models using seasonal daily means and the more complex models using hourly temperature bins. The regression output displaying the impact of growing season daily mean temperature and precipitation on net farm income per acre, crop revenue (gross income) per acre, and agricultural GDP is available in Table 2.3.²⁰ For the linear and most naive specification, a 1°C increase in daily mean temperature is significantly correlated with a −\$3.14 impact on net farm income per acre, a −\$7.97 impact on crop revenue per acre, and a −6.7% point impact on agricultural GDP. For reference, mean net farm income and crop revenue per acre across the entire sample are \$74.44 and \$228.24, respectively. The difference in impact between net and gross income is may be partially explained by monetary payments from insurance and subsidies. For all three outcomes, including squared terms substantially improves model fit, the linear term switches sign, and the squared temperature variable is highly significant.

The coefficient estimates for the linear hourly temperature bin model ($L = 1$) are available in Figure 2.4, including 95% confidence intervals. The omitted categories for temperature and precipitation are 18 to 21°C (64.5 to 70°F) and 0mm, respectively. One benefit of this model specification is that it identifies specific temperature and precipitation ranges that adversely affect net farm income per acre. This is particularly relevant for farmers who may look at weather forecasts to

²⁰Agricultural GDP is obtained from the BEA using NAICS code 11 for Agriculture, Forestry, Fishing, and Hunting.

inform short-term damage mitigation strategies. Results suggest a steady increase in farm income with temperatures rising from below 0 to 29°C, after which income declines rapidly. Statistically significant findings indicate that each additional hour in temperatures above 33°C (91.5°F) reduces farm profits by approximately \$0.05 per acre, compared to the baseline. Considering the average farm size in the US in 2021 was 445 acres, this translates to a loss of \$22.25 *per hour* when temperatures exceed 33°C. The effects of different precipitation levels are less conclusive, though there is suggestive evidence that daily precipitation above 5mm can be detrimental relative to no rain.

When $L = 1$, $\hat{\beta}_{k,1}$ denotes the marginal effect of an extra hour in the k^{th} bin. Choosing $L = 1$ offers the advantage of straightforward interpretation, with $\hat{\beta}_{k,1}$ reflecting the average effect across all counties. However, setting $L > 1$ allows for marginal effects to vary based on each county's temperature profile, which might be relevant if a county accustomed to high temperatures reacts differently to additional temperature increases. For instance, with $L = 2$, the marginal effect in county i would be $\hat{\beta}_{k,1} + 2\hat{\beta}_{k,2}HT_{it,k}$, and for $L = 3$, it extends to $\hat{\beta}_{k,1} + 2\hat{\beta}_{k,2}HT_{it,k} + 3\hat{\beta}_{k,3}HT_{it,k}^2$. As reported in Figure 2.4, $\hat{\beta}_{[33,36),1} = -\0.051 . The calculated marginal effects for $L = 2$ and $L = 3$, using the sample average of 124 hours spent in the $[33, 36)^\circ\text{C}$ bin, are approximately $-\$0.079$ (calculated as $-0.098 + 2(7.56 \times 10^{-5}) \cdot 124$) and $-\$0.063$ (calculated as $-0.047 + 2(-0.0001) \cdot 124 + 3(1.96 \times 10^{-7}) \cdot 124^2$), respectively. At the sample average, it appears the linear model underestimates the magnitude of the negative effect assuming the nonlinear models better reflect the true data generating process. Conversely, the linear model overestimates the negative impact in warmer climates. The marginal temperature curves presented in Figure 2.5 effectively demonstrate that as a county experiences more time in hotter temperatures on average, the added adverse effect of an extra hot hour is less pronounced than in cooler counties. The dashed lines in the figure represent sample averages for Robeson County, North Carolina, and Pima County, Arizona, highlighting how the nonlinear models capture differing impacts in Arizona's desert climate compared to North Carolina's humid subtropical climate. I also present marginal effect curves for the three agricultural outcomes from the simple seasonal mean models in Figure 2.6. For each outcome, the linear model estimates a constant

negative marginal effect for a 1°C increase in temperature. Once the squared and cubed terms are included, the marginal effect is positive until approximately 16°C for crop revenue per acre and approximately 18°C for net farm income per acre. For agricultural GDP, temperature increases appear to be beneficial until approximately 15°C for the squared model. Assuming this 15-18°C daily mean temperature range is plausibly optimal, it is evident from Figure 2.3 that substantially fewer counties will lie in that range in 2060-2100.

Without considering formal model selection criteria, these findings suggest that flexible, non-linear models offer a more nuanced understanding of the impacts of temperature on agricultural outcomes. Compared to the linear models, which may oversimplify the relationship between weather and agricultural productivity, nonlinear approaches are better capable of capturing heterogeneity in responses across space, which one would expect to exist across landscapes in a country as geographically and climatically diverse as the US. Nonlinear models reveal that the effect of temperature on net farm income, crop revenue, and agricultural GDP is not uniformly negative or positive but varies across a temperature spectrum. By incorporating squared and higher-order terms, nonlinear models identify an optimal temperature range, beyond which agricultural outcomes begin to suffer. This capacity to delineate specific temperature thresholds for detrimental or beneficial impacts provides a more intuitive and accurate reflection of the true data generating process.

2.5.2 Climate Change Impacts on Agricultural Productivity

Using Equation 2.9, I estimate the impacts of future weather variability due to climate change on net farm income per acre for every county in the conterminous US. For simplicity, all estimates are generated using the ACCESS-CM2 model projection. Results are qualitatively similar for six additional model projections and are available in the Supplementary Analysis. Predicted future impacts with bootstrapped 95% confidence intervals are presented in Figure 2.7. The bottom panel reports the median impact per acre across all counties for each of the 66 model specifications. Generally, as region-by-year fixed effects become more disaggregated, the median impacts become

less negative. For the preferred models with FRR- and LRR-year fixed effects, median impact per acre ranges from $-\$13$ to $-\$21$ in net income, although the confidence intervals for each model prediction overlap substantially. The top panel reports the aggregated impact across all farmland in the conterminous US using 2017 acreage. Annual impacts range from $-\$13$ to $-\$18$ billion in agricultural profits for the preferred models, again with substantial overlap in confidence intervals across models. Overall, irrespective of the model specification, each set of fixed effects consistently forecasts a statistically similar and negative median or average aggregate impact.²¹

Thus far, choosing one functional form of weather over another appears fairly inconsequential in terms of estimating the *average* impacts of future weather variability. However, this is not the case for the *distributional* impacts. Figure 2.8 displays smoothed density curves of the predicted impacts per acre across all counties for each model specification with FRR-year fixed effects. Across all models, county-level estimates range from just below $-\$75$ to above $\$25$ per acre. While there is some overlap across the center of each distribution, the height and location of peaks and the width of tails vary considerably. This is perhaps unsurprising given that the nonlinear models allow for greater heterogeneity in the marginal effects of temperature, as displayed in Figure 2.5 and Figure 2.6.

Suppose a policymaker were to use the estimates from one of these models to provide targeted climate adaptation payments to vulnerable counties, i.e., those with predicted large and negative impacts. Choosing one model versus another will result in a different allocation of resources if models are producing systemically different estimates across space. Spatial disparities in estimates could then lead to an inefficient allocation of payments, particularly if the chosen model has less support in terms of theoretical backing, fit statistics, and/or predictive power than another. Table 2.4 reports the regression output from Equation 2.11, which shows how predicted impacts of climate change are correlated across census region indicators and measures of agricultural dependency.

²¹The exception to this are the models with state-by-year fixed effects, which produce estimates not statistically different from zero. However, given the notably low adjusted R^2 values in Figure 2.12 for these models, it is more plausible that these fixed effects are excessively absorbing variation, making it difficult to discern the impact of weather, rather than suggesting that weather has no impact after accounting for localized shocks.

Excluding the linear seasonal means model, the ranges of average impacts across models to net farm income per acre within each census region are as follows:

West: −\$6.48 to \$2.80 per acre

South: −\$35.62 to −\$16.88 per acre

Midwest: −\$12.50 to −\$7.90 per acre

Northeast: −\$4.39 to \$4.05 per acre.

Predicted impacts range similarly across measures of agricultural dependency, switching from negatively to positively correlated depending on model specification.

The variability in predicted impacts within and across counties highlights the need for prudent modelling choices when assessing the distributional impacts of climate change on agricultural productivity.²² In the worst-case targeted policy scenario, relying on an unsuitable model could overlook areas that are not only vulnerable to climate change but also heavily reliant on agriculture for their economic productivity. Such an oversight could exacerbate regional disparities and misdirect funds away from vulnerable communities.

2.5.3 Model Performance

As highlighted above, the disparities in model estimates across space warrants further scrutiny of model selection beyond economic intuition. The results reported below correspond to the case where *net farm income per acre* is the dependent variable. Equation 2.8 is estimated for 11 different functional forms of weather with 6 different levels of fixed effects for a total of 66 model specifications. The models are first judged by in-sample fit criteria. Figure 2.12 displays each model's AIC, BIC, and within adjusted R^2 values. The minimum AIC and BIC values across all models are subtracted from each model's reported value for easier interpretation and viewing. Per Burnham & Anderson (2004), models with a difference greater than 10 between their AIC or BIC value and the minimum value have no support as being the best model in the candidate set. Using

²²To further assess distributional impacts, I estimate an ensemble of fully-interacted regional models to discern how the impacts of temperature and precipitation vary across different climates and agricultural systems within the US. Results from this exercise are available in the Supplementary Analysis.

AIC, the hourly binned model with $L = 3$ (cubic) is preferred across all levels of fixed effects. This is also the most complex model with 57 covariates. Using BIC, which incorporates a larger penalty for model complexity, there is support for both hourly binned and seasonal mean models depending on the included fixed effects. With more disaggregated region-by-year fixed effects, the seasonal mean models with $L > 1$ are the best candidates. In terms of within adjusted R^2 , the cubic hourly binned model appears to explain the most in-sample variation in net farm income per acre. However, it is important to note that the values are exceptionally small, indicating that there is a minimal amount of additional variation explained by the covariates after saturating the models with fixed effects. In general, in-sample fit statistics favor the flexible hourly binned or seasonal mean models.

For the CV exercises, a root mean squared error (RMSE) is calculated for each model prediction, where $RMSE = \sqrt{\frac{1}{n} \sum_{i=1}^n (y_{it} - \hat{y}_{it})^2}$. Models with the lowest RMSE values display the best out-of-sample prediction. Forecast, backcast, and K-fold results are presented in Figure 2.13. In terms of backcast prediction, the top performing models are the daily and hourly binned models with squared terms and demeaned by the FRR-by-year fixed effects. However, several other model specifications perform similarly—the top ten models have RMSE values that range from \$107 to \$117. For forecast prediction, the hourly binned model with linear terms demeaned by meridian-by-year, year, and FRR-by-year fixed effects have the three lowest RMSE, respectively, followed by the hourly binned model with quadratic terms and demeaned by FRR-by-year. The RMSE values for these top four models range from \$88 to \$107, and the next lowest RMSE is \$117. For both methods that involve demeaning the data by fixed effects, the cubic models have a tendency to overfit the training data and provide worse out-of-sample predictions. For the K-fold CV method that can handle both county and time fixed effects, this is not the case, and the hourly binned models with squared and cubed terms result in the lowest RMSE. However, it is important to note that RMSE values are similar across all model specifications when using this method. This is perhaps unsurprising given the low within adjusted R^2 values presented in Figure 2.12, indicating that the within-county and -year averages explain most of the variation in the data.

To summarize the results above, Figure 2.9 displays the best performing weather specification and runner up across all model selection criteria, delineated by the level of fixed-effects used. Tiles are colored according to which model resulted in the lowest RMSE for out-of-sample prediction, lowest AIC and BIC values, and highest within-adjusted R^2 . Tiles with black stripes correspond to models that use the novel hourly temperature data developed for this study. The hourly binned temperature specifications perform best in 24 of the 36 possible outcomes, and 19 of 36 for second best. In stark contrast, the degree day specification, exceptionally prevalent in the contemporary literature, performs best 1 of 36 times and is never a runner up. These results perhaps suggest that with new opportunities to exploit variation in large datasets, models should in turn be more flexible in identifying relationships between weather and economic outcomes.

2.5.4 Consequences of Modeling Choices for Targeted Policy and the Social Cost of Carbon

Given the intuitive marginal effects outlined in Figure 2.5 and model performances outlined in Figure 2.9, it appears that the nonlinear hourly binned models are the best candidates for identifying the impact of temperature and precipitation on net farm income. Figure 2.10 displays the county-level differences in projected impacts between the Hour Bins, Cubic and Growing Degree Days, 10-32 model specifications under six CMIP6 models. Assuming the hourly binned models is a better reflection of the true data generating process, these maps effectively demonstrate the misallocation of resources if a policymaker were to design targeted climate policy based on the literature-prevalent growing degree days model. This hypothetical misallocation would result in insufficient resources allocated to counties in the Intermountain West (particularly Southern Arizona and California), central and southwest Texas, and much of the Southeast and eastern Heartland. Conversely, resources would be overallocated to northern counties and counties along the Pacific coast. These results are consistent across all six climate projections.

The implications of modeling choices also extend beyond geographic discrepancies. Figure 2.11 illustrates the variability of climate impact estimates across income deciles, showing

the potential for distributional inequalities when policy is informed by different models. This heterogeneity of impact estimates is particularly pronounced for the lowest income decile, where the range of negative impacts is consistently wide across all climate projections. Such disparities in projected impacts across models may have implications for equity weighting in the social cost of carbon (SCC). Policymakers advocating for a more equitable approach to climate change mitigation have suggested placing higher equity weights on more vulnerable groups, such as lower-income populations, to reflect their reduced capacity to adapt and higher relative suffering. This adjustment aims to ensure that the social cost of carbon fully captures the increased burdens borne by these groups, potentially leading to a more just climate policy framework. Any consistent over- or underestimation of damages to more or less vulnerable groups by certain models can skew the effectiveness of equity weights. If a model systematically underestimates the damages to lower-income deciles, it could lead to a miscalculation of the SCC, underrepresenting the true economic and social burdens of climate change on these communities. Conversely, overestimation could lead to an inflated SCC, potentially resulting in unnecessarily stringent measures that could have unintended economic repercussions.

