Department of Economics and Business
Working Paper Series

What Drives the Value of Water Markets Under Uncertainty? The Economics of Water Supply Forecasts Across Snow and Rain-fed Basins

Hannah Kamen
Jared C. Carbone
Ben Livneh
Parthkumar Modi
Eric Small
Bill Szafranski
Cameron Wobus

Working Paper 2026-01
http://econbus-papers.mines.edu/working-papers/wp202601.pdf

Colorado School of Mines
Department of Economics and Business
1500 Illinois Street
Golden, CO 80401

April 2026

Colorado School of Mines
Department of Economics and Business
Working Paper No. **2026-01**
April 2026

Title:

What Drives the Value of Water Markets Under Uncertainty? The Economics of Water Supply Forecasts Across Snow and Rain-fed Basins*

Author(s):

Hannah Kamen, Corresponding author: Department of Economics and Business, Colorado School of Mines, Golden, CO 80401, hkamen@mines.edu

Jared C. Carbone, Department of Economics and Business, Colorado School of Mines

Ben Livneh, Department of Civil, Environmental and Architectural Engineering, University of Colorado and CIRES, University of Colorado

Parthkumar Modi, Alabama Water Institute

Eric Small, Department of Geological Sciences, University of Colorado

Bill Szafranski, Lynker Corporation

Cameron Wobus, CK Blueshift, LLC

ABSTRACT

Climate change is shifting Western US basins from snow-fed toward rain-fed systems, weakening the precision of water supply forecasts. We investigate the value of water markets in reducing economic damages from supply uncertainty. We produce hydrologic supply forecasts from data across 889 SNOTEL stations and show that forecast uncertainty is higher in rain-fed basins. We embed these empirical uncertainty measures into a water-trade model and find that when some production decisions are irreversible (e.g., planting in agriculture), and made before water availability is realized, trade helps offset losses from sub-optimal investments. We find that gains are 7% larger in rain-fed basins on average, with offsetting effects increasing in both forecast bias and variance. As rising temperatures are expected to reduce the fraction of precipitation falling as snow, we project that water trade will become a more valuable mechanism for rationing water.

JEL **classifications: D81, E27, F17, Q25, Q54**
Keywords: climate change, drought, forecast uncertainty, water trade

*We gratefully acknowledge funding and support from the National Science Foundation's Dynamics of Integrated Socio-Environmental Systems (CHN2) program. We are thankful for comments received at the American Geophysical Union (AGU) Fall Meeting, the AGU Chapman Conference, the University of Colorado Environmental and Resource Economics Workshop, and the Colorado School of Mines. We also thank Steven Smith for valuable advice and feedback. Any errors are our own.

1. Introduction

In much of the western United States, recurring drought and rising variability in water supplies have intensified interest in policies that can reallocate scarce water toward higher value uses. Water markets are often proposed as a key adaptation tool because they can, in principle, move water to producers with the highest willingness to pay during shortages. Yet in many basins, allocation remains largely determined by rights seniority rather than value, so water can remain tied to relatively low value agricultural uses even when scarcity is severe (MacDonnell 2014; Zellmer and Amos 2021). This institutional setting is shaped by the prior appropriation doctrine, which assigns rights based on beneficial use and developed in contrast to riparian doctrines more common in the eastern United States (Burness and Quirk 1979). Prior appropriation also provides the governance foundation that makes water transfers contractible by establishing and enforcing property rights, without which there is little incentive to enter water rights contracts (Greif et al. 1994; Dixit 2009; Leonard and Libecap 2015). At the same time, because western rights were historically allocated to agriculture and remain concentrated in relatively low value crops, prior appropriation alone does not ensure allocation to the highest value use (Brewer et al. 2008). A large literature therefore studies water markets as a mechanism to reduce economic losses from drought (Vaux and Howitt 1984; Grafton et al. 2011; Goemans and Pritchett 2014; Rafey 2023) and under climate change (Gohar and Ward 2010; Hagerty 2019; Arellano Gonzalez et al. 2021). However, much less is known about how producers respond when water supplies are uncertain (Marques et al. 2005; Calatrava and Garrido 2005; Rigby et al. 2010; Zuo et al. 2014), and more specifically how the economic value of water markets depends on the degree and structure of that uncertainty.

This paper brings the research question to the foreground: how do gains from allowing water trade change as seasonal water supply forecasts become more or less accurate, and which features of forecast performance matter most for market value? Related work that explicitly links uncertainty to the value of trade is limited. Calatrava and Garrido (2005), for example, show in a simulation of an agricultural economy in a Spanish river basin that water trading can offset economic losses associated with uncertain supplies. What remains unresolved is how the value of water markets varies with the specific parameters that determine hydrologic forecast accuracy, and therefore how the gains from implementing or improving markets change when forecast conditions shift. This question is increasingly relevant given hydrologic evidence that climate change is degrading seasonal water supply forecasts in historically snow fed basins by shifting key forecast drivers, increasing the risk of larger and more frequent forecast errors (Livneh and Badger 2020; Modi et al. 2022). In the Western United States, our setting, this matters both because many basins remain snow fed and because allocation under prior appropriation has historically been relatively inflexible. Together, these features make it critical to understand how the benefits of water markets depend on water supply uncertainty, and how those benefits evolve as forecast accuracy changes.

The concern that water supply forecast accuracy will degrade in historically drought prone regions such as the Western US is well supported in the hydrological literature. Here we define “water supply forecasts” broadly as season-ahead predictions of total water supply volume available for use in agricultural production. Over half of annual runoff across the Western US historically originates as snowmelt (Li and Zhang., 2017), making knowledge of snowpack crucial to the predictability of seasonal water supply forecasts (Page and Dilling, 2019). However, ongoing and projected climate warming is expected to reduce spring snow water equivalent (SWE) across many parts of the Western United States (e.g., Li and Zhang., 2017; Siirila-Woodburn et al., 2021) and alter snowmelt timing (Dettinger and Cayan ,1995; Stewart et al., 2004). As a result, the reliability and performance of these forecasts in snow-dominated systems will be compromised (e.g., Livneh and Badger, 2020; Modi et al., 2022), and total water availability diminished (Berghuijs et al. 2014, Xiao et al., 2019, Williams et al., 2020). These changes, attributable to changes in global atmospheric and oceanic temperature and circulations result in regional changes in precipitation and temperature (Zhang et al., 2021). It is expected that such changes will impact seasonal forecasting systems, especially in rain-dominated systems, which have been shown to be less skillful than in snow-dominated systems (Pagano et al., 2004). Rising temperatures may reduce the fraction of precipitation falling as snow by up to 30% (Figure 6) in intermountain and continental regions of the Rockies in the future, fundamentally altering the regional water cycle.

These factors combine to exacerbate the financial risk faced by producers insofar as they must use increasingly uncertain forecasting information to make season-ahead investment decisions. In the agricultural context that we reference in our model, season-ahead investment decisions are highly applicable; for example, farmers typically hire farm-hands or purchase harvest materials prior to the growing season (Calatrava and Garrido 2005) or make decisions regarding irrigation acreage (Li et al. 2017) or cropping patterns (Marques 2005), though applicability is certainly not limited to this sector. Given the importance of season-ahead investments in water intensive sectors, understanding the role of water trade in ameliorating losses from supply uncertainty is critical due to the likelihood that climate change will increase forecast uncertainty.