2.6 Conclusion

In recent years, the pursuit to econometrically determine the impacts of weather and climate on economic outcomes has intensified. A variety of empirical methodologies have been used in this endeavor, leading to continued discourse on the suitability of different aspects of these approaches. Although considerable attention has been paid to the appropriate use of lags, the choice between growth and level outcomes, and the application of fixed effects and time trends, the functional form of weather variables has received comparatively less focus. The emergence of rich meteorological and economic datasets has enabled the creation large panels capable of supporting complex models without compromising statistical power, and as such, specifications that require a large number of covariates warrant exploration.

This study examines the impacts of climate change on agricultural productivity, with a focus on *net farm income* in the United States, and I employ a suite of model specifications to assess the sensitivity of the estimated effects of temperature and precipitation. Congruent with the existing literature, nonlinear specifications that allow marginal effects to vary across counties have the most support in terms of model selection and theoretical backing. Although the degree day functional form is the most prevalent in the climate-agriculture economics literature, in this analysis I find little to no support for that specification. This is likely due to dependence on static threshold values to define heat accumulation, a simplification that may be less appropriate for more comprehensive outcomes like profit versus crop-specific yields.

A key observation is that no single model consistently excels across all selection criteria, although all reasonably defensible models predict a consistent aggregate negative impact on US agricultural productivity due to climate change: $-\$10$ to $-\$20$ billion in annual agricultural profits (15-30% of current levels) by 2060-2100. Although I find that the choice of the functional form of weather appears to be relatively inconsequential in estimating *average* impacts, another notable observation is that the estimated *distributional* impacts, i.e., within and across counties, vary substantially across models. This is an important finding if the end use goal of these models is to inform targeted policy endeavors.

Additionally, climate-economy relationships are likely highly nonlinear nature due to the interplay between both biophysical and behavioral responses, which is supported by the intuitive marginal effects and model selection results for the more flexible model specifications presented in this study. The most complex models, incorporating over 60 covariates and thousands of fixed effects, could not have been feasibly or meaningfully estimated with prior data sources purely due to lack of sample size. For the present analysis, the use of new agricultural data from the BEA Regional Economic Accounts, in place of the traditionally used quinquennial Census of Agriculture, alongside advanced meteorological datasets, allows over five times as much variability in weather and economic outcomes to be leveraged for identification compared to previous, comparable studies. Furthermore, similar economic data sources are not unique to the BEA, and new

satellite imagery and geospatial data products are in continual development. As such, the current wealth of data products now allow the creation of large panels that permit the estimation of more sophisticated statistical models that arguably better reflect the true data generating process.

In summary, I find that increases in temperature positively affect farm income up to a certain threshold, after which the effects become negative. Specifically, temperatures above 33°C significantly reduce farm profits. The impact of precipitation is less definitive, but excessive rainfall appears to have adverse effects. I also find notable regional variation in the impacts of climate change, particularly for the specifications that allow coefficients to vary across space. Given some uncertainty about which model most accurately reflects the true data generating process, I develop a composite damage measure in the Supplementary Analysis to provide a comprehensive measure of risk, identifying counties consistently and severely affected by climate change.

The regional heterogeneity in climate change impacts highlights the need for localized policy and resource management strategies. Using the county as the unit of analysis has the benefit of communicating information already conformed to a common organizational unit for existing institutions. For example, the US Department of Agriculture's (USDA) Farm Service Agency, Natural Resource Conservation Service, and the Rural Development agency already tailor program characteristics to meet county-specific needs, and USDA Service Centers designed for community access are available in most counties.²³ Counties also typically have emergency managers who oversee planning and responses to disasters. Results and methods presented in this study can assist in designing more effective and targeted policies.

²³See <https://offices.sc.egov.usda.gov/locator/app>.

2.7 Tables

Table 2.1: Econometric specifications of $f(\mathbf{X}_{it}; \beta)$

Seasonal Means	$\sum_{l=1}^L \left[\beta_l (temp_{it})^l + \alpha_l (prcp_{it})^l \right]$
Degree Days	$\beta_1 DD_{it,(-\infty, 10]} + \beta_2 DD_{it,(10, H^U]} + \beta_3 DD_{it,(H^U, \infty]} + \sum_{l=1}^L \left[\alpha_{m,l} (P_{it,m})^l \right]$
Day Bins	$\sum_{l=1}^L \left[\sum_{k=1}^K \beta_{k,l} (DT_{it,k})^l + \sum_{m=1}^M \alpha_{m,l} (P_{it,m})^l \right]$
Hour Bins	$\sum_{l=1}^L \left[\sum_{k=1}^K \beta_{k,l} (HT_{it,k})^l + \sum_{m=1}^M \alpha_{m,l} (P_{it,m})^l \right]$

L : degree of nonlinearity allowed across variable dimensions

$temp$ and $prcp$: growing season daily mean temperature and precipitation

$DD_{(H^L, H^U]}$: degree days accumulated between thresholds H^L and H^U

DT_k : days spent in k^{th} 3°C-wide temperature bin (12 to 15°C omitted)

P_m : days spent in m^{th} 1mm-wide precipitation bin (0mm omitted)

HT_k : hours spent in k^{th} 3°C-wide temperature bin (18 to 21°C omitted)

Table 2.2: Common functional forms of weather in the economics literature

Specification of $f(\mathbf{X})$	Agricultural Outcomes		
	Crop Yields	Land Values	Profit/Revenue/GDP
Seasonal/Monthly Means	Mérel & Gammans (2021)	Mendelsohn et al. (1994)	Olper et al. (2021)
Growing Degree Days	Schlenker & Roberts (2009)	Schlenker et al. (2005, 2006)	Deschênes & Greenstone (2007, 2012)
Daily Means (binned)	×	×	Deryugina & Hsiang (2017)
Hourly Measures (binned)	Schlenker & Roberts (2009)	×	×

Table 2.3: The impact of growing season daily mean temperature and precipitation on agricultural productivity

Model:	Net Farm Income per Acre			Crop Revenue per Acre			log(Ag. GDP)		
	$L = 1$	$L = 2$	$L = 3$	$L = 1$	$L = 2$	$L = 3$	$L = 1$	$L = 2$	$L = 3$
Temperature	-3.14** (1.48)	22.2*** (5.19)	3.05 (12.2)	-7.97*** (2.03)	35.4*** (7.89)	-11.2 (20.8)	-0.067*** (0.018)	0.135* (0.072)	-0.311 (0.231)
Temperature ²		-0.615*** (0.130)	0.418 (0.707)		-1.07*** (0.202)	1.44 (1.20)		-0.005*** (0.002)	0.019 (0.013)
Temperature ³			-0.018 (0.013)			-0.043** (0.022)			-0.0004* (0.0002)
Precipitation	-1.05 (0.943)	11.7*** (2.92)	5.58 (5.17)	-5.43*** (1.24)	8.95** (4.04)	0.326 (7.64)	-0.006 (0.010)	0.040 (0.033)	0.080 (0.070)
Precipitation ²		-1.74*** (0.382)	-0.114 (1.27)		-1.99*** (0.531)	0.227 (1.92)		-0.006 (0.004)	-0.017 (0.015)
Precipitation ³			-0.132 (0.104)			-0.174 (0.160)			0.0008 (0.001)
<i>Fixed effects</i>									
FIPS	✓	✓	✓	✓	✓	✓	✓	✓	✓
FRR-by-year	✓	✓	✓	✓	✓	✓	✓	✓	✓
Observations	121,910	121,910	121,910	122,284	122,284	122,284	51,400	51,400	51,400
Within R ²	0.00028	0.00283	0.00297	0.00159	0.00484	0.00516	0.00126	0.00222	0.00260

Two-way clustered (FIPS & State-by-year) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table 2.4: Cross-sectional regression results assessing the distributional impacts of predicted climate change impacts on county net farm income per acre by model

Model:	Seasonal Means			Day Bins			Hour Bins			Degree Days	
	$L = 1$	$L = 2$	$L = 3$	$L = 1$	$L = 2$	$L = 3$	$L = 1$	$L = 2$	$L = 3$	10-32	10-35
Constant	-12.8*** (0.133)	1.62*** (0.609)	1.28* (0.729)	-0.803** (0.402)	2.80*** (0.528)	1.59*** (0.553)	-2.20*** (0.433)	-1.48** (0.578)	-3.52*** (0.701)	-5.03*** (0.468)	-6.48*** (0.402)
Ag./Tot. GDP	-0.009 (0.012)	0.067 (0.054)	0.085 (0.065)	-0.074** (0.036)	-0.087* (0.047)	-0.130*** (0.049)	-0.132*** (0.039)	0.102** (0.051)	0.180*** (0.062)	-0.010 (0.042)	0.051 (0.036)
Farm/Tot. Inc.	-0.021 (0.023)	-0.160 (0.107)	-0.148 (0.128)	0.017 (0.071)	0.053 (0.093)	0.109 (0.097)	0.221*** (0.076)	-0.080 (0.102)	-0.182 (0.123)	-0.127 (0.082)	-0.174** (0.071)
Farm/Tot. Emp.	-0.062*** (0.010)	-0.222*** (0.044)	-0.271*** (0.053)	-0.036 (0.029)	-0.058 (0.038)	-0.011 (0.040)	0.025 (0.031)	-0.331*** (0.042)	-0.466*** (0.051)	-0.291*** (0.034)	-0.313*** (0.029)
South	0.665*** (0.143)	-32.5*** (0.654)	-36.9*** (0.783)	-25.1*** (0.432)	-32.2*** (0.567)	-32.3*** (0.594)	-25.3*** (0.466)	-20.4*** (0.621)	-18.0*** (0.753)	-14.6*** (0.503)	-10.4*** (0.432)
Midwest	-0.415*** (0.147)	-13.7*** (0.676)	-12.1*** (0.809)	-11.7*** (0.446)	-13.4*** (0.586)	-12.7*** (0.613)	-10.1*** (0.481)	-6.42*** (0.641)	-5.67*** (0.778)	-3.77*** (0.519)	-3.09*** (0.446)
Northeast	-0.973*** (0.214)	-4.87*** (0.981)	-2.50** (1.17)	-3.59*** (0.647)	-4.53*** (0.851)	-2.79*** (0.890)	1.09 (0.698)	4.47*** (0.931)	7.57*** (1.13)	4.74*** (0.754)	2.27*** (0.648)
Observations	3,068	3,068	3,068	3,068	3,068	3,068	3,068	3,068	3,068	3,068	3,068
R ²	0.089	0.542	0.546	0.605	0.604	0.591	0.605	0.415	0.311	0.389	0.305
Adjusted R ²	0.086	0.541	0.545	0.604	0.602	0.590	0.605	0.414	0.310	0.387	0.303

Signif. Codes: ***: 0.01, **: 0.05, *: 0.1

2.8 Figures




Figure 2.1: Differences in temperature measures between Orleans County, Louisiana and Pima County, Arizona. Each plot shows the mean of each weather variable across all years in the sample for each county.




Figure 2.2: Estimated upper temperature thresholds for harmful degree days for different crop types. The curves represent estimated idealized daily temperature curves for an average summer solstice for each county. The cross marks correspond to daily mean temperatures.




Figure 2.3: Changes in the distribution of daily mean temperature and precipitation during the growing season from 1980-2021 to 2060-2100 across all counties. Future climate data is from CMIP6 model projections under the SSP245 scenario.




Figure 2.4: Estimated marginal effects of temperature and precipitation on net farm income per acre. The coefficient point estimates with 95% confidence intervals correspond to the linear ($L = 1$) hourly bin specification with ERS FRR-by-year fixed effects.




Figure 2.5: How the marginal effect of an extra hour spent in the $[33, 36)^{\circ}\text{C}$ temperature bin changes as a county's climate becomes warmer by model specification.




Figure 2.6: How the marginal effect of a 1°C mean growing season temperature increase changes as a county's climate becomes warmer by model specification.




Figure 2.7: Comparison of predicted climate change impacts on US agricultural profits by model specification over 2060-2100 with bootstrapped 95% confidence intervals. Future climate data is from the CMIP6 ACCESS-CM2 model projection under the SSP245 scenario.




Figure 2.8: Comparison of predicted distributional impacts of climate change on county net farm income per acre by model specification with county and FRR-year fixed effects. Each density curve represents all point estimates of the average annual change in net farm income every county in the sample over 2060-2100. Future climate data is from the CMIP6 ACCESS-CM2 model projection under the SSP245 scenario.




Figure 2.9: Comparison of model performance by model fit and out-of-sample prediction. Tiles are colored by the best performing model by the following criteria: lowest RMSE for all out-of-sample prediction methods, lowest AIC and BIC, and highest within-adjusted R^2 .




Figure 2.10: Differences in projected impacts to net farm income per acre between the Hour Bins, Cubic and Growing Degree Days, 10-32 model specifications under six CMIP6 model projections. All future climate data assumes the SSP245 scenario.




Figure 2.11: Range of projected impacts across models to total net farm income by net farm income decile under six CMIP6 model projections. All future climate data assumes the SSP245 scenario. The range of impacts correspond only to models estimated with ERS Farm Resource Region-by-year fixed effects and omits the Seasonal Means, Linear specification.

2.9 Supplementary Analysis

2.9.1 Model Selection Criteria




Figure 2.12: Comparison of fit statistics by model specification for net farm income per acre.

2.9.2 Crop Revenue and Agricultural GDP

Figure 2.14 and Figure 2.15 display the aggregate projected impacts of climate change on crop revenue and agricultural GDP using estimated coefficients from the econometric specifications outlined in Table 2.1. In general, models produce qualitatively similar estimates. Interestingly, negative impacts to agricultural GDP are estimated to be the largest for models using ERS FRR-by-year and USDA LRR-by-year fixed effects compared to others.

2.9.3 Climate Model Sensitivity

To investigate the sensitivity of my results to the choice of CMIP6 model, Figure 2.16, Figure 2.17, and Figure 2.18 display projected impacts to net farm income per acre, aggregate agri-




Figure 2.13: Comparison of cross validation performance by by model specification for net farm income per acre.

cultural profits, and distributional impacts for six additional climate models, respectively. There is slight variation in the magnitude of impacts due to different warming conditions embedded within each climate projection, however results are qualitatively similar to the main findings presented in this paper.

2.9.4 Regional Models

Although the preferred specification of Equation 2.8 controls for local shocks through FRR-by-year fixed effects, the model still produces a single set of coefficient estimates for all counties. Given the heterogeneity in climate and agricultural practices across the US, it may be insufficient to assume that these coefficients accurately reflect the varied responses to temperature and precipitation changes in different regions. For example, agricultural systems in counties west of the 100th meridian tend to use irrigation, which can mitigate the impact of extreme temperatures through evaporative cooling and maintaining soil moisture. Therefore, the impact of an additional hour in the 33-36°C temperature bin will likely differ between counties in Southern California’s Imperial Valley and rain-fed counties in the Corn Belt.

I extend Equation 2.8 by allowing β and α coefficients to vary across diverse regions through a series of interacted models that take the form:

$$y_{it}^{(k)} = \sum_{j=1}^J \left[\mathbf{1}(R_i^{(k)} = R_j^{(k)}) \times \mathbf{f}(\mathbf{X}_{it}; \beta_j^{(k)}, \alpha_j^{(k)}) \right] + \theta_i + \gamma_t + \varepsilon_{it}, \quad k = 1, \dots, K, \quad (2.12)$$

where $y_{it}^{(k)}$ represents the economic outcome for county i in year t under the k -th regional aggregation, and $\mathbf{1}(R_i^{(k)} = R_j^{(k)})$ is an indicator function that takes the value 1 when the county i belongs to the j -th region denoted by $R_j^{(k)}$, and 0 otherwise. This allows the function $\mathbf{f}(\mathbf{X}_{it}; \beta_j^{(k)}, \alpha_j^{(k)})$ to capture the effects of weather outcomes that are region-specific, with $\beta_j^{(k)}$ and $\alpha_j^{(k)}$ representing the set of parameters that vary by region and aggregation level k . The term θ_i continues to control for county fixed effects and γ_t is a year fixed effect. Finally, ε_{it} is the idiosyncratic error term that captures unexplained variation in the economic outcome after controlling for the fixed effects and the region-specific weather impacts. By interacting the region dummies with the functional form of weather outcomes, Equation 2.12 allows for a flexible assessment of how climate impacts on the economic outcome might differ across various regional climates and geographies. Equation 2.12 is estimated across 11 different functional forms of $\mathbf{f}(\mathbf{X}_{it})$ for $K = 4$ regional aggregations. The regional aggregations mirror the variations in region-by-year fixed effects in Equation 2.8, excluding states.

Regional Heterogeneity in Weather Effects

The regional models outlined in Equation 2.12 permit weather parameters to vary across regions, allowing for more heterogeneity in temperature responses across space. As evidenced in Figure 2.6, the “optimal” mean growing season temperature for agricultural productivity is somewhere from 15-18°C in the model that estimates a single set of coefficients for all counties. When coefficients are able to vary across regions, that optimal range may vary as well. Figure 2.19 shows how the marginal effect of temperature differs for counties east and west of the 100th meridian line, the historical dividing line for rain-fed (east) and irrigated (west) agriculture. The two panels on

the left correspond to linear hourly bin specification. Consistent with the notion that irrigation can mitigate the adverse affects of extreme temperatures, net farm income in western counties does not suffer until temperatures exceed 36°C (96.8°F). Conversely, eastern counties experience decreases in net farm income at temperatures exceeding 29°C (84°F). Compared to the results presented in Figure 2.4, the effects of hotter temperature are overestimated for western counties and underestimated for eastern counties. The panel on the right of Figure 2.19 displays the marginal effect curves generated from coefficients from the cubic growing season daily mean specification, and it suggests that the optimal temperatures for eastern and western counties are 17°C and 21°C , respectively.²⁴

Equation 2.12 attempts to address heterogeneity in climate and agricultural practices by capturing different possible responses across space. This creates the potential for the negative effects of hot temperatures to be less severe in counties that have better adaptation measures (e.g., irrigation in the west) and more severe elsewhere. Figure 2.20 displays the aggregate annual impacts on agricultural profits estimated from the regional models, and each column corresponds to the regions over which coefficients can vary. Remarkably, the overall estimated impact of climate change on agricultural profits in the US when using Equation 2.12 remains consistent with the results obtained from Equation 2.8. Most models produce estimates still ranging from $-\$15$ to $-\$25$ billion. However, and perhaps unsurprisingly, Equation 2.12 reveals far more heterogeneity in climate responses across counties. Figure 2.21 displays smoothed density curves of the predicted impacts per acre across all counties for four different functional forms of whether across different regional aggregations. In each case, the regionally interacted models produce multi-modal distributions, often with a larger portion of counties benefiting from climate change when compared to the estimates from Equation 2.8.

The projected impacts in each county for the linear hourly temperature bin model are also presented on maps in Figure 2.22. The top left map presents impacts from the specification that

²⁴The curve for western counties may also suggest an optimal mean growing season temperature at 13°C , however this would require a county's climate to cool to most cases.

produces a single set of coefficients for all counties, most of which are non-positive. Moving to the regional models, there is an increased prevalence of counties that may potentially benefit from climate change. Across the regional models, counties located in the northeast, northwest, and the Intermountain West (excluding the Southwest) are consistently predicted to see marginal benefits in net farm income per acre.

2.9.5 Characterizing Vulnerability

A salient goal of scrutinizing the distributional impacts of climate change is to identify which regions or counties are potentially most vulnerable to climate change. Unfortunately, one county may be identified as vulnerable in one model but unaffected in another, and as evidenced in Figure 2.12 and Figure 2.13, no single model performs best in all model selection criteria. To address this, I develop a composite vulnerability index that integrates the magnitude and consistency of impacts across all models, excluding the linear seasonal means model, and seven climate projections. This ensemble approach offers a comprehensive measure of risk, identifying counties consistently and severely estimated to be affected by climate change. The development and application of this index are as follows:

1. For each county i , model j , and climate projection k , compute an impact score S_{ijk} that reflects the sign and magnitude of the estimated impact $\hat{y}_{ijk} = \mathbf{f}_j(\Delta\bar{\mathbf{X}}_{ik}; \hat{\beta}, \hat{\alpha})$ from Equation 2.9. The score is calculated as the estimated impact multiplied by -1:

$$S_{ijk} = -\hat{y}_{ijk}.$$

2. Sum these impact scores across all models and projections for each county to obtain a cumulative impact score:

$$C_i = \sum_j \sum_k S_{ijk}.$$

3. Normalize the cumulative impact scores between 0 and 1 to facilitate comparison across counties:

$$N_i = \frac{C_i - \min(C)}{\max(C) - \min(C)},$$

where N_i is the normalized impact score for county i , and $\min(C)$ and $\max(C)$ are the minimum and maximum cumulative impact scores across all counties, respectively.

4. Adjust the normalized score by the proportion of negative estimates (P_i) to incorporate the frequency of negative impacts. This results in the final vulnerability index for county i :

$$V_i = N_i \times P_i.$$

The index V_i captures both the magnitude of all impacts and the consistency of negative impacts. A higher value of V_i indicates a county that not only experiences larger overall impacts (as reflected by the cumulative score) but also a higher frequency of negative impacts, signaling a higher degree of vulnerability. This composite index offers a nuanced perspective on risk, highlighting counties with significant potential impacts and a pattern of consistently negative projections across models.

This index is calculated twice, once for the base specifications (Equation 2.8) and again for the regional specifications (Equation 2.12). A valid criticism of the regional models is the discrete change in coefficients for adjacent counties imposed by the interaction terms. In the top right panel of Figure 2.22, many counties just west of the 100th meridian are projected to experience climate change benefits, whereas many neighboring counties just east of the 100th meridian have negative projected changes. The vulnerability index above serves to smooth out estimates across different regional boundaries, and it is displayed in Figure 2.23.

For both indices, the counties with the largest and most consistent negative estimates across models lie in Southern Arizona and Southern California. Although these counties heavily utilize irrigation, they are also projected to see substantial increases in time spent above 36°C (96.8°F). This area of the US is already experiencing such extreme heat events. In the summer of 2023, Phoenix,

Arizona, located in Maricopa County, experienced a record-breaking 31-day streak of daily high temperatures exceeding 43°C (110°F).²⁵ Additionally, both maps indicate counties along the Pacific coast, in high elevations in the Rocky and Appalachian Mountains, and in the Northeast near the border of Canada are not particularly vulnerable in terms of agricultural profitability. However, these regions do not contain the largest counties in terms of agricultural output. The largest discrepancy between the two indices is the projected vulnerability of western counties (excluding the Southwest).




Figure 2.14: Comparison of predicted climate change impacts on US gross crop revenue by model specification over 2060-2100 with bootstrapped 95% confidence intervals. Future climate data is from the CMIP6 ACCESS-CM2 model projection under the SSP245 scenario.

²⁵<https://www.weather.gov/psr/ExtremeTemps>




Figure 2.15: Comparison of median predicted climate change impacts on US agricultural GDP by model specification over 2060-2100 with bootstrapped 95% confidence intervals. Future climate data is from the CMIP6 ACCESS-CM2 model projection under the SSP245 scenario.




Figure 2.16: Comparison of median predicted climate change impacts on net farm income per acre by model specification and CMIP6 model over 2060-2100 with bootstrapped 95% confidence intervals. All future climate data assumes the SSP245 scenario.




Figure 2.17: Comparison of predicted climate change impacts on aggregate US agricultural profits by model specification and CMIP6 model over 2060-2100 with bootstrapped 95% confidence intervals. All future climate data assumes the SSP245 scenario.




Figure 2.18: Comparison of predicted distributional impacts of climate change on county net farm income per acre by model specification with county and FRR-year fixed effects and CMIP6 model. Each density curve represents all point estimates of the average annual change in net farm income every county in the sample over 2060-2100. All future climate data uses the SSP245 scenario.




Figure 2.19: Comparison of estimated marginal effects of temperature on net farm income per acre on counties east and west of the 100th meridian. On the left, coefficient point estimates with 95% confidence intervals correspond to the linear ($L = 1$) hourly bin specification. On the right, marginal effects curves correspond to the cubic ($L = 3$) growing season daily mean specification.




Figure 2.20: Comparison of predicted average annual impacts of climate change on aggregate US agricultural profits by model specification and regional aggregation over 2060-2100 with bootstrapped 95% confidence intervals. Future climate data is from the CMIP6 ACCESS-CM2 model projection under the SSP245 scenario. Cubic specifications for the binned models are omitted due to exceptionally large confidence intervals.




Figure 2.21: Comparison of predicted distributional impacts of climate change on county agricultural profits per acre by model specification and regional aggregation. Each density curve represents all point estimates of the average annual change in net farm income per acre for every county in the sample over 2060-2100. Future climate data is from the CMIP6 ACCESS-CM2 model projection under the SSP245 scenario.




Figure 2.22: Predicted impacts of climate change on county agricultural profits per acre by regional aggregation over 2060-2100 for the hour bins, linear model. Future climate data is from the CMIP6 ACCESS-CM2 model projection under the SSP245 scenario.




Figure 2.23: Vulnerability of county-level agricultural profits to climate change

Chapter 3

A Comparative Analysis of Water Allocation Institutions and Economic Efficiency

3.1 Introduction

The western United States (US), one of the fastest-growing regions in the country, encompasses the driest areas of North America. New and competing interests for water continue to emerge, yet surface water shortages are expected to worsen due to increasing frequency and severity of droughts induced by climate change (Gutzler & Robbins 2011). The extent to which increasing water scarcity affects western river basins will depend on the ability of human systems to adapt. Although technological innovations have led to marked improvements in water use efficiency, institutional changes in the governance of western water are occurring at a slower pace (Brewer et al. 2006, Libecap 2011, Leonard et al. 2019). Achieving outcomes close to a social optimum through technological improvement alone is unlikely if the mechanism governing water use inherently introduces inefficiencies. The present research aims to compare the efficiency of different water allocation institutions and evaluate the implications of return flows, heterogeneous users, and user

locations (i.e., upstream or downstream) on allocative performance and market efficiency. I focus on comparing prior appropriation and alternative, share-based mechanisms—all with and without water markets. While it is well established in the economics literature that a system of prior appropriation creates inefficiencies due to asymmetry in risk across water users, less attention has been given to how the initial allocation of water rights can exacerbate or alleviate these inefficiencies. Although perfectly competitive markets can reallocate resources to an efficient outcome, the initial distribution of property rights has implications for how much reallocation must occur within the market. This is particularly relevant for contemporary water trading, which are typically associated with high transaction costs. Additionally, I find that due to return flows—the positive externalities associated with water use—the initial location of prior appropriative water rights along a stream can affect the total amount of water available for use in the onset.