Our contribution is two-fold. First, we develop a theoretical foundation that demonstrates how water supply forecast accuracy interacts with the value of water markets. Our theoretical findings demonstrate that trade increases the degree to which producers can offset losses from water supply uncertainty relative to a water allocation based on relative seniority of water rights (e.g., prior appropriation). Second, we examine the relationship between basin precipitation from snowpack (versus rainfall), forecast accuracy, and the value of water markets using a water forecast dataset comprised of snow-observing basins located in states within the US Census-defined Pacific and Mountain divisions of the United States, referred to in the remainder of this paper as the Western US (Census Bureau 2024). We find that forecast uncertainty is substantially higher in rain-fed basins: median normalized forecast standard deviation rises from 2% in predominantly snow-fed basins to 6.8% in rain-fed basins, and median absolute forecast bias ranges from 3% to 15% across these categories. We then embed these empirical uncertainty measures into a two-period water-trade model and find that profit gains from trade are on average 7% larger in predominantly rain-fed basins. Unlike much of the existing water market literature, we show that the gains from trade depend on drought severity as well as uncertainty in season ahead water forecasts. When producers make irreversible investments before water supplies are realized, trade can partially undo the efficiency losses from these ex ante, and often suboptimal, investment choices. We show that as forecast bias and variance increase, so does the value of allowing trade relative to rigid allocation.

Our work suggests that water markets serve as a substitute for forecast accuracy. This is particularly important due to the likelihood that forecast accuracy will decrease as climate change results in a lower proportion of precipitation from snowfall (Li et al., 2017; Siirila-Woodburn et al., 2021). Taken together, these findings suggest that water markets may become a more valuable policy tool not only due to increased water scarcity, but also due to increased water supply uncertainty. We proceed in the following manner. Section 2 describes a simple model of water trade and discusses its theoretical implications. Section 3 describes the process by which we link our economic model to our forecasting dataset to generate estimates of returns to trade across forecasts with varying levels of uncertainty. Section 4 discusses our results, and Section 5 concludes.

2. Water Trade under Uncertainty

In this section, we define key forecast parameters and describe a simple model of water trade between agricultural producers in a prior appropriation rights system when future water supply is uncertain. This framework helps us establish an intuition for understanding the role of uncertainty in determining water market value, which we use to inform our quantitative results in section 4. For tractability, we abstract from details often incorporated into water rights contracts in our model, including transaction costs, timing of use, and return flow externalities (Womble and Hanneman 2020a).

In our analysis, we define season-ahead water supply forecasts as information received by agricultural producers pertaining to expected water supply volume during the growing season that same year, prior to the growing season and coincident with the point in time at which pre-season investment decisions must be made by the producer. While our methodology for generating hydrological forecasts implies a zero-lead time, which means producers would have access to such information on April 1st of the year they intend to plant, we abstract from examining how the date of forecast dispersal relative to the growing season may implicate water market value, however this would be an interesting topic for future research.

Our forecast parameters of interest are bias and variance. Forecast bias is defined as the sign and magnitude of the percent difference between the mean of the true water supply distribution and the forecast supply distribution. Forecast variance represents the variance of the distribution of expected water supply outcomes. Table 1 defines the key terms and parameters related to these concepts that we use throughout the paper.

Table 1. Important terms and parameters

Term (parameter)	**Definition**
$Forecast\ mean\ (\mu_f)$	Mean forecasted water supply
$Realization\ mean\ (\mu_o)$	Mean observed water supply

Forecast variance (σ_f^2)	Variance of water forecast distribution (equivalent to distribution of observed outcomes).
Forecast bias $[(\mu_f - \mu_o)/\ \mu_o]\times 100$	Percent difference between realization mean and forecast mean

2.1 Water Trade Model Overview

In this section we discuss a general overview of the water trade model we use to represent water allocation under prior appropriation law. Sections 2.2 and 2.3 discuss the specific functional forms driving producer decision-making, and how the value of trade is determined in our model. In this context, we interpret a prior appropriation system as a rule for water allocation in which a senior producer receives water in priority to junior producers in the event of drought, capturing the fundamental allocation feature of prior appropriation law (Johnson et al. 1981). Consider two agricultural producers with identical production technologies. The senior producer, denoted s, receives water with higher priority than the junior producer, denoted j, such that in the event of a drought the senior producer receives the total volume of its water right, $\breve{w}_s$, and the junior producer receives the remainder, up until the total volume delineated by its water right, $\breve{w}_j$. We assume that the senior and junior producer's maximum allocations are identical i.e. $\breve{w}_s = \breve{w}_j$. In the event of total water supply realization W, water allocations to the senior and junior producers ($\overline{w_s}$ and $\overline{w_j}$, respectively), take the form of equation (1). For simplicity, we assume $W > \breve{w}_s$ always holds, such that the senior producer never receives less water than the total volume delineated by its water right, so that $\breve{w}_s = \overline{w_s}$ in all cases.

$$\begin{cases} \overline{w_s} = \breve{w}_s , \overline{w_j} = \breve{w}_j \ \ if \ \ W \geq \breve{w}_s + \breve{w}_j \\ \overline{w_s} = \breve{w}_s , \overline{w_j} = W - \breve{w}_s \ \ if \ \ W < \breve{w}_s + \breve{w}_j \ and \ W > \breve{w}_s \end{cases} \tag{1}$$

We model producer decision making over two time periods and two policy scenarios. In the first period, referred to hereafter as the "forecast period", each uses a season-ahead water forecast to make a labor investment decision that is irreversible in the second period.[1] For example, agricultural producers typically must make farm-hand labor and harvesting material purchases prior to the growing season (Calatrava and Garrido 2005). In the second period, referred to hereafter as the "realization period", the water supply is realized and production occurs. We model the realization period over two policy scenarios. In our benchmark "rigid" scenario, producers cannot transfer water and water allocations to each of the producers remain at the endowment level defined in equation (1). In our counterfactual "trade" scenario producers trade water in the event of a shortage, i.e. when the inequalities in the second line of equation (1) hold.

[1] We use labor as an example. Generically, these inputs can be thought of a composite of all variable inputs other than water.

In trade scenarios, trade occurs in the realization period prior to production. This sequential process is outlined in Figure 1.

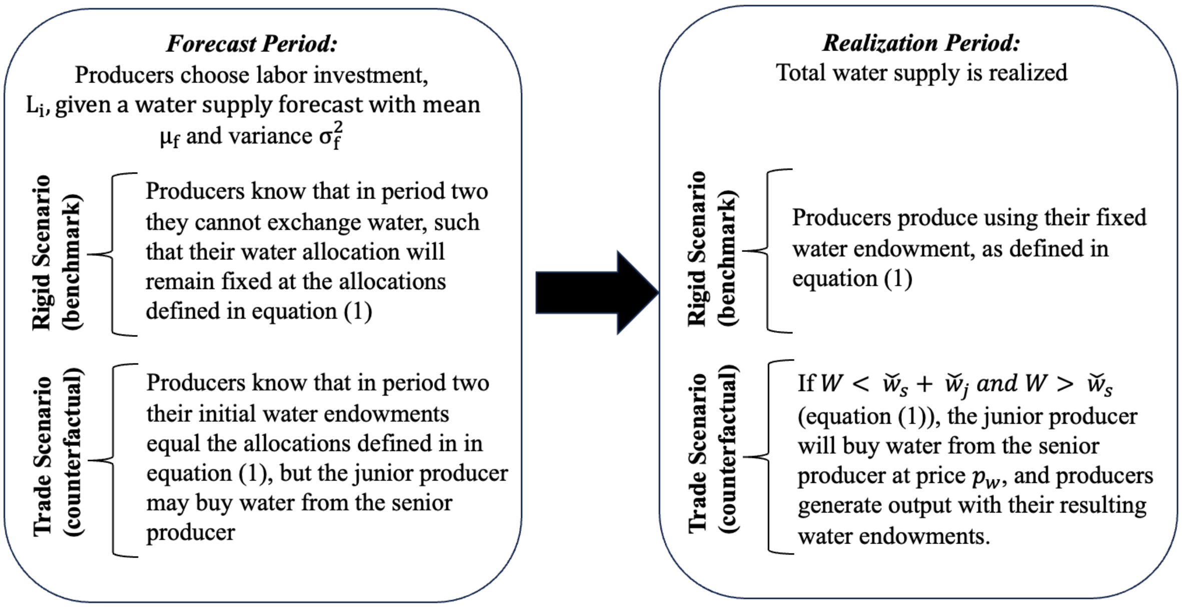

Figure 1. Summary of rigid and trade scenarios the in two- period model

2.2 Producer Behavior

In this section we discuss producer decision-making in the forecast and realization periods of our model. Central to our model is the assumption that producers seek to maximize profits in each period, in accordance with standard micro-economic theory (Weston and Fred 1950). We proceed by first describing the profit maximization problem faced by producers in the forecast and realization periods, for both the benchmark rigid and counterfactual trade scenarios. We then discuss how forecast bias and variance impact gains from trade when water is realized in the realization period. This logic helps motivate our results discussion in section 3.

In the forecast period, each producer i observes a probabilistic forecast for total basin water supply, W, which is distributed normally with mean $\mu_f$and variance σ_f^2. Producers take the forecast as the relevant information set when making an irreversible labor choice, L_i^*. In the realization period, total water supply W is realized and each producer receives a realized endowment $\bar{w}_i(W)$implied by the prior appropriation allocation rule in equation (1). In the rigid scenario, producers cannot trade and therefore use their endowment, $w_i = \bar{w}_i(W)$. In the trade scenario, producers can buy and sell water in a spot market at an endogenously determined price $p_w(W)$, choosing water use w_i subject to market clearing. In both scenarios, producers' labor inputs in the realization period are constrained to their forecast-period choice, L_i^*.

In both periods each producer's output, Y_i, is represented using a constant elasticity of substitution (CES) production function, which defines the relationship between producers' water and labor inputs and the quantity of output produced. The CES representation of production is widely used across the economics literature due to both its empirical reproducibility and convenient form (Miller, 2008). Producers take the output price and wage as exogenous, p_Y and p_ℓ, respectively. Realized output is given by:

$$Y_i = \left(\theta_w w_i^\rho + (1-\theta_w) L_i^\rho\right)^{1/\rho}. \qquad (2)$$

Here θ_w represents producers' identical benchmark expenditure share of water, and ρ is related to the elasticity of substitution between water and labor, σ, through $\rho \equiv \frac{\sigma-1}{\sigma}$ (equivalently, $\sigma = 1/(1-\rho)$). We set $\rho = -1$(so $\sigma = 0.5$) based on the 0.05–0.9 range of substitution elasticities between water and other inputs for agricultural producers reviewed in Graveline (2016). For tractability, we assume σ, and therefore ρ, are identical between producers.

2.2.1 Forecast-period problem (labor choice).

In the forecast period, each producer chooses labor to maximize expected profits, taking as given (i) the distribution of realized basin water W, and (ii) the subsequent realization-period allocation rule and (potential) trade described above. Let $w_i^k(W; L)$ denote producer i's realization-period water use as a function of realized $W$and the vector of labor choices L, under scenario $k \in \{R, T\}$, where R denotes the rigid scenario and Tdenotes the trade scenario. Under rigidity, $w_i^R(W; L) = \bar{w}_i(W)$. Under trade, $w_i^T(W; L)$ is determined in equilibrium by the spot market described below in section 2.2.2. Producer i's forecast-period problem under scenario $k$is:

$$\max_{L_i \geq 0} \ \mathbb{E}_f \left[p_Y \, Y_i \left(L_i, \, w_i^k(W; L)\right) - p_\ell L_i\right], \qquad (3)$$

where $\mathbb{E}_f[\cdot]$denotes expectations taken with respect to the forecast distribution of W. The first order condition characterizing optimal labor demand is:

$$p_\ell = p_Y \, \mathbb{E}_f \left[\mathrm{MPL}_i \left(L_i, \, w_i^k(W; L)\right)\right], \qquad (4)$$

where $\mathrm{MPL}_i$is the marginal product of labor in the realization-period production function. Differentiating (2) yields closed-form expressions for marginal products:

$$\mathrm{MPL}_i \equiv \frac{\partial Y_i}{\partial L_i} = (1-\theta_w) L_i^{\rho-1} Y_i^{\,1-\rho}, \qquad (5)$$

$$\mathrm{MPW}_i \equiv \frac{\partial Y_i}{\partial w_i} = \theta_w w_i^{\rho-1} Y_i^{\,1-\rho}. \qquad (6)$$

Equation (4) states that each producer chooses L_i^* such that the wage is equal to the expected marginal value of labor under the forecast distribution. Purchasing any more or less labor than implied by (4) would imply that the expected marginal value of labor is either smaller or larger than its price, respectively, and therefore that expected profits are not maximized.

Equation (4) also provides intuition for how forecast variance and forecast mean affect labor choices. Under CES production with $\rho < 1$, output is concave in inputs and marginal products exhibit diminishing returns. As a result, the expected marginal value of labor can be sensitive to forecast variance through Jensen-type effects. Intuitively, when water realizations are uncertain, low realizations of W relative to expectation generate larger reductions in realized labor productivity than gains from relatively high realizations. Producers therefore hedge in response to greater forecast uncertainty by choosing a smaller labor investment, holding the forecast mean μ_f constant (Just, 1975). Figure 2a illustrates this logic: because expected labor productivity is concave in realized water availability over the relevant range, increases in σ_f^2 increase downside risk in realized productivity and reduce the optimal labor choice L_i^*.

Another implication of (4) is that producers hire less labor in response to a lower forecast mean μ_f. Figure 2b illustrates producers' expected labor productivity based on two forecasts in which the expected value differs but the variance is constant. In the case of the relatively "low mean" forecast (red), producers invest less relative to the "high mean" forecast (blue) due to the lower expected water availability implied by μ_f. Thus, holding the realization mean μ_o constant, producers facing a negatively biased forecast (i.e., $\mu_f < \mu_o$) will hire less labor than when the forecast mean is positively biased. This effect is reinforced as μ_f becomes lower due to the curvature of the expected marginal product of labor with respect to expected water quantity (Figure 2b).

2.2.2 Realization-period problem (water use and trade)

In the realization period, basin water supply W is realized and each producer receives a realized endowment $\bar{w}_i(W)$ (equation (1)). In the rigid scenario, producers make no decision in the realization period and simply use their endowment: $w_i = \bar{w}_i(W)$. In the trade scenario, producers can buy and sell water in a spot market at price $p_w(W)$, choosing w_i to maximize realized profits subject to the labor constraint $L_i = L_i^*$:

$$\max_{w_i \geq 0} \; p_Y Y_i - p_w(W)(w_i - \bar{w}_i(W)) - p_\ell L_i^*, \qquad (7)$$

subject to $L_i = L_i^*$ and the market clearing condition $\sum_i w_i = \sum_i \bar{w}_i\,(W)$. The first order condition for an interior optimum is:

$$p_w(W) = p_Y \,\mathrm{MPW}_i. \qquad (8)$$

Equation (8) states that in the trade scenario, each producer buys or sells water until the spot price equals the marginal value product of water in production. In drought realizations for which the junior producer is shorted under the allocation rule (i.e., realizations in which $\bar{w}_j(W)$is low relative to desired use), the junior producer's marginal product of water is high. Because MPW_i is increasing in the degree of shortage and increasing in L_j^* through complementarity between inputs (for $\rho < 1$), water transferred to the junior producer becomes especially valuable in severe drought states. Taken together with the effects of μ_f and σ_f^2on the labor choice L_i^*, this mechanism forms the basis of our discussion in the next section regarding how forecast bias and variance shape ex post gains from trade.