The riparian water rights doctrine, rooted in English common law, is prevalent in the humid, eastern US. It grants landowners abutting water bodies the right to use adjacent water “reasonably,” without significantly diminishing its supply for downstream users. However, in the arid regions of the western US, where the prevalence of irrigated agriculture requires large diversions and sometimes low return flows, a riparian doctrine is insufficient for ensuring the right of downstream users to divert water. Consequently, the western US has adopted a doctrine of appropriative water rights (Jones & Cech 2009). This system awards water rights based on the actual physical diversion and beneficial use of water, adhering to a “first in time, first in right” principle. The adoption of beneficial consumptive use as the criterion for water rights aims to encourage investment in water diversion infrastructure and limit excessive speculation on unused waters. There is also evidence that appropriative water rights incentivized *collective* investment in western irrigation infrastructure due to decreased risk for senior users (Libecap 2007, Leonard & Libecap 2019).

Under a doctrine of prior appropriation, users must obtain a water right to divert water from any natural system. Once a water right is administered, it is placed within a hierarchy of priority. When basin water supplies are insufficient to fulfill all allotments, older (senior) water rights have priority, potentially leaving newer (junior) rights with nothing. This system introduces uneven risk across

water users, making prior appropriation widely considered inefficient in the absence of a competitive water rights market (Burness & Quirk 1979, 1980, Huffaker et al. 2000). In their seminal work, Burness & Quirk (1979, 1980) analyze the efficiency of prior appropriation with simplifying assumptions, focusing on the long-run equilibrium of expected diversions and profits when firms with identical, separable profit functions share a waterway. They acknowledge that an efficient use of resources under an appropriative framework would require a competitive leasing market contingent on significant investment in the construction of diversion facilities to capture overall water capacity. However, over 40 years later, many western basins are now over-appropriated, and several assumptions in Burness & Quirk (1979, 1980) may not align with contemporary conditions. Firstly, although modern water markets do exist, they are typically constrained to localized trading with high transaction costs encompassing legal fees, negotiation costs, and the expenses associated with transferring rights (Brewer et al. 2006, Libecap 2011, Leonard et al. 2019). This renders current water allocations relatively inflexible, which can hinder the transfer of water to more beneficial uses as new and potentially more valuable uses emerge (Brewer et al. 2008). Consequently, the initial allocation of water rights becomes an important determinant of water use efficiency, as high transaction costs may prevent the market from correcting inefficiencies in allocation. Secondly, the authors do not address scenarios where average water availability is insufficient to satisfy all administered water rights. This situation is particularly pertinent in regions experiencing water scarcity, where the hierarchy of water rights can result in junior users being deprived of water despite holding legal entitlements. Not only does this prevent users who potentially have a high marginal benefit of water from generating value, there are implications for return flows when curtailed users are located upstream. Lastly, the authors focus on the case in which value is only generated from consumptive use. Many water uses generate some value non-consumptively, such as hydroelectric power generation, recreational activities, ecosystem services, and many household uses. Under a strict prior appropriative regime, it is possible that no water may be diverted for such uses during times of shortage even if the consumptive use is relatively low.

One alternative institution to prior appropriation are water sharing mechanisms, which aim to redistribute risk. For example, some farmers have formed cooperative organizations called “irrigation districts,” pooling the priority of their water rights and redistributing risk evenly across shares. Water is then allocated among members according to the number of shares they own. This system ensures nonzero water allocations each year for each shareholder, which creates more certainty for users to plan for the future. Ji & Cobourn (2018) find that irrigation districts allocate larger portions of land to drought-sensitive, high-value crops compared to areas with similar overall priority without a share-based system, resulting in a 6% average increase in profit. Water sharing plans have also been adopted at larger scales in Australia and Chile (Grafton et al. 2011, Tingey-Holyoak 2014).

Although shared-based mechanisms can remedy the uneven risk introduced by prior appropriation, efficiency ultimately depends on the benefits generated by individual water users (Bennett et al. 2000). Return flows also require consideration, as water diverted is not a perfectly rivalrous good. Most water uses necessitate diverting more water than needed for consumption, with return flows ranging from 5-90% depending on the use (Burness & Quirk 1980). In the South Platte Basin of Colorado, annual surface water diversions double the average natural inflow (CWCB 2015). Efficiency, therefore, depends not only on the heterogeneous value generation across users but also on their physical location within a river basin. For instance, if an upstream user with high return flows diverts one unit of water, a large portion of that unit can continue to generate value as it reenters the stream for the next user. The allocative efficiency of priority based mechanisms can further change depending on whether the upstream user is junior or senior and high value or low value. Given the complexities arising from return flows and user heterogeneity, no single water allocation institution is likely to be preferable across all western basins. The most efficient institution for a given basin will depend on its specific composition. Figure 3.1 highlights the heterogeneity of watersheds in the western US, displaying HUC8 classifications from Duan et al. (2018) and the location of relatively big cities. Large population centers are located within both headwater (upstream) and terminus (downstream) HUC8 basins and across states with varying water uses and

needs. For example, California's water demands are remarkably diverse. Large urban areas, such as Los Angeles and San Francisco along the state's central and southern coast, rely heavily on water for domestic use, industry, and maintaining green spaces in cities. Agriculture, a significant sector in California's economy, demands large quantities of water for irrigation. The state's water sources are equally varied, encompassing snowmelt from the Sierra Nevada mountains in northern California and major river systems like the Colorado River and the Sacramento-San Joaquin River Delta.

In general, the current body of economic research offers important perspectives on optimizing water resource use under institutional constraints, particularly emphasizing the role of market mechanisms. The Coase Theorem (Coase 1960), a cornerstone in the field of environmental economics, posits that if property rights are clearly defined and transaction costs are negligible, the allocation of resources will be efficient regardless of the initial distribution of these rights. This theorem stresses the potential for private negotiations to address externalities and achieve optimal resource allocation. However, the application of Coasean bargaining to water rights governed by a doctrine of prior appropriation is not clear cut. When high priority water rights are located downstream, the ability of upstream junior users to divert water and generate return flows can be stifled, potentially distorting the total amount of water that can be reallocated through a market. This challenges the assumption that initial allocations do not affect efficiency, particularly under conditions where the physical characteristics of the resource impose unique constraints.

In contrast to fixed-allocation mechanisms like prior appropriation, share-based mechanisms, where water rights are allocated as proportional shares of the available supply rather than as fixed quantities, offer a different paradigm. These mechanisms provide two possible improvements to prior appropriation. First, they always ensure some amount of upstream diversions, which can generate more return flows in some cases. Second, water sharing plans inherently reduce the asymmetry in risk across users, potentially leading to an initial allocation that reflects more efficient use from the outset. Share-based systems may therefore reduce the subsequent reliance on markets to achieve social optimality. The key insight here is that the initial allocation of property

rights not only potentially affects the availability of the resource due to the positive externalities of return flows, but also determines the degree of market participation necessary to reach an economically efficient use of resources. In scenarios where the initial allocation generates a more efficient distribution, the market's role in adjusting allocations to meet changing demands and values is less critical.

The initial allocation of property rights and its impact on market efficiency is a theme not unique to water rights. In the context of land demarcation, (Libecap 2007) show that different property allocation regimes have enduring impacts on land values, transaction costs, and the structure of land markets. They find that the Rectangular System, by providing standardized, secure property rights, enhances economic outcomes compared to the more flexible Metes and Bounds system. In the context of sulfur dioxide cap-and-trade systems, (Henry III et al. 2011) demonstrate that the initial allocation of pollution permits can lead to increased social costs by shifting emissions to areas of higher marginal damage. Such examples in the broader economics literature highlight the principle that the initial distribution of property rights can influence the efficiency of resource allocation, even in the presence of well-functioning markets.

The present research builds upon the current literature by exploring the nuanced aspects of water allocation that remain relatively unexplored, examining how the initial allocation of water rights and physical characteristics of a river basin can either aggravate or mitigate potential inefficiencies. I develop a series of optimization models that emulate different institutions that exist throughout the world and use Monte Carlo simulations of water supply to examine their sensitivity to randomly generated physical characteristics of a river basin. I find that system-wide value generation under prior appropriation is particularly sensitive to the location of priority along a stream and the utilization of return flows during times of shortage, even in the presence of a perfectly competitive leasing market. For simulations in which there is no water shortage, markets unsurprisingly and unambiguously generate the most value. Ideal conditions for priority-based mechanisms include relatively senior upstream water users that either generate high value or high return flows. However, as the number of water users in the model increases, these ideal conditions become less likely,

and share-based mechanisms can generate more value than a market under certain conditions. By examining the physical factors that impact institutional efficiency, this study provides valuable insights for policymakers and stakeholders in the development and implementation of effective and efficient water allocation strategies.

The remainder of the article is organized as follows. First, I present a series of modeling frameworks that emulate different water allocation mechanisms for a hypothetical river basin. I begin by outlining a simple two-user case to emphasize the impact of basin composition on economic efficiency and extend the models to accommodate n -users for generalizability. Next, I employ Monte Carlo simulations to evaluate efficiency outcomes of each mechanism under various basin compositions and random water supply conditions. I outline the conditions that have the greatest impact on market efficiency, and regression models are applied to the simulation results to characterize basin compositions that bode well for specific water allocation mechanisms. Lastly, I discuss potential policy considerations and implications.

3.2 Theoretical Models of Water Allocation Institutions

3.2.1 Two Users

Social Planner

Before examining real-world water allocation institutions, consider the social planner (SP), whose objective is to maximize the total value of production within a river basin by allocating water w_i to two users with value functions $v_i(w_i)$, where $i = \text{upstream } (u) \text{ or downstream } (d)$. Available water in a given year is a random draw $\tilde{w} \in \mathbf{W}$, where $\mathbf{W}$ is the distribution of possible water supply conditions. Some proportion of water $\theta \in (0, 1)$ diverted to the upstream user returns to the stream and becomes available for the downstream user as a return flow. The SP's deterministic

maximization problem with two water users in a given year is

$$\begin{aligned}
& \max_{w_u, w_d \geq 0} && v_u(w_u) + v_d(w_d) \\
& \text{s.t.} && w_u \leq \tilde{w} \\
& && w_d \leq \tilde{w} - (1 - \theta)w_u.
\end{aligned} \tag{3.1}$$

If the two users are homogeneous, i.e., $v_u = v_d = v$, it can be easily shown that $w_u^* > w_d^*$ and $\partial v / \partial w_d^* > \partial v / \partial w_u^*$ at the optimum. Due to return flows, more water is diverted to the upstream user, resulting in a lower marginal benefit of water compared to the downstream user.²⁶ When $v_u \neq v_d$, optimal allocations will depend on the relative value functions.

Prior Appropriation

Now consider the same basin layout, but water allocation is determined by PA. Each user owns a water right R_i^m that determines a maximum amount of water allowable for diversion, where m = senior (s) or junior (j). Under PA, R_i^s must be met before the junior user can divert any water. In practice, any water use by a junior user that prevents any senior user from receiving their full allotment is prohibited. Water allocation is therefore no longer determined by a maximization problem, but by the location of the junior and senior user, their respective water rights, and the total water supply. There are two possible cases for a given level of water supply: (i.) the senior user is upstream or (ii.) the senior user is downstream.

²⁶The role of return flows in determining upstream versus downstream diversions is further demonstrated in the Supplementary Analysis.

- i. When the upstream user is senior and the downstream user is junior, the resulting water allocations are:

$$\begin{aligned} w_u^s &= \begin{cases} \tilde{w} & \text{if } \tilde{w} < R_u^s \\ R_u^s & \text{if } \tilde{w} \geq R_u^s \end{cases} \\ w_d^j &= \begin{cases} \theta \tilde{w} & \text{if } \tilde{w} \leq R_u^s \\ \tilde{w} - (1 - \theta)R_u^s & \text{if } R_u^s < \tilde{w} < (1 - \theta)R_u^s + R_d^j \\ R_d^j & \text{if } \tilde{w} \geq (1 - \theta)R_u^s + R_d^j. \end{cases} \end{aligned} \quad (3.2)$$

- ii. When the upstream user is junior and the downstream user is senior, the resulting water allocations are:

$$\begin{aligned} w_d^s &= \begin{cases} \tilde{w} & \text{if } \tilde{w} < R_d^s \\ R_d^s & \text{if } \tilde{w} \geq R_d^s \end{cases} \\ w_u^j &= \begin{cases} 0 & \text{if } \tilde{w} \leq R_d^s \\ \frac{\tilde{w} - R_d^s}{1 - \theta} & \text{if } R_d^s < \tilde{w} < (1 - \theta)R_u^j + R_d^s \\ R_u^j & \text{if } \tilde{w} \geq (1 - \theta)R_u^j + R_d^s. \end{cases} \end{aligned} \quad (3.3)$$

In each case, $(1 - \theta)R_u^m + R_d^m$ characterizes the total amount of water necessary to fulfill the maximum allowable diversion for both users. For the senior user, location is irrelevant and $w_u^s = w_d^s$ in any water supply scenario. However, the junior user faces different outcomes in times of scarcity, i.e., $\tilde{w} < (1 - \theta)R_u^m + R_d^m$, depending on location. If downstream, the junior user will always receive at least the return flows from the upstream user, $\theta \tilde{w}$. If upstream, the junior user faces the possibility of receiving nothing if the basin water supply is insufficient to fulfill even the senior right, i.e., $\tilde{w} < R_d^s$. Although simple, the two user case clearly conveys how risk is disproportionately assumed by the junior user. Table 3.1 provides a stylized example of how water allocations differ for the junior user under different stream location and supply scenarios. If the

users have the same value functions in both cases, case (ii.) unambiguously results in lower basin-wide value during shortages due to the upstream, junior user receiving less water.