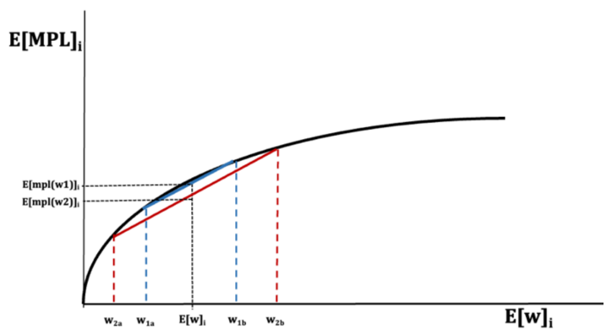

Figure 2a. Expected marginal product of labor faced by farmers when season-ahead forecast is relatively "low variance" (blue) and "high variance" (red)

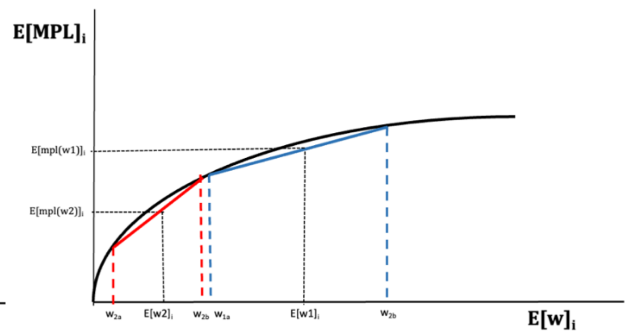

Figure 2b. Expected marginal product of labor faced by farmers when season-ahead forecast is relatively "low mean" (red) and "high mean" (blue)

2.3 Value of Trade

In this section we define how we calculate the value of trade in our model and explore how we can expect the value of trade to behave across the parameter space, given our discussion of producer behavior in the previous section. We define the value of trade as the difference between total realized profits across producers in the trade scenario and total realized profits in the rigid scenario, for a given realized basin water supply W:

$$\Delta\Pi(W) \equiv \sum_i \pi_i^T(W) - \sum_i \pi_i^R(W). \qquad (11)$$

Here $T$and Rdenote the trade and rigid scenarios, respectively. Realized profits in each scenario are:

$$\pi_i^T(W) = p_Y\, Y_i\left(L_i^{T*}, w_i^{T*}(W)\right) - p_w(W)\left(w_i^{T*}(W) - \bar{w}_i(W)\right) - p_\ell L_i^{T*} \qquad (12)$$

$$\pi_i^R(W) = p_Y\, Y_i\, \left(L_i^{R*}, \bar{w}_i(W)\right) - p_\ell L_i^{R*}, \qquad (13)$$

where L_i^{k*} is the forecast-period labor choice under scenario $k \in \{T, R\}$, $\bar{w}_i(W)$ is the realized endowment implied by the allocation rule in equation (1), and $w_i^{T*}(W)$ is the realization-period water use chosen in the spot market under trade.

A key implication of this definition is that water payments are transfers between producers and therefore do not affect aggregate profits. In particular, the trade scenario satisfies market clearing, $\sum_i w_i^{T*}(W) = \sum_i \bar{w}_i(W)$, implying:

$$\sum_i p_w(W)\left(w_i^{T*}(W) - \bar{w}_i(W)\right) = p_w(W)\left(\sum_i w_i^{T*}(W) - \sum_i \bar{w}_i(W)\right) = 0 \qquad (14)$$

Substituting (12) , (13) and (14) into (11) yields this expression for the value of trade:

$$\Delta\Pi(W) = p_Y[\textstyle\sum_i Y_i\ (L_i^{T*}, w_i^{T*}(W)) - \sum_i Y_i\ (L_i^{R*}, \bar{w}_i(W))] - p_\ell[\sum_i L_i^{T*} - \sum_i L_i^{R*}]. \qquad (15)$$

Thus, the value of trade is driven by two forces: (i) how much additional output is generated when water can be reallocated across producers in the realization period (holding labor fixed within each scenario), and (ii) how the anticipation of trade changes labor investment in the forecast period (i.e., differences between L_i^{T*} and L_i^{R*}).

Because producers choose labor in the forecast period under uncertainty and cannot adjust labor once W is realized, the value of trade in this context rests in how much producers can use traded water to offset the sub-optimality of their labor decision ex post. When realized water is lower than expected, producers' labor decisions tend to be "too high" relative to realized water availability, implying a high marginal product of water, especially for the junior producer when they are shorted under the allocation rule. Trade increases total profits by transferring water toward the producer with the highest marginal value product of water, until marginal values are equalized across producers, equation (8). When realized water is higher than expected, producers' labor decisions tend to be "too low" relative to realized water availability; trade then reallocates water to reduce the opportunity cost that arises from being unable to expand labor in the realization period, again by moving water toward its highest marginal value in production.

This logic suggests that gains from trade increase in the degree to which realized water differs from what producers planned for when choosing labor. In other words, $\Delta\Pi(W)$ will tend to be larger when deviations between realized and expected water volumes generate a larger mismatch between realized and expected input proportions. Moreover, the model predicts that the value of trade is higher for water realizations below expectation than for realizations above expectation. The reason is that when water becomes scarce, the junior producer's marginal product of water rises sharply (due to diminishing returns and complementarity between inputs for $\rho < 1$), and the prior appropriation allocation rule can create large wedges in marginal products across producers, relative to periods of water abundance. In drought states, trade therefore produces relatively large efficiency gains by resolving a larger wedge in the marginal product of water between producers.

Forecast bias and forecast variance affect $\Delta\Pi(W)$ through their impacts on the forecast-period labor choices $L_i^{T*}$and L_i^{R*}(Section 2.2). Forecast variance induces hedging behavior: when $\sigma_f^2$is higher, producers reduce labor investment to limit downside exposure to low-water realizations. This hedging effect reduces the extent of ex post mismatch between labor and realized water in drought states, reducing the scope for profitable trade when $W$is lower than expected. In contrast, when $W$is higher than expected, this same hedging behavior can exacerbate the extent to which labor is "too low," increasing the value of trade by making the reallocation of water more consequential in alleviating the opportunity cost of fixed labor. These mechanisms are illustrated in Figure 3, shows how the the relationship between forecast variance and gains from trade can differ depending on whether realized water is above or below expectation. In the analysis that follows, we quantify how economic gains from trade; ΔΠ(W), vary with forecast bias and forecast variance using hydrologic water forecast data drawn from basins across the Western U.S.

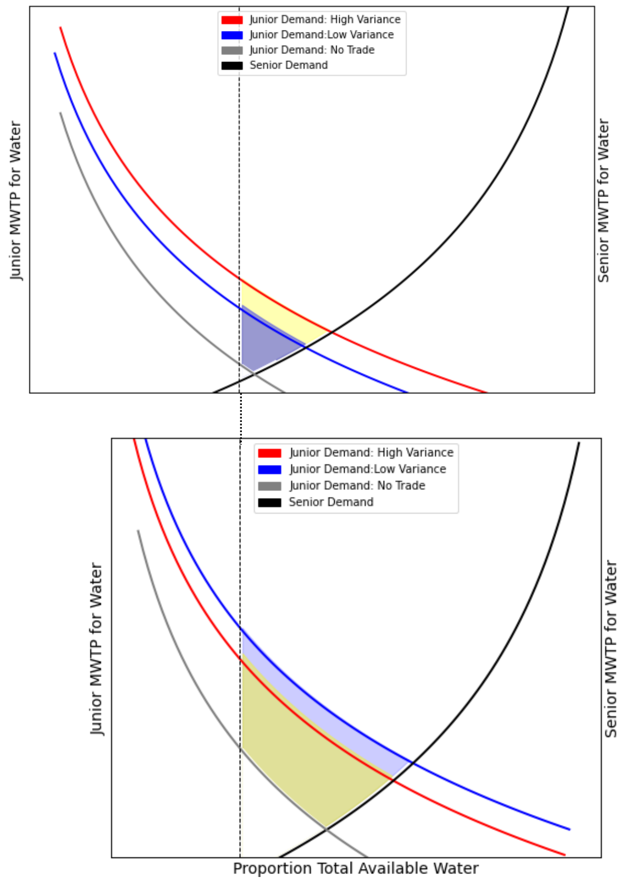

Figure 3. Gains from trade (yellow and blue shaded regions) when realized water volume in the realization period is higher than expected (top panel) vs lower than expected (bottom panel). The width of the x-axis in each panel represents the total volume of water supply realized in the realization period. The blue shaded region in the bottom panel represents the amount by which the junior producer under a low variance forecast (whose water demand is represented by the blue series) benefits more from trade than the junior producer under a high variance forecast (whose water demand is represented by the red series) when the total volume of water realized in the realization period is lower than expected. The yellow shaded region in the top panel represents a reversed relationship between gains from trade in the low and high variance forecast scenarios when total water realized in the realization period is higher than expected. The change in relative gains is due to shifts in the suboptimality of the junior user's labor decision once actual water volume is realized in the realization period for each forecast type.

3 Empirical analysis

In this section we describe the application of our water trade model to water supply forecasts in the Western US. In Section 3.1 we summarize our forecasting methodology and motivate our empirical analysis by examining the relationship between forecast bias and variance in relation to the proportion of precipitation falling in the basin from snow (SWE/P). In section 3.2 we discuss our process for linking the forecasts in our hydrological dataset to our producer model. This comprises the experimental setup we use to link our theoretical model outlined in section 2 to the differences in forecast parameters observed across low and high SWE/P basins . In section 3.3. we discuss our sensitivity analysis, in which we examine the relationship between economic outcomes forecasted by our model and the ability for producers to substitute between water and other inputs. Finally, in section 3.4 we describe our station level analysis, in which we examine gains from trade utilizing the full range of forecast variability in our dataset.

3.1 Forecast Data Summary Statistics

Our forecasting dataset is based on the methods developed by Modi et al. (2022). We employ a modified Natural Resources Conservation Service (NRCS) statistical procedure to generate water supply forecasts for 889 snow-observing locations. We trained a Principal Component Regression (PCR) using ERA5 land hourly (Munoz et al., 2021) that was aggregated to a daily timestep and extrapolated from the proximal grid cell at each location. ERA5 predictors of SWE and accumulated precipitation are transformed into standardized anomalies (i.e., subtraction of mean and division by standard deviation based on training years), and April-July runoff volume is semi-normalized via a square root transformation (Lehner et al., 2017; Garen 1992; Modi et al., 2022). The modification to the NRCS's PCR procedure is undertaken in the process of retaining the principal components. While the NRCS procedure uses a significance and sign test on regression coefficients to retain the number of principal components via an iterative process, a cross-validation approach is used here to retain the number of principal components. Specifically, a 10-fold cross-validation, i.e., a 'test' of model on ten different samples, calculates the model's mean squared error, with the addition of the principal component one at a time. The number of principal components corresponding to the best model skill score is retained (Modi et al., 2022). The forecasts using this modified PCR are generated across the 889 locations from WY1986-2020.

We calculate prediction bias in each station-year as the percent difference between the mean predicted streamflow and the observation mean. To construct a measure of forecast variance we use a jackknife sampling process to estimate the variance in predictions on each station-year. We categorize basins into five evenly spaced categories between 0-1 based on SWE/P value. After calculating median bias and variance for each station over the 34-year period we eliminate 15 stations in which forecast bias and variance fall outside of three standard deviations from the mean within each SWE/P category. Median absolute bias, normalized standard deviation and station count for each SWE/P category are described in Table 2. Median absolute bias is used to proxy the severity of bias in each bin, though we explore implications of bias sign by conducting our analysis in section 3.2 with both median absolute bias and negated absolute bias in the lowest and highest SWE/P categories. Figure 4 displays box plots of bias and normalized standard deviation across each SWE/P category, Figure 5 displays the distribution of these parameters across our forecast dataset for each SWE/P bucket.

Table 2. Forecast data summary statistics

SWE/P Category	Median Normalized Std. Dev. (% Observed Streamflow)	Median Percent Abs. Bias (%)	Station Count
0.0-0.2	6.8%	(+-) 14.8%	137
0.2-0.4	4.9%	(+-) 11.9%	135
0.4-0.6	4.7%	(+-) 9.0%	168
0.6-0.8	3.5%	(+-) 8.3%	287

0.8-1.0	2.0%	(+-) 2.7%	45

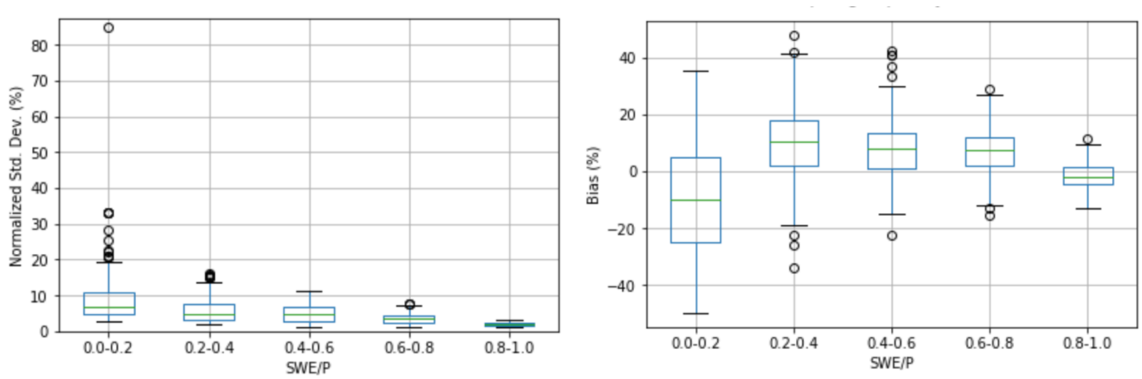

Figure 4. Boxplots of normalized forecast standard deviation (left) and bias (right) by SWE/P category The middle line in each plot represents the category median. Circular markers represent observations outside of each category's interquartile range.