One criticism of the doctrine of prior appropriation is the economic inefficiency introduced through an uneven distribution of risk. During times of shortage, some users who could generate significant value per unit of water receive nothing, while others with higher priority rights may not put all of the resource to its most productive use. This uneven risk can lead to a suboptimal allocation of resources. Alternatively, share-based institutions can partly remedy this inefficiency by ensuring that all users receive a non-zero allocation of water, even during times of scarcity. Under a PS system, water is allocated proportionally according to shares until all endowments are satisfied in a descending order from highest to lowest number of shares. If a basin were to transition from PA to PS, shares may be determined by users' previous priorities or endowments or both. Share-based institutions aim to redistribute risk more equitably among water users, hopefully promoting a more efficient use of resources.

Proportional Sharing

Before addressing PS, here I characterize how water rights may have formed along a stream. This intermediate step has the benefit of ensuring return flows are accounted for when developing a stylized PS system. The senior user first appropriates some portion $\alpha_i^m \in (0, 1)$ of average water supply from a completely unappropriated basin, and then the junior user appropriates some portion of the remaining available water. The characterization again depends on the location of each user, where $\bar{w}$ is the mean of $\mathbf{W}$.

- i. The upstream user is senior, the downstream user is junior, and the resulting water rights are:

$$\begin{aligned} R_u^s &= \alpha_u^s \bar{w} \\ R_d^j &= \alpha_d^j \bar{w} [1 - (1 - \theta) \alpha_u^s]. \end{aligned} \tag{3.4}$$

ii. The upstream user is junior, the downstream user is senior, and the resulting water rights are:

$$\begin{aligned} R_d^s &= \alpha_d^s \bar{w} \\ R_u^j &= \alpha_u^j \bar{w} \left(\frac{1 - \alpha_d^s}{1 - \theta} \right). \end{aligned} \quad (3.5)$$

Again consider the same basin layout as in the SP and PA examples, but water allocation is now determined by PS. Assume this institution has replaced PA and the existing water users trade in their water rights for a share

$$\delta_i^m = \frac{R_i^m}{\sum_i R_i^m} \in (0, 1). \quad (3.6)$$

A more realistic formulation of shares that is driven by seniority is introduced in the n-user framework, however this simple version is sufficient for examining how allocations change under a share-based mechanism. Under PS, if $\tilde{w}$ is not sufficient to fulfill the sum of all previous water rights, the shortage is distributed equally according to shares. Again, there are two cases depending on the location of the users.

i. The upstream user was senior, the downstream user was junior, and the resulting shares are:

$$\begin{aligned} \delta_u^s &= \frac{\alpha_u^s}{\alpha_u^s + \alpha_d^j [1 - (1 - \theta)\alpha_u^s]} \\ \delta_d^j &= \frac{\alpha_d^j}{\alpha_u^s + \alpha_d^j [1 - (1 - \theta)\alpha_u^s]}. \end{aligned} \quad (3.7)$$

ii. The upstream user was junior, the downstream user was senior, and the resulting shares are:

$$\begin{aligned} \delta_d^s &= \frac{\alpha_d^s}{\alpha_d^s + \alpha_u^j \left(\frac{1 - \alpha_d^s}{1 - \theta} \right)} \\ \delta_u^j &= \frac{\alpha_u^j}{\alpha_d^s + \alpha_u^j \left(\frac{1 - \alpha_d^s}{1 - \theta} \right)}. \end{aligned} \quad (3.8)$$

In each case, the junior user is now guaranteed a nonzero allocation of water in a given period.

3.2.2 N Users

In this section I outline four institutional mechanisms for allocating water relative to the social planner. The first is a strict doctrine of prior appropriation, where no water can be traded among users. The second and third mechanism also do not have markets, and water is allocated equally or proportionally according to shares. The final mechanism emulates a leasing market, in which water allocations from any of the prior institutions can be traded.

Social Planner

The two-user case provides helpful intuition for how stream position and seniority can affect the overall value of water within a basin. The SP problem can be easily extended to accommodate an n number of users, where i indexes stream position from furthest upstream (1) to furthest downstream (n):

$$\begin{aligned}
 & \max_{w_1, \dots, w_n \geq 0} \quad \sum_{i=1}^n v_i(w_i) \\
 & \text{s.t.} \quad \left. \begin{aligned}
 w_1 &\leq \tilde{w} \\
 w_2 &\leq \tilde{w} - (1 - \theta_1)w_1 \\
 &\vdots \\
 w_n &\leq \tilde{w} - \sum_{i=1}^{n-1} (1 - \theta_i)w_i
 \end{aligned} \right\} \text{Physical constraints.}
 \end{aligned} \tag{3.9}$$

There are now n number of constraints, where each imposes the maximum amount of water that can be physically delivered to the i^{th} user. Again, the SP problem conveys the upper bound on the total value that can be generated from the available water in a given period. Figure 3.2 provides a visual representation of the hypothetical river basin outlined in equation (9) with $n = 3$ users. Notably, return flows generated from user i are available to all downstream users.²⁷

²⁷All else equal, the Social Planner will always divert more water to upstream users, as demonstrated in the Supplementary Analysis.

Prior Appropriation

With n users, it is not feasible to discretely characterize every possible stream position and priority hierarchy as in the two-user case. Instead, the institution of prior appropriation can be emulated with a constrained linear programming problem. The following maximization problem will hereinafter be referred to as the PA allocation function,

$$\begin{aligned}
 & \max_{w_1^m, \dots, w_n^m \geq 0} \quad \sum_{i=1}^n \beta_i^m w_i^m \\
 & \text{s.t.} \quad \left. \begin{aligned}
 & w_1^m \leq \tilde{w} \\
 & w_2^m \leq \tilde{w} - (1 - \theta_1)w_1^m \\
 & \vdots \\
 & w_n^m \leq \tilde{w} - \sum_{i=1}^{n-1} (1 - \theta_i)w_i^m
 \end{aligned} \right\} \text{Physical constraints} \\
 & \quad \left. w_i^m \leq R_i^m \quad \forall i \in \{1, \dots, n\} \right\} \text{Institutional constraints}
 \end{aligned} \tag{3.10}$$

where m indexes priority from 1 to M with 1 being the highest priority user, and β_i^m and R_i^m correspond to a priority coefficient and maximum allowable quantity for user i , respectively. For the highest priority user, β_i^m must be sufficiently large enough such that the PA allocation function always gives the next unit of water to the user with $m = 1$ until $w_i^m = R_i^m$, and so on and so forth down the priority ranking. The n institutional constraints are imposed to ensure no user is allocated more than R and the physical constraints are identical to the SP problem, resulting in a total of $2n$ constraints.

For the PA allocation function, i still indexes stream position, but it does not reflect the hierarchy of priority. It is not necessarily the case that β_1^m corresponds to the highest or lowest priority user. In the following water supply simulations section, both i and m are randomly assigned to each user to analyze how efficiency changes under different hierarchies and stream positions. The solution to the PA allocation function yields a vector of water delivered to each user. Once the

allocations are determined, they can be plugged into the function $\sum_{i=1}^n v_i(\hat{w}_i^m)$ to determine the total value generated under prior appropriation.

Proportional Sharing

Suppose users under a prior appropriation system elect to pool their risk and assign shares according to priority levels. Under a proportional sharing (PS) system, water is allocated proportionally according to each user's number of shares until all endowments are satisfied in a descending order from highest to lowest number of shares. This ensures non-zero allocations but still maintains some manner of hierarchy. Let γ_i^m represent a user's shares and δ_i^m their proportional allocation:

$$\gamma_i^m = \sum_{j=m}^n \frac{1}{j} \quad (3.11)$$

$$\delta_i^m = \frac{\gamma_i^m}{\sum_{m=1}^n \gamma_i^m}. \quad (3.12)$$

Different formulations of γ_i^m can be used to tighten or loosen the level of hierarchy.²⁸ Essentially, all users are allocated water at a rate proportional to δ_i^m .

The PS mechanism can be modeled with an iterative linear optimization approach. Again, let R_i^m be the water right that determines the maximum amount of water allowable for user i with priority m . PS can be emulated as follows:

1. Calculate the minimum amount of water necessary, $\hat{w}_1$, to satisfy the maximum allowable quantity for the highest priority user (R_i^1) while abiding by the physical and institutional

²⁸For $n = 4$, $\delta_i^1 = 0.520833$, $\delta_i^2 = 0.276833$, $\delta_i^3 = 0.145833$, and $\delta_i^4 = 0.0625$ using equation (12).

constraints below.

$$\begin{aligned}
& \min_{\dot{w}} \quad \dot{w} \\
& \text{s.t.} \quad \left. \begin{aligned} w_1^m &\leq \dot{w} \\ w_2^m &\leq \dot{w} - (1 - \theta_1)w_1^m \\ &\vdots \\ w_n^m &\leq \dot{w} - \sum_{i=1}^{n-1} (1 - \theta_i)w_i^m \end{aligned} \right\} \text{Physical constraints} \\
& \quad \quad \quad \left. \begin{aligned} w_i^1 &= R_i^1 \\ w_i^m &= \delta_i^m (\sum_{j=1}^n w_j) \end{aligned} \right\} \text{Institutional constraints}
\end{aligned} \tag{3.13}$$

2. If $\tilde{w} \leq \dot{w}_1$, use the following maximization problem to determine allocations.

$$\begin{aligned}
& \max_{w_1^m, \dots, w_n^m \geq 0} \quad \sum_{i=1}^n w_i^m \\
& \text{s.t.} \quad \left. \begin{aligned} w_1^m &\leq \tilde{w} \\ w_2^m &\leq \tilde{w} - (1 - \theta_1)w_1^m \\ &\vdots \\ w_n^m &\leq \tilde{w} - \sum_{i=1}^{n-1} (1 - \theta_i)w_i^m \end{aligned} \right\} \text{Physical constraints} \\
& \quad \quad \quad \left. \begin{aligned} w_i^m &\leq R_i^m \\ w_i^m &= \delta_i^m (\sum_{j=1}^n w_j) \end{aligned} \right\} \text{Institutional constraints}
\end{aligned} \tag{3.14}$$

If $\tilde{w} > \dot{w}_1$, store the allocations from the minimization problem in step 1 in vector $\hat{\mathbf{w}}_1 = (\hat{w}_1^m, \hat{w}_2^m, \dots, \hat{w}_n^m)$ and proceed to step 3.

3. Subtract $\hat{\mathbf{w}}_1$ from the vector of maximum allowable quantities $\mathbf{R}$ to get the remaining amount of water each user requires to satisfy their water right, $\hat{\mathbf{R}}_1 = (R_1^m - \hat{w}_1^m, R_1^2 - \hat{w}_2^m, \dots, R_n^m - \hat{w}_n^m)$. Subtract $\dot{w}_1$ from $\tilde{w}$ to get the remaining amount of water in the stream, $\tilde{w}_1$. Set $\delta_i^1 = 0$ since $R_i^1 - \hat{w}_i^1 = 0$, and generate new shares and proportions using equation (14) for $i - 1$ users. Calculate the minimum amount of water necessary, $\dot{w}_2$, to satisfy

the maximum allowable quantity for the second highest priority user ($R_i^2 - \hat{w}_i^2$), who now holds the largest share, using the same minimization problem from step 1 with the updated parameters.

4. If $\tilde{w}_1 \leq \dot{w}_2$, use the maximization problem from step 2 with updated parameters to determine second stage allocations and add them to $\hat{w}_1$ to get final allocations. If $\tilde{w}_1 > \dot{w}_2$, repeat step 3 with updated parameters.
5. Repeat the process until there is no water remaining or all water rights are satisfied.

The iterative approach to PS allocation provides the vector of water delivered to each user. This vector is the sum of all water allocated through each iteration, culminating in the the maximization problem in step 2 when there is not enough water left to satisfy the next highest priority user. In essence, the PS allocation process always diverts the next available unit of water in proportion to each user's number of shares while abiding by the physical constraints of the basin.

Equal Sharing

A simple, and naive, share-based mechanism for water allocation is one where all users are required to receive the same amount of water. An institution where water is equally shared (ES) across all water users can also be imitated with a linear programming problem. The following is the ES allocation function:

$$\begin{array}{ll}
 \max_{w_1, \dots, w_n \geq 0} & \sum_{i=1}^n v_i(w_i) \\
 \text{s.t.} & \left. \begin{array}{l}
 w_1 \leq \tilde{w} \\
 w_2 \leq \tilde{w} - (1 - \theta_1)w_1 \\
 \vdots \\
 w_n \leq \tilde{w} - \sum_{i=1}^{n-1} (1 - \theta_i)w_i
 \end{array} \right\} \text{Physical constraints} \\
 & \left. \begin{array}{l}
 w_i \leq R_i \quad \forall i \in \{1, \dots, n\} \\
 w_i = w_j \quad \forall i \neq j
 \end{array} \right\} \text{Institutional constraints}
 \end{array} \tag{3.15}$$

where additional institutional constraints ensure all users receive the same amount of water.

Leasing Markets without Transaction Costs

The mechanisms above each allocate water according to a set of institutional constraints that do not consider the value generated by individual users and therefore do not attempt to equate marginal value across users. A leasing market can remedy this introduced inefficiency, which allows the trading of water among users. Users with lower-valued uses can sell water to those who can generate higher value from the water, leading a more efficient reallocation of resources. The following model emulates a leasing market in which users are allowed to trade water *after* allocations are determined from the PA, PS, or ES allocation function. Annual water supplies are often estimated from the amount of winter snow pack, so the initial allocations can be interpreted as expected allocations for the given period.