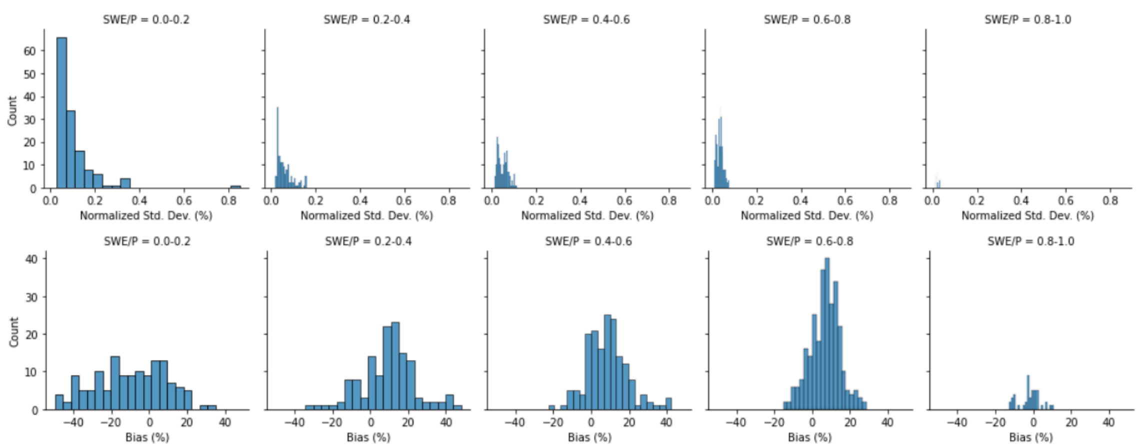

Figure 5. Histogram plots of Percent forecast bias and normalized standard deviation for each SWE/P category

Our summary statistics from Table 2 indicate that both median forecast bias and variance decrease as SWE/P increases. Figure 5 shows histogram plots of median normalized standard deviation and median percent bias for each station in each SWE/P bin. These plots show that the frequency of relatively extreme observations of each of these forecast parameters is higher overall in rain-dominant basins (basins with low SWE/P). Taken together with the theoretical results from section 2, this suggests that water trade is likely particularly important in basins with a lower proportion of streamflow from snowpack. As discussed in the introduction, this has important implications in the context of climate change as warming is projected to decrease SWE/P in basins that have historically been snow dominated, as the majority are in the Western US (Figure 6).

The theoretical model outlined in section 2 shows that gains to trade increase as forecast accuracy declines i.e. as forecast bias and variance increase. In this section we motivate our empirical analysis by showing that both forecast bias and variance increase with decreased SWE/P using data from SNOTEL stations in the Western United States. We proceed by using estimates of forecast parameter differences from these summary statistics to inform our water market value analysis, which is described in the following sections.

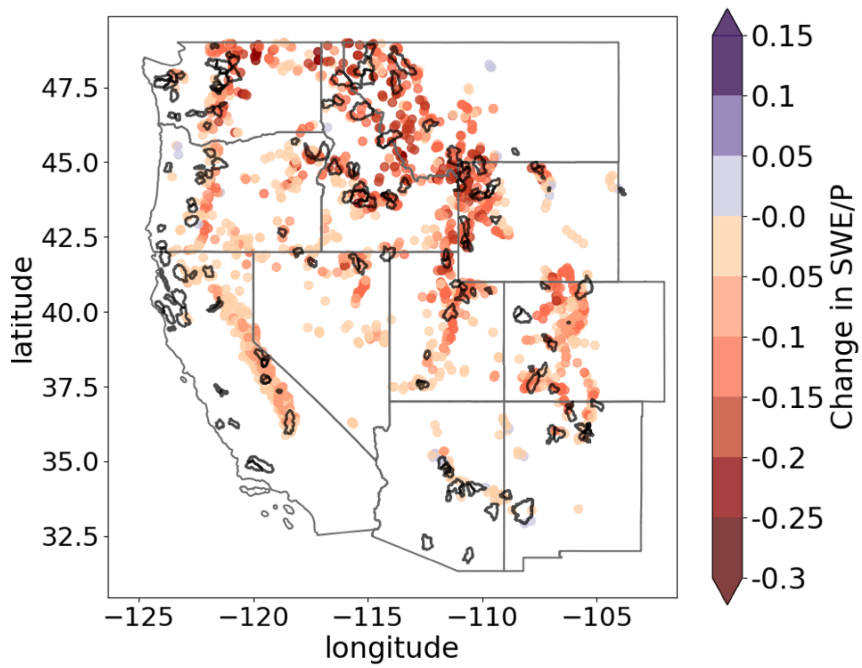

Figure 6 Reductions in future snowpack: Projected changes in the ratio of annual peak snow water equivalent (SWE) to accumulated precipitation (SWE/P) for mid-century (2030-2059) relative to historical (1981-2010) years in ssp5-8.5 scenario. Each point represents a snow-observing location, and the black polygon is a basin boundary representing flow at a USGS stream gage. Analysis shown for snow-dominated regions using the six climate models included in the ORNL CMIP6 dataset (Kao et al., 2022).

3.2 Experimental Setup

In order for our experimental design and preceding water market value analysis to reflect differences in key forecast parameters across observed Western US stations with substantial variation in SWE/P, we use summary statistics from the previous subsection from the lowest (0.0-0.2) and highest (0.8-1.0) to inform the forecasts faced by each of the senior and junior producers in our water trade model. We run two experiments. In the first experiment, hereafter referred to as the “variance experiment” we examine differences in gains from trade between a “low variance” and “high variance” scenarios. The “low variance” scenario reflects a forecast variance scaled to reflect median normalized standard deviation in the lowest SWE/P (0.0-0.2) (most rain-fed) category from our empirical forecast dataset (Table 2). The “high variance” scenario reflects a forecast variance scaled to reflect median normalized standard deviation in the highest SWE/P (0.8-1.0) (most snow-fed) category from Table 2. For tractability, both forecasts in this experiment exhibit zero-bias.

In the second experiment, hereafter referred to as the "variance and bias experiment" we allow both forecast variance and bias to vary and reflect median normalized standard deviation and both positive and negative percent absolute bias for the lowest (0.0-0.2) and highest (0.8-1.0) SWE/P category in Table 2. Thus, gains from trade are examined across four forecasts in the variance and bias experiment and two in the variance experiment. We examine water market value across bias sign in the variance and bias experiment due to the sensitivity of water trade to drought conditions, motivated by our theoretical model. In this experiment we investigate the sensitivity of trade gains to the amount by which the junior producer is shorted, in addition to producers' hedging behavior in response to downside risk.

For each forecast in each experiment, we run the rigid and trade scenarios in our model of water trade over 100 water realization draws from the true water realization distribution. This is done to evaluate how bias and variance in each forecast impact economic returns from trade *on average* across production years. To control for basin scale, we normalize differences in basin average streamflow volume by assigning the median average observed streamflow to the mean of the realization distribution across all SNOTEL stations. An outline of our experimental setup is presented in Table 3.

3.3 Sensitivity Analysis

The relative substitutability between water and other inputs to agricultural production can vary depending upon producers ability to adapt to drought conditions via modalities other than water trade. For example, producers may be able to substitute water intensive crops to crops that are relatively more intensive in other inputs (Lund et al. 2018) or into irrigation efficiency investments (Medellín-Azuara et. al 2012). We examine the sensitivity of economic gains to trade in our experimental setup by repeating our gains from trade analysis described in section 3.2 when the substitution elasticity between water and labor is "low" and again when the substitution elasticity between water and labor is "high". We select low (0.05) and high (0.9) values for the value of σ based on the lower and upper bounds for this parameter discussed in Graveline (2016).

3.4 Station Level Analysis

In the experiments described above we evaluate gains from trade over the median forecast bias and variance within in each SWE/P category. This process allows us to link the experiment results to our theoretical discussion in section 2 in a tractable manner. However, as indicated by the boxplots and histograms in Figures 4 and 5, SWE/P groups differ considerably by range and distribution of forecast parameters. To reconcile differences in the frequency and range of forecast parameters between categories, we conduct a complementary analysis in which we run our producer model for each station in the dataset.

In this analysis, we run our producer model once for the trade and rigid scenarios for each station. For each station we input a single observed water volume, forecast variance, and bias, which represent the median normalized values of these parameters for each station over our 1986-2020 sample period. Just as in the experimental setup, we exclude stations that demonstrate a variance and bias over three standard deviations away from the mean of each SWE/P category.

Table 3. Outline of experimental setup. In order for our experiments to reflect the median difference in key forecast parameters across observed Western US stations with varying proportion precipitation falling as snow, differences in forecast variance in both experiments, and differences in forecast bias in the Bias and Variance Experiment reflect median values from the lowest (0.0-0.2) and highest (0.8-1.0) SWE/P bins in Table 2, rounded to the nearest percent. $\overline{\mu}$ *reflects a benchmark parameter set to the median water volume observed in the forecast dataset.*

<table>
<tr><td colspan="2"><u>Variance Experiment</u></td></tr>
<tr><td colspan="2">Forecast 1: High Variance, Low SWE/P (0.0-0.2)
$\sigma_f = 0.07 \times \mu_o$
$\mu_f = \mu_o = \overline{\mu}$</td></tr>
<tr><td colspan="2">Forecast 2: Low Variance, High SWE/P (0.8-1.0)
$\sigma_f = 0.02 \times \mu_o$
$\mu_f = \mu_o = \overline{\mu}$</td></tr>
<tr><td colspan="2"><u>Bias and Variance Experiment</u></td></tr>
<tr><td>Positive Bias</td><td>Negative Bias</td></tr>
<tr><td>Forecast 1: Low SWE/P (0.0-0.2)
$\sigma_f = 0.07 \times \mu_o$
$\mu_f = \mu_o \times 0.15 + \mu_o$</td><td>Forecast 1: Low SWE/P (0.0-0.2)
$\sigma_f = 0.07 \times \mu_o$
$\mu_f = \mu_o \times -0.15 + \mu_o$</td></tr>
<tr><td>Forecast 2: High SWE/P (0.8-1.0)
$\sigma_f = 0.02 \times \mu_o$
$\mu_f = \mu_o \times 0.03 + \mu_o$</td><td>Forecast 2: High SWE/P (0.8-1.0)
$\sigma_f = 0.02 \times \mu_o$
$\mu_f = \mu_o \times -0.03 + \mu_o$</td></tr>
</table>

4 Results

4.1 Variance experiment

Figure 7.i displays profit gains from trade as a percent of profit in the benchmark rigid case for Forecasts 1 and 2 in the Variance experiment. Each marker represents gains from trade for an observed water volume drawn from the true water supply distribution (n=100). Aligned with our theoretical discussion in section 2, gains are increasing as water realizations move away from the expected total water volume and are zero at the point at which realized and expected marginal product of labor are equated. Because the high variance (low SWE/P) basin (Forecast 1) yields water realizations that are farther away from expectation in both directions than realizations in the low variance basin (Forecast 2), overall gains to trade are higher in the former over our sample realizations.

As discussed in section 2, the value of trade increases as producers' labor decisions diverge relative to labor decisions under certainty. Illustrating this result, Figure 7.ii plots the percent decrease in profit losses from uncertainty that attributable to trade. This metric reflects the difference between profit losses relative to certainty in the rigid and trade cases for each forecast, which is increasing in the magnitude of the labor discrepancy (Figure 7.iii). This is because trade allows producers to ameliorate losses from hiring too much or too little labor in the forecast period. As the magnitude of producers' labor discrepancy increases, the margin by which trade can help offset mistakes due to forecast uncertainty in the forecast period increases.

Over the range of water realizations in which water volumes are commensurate between Forecasts 1 and 2 in Figure 7.i (96-107% of expected water volume), gains are relatively higher for the low variance (high SWE/P) case as realized water falls below expectation, while gains are relatively lower than in the high variance case as water realized increases past expected water. This is because producers facing the low variance forecast hedge relatively less than those facing the high variance forecast, thus at lower water realizations the magnitude of the labor discrepancy from decisions made in the forecast period is larger than it is in the high variance forecast, while the reverse is true as water realizations increase past expectation.

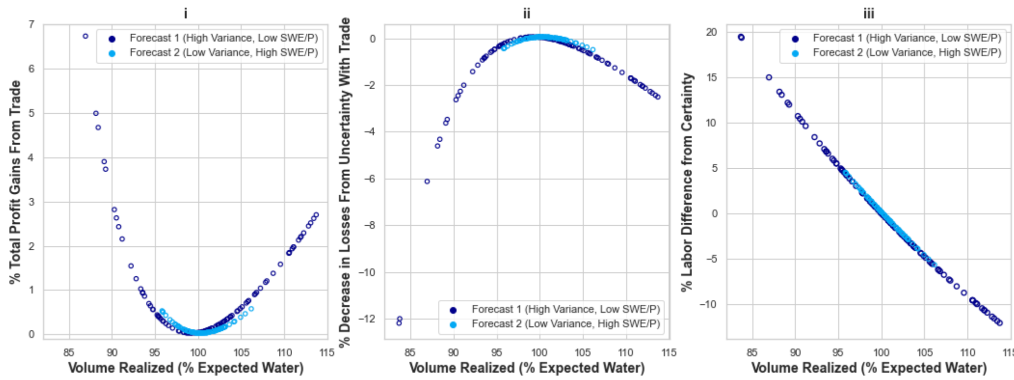

Figure 7. Panels i-iii display results from the Variance experiments in which we vary forecast variance according to the forecast summary data in section 3.1 (Table 2). Each panel displays water volume realized as percent of expected water on the x-axis, and moving from left to right the y-axis reflects percent profit gains from trade, percent decrease in losses from uncertainty due to utilizing trade relative to the rigid scenarios, and percent difference between labor hired under uncertainty and certainty across each water realization.

4.2 Bias and Variance Experiment

Figures 8a and 8b display results from our Bias and Variance experiment, in which we vary median forecast bias magnitude for high and low SWE/P categories from Table 2. Figures 8a.i and 8b.i display gains from trade across water realizations for the positive and negative bias values from Table 2, respectively. Note that the scale of the y-axis in both charts is substantially larger than the scale of the y-axis in the trade gains chart for the variance experiment (Figure 7.i). This is because holding forecast variance constant, forecast bias offsets water realizations further from expectation than the Variance experiment alone, reflected by the magnitude of water realizations recorded on the x-axes in Figures 8a and 8b relative to Figure 7. Specifically, in Figure 8b.i , the lowest water realization is 72% of expected value, which is when the largest gains from trade occur. In the Variance experiment (Figure 7) water realizations range from 80-115% of expected value. As a result, producers can offset much higher potential profit losses due to uncertainty by trading water in the second period of the model. This is because labor discrepancies relative to certainty increase with distance from expected water (Figures 8a.ii and 8bii).

Comparing low SWE/P basins to high SWE/P basins in each experiment shows that gains to trade are significantly higher in the low SWE/P. This is not surprising since precipitation dominated basins are more difficult to forecast than snow dominated basins. The resulting higher positive and negative bias cause water realizations to occur much farther away from expected volumes than in the high SWE/P basin. This means that producers' expected marginal product of labor is much more likely to differ substantially from that which is realized (Figure 2b), increasing the role of trade in correcting profit losses relative to certainty (Figures 8a.ii and 8bii).