To proxy the outcome of a perfectly competitive market with no transaction costs, the model is set up as a Social Planner's problem with additional institutional constraints. For the leasing market, let $\hat{w}_i$ be the initial allocation for user i who can choose to buy or sell an amount t_i :

$$\begin{aligned}
 & \max_{t_1, \dots, t_n \geq 0} \quad \sum_{i=1}^n v_i(\hat{w}_i + t_i) \\
 & \text{s.t.} \quad \left. \begin{aligned}
 & \hat{w}_1 + t_1 \leq \tilde{w} \\
 & \hat{w}_2 + t_2 \leq \tilde{w} - (1 - \theta_1)(\hat{w}_1 + t_1) \\
 & \vdots \\
 & \hat{w}_n + t_n \leq \tilde{w} - \sum_{i=1}^{n-1} (1 - \theta_i)(\hat{w}_i + t_i)
 \end{aligned} \right\} \text{Physical constraints} \\
 & \quad \left. \begin{aligned}
 & \hat{w}_i + t_i \geq 0 \quad \forall i \in \{1, \dots, n\}
 \end{aligned} \right\} \text{Non-negative allocations} \\
 & \quad \left. \begin{aligned}
 & \sum_{i=1}^n t_i = 0
 \end{aligned} \right\} \text{Market clearing condition.}
 \end{aligned} \tag{3.16}$$

If $t_i > 0$, user i can be considered a net buyer and if $t_i < 0$, a net seller. The first set of constraints ensure that user i can only buy water up to an amount such that $\hat{w}_i + t_i$ is less than or equal to the

total amount of water in the river at their point of diversion. The second set of constraints ensures that no user can sell more water than they were initially allocated. The last constraint is a market clearing condition that ensures that sum of all water bought is equal to the sum of all water sold. The solution to the water market optimization problem, $\hat{t}$, can then be added to the initial vector of allocations $\hat{w}$ to get the final allocations.

3.3 Results

3.3.1 Water Supply Simulations

In this section I employ a Monte Carlo simulation to randomly generate water supply conditions to compare allocation and value outcomes across the institutional mechanisms outlined above. For the two-user case, I compare four simple basin compositions to highlight the sensitivity of prior appropriation to the location of priority and value along a waterway. For the n -user case, basin compositions are randomly generated. Users are randomly assigned a priority number, stream position, return flow parameter, and value function to assess how robust water allocation mechanisms are to varying institutional and physical constraints.

Two Users

To compare the efficiency of each institution, I allow users to be heterogeneous in terms of their value functions and use a Monte Carlo simulation to randomly sample 5,000 draws of $\tilde{w}$ from a log-normal distribution. Four cases are compared across the three institutions (SP, PA, PS): (i.) the senior user is upstream and high value, (ii.) the senior user is downstream and high value, (iii.) the senior user is upstream and low value, and (iv.) the senior user is downstream and low value. In each scenario, the junior user assumes the opposite characteristics. The outcomes of PA and PS would be identical when there is enough water to fulfill the endowments of both users, i.e., $\tilde{w} \geq (1 - \theta)R_u^m + R_d^m$, so the following results examine only the iterations where there was a water

shortage. The models are parameterized such that: $v_{high}(w) = 20w^{\frac{1}{2}}$, $v_{low}(w) = 10w^{\frac{1}{2}}$, $\theta = 0.25$, $\alpha_i^s = 0.6$, $\alpha_i^j = 0.75$, and $\bar{w} = 500$.

The results of the simulations are presented in Figure 3.2, which displays distributions of total value generated for each institution across all random water supply draws. Distributions with significant area on the right side of each panel indicate a larger proportion of simulations where high value is generated. In each scenario, the SP's distribution characterizes the upper bound on efficiency. In the scenario where the senior user is upstream and high value, PA generates more basin-wide value from water than PS on average, as evidenced by the higher peak on the right hand side of the panel. Under PA, water is first allocated to the higher value user upstream, and return flows are then available for the less valuable use downstream. Moving from PA to PS reallocates some water away from a high value use. In the scenarios where the senior user is low value, PS unambiguously generates more value than PA. When the senior user is downstream and low value (bottom right panel), there is potential for PA to generate very little relative value during extreme shortages, as shown in the peak in the left tail. An ambiguous finding results from when the senior user is downstream and high value. For minor shortages, PA may be preferred to PS, but during extreme shortages where the junior user receives no water, the senior uses all available water to a point of low marginal benefit. In general, Figure 3.2 conveys how PA is more sensitive to basin composition than PS in terms of risk. Moving from a desirable to less desirable composition causes the distribution to flatten, increasing the risk of low-value outcomes in the left tail.

N Users

Similar to the two-user simulations, the allocative efficiency of the following mechanisms with and without a subsequent leasing market are compared using a Monte Carlo simulation: 1. prior appropriation (PA), 2. proportional sharing (PS), and 3. equal sharing (ES). For the leasing market scenarios (PAM, PSM, ESM), transaction costs are set to zero to capture a perfectly efficient market. Water supplies are randomly sampled from a log-normal distribution and the maximum amount any user can divert is 500 units of water. For each water supply draw, users along the

stream are randomly assigned a stream position, a priority (and a subsequent number of shares), a value function, and a return flow parameter. The value functions are all non-negative, monotonically increasing, and display diminishing marginal returns. Detailed information on all parameters of the hypothetical river basin are outlined in Table 3.2. The optimization problems for each of the mechanisms are then solved for that particular composition of users. This process is repeated for a system with three to 36 users, with 15,000 water supply draws for each case.²⁹ Table 3.3 displays an example of a simulation with $n = 10$ users and a randomly sampled water supply. In this example, the optimization models are each parameterized according to their relevant values in this table, and the outcomes are compared for $\tilde{w} = 2, 252$.

For each institution and simulation, I also compute the absolute difference between the water allocations without the market (PA, PS, ES) and with the market (PAM, PSM, ESM), to generate a percentage that provides a total measure of how much water had to be shifted under each mechanism to reach the efficient (market) outcome. A higher percentage indicates a greater shift in allocations due to the introduction of the market, and, all else equal, plausibly higher participation and transaction costs required to reach that optimal outcome. The calculation for PA would be:

$$\% \text{ Reallocation}_{\text{PA}} = \frac{\sum_{i=1}^n |w_i^{\text{PA}} - w_i^{\text{PAM}}|}{2 \times \sum_{i=1}^n w_i^{\text{PA}}} \times 100 \quad (3.17)$$

where w_i^{PA} and w_i^{PAM} are the water allocations to the i^{th} user under PA without and with the perfectly competitive leasing market, respectively. The division by 2 is included to account for the fact that reallocating water from one user to another is counted in both the reduction for one and the increase for the other, thus doubling the total apparent reallocation when summed absolutely.

Figure 3.4 displays the average value generated (left panel) and average water reallocated (right panel) across all simulations by institution and number of users. The left panel of the figure presents a clear trend: in the absence of a leasing market, share-based mechanisms outperform PA in terms of average total value created, with the disparity widening as the user base expands.

²⁹Due to the complex and nonlinear optimization problems, the solver fails to find a solution for at least one institution in roughly 9.1% of the draws. When a solution cannot be found, that draw is omitted.

This suggests that share-based mechanisms have the capacity to allocate water more efficiently in the absence of an active market. With a leasing market in place, as indicated by the overlapping red, yellow, and orange curves, total average value is effectively equivalent across institutions, which aligns with the principle that the efficiency of resource allocation should not depend on the initial allocation of resources. Nonetheless, the right panel provides a nuanced perspective on this narrative. Despite achieving similar levels of value creation, the proportion of water that necessitates reallocation under the PA market (PAM) system is substantially higher—almost double—compared to the equal sharing market (ESM) system. The need for significant reallocation in the PAM system to attain efficiency relative to the ESM system highlights an often-overlooked dimension of Coasean bargaining: the initial distribution of rights has substantial implications for the volume of resources that must be reallocated to achieve efficiency. This suggests that more efficient initial allocations, as evidenced in the share-based systems, could reduce the need for market reallocations, thereby minimizing transaction costs and potential barriers to achieving an optimal distribution of resources.

Due to return flows, as demonstrated in the Supplementary Analysis, the Social Planner allocates more water upstream to balance the marginal benefits of the last units of water diverted to each user. However, as shown in Section 2.1.2 and Table 3.1, a senior user downstream of a junior user can result in less overall diversions. So, for each randomly generated basin composition, I calculate a Pearson correlation coefficient ρ to examine the strength and direction of the relationship between stream position and user priority under PA. For example, consider a case with $n = 3$ users where the priorities assigned (in order of stream position) are 1, 2, 3. Here, the stream position correlates perfectly with the user priorities, yielding a correlation coefficient of $\rho = 1$. In contrast, if the user priorities are assigned as 3, 2, 1, the stream position correlates perfectly negatively with the user priorities, resulting in a correlation coefficient of $\rho = -1$. In essence, a positive correlation would suggest that users' priorities are aligned with the natural flow of the stream, i.e., upstream users have higher priority. Conversely, a negative correlation implies that high-priority rights are downstream.

Figure 3.5 displays the relationship between stream-priority location and the total amount of water diverted to all users. This relationship is calculated simply by regressing ρ on the percentage of total water that is diverted in excess of the natural inflow $\tilde{w}$ across all simulations. For the ES mechanism, stream-priority location has no impact on the amount of water diverted, because the institutional constraint in equation (15) simply forces everyone to diverted the same amount of water. The PA mechanism displays the steepest slope, indicating that when downstream users have senior rights, there is a constraint on the ability of upstream users to divert water, which could lead to decreased overall utilization of the resource. However, a higher correlation mimics the strategy of the Social Planner to divert more water upstream, so there are then some instances when PA with an active market is the most efficient mechanism. The magnitude of the slope is reduced for the PS mechanism, as it can be interpreted as a middle ground institution between ES and PA. In summary, the findings presented in Figure 3.4 and Figure 3.5 convey that the distribution of user priorities across the stream not only determines the level of market transactions required to achieve social optimality, but it also has a direct impact on the total volume of water that can be harnessed by the system.

Figure 3.4 demonstrates that each of the market mechanisms generates the same amount of total value on average, but this is not the case for all stream compositions. As shown in Figure 3.5, the total amount of water diverted is impacted by the location of users. As such, Figure 3.6 shows the distribution of total value generated (relative to the Social Planner) across all simulations that fall into four categories: the combinations of high priority downstream ($\rho \leq -0.5$) and high priority upstream ($\rho \geq 0.5$), and high and low shortage. The high and low shortage scenarios are defined as the top and bottom 25th percentile of shortage per user, which is calculated as the minimum amount of water required to satisfy all user endowments minus the natural inflow $\tilde{w}$ and divided by the number of users. In Figure 3.6, the top four panels correspond to the mechanisms without a market, and it is evident that PA is particularly sensitive to both the location of users and the level of shortage. The peaks of each distribution move the left as high priority users move downstream, indicating greater efficiency losses, and the distributions flatten as the magnitude of

shortage increases. The share-based mechanisms display less sensitivity in these dimensions. The bottom four panels of Figure 3.6 correspond to the mechanisms with a market, and unsurprisingly all distributions display a rightward shift relative to no market. However, PAM still displays a sensitivity to the four scenarios. Given that these markets have zero transaction costs and the same value functions across users, the only explanation for differences in efficiency is that different amounts of water are being traded (despite the same natural inflow). This further demonstrates that the total volume of water harnessed by the system can be impacted by the location of priority along the stream and potentially further exacerbated during times of extreme shortage. In fact, ES *without* a water market still generates more value than PAM in 17.5% of simulations. Of these simulations, the average stream position-priority correlation coefficient is -0.2 and the average shortage per user is 180 units, which corresponds to the 27th percentile across all simulations.

3.4 Conclusion and Policy Discussion

This research examines the impact of water allocation institutions on economic efficiency by comparing prior appropriation and alternative, share-based mechanisms. As water scarcity intensifies due to climate change, this study underscores the need for adaptive institutional changes beyond technological improvements. I demonstrate how the initial allocations of water rights, influenced by return flows and the location of users, affect both market efficiency and the total availability of water. Prior appropriation, the predominant water allocation mechanism in the western US, is particularly sensitive to basin composition. Additionally, share-based water allocation mechanisms can provide more equitable risk distribution and efficiency compared to prior appropriation in many instances, even in the absence of a water market. The findings also suggest that market mechanisms, while important, are not a panacea; institutional configurations must be tailored to basin-specific characteristics and the physical realities of water resources.

This analysis shortcomings that must best addressed. First, the parameterization of the theoretical models in this analysis is somewhat arbitrary. The purpose of these models is solely to

highlight the costs and benefits of different institutional constraints, not to adequately mimic reality. To apply the findings of this article to an existing river basin, it would be necessary to generate realistic value functions and return flow dynamics, which can be challenging. Secondly, the models presented in this research do not account for decreases in water *quality*. I claim that, in some instances, water sharing mechanisms generate more value due to a higher capacity to capitalize on return flows. However, return flows tend to be associated with decreases in water quality, which can impact downstream value generation. Lastly, to consider the feasibility of policy action, moving away from a property rights system like PA can be exceptionally difficult.

Despite some limitations, the results from this paper have important policy implications. First, it is a poignant reminder for policymakers that the design of water allocation systems must account for the geophysical context to ensure both the efficient allocation and the maximal utilization of water resources. The composition of a river basin and the initial allocation of property rights have implications for both the efficiency of markets and the total volume of water that can be used. The presence of return flows and the location of users need to be considered when designing rules for governance.

For future work, the modeling framework presented in this study could be modified to further investigate the availability and timing of return flows, which could also impact institutional efficiency. In reality, the return flows from one user may not be available for all downstream users. There may also be uncertainty about *when* that water reenters the system. Additional constraints could also be included to emulate instream flow water rights. States such as Arizona and Colorado have incorporated the use of water rights to ensure certain sections of rivers maintain a minimum flow for ecological purposes. These minimum flow requirements would generate ecological benefits, though there may be trade offs for the location and seniority of instream flow water rights.

3.4.1 The Colorado River

The findings presented in this paper hold significant importance in the context of ongoing negotiations among Western states for reducing withdrawals of Colorado River water. The US

Bureau of Reclamation has issued warnings of imposing more stringent regulations on water release from reservoirs if a total demand reduction of 4 million acre-feet is not achieved (Flavelle 2023a). California, Arizona and Nevada came to a temporary agreement to reduce water use, however the federal government will pay \$1.2 billion to various stakeholders to achieve that reduction (Flavelle 2023b). As climate change exacerbates water scarcity in the western United States, there is an urgent need to develop and implement efficient water allocation institutions that can adapt to evolving environmental and socioeconomic conditions.

The Colorado River is a crucial water source for seven states in the US and two states in Mexico, supporting over 40 million people and extensive agricultural activities. Given the river's importance, the ongoing negotiations between states aim to address the increasing demand for its water resources and balance the competing interests of various stakeholders. This research can inform these negotiations by providing insights into the relative efficiency of different water allocation mechanisms under different basin compositions and water supply conditions. Understanding the strengths and weaknesses of existing mechanisms, such as prior appropriation and share-based systems, can help policymakers identify potential improvements in water governance that foster more equitable and sustainable water use.

Moreover, the US Bureau of Reclamation's warning highlights the need for effective demand management strategies to avoid severe restrictions on water releases from reservoirs, which could have significant consequences for water-dependent communities and ecosystems. The results from this analysis can guide decision-makers in identifying water allocation mechanisms that promote more efficient water use. By examining the impacts of heterogeneity in value generation and user location on allocative efficiency, this study can also inform targeted policy interventions that address the specific needs and constraints of different user groups within a river basin. If water markets are the preferred intervention, this article provides insights into how much participation is necessary to outperform other mechanisms. Ultimately, this research has the potential to contribute to more sustainable water management practices in the western United States and other arid climates. By identifying the most efficient water allocation mechanisms under various basin

compositions, it can support the development of policies that promote more equitable and efficient water use in the face of growing demands and climate change-induced water scarcity.

3.5 Tables

Table 3.1: Water allocation outcomes under different water supply scenarios; $R_i^s = 250$, $R_i^j = 200$, $\theta = 0.25$

$\tilde{w}$	(i.)			(ii.)		
	w_u^s	w_d^j	Total	w_d^s	w_u^j	Total
100	100	25	125	100	0	100
300	250	112.5	362.5	250	66.7	316.7
500	250	200	450	250	200	450

Table 3.2: Detailed parameters for water supply simulations

Parameter	Sample Set	Replacement
Water Supply ($\tilde{w}$)	Log-normal distribution	With
User Endowments (R)	500	N/A
Number of Users (n)	3 to 36 in increments of 1	N/A
Stream Position (i)	1 to n	Without
Priority Number (m)	1 to n	Without
Priority Coefficient (β)	20^{-m}	N/A
Return Flow Coefficient (θ)	0.1, 0.2, 0.3, ..., 0.9	With
Value Coefficient (v)	Normal distribution	With
Value Function	$v \cdot w^{\frac{1}{2}}$	N/A

Table 3.3: Parameters for a sample simulation with ten users and $\tilde{w} = 2, 252$

Parameter	User 1	User 2	User 3	User 4	User 5	User 6	User 7	User 8	User 9	User 10
Stream position (i)	1	2	3	4	5	6	7	8	9	10
Priority no. (m)	5	6	4	10	7	8	2	1	3	9
Share coef. (δ)	0.085	0.065	0.11	0.01	0.048	0.034	0.193	0.293	0.143	0.021
Value coef. (v)	68	738	435	225	950	454	121	592	321	534
Return flow coef. (θ)	0.4	0.1	0.8	0.6	0.7	0.8	0.8	0.3	0.7	0.3
Endowment (R)	500	500	500	500	500	500	500	500	500	500

3.6 Figures




Figure 3.1: Spatial heterogeneity in hydraulic connectivity of western HUC8 basins per Duan et al. (2018)




Figure 3.2: Visual representation of a hypothetical waterway with three users in which return flows from an upstream user are available to all downstream users.




Figure 3.3: Monte Carlo generated total value distributions under different water allocation mechanisms with two water users




Figure 3.4: Average total value generated and percentage of reallocated water by institution and number of water users




Figure 3.5: How priority location impacts the total amount of water available for trade




Figure 3.6: Distributions of total value losses relative to the Social Planner by institution, priority composition, and water shortage scenarios

3.7 Supplementary Analysis

3.7.1 The Role of Return Flows in Water Allocation

Consider n users along a river, with a uniform return rate θ for water used by each user. Each user i has a concave, differentiable value function $v(w_i)$, representing the total value generated from water diverted w_i . The objective of the Social Planner is to maximize the total value from water allocations subject to the total available water $\tilde{w}$. The simplified problem is formulated as:

$$\max_{w_1, \dots, w_n \geq 0} \sum_{i=1}^n v(w_i) \quad \text{s.t.} \quad w_1 + (1 - \theta)w_2 + \dots + (1 - \theta)^{n-1}w_n \leq \tilde{w} \quad (3.18)$$

The Lagrangian for this optimization problem is:

$$\mathcal{L} = \sum_{i=1}^n v(w_i) - \lambda \left(\sum_{i=1}^n (1 - \theta)^{i-1} w_i - \tilde{w} \right), \quad (3.19)$$

where λ is the Lagrange multiplier for the water availability constraint. Differentiating $\mathcal{L}$ with respect to w_i yields the first-order conditions:

$$\frac{dv}{dw_i} = \lambda(1 - \theta)^{i-1}. \quad (3.20)$$

Since $(1 - \theta) \in (0, 1)$ and i indexes stream position from upstream to downstream, this condition shows that the marginal value of water allocation decreases as we move downstream, adjusting for the return flow effect. Thus, the allocation w_i must decrease for increasing i to balance marginal value across users, leading to $w_1 > w_2 > \dots > w_n$. This pattern indicates that, due to return flows, more water is allocated upstream *ceteris paribus*.

3.7.2 Regression Exercises

In the above analysis, institutional performance is compared across hundreds of thousands of simulations under randomly generated basin compositions and different supply scenarios. While

the mechanisms that consistently generate a large amount of value relative to others show robustness to varying basin characteristics, it is perhaps still not entirely clear how the institutional and physical constraints may mitigate or exacerbate inefficiencies. As a final exercise to try and characterize why certain mechanisms perform better or worse under different circumstances, I employ a series of regressions to explain why a mechanism generates the most value for a given simulation. Specifically, I run separate linear regression models for each water allocation mechanism:

$$\frac{y_k^j - y_k^{SP}}{y_k^{SP}} = \alpha^j + \beta_1^j X_k^j + \beta_2^j C_k + \varepsilon_k^j, \quad j = 1, \dots, J, \quad (3.21)$$

where $\frac{y_k^j - y_k^{SP}}{y_k^{SP}}$ is the proportional decrease in value generated for mechanism j relative to the social planner in simulation iteration k . Matrix X_k^j contains variables characterizing institution-specific constraints, such as the correlation between priority and stream position and the correlation between priority and value. Matrix C_k contains information characterizing the physical characteristics of the basin such as the number of users, the level of shortage, and the correlation between stream position and return flows. The goal of this exercise is to identify which mechanisms are most sensitive to institutional constraints and basin compositions. Regression output from equation (18) is available in Table 3.4.

3.7.3 Other Considerations

In addition to system-wide value, there are other outcomes from allocating water that warrant attention. For a 36 user system, Figure 3.7 displays distributions for total value, average value per unit of diverted water, total diverted water, and basin outflow for each institution.

Table 3.4: Cross-sectional regression results assessing the impacts of basin composition and water shortage on institutional efficiency relative to the Social Planner

Model:	PA (1)	PAM (2)	PS (3)	PSM (4)	ES (5)	ESM (6)
Constant	-0.034*** (0.0008)	-0.078*** (0.0005)	-0.141*** (0.0004)	-0.095*** (0.0004)	-0.086*** (0.0003)	-0.037*** (0.0003)
# Users	-0.017*** (7.6×10^{-5})	-0.005*** (5.02×10^{-5})	-0.009*** (4.15×10^{-5})	-0.005*** (3.96×10^{-5})	-0.010*** (3.46×10^{-5})	-0.008*** (3.27×10^{-5})
# Users Sq.	0.0003*** (1.94×10^{-6})	5.35×10^{-5} *** (1.28×10^{-6})	0.0001*** (1.06×10^{-6})	7.94×10^{-5} *** (1.01×10^{-6})	0.0002*** (8.82×10^{-7})	0.0001*** (8.32×10^{-7})
Shortage	-0.001*** (6.86×10^{-6})	0.0003*** (4.53×10^{-6})	-0.0004*** (3.75×10^{-6})	0.0001*** (3.58×10^{-6})	0.0002*** (3.13×10^{-6})	0.0002*** (2.95×10^{-6})
Shortage Sq.	-9.13×10^{-7} *** (2.53×10^{-8})	-8.67×10^{-7} *** (1.67×10^{-8})	1.22×10^{-6} *** (1.38×10^{-8})	2.34×10^{-7} *** (1.32×10^{-8})	-4.02×10^{-7} *** (1.15×10^{-8})	-3.71×10^{-7} *** (1.09×10^{-8})
Prio. SP Corr.	0.080*** (0.0005)	0.035*** (0.0003)	0.025*** (0.0003)	0.023*** (0.0003)	-0.0005** (0.0002)	-0.0006*** (0.0002)
Prio. Value Corr.	0.148*** (0.0005)	0.004*** (0.0003)	0.106*** (0.0003)	0.002*** (0.0003)	0.0004 (0.0002)	0.0002 (0.0002)
RF SP Corr.	-0.044*** (0.0005)	-0.033*** (0.0003)	-0.023*** (0.0003)	-0.025*** (0.0003)	-0.032*** (0.0002)	-0.033*** (0.0002)
<i>Fit statistics</i>						
Observations	467,773	467,773	467,773	467,773	467,773	467,773
R ²	0.59827	0.18693	0.38230	0.21444	0.40189	0.37203

IID standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*




Figure 3.7: Monte Carlo generated total value, average value, total water, and basin outflow distributions with 36 water users

Bibliography

- Abatzoglou, J. T. (2013), ‘Development of gridded surface meteorological data for ecological applications and modelling’, *International Journal of Climatology* **33**(1), 121–131.
- Adam, J. C., Hamlet, A. F. & Lettenmaier, D. P. (2009), ‘Implications of global climate change for snowmelt hydrology in the twenty-first century’, *Hydrological Processes: An International Journal* **23**(7), 962–972.
- Auffhammer, M. & Schlenker, W. (2014), ‘Empirical studies on agricultural impacts and adaptation’, *Energy Economics* **46**, 555–561.
- Bauder, T. A., Waskom, R. M. & Andales, A. (2014), ‘Nitrogen and irrigation management’, *Colorado State University Extension Fact Sheet* (0.514).
- Bennett, L. L., Howe, C. W. & Shope, J. (2000), ‘The interstate river compact as a water allocation mechanism: efficiency aspects’, *American Journal of Agricultural Economics* **82**(4), 1006–1015.
- Blanc, E. & Schlenker, W. (2017), ‘The use of panel models in assessments of climate impacts on agriculture’, *Review of Environmental Economics and Policy* .
- Blumberg, J., Goemans, C. & Manning, D. (2024), ‘Perceived water scarcity and irrigation technology adoption’, *American Agriculture, Water Resources, and Climate Change* p. 173.
- Brewer, J., Glennon, R., Ker, A. & Libecap, G. (2006), ‘Transferring water in the american west: 1987-2005’, *U. Mich. JL Reform* **40**, 1021.

- Brewer, J., Glennon, R., Ker, A. & Libecap, G. (2008), '2006 presidential address water markets in the west: prices, trading, and contractual forms', *Economic Inquiry* **46**(2), 91–112.
- Burke, M., Hsiang, S. M. & Miguel, E. (2015), 'Global non-linear effect of temperature on economic production', *Nature* **527**(7577), 235–239.
- Burness, H. S. & Quirk, J. P. (1979), 'Appropriative water rights and the efficient allocation of resources', *The American Economic Review* **69**(1), 25–37.
- Burness, H. S. & Quirk, J. P. (1980), 'Water law, water transfers, and economic efficiency: The colorado river', *The Journal of Law and Economics* **23**(1), 111–134.
- Burnham, K. P. & Anderson, D. R. (2004), 'Multimodel inference: understanding aic and bic in model selection', *Sociological methods & research* **33**(2), 261–304.
- Butler, C. D. (2018), 'Climate change, health and existential risks to civilization: A comprehensive review (1989–2013)', *International journal of environmental research and public health* **15**(10).
- Callaway, B. & Sant'Anna, P. H. (2021), 'Difference-in-differences with multiple time periods', *Journal of Econometrics* **225**(2), 200–230.
- Carey, J. M. & Zilberman, D. (2002), 'A model of investment under uncertainty: modern irrigation technology and emerging markets in water', *American Journal of Agricultural Economics* **84**(1), 171–183.
- Case, A. (1992), 'Neighborhood influence and technological change', *Regional science and urban economics* **22**(3), 491–508.
- Cline, W. R. et al. (1992), *The economics of global warming*, Institute for International Economics.
- Coase, R. (1960), 'The problem of social cost', *Journal of Law and Economics* pp. 1–44.
- Cobourn, K. M., Ji, X., Mooney, S. & Crescenti, N. F. (2021), 'The effect of prior appropriation water rights on land-allocation decisions in irrigated agriculture', *American Journal of Agricultural Economics* .

- CWCB (2015), 'Colorado's water plan', *Colorado Division of Natural Resources* .
- Dell, M., Jones, B. F. & Olken, B. A. (2012), 'Temperature shocks and economic growth: Evidence from the last half century', *American Economic Journal: Macroeconomics* **4**(3), 66–95.
- Deryugina, T. & Hsiang, S. (2017), The marginal product of climate, Technical report, National Bureau of Economic Research.
- Deschênes, O. & Greenstone, M. (2007), 'The economic impacts of climate change: evidence from agricultural output and random fluctuations in weather', *American economic review* **97**(1), 354–385.
- Deschênes, O. & Greenstone, M. (2012), 'The economic impacts of climate change: evidence from agricultural output and random fluctuations in weather: reply', *American Economic Review* **102**(7), 3761–3773.
- Dieter, C., Maupin, M., Caldwell, R., Harris, M., Ivahnenko, T., Lovelace, J., Barber, N. & Linsey, K. (2018), 'Estimated use of water in the United States in 2015', *US Geological Survey* .
- Diffenbaugh, N. S. & Field, C. B. (2013), 'Changes in ecologically critical terrestrial climate conditions', *Science* **341**(6145), 486–492.
- Duan, K., Sun, G., Caldwell, P. V., McNulty, S. G. & Zhang, Y. (2018), 'Implications of upstream flow availability for watershed surface water supply across the conterminous united states', *JAWRA Journal of the American Water Resources Association* **54**(3), 694–707.
- EPA (2016), 'Climate change indicators in the United States', *United States Environmental Protection Agency* .
- Fisher, A. C., Hanemann, W. M., Roberts, M. J. & Schlenker, W. (2012), 'The economic impacts of climate change: evidence from agricultural output and random fluctuations in weather: comment', *American Economic Review* **102**(7), 3749–3760.

- Flavelle, C. (2023a), ‘Biden administration proposes evenly cutting water allotments from colorado river’, *The New York Times* .
- Flavelle, C. (2023b), ‘A breakthrough deal to keep the colorado river from going dry, for now’, *The New York Times* .
- Foster, T. & Brozović, N. (2018), ‘Simulating crop-water production functions using crop growth models to support water policy assessments’, *Ecological Economics* **152**, 9–21.
- Frisvold, G. & Bai, T. (2016), ‘Irrigation technology choice as adaptation to climate change in the western United States’, *Journal of Contemporary Water Research & Education* **158**(1), 62–77.
- Getches, D. H. (2009), *Water Law in a Nutshell*, 4th edn, Thomson West.
- Graff Zivin, J. & Neidell, M. (2014), ‘Temperature and the allocation of time: Implications for climate change’, *Journal of Labor Economics* **32**(1), 1–26.
- Grafton, R. Q., Libecap, G., McGlennon, S., Landry, C. & O’Brien, B. (2011), ‘An integrated assessment of water markets: a cross-country comparison’, *Review of Environmental Economics and Policy* .
- Gupta, A., Rico-Medina, A. & Caño-Delgado, A. I. (2020), ‘The physiology of plant responses to drought’, *Science* **368**(6488), 266–269.
- Gutzler, D. S. & Robbins, T. O. (2011), ‘Climate variability and projected change in the western United States: regional downscaling and drought statistics’, *Climate Dynamics* **37**(5), 835–849.
- Hamlet, A. F., Mote, P. W., Clark, M. P. & Lettenmaier, D. P. (2005), ‘Effects of temperature and precipitation variability on snowpack trends in the western United States’, *Journal of Climate* **18**(21), 4545–4561.
- Heal, G. & Park, J. (2020), ‘Reflections—temperature stress and the direct impact of climate change: a review of an emerging literature’, *Review of Environmental Economics and Policy* .

- Henry III, D. D., Muller, N. Z. & Mendelsohn, R. O. (2011), 'The social cost of trading: Measuring the increased damages from sulfur dioxide trading in the united states', *Journal of Policy Analysis and Management* **30**(3), 598–612.
- Howden, S. M., Soussana, J.-F., Tubiello, F. N., Chhetri, N., Dunlop, M. & Meinke, H. (2007), 'Adapting agriculture to climate change', *Proceedings of the National Academy of Sciences* **104**(50), 19691–19696.
- Hsiang, S. (2016), 'Climate econometrics', *Annual Review of Resource Economics* **8**, 43–75.
- Huffaker, R., Whittlesey, N. & Hamilton, J. R. (2000), 'The role of prior appropriation in allocating water resources into the 21st century', *International Journal of water resources development* **16**(2), 265–273.
- Ji, X. & Cobourn, K. M. (2018), 'The economic benefits of irrigation districts under prior appropriation doctrine: An econometric analysis of agricultural land-allocation decisions', *Canadian Journal of Agricultural Economics/Revue canadienne d'agroeconomie* **66**(3), 441–467.
- Ji, X. & Cobourn, K. M. (2021), 'Weather fluctuations, expectation formation, and short-run behavioral responses to climate change', *Environmental and Resource Economics* **78**(1), 77–119.
- Johnston, J. & DiNardo, J. (2009), *Econometric methods*, 4th edn, McGraw-Hill.
- Jones, P. A. & Cech, T. (2009), *Colorado water law for non-lawyers*, University Press of Colorado.
- Kassel, K. & Martin, A. (2023), 'Ag and food sectors and the economy', *Economic Research Service, Ag and Food Statistics: Charting the Essentials* .
- Kolstad, C. D. & Moore, F. C. (2020), 'Estimating the economic impacts of climate change using weather observations', *Review of Environmental Economics and Policy* .
- Koundouri, P., Nauges, C. & Tzouvelekas, V. (2006), 'Technology adoption under production uncertainty: theory and application to irrigation technology', *American Journal of Agricultural Economics* **88**(3), 657–670.

- Lemoine, D. (2018), Estimating the consequences of climate change from variation in weather, Technical report, National Bureau of Economic Research.
- Leonard, B., Costello, C. & Libecap, G. D. (2019), 'Expanding water markets in the western united states: Barriers and lessons from other natural resource markets', *Review of Environmental Economics and Policy*.
- Leonard, B. & Libecap, G. D. (2019), 'Collective action by contract: prior appropriation and the development of irrigation in the western United States', *The Journal of Law and Economics* **62**(1), 67–115.
- Li, M., Xu, W. & Zhu, T. (2019), 'Agricultural water allocation under uncertainty: redistribution of water shortage risk', *American Journal of Agricultural Economics* **101**(1), 134–153.
- Libecap, G. D. (2007), 'The assignment of property rights on the western frontier: lessons for contemporary environmental and resource policy', *The Journal of Economic History* **67**(2), 257–291.
- Libecap, G. D. (2011), 'Institutional path dependence in climate adaptation: Coman's "some unsettled problems of irrigation"', *American Economic Review* **101**(1), 64–80.
- Linville, D. E. (1990), 'Calculating chilling hours and chill units from daily maximum and minimum temperature observations', *HortScience* **25**(1), 14–16.
- Lioubimtseva, E. (2004), 'Climate change in arid environments: revisiting the past to understand the future', *Progress in Physical Geography* **28**(4), 502–530.
- Lobell, D. B., Roberts, M. J., Schlenker, W., Braun, N., Little, B. B., Rejesus, R. M. & Hammer, G. L. (2014), 'Greater sensitivity to drought accompanies maize yield increase in the us midwest', *Science* **344**(6183), 516–519.
- Luedeling, E. & Fernandex, E. (2022), *chillR: Statistical Methods for Phenology Analysis in Temperate Fruit Trees*. R package version 0.72.7.

- Maddison, D. (2007), *The perception of and adaptation to climate change in Africa*, Vol. 4308, World Bank Publications.
- Mancosu, N., Snyder, R. L., Kyriakakis, G. & Spano, D. (2015), 'Water scarcity and future challenges for food production', *Water* **7**(3), 975–992.
- Masson-Delmonte, V., Zhai, P., Pirani, A., Connors, S., Péan, C., Berger, S., Caud, N., Chen, Y., Goldfarb, L., Gomis, M. et al. (2021), 'IPCC 2021: Climate Change 2021: The physical science basis: Contribution of working group i to the sixth assessment report of the intergovernmental panel on climate change'.
- McIntosh, C. T. & Schlenker, W. (2006), 'Identifying non-linearities in fixed effects models', *UC-San Diego Working Paper*.
- McKee, T. B., Doesken, N. J., Kleist, J., Shrier, C. J. & Stanton, W. P. (2000), 'A history of drought in colorado: Lessons learned and what lies ahead'.
- Mendelsohn, R., Nordhaus, W. D. & Shaw, D. (1994), 'The impact of global warming on agriculture: a ricardian analysis', *The American economic review* pp. 753–771.
- Mérel, P. & Gammans, M. (2021), 'Climate econometrics: Can the panel approach account for long-run adaptation?', *American Journal of Agricultural Economics* **103**(4), 1207–1238.
- Michler, J. D., Josephson, A., Kilic, T. & Murray, S. (2020), 'Estimating the impact of weather on agriculture', *arXiv preprint arXiv:2012.11768*.
- Mote, P. W. (2006), 'Climate-driven variability and trends in mountain snowpack in western north america', *Journal of Climate* **19**(23), 6209–6220.
- Mote, P. W., Hamlet, A. F., Clark, M. P. & Lettenmaier, D. P. (2005), 'Declining mountain snowpack in western north america', *Bulletin of the American Meteorological Society* **86**(1), 39–50.
- National Academies of Sciences, E., Medicine et al. (2017), *Valuing climate damages: updating estimation of the social cost of carbon dioxide*, National Academies Press.

- Newell, R. G., Prest, B. C. & Sexton, S. E. (2021), ‘The gdp-temperature relationship: implications for climate change damages’, *Journal of Environmental Economics and Management* **108**, 102445.
- Nordhaus, W. D. (1992), ‘An optimal transition path for controlling greenhouse gases’, *Science* **258**(5086), 1315–1319.
- Nordhaus, W. D. (1994), *Managing the global commons: the economics of climate change*, Vol. 31, MIT press Cambridge, MA.
- Olper, A., Maugeri, M., Manara, V. & Raimondi, V. (2021), ‘Weather, climate and economic outcomes: Evidence from Italy’, *Ecological Economics* **189**, 107156.
- Pielke, R. A., Doesken, N., Bliss, O., Green, T., Chaffin, C., Salas, J. D., Woodhouse, C. A., Lukas, J. J. & Wolter, K. (2005), ‘Drought 2002 in Colorado: an unprecedented drought or a routine drought?’, *Pure and Applied Geophysics* **162**(8), 1455–1479.
- Pritchett, J., Thorvaldson, J. & Frasier, M. (2008), ‘Water as a crop: limited irrigation and water leasing in Colorado’, *Review of Agricultural Economics* **30**(3), 435–444.
- Schauberger, B., Archontoulis, S., Arneth, A., Balkovic, J., Ciais, P., Deryng, D., Elliott, J., Folberth, C., Khabarov, N., Müller, C. et al. (2017), ‘Consistent negative response of US crops to high temperatures in observations and crop models’, *Nature Communications* **8**(1), 1–9.
- Schlenker, W., Hanemann, W. M. & Fisher, A. C. (2006), ‘The impact of global warming on US agriculture: an econometric analysis of optimal growing conditions’, *Review of Economics and Statistics* **88**(1), 113–125.
- Schlenker, W. & Lobell, D. B. (2010), ‘Robust negative impacts of climate change on African agriculture’, *Environmental Research Letters* **5**(1), 014010.

- Schlenker, W., Michael Hanemann, W. & Fisher, A. C. (2005), 'Will us agriculture really benefit from global warming? accounting for irrigation in the hedonic approach', *American Economic Review* **95**(1), 395–406.
- Schlenker, W. & Roberts, M. J. (2009), 'Nonlinear temperature effects indicate severe damages to us crop yields under climate change', *Proceedings of the National Academy of sciences* **106**(37), 15594–15598.
- Schneekloth, J. & Andales, A. (2017), 'Seasonal water needs and opportunities for limited irrigation for colorado crops', *Colorado State University Extension Fact Sheet* (4.718).
- Spence, A., Poortinga, W., Butler, C. & Pidgeon, N. F. (2011), 'Perceptions of climate change and willingness to save energy related to flood experience', *Nature climate change* **1**(1), 46–49.
- Stevens, A. (2019), Temperature, wages, and agricultural labor productivity, Technical report, UC Berkeley Working Paper, Accessible on UC Berkeley website.
- Stott, P. (2016), 'How climate change affects extreme weather events', *Science* **352**(6293), 1517–1518.
- Tingey-Holyoak, J. L. (2014), 'Water sharing risk in agriculture: Perceptions of farm dam management accountability in australia', *Agricultural Water Management* **145**, 123–133.
- US Census Bureau (2021), 'New vintage 2021 population estimates available for the nation, states and Puerto Rico', *RELEASE NUMBER CB21-208* .
- Wallander, S., Aillery, M., Hellerstein, D. & Hand, M. (2013), 'The role of conservation programs in drought risk adaptation', *Economic Research Service ERR* **148**.
- Walthall, C. L., Anderson, C. J., Baumgard, L. H., Takle, E., Wright-Morton, L. et al. (2013), 'Climate change and agriculture in the united states: Effects and adaptation'.
- Wang, X. (2017), 'Understanding climate change risk perceptions in china: Media use, personal experience, and cultural worldviews', *Science Communication* **39**(3), 291–312.

- Wang, X. & Lin, L. (2018), 'The relationships among actual weather events, perceived unusual weather, media use, and global warming belief certainty in china', *Weather, climate, and society* **10**(1), 137–144.
- Waskom, R. M. (2013), 'HB12-1278 Study of the South Platte River Alluvial Aquifer', *Colorado Water Institute* .
- Waskom, R. & Neibauer, M. (2014), 'Water conservation in and around the home', *Colorado State University Extension* **9.952**.
- Waskom, R., Rein, K., Wolfe, D. & Smith, M. (2016), 'How diversion and beneficial use of water affect the value and measure of a water right', *Fort Collins, CO, USA* .