Again, hedging causes slight differences in returns over comparable water volumes (Figures 8a.i and 8b.i). The high SWE/P forecast (Forecast 2) yields higher (lower) returns to trade than low SWE/P (Forecast 1) in this range (92-102% of expected water volume) when forecast bias is positive (negative). This is because for a given volume of water below (above) expectation, producers in a high (low) SWE/P basin hedged less (more) in the forecast period, as a result, their labor discrepancy is relatively higher (lower) when low (high) water volumes are realized (Figures 8a.iii and 8b.iii).

As expected, positive bias yields larger returns to trade than negative bias. This is because as the likelihood of water realizations that are lower than expected increase, the junior producer's marginal product of water in the rigid case increases. As a result, a transfer of water to the junior producer offsets profit losses from uncertainty more than a transfer when water is less scarce. This is demonstrated by the larger magnitude of profit losses relative to certainty for the rigid case relative to the trade in rigid case in Figure 8a.ii (positive bias) versus in 8b.ii (negative bias), across low and high SWE/P values. In general, when profit losses due to uncertainty are smaller, gains from trade are relatively smaller. Ameliorating losses due to uncertainty is the main avenue through which trade produces value in our model.

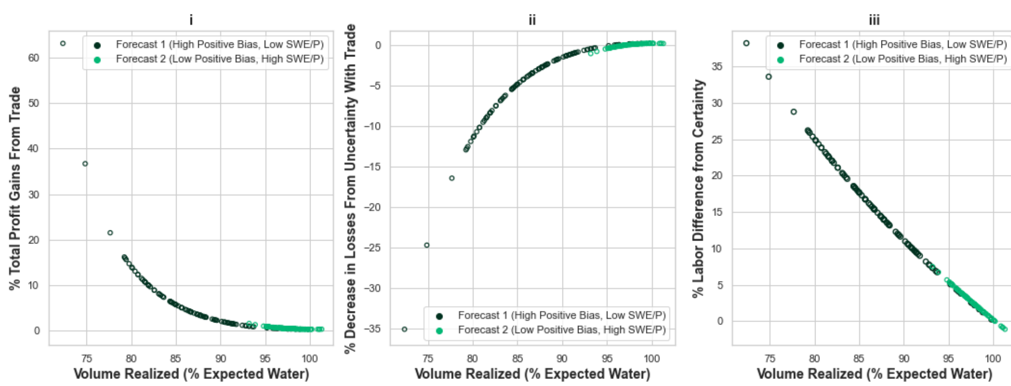

Figure 8a. Panels i-iii display results from the positive Variance and Bias experiment in which we vary forecast parameters according to the forecast summary data in section 3.1 (Table 2), limited to positive forecast bias. Each panel displays water volume realized as percent of expected water on the x-axis, and moving from left to right each y-axis reflects percent profit gains from trade, percent decrease in losses from uncertainty due to utilizing trade (relative to the rigid scenarios, and percent difference between labor hired under uncertainty and certainty across each water realization.

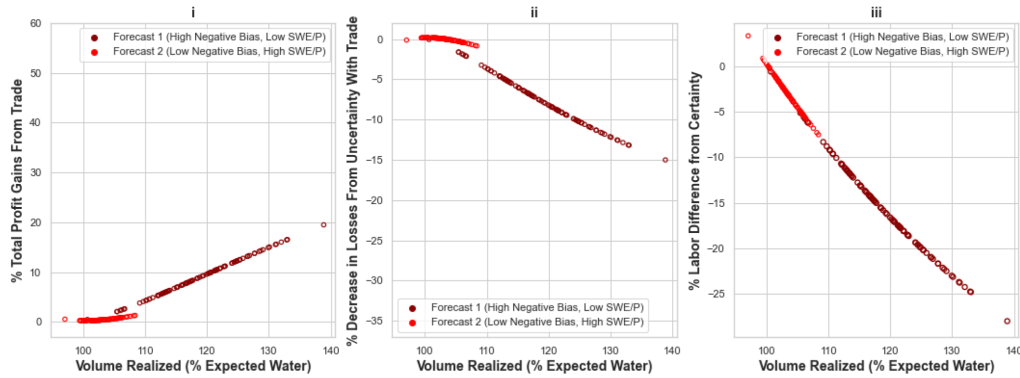

Figure 8b. Panels i-iii display results from the negative Variance and Bias experiment in which we vary forecast parameters according to the forecast summary data in section 3.1 (Table 2), limited to negative forecast bias. Each panel displays water volume realized as percent of expected water on the x-axis, and moving from left to right each y-axis reflects percent profit gains from trade, percent decrease in losses from uncertainty due to utilizing trade (relative to the rigid scenarios, and percent difference between labor hired under uncertainty and certainty across each water realization.

4.3 Sensitivity Analysis

Figures 9-12 comprise results from our sensitivity analysis. Each chart represents the percent difference in gains from trade between low (0.05) and high (0.9) values specified for the junior producer's ability to substitute between water and labor (σ_j), and our baseline σ_j specification of 0.5, which is identical between junior and senior producers. Values for our low and high values were chosen based on the lower and upper bounds of the range of substitution parameter values discussed in Graveline (2016). Figure 9 shows these percent differences across water realizations from our "High Variance, Low SWE/P" run of the variance experiment. We omit comparisons for the "Low Variance, High SWE/P" run of the variance experiment for this and all other experiments in the sensitivity analysis. This is because variation in differences in gains from trade between our specifications for the junior user's substitution elasticity parameter are due primarily to differences across the range of realized water volume in each experiment, as opposed to differences in hedging between the low and high variance forecasts under comparable water realizations. Because we observe the highest range in realized water volume for the high variance forecasts, we make these the focus of our results exposition for this analysis.

In Figure 9, gains from trade are markedly higher than in the baseline specification for σ_j when the junior user's substitution parameter is increased to 0.9 from 0.5. The reasoning for this increase is as follows; when water and labor become relatively more substitutable, the junior user hires more labor in the forecast period. Intuitively, this is because an increased ability to utilize water and labor interchangeably means that labor can better serve to increase output in the realization period, holding expected water volume constant. Due to the convexity of the production technology, a higher labor decision from the forecast period increases the marginal product of water in the realization period, shifting junior water demand upward. The upward shift in the junior user's water demand curve is especially high as the junior users' labor decision in the forecast period is more sub-optimal, this is because as water becomes more substitutable with labor, the junior user may now better "make-up" for a suboptimal labor choice in the forecast period by way of water transfers.

The mechanics of these effects in our model are illustrated in Figure 10. Focusing on the comparison between "High σ_j" and "Moderate σ_j", with "Moderate σ_j", representing our baseline case, the left panel of Figure 10 depicts a larger labor decision made by the junior user in the forecast period in the "High σ_j" case. Accordingly, the marginal product of water for the junior user is much higher in the realization period (middle panel). This is due both to a higher labor investment, as well as the increased capability for the junior user to offset a poor labor decision in the forecast period with the water market. The right panel of Figure 10 further supports this reasoning; the junior user buys the entirety of the senior user's endowment across all water realizations due to their large upward shift in marginal product of water relative to the senior user.

Figure 10 further illustrates that when the junior user's substitution parameter becomes lower than that of the senior user, gains from trade are *also* higher than in our baseline case when each user faces the same substitution parameter. As can be seen from the negative range of the gray series in the right panel of Figure 10, this is because the downward shift in the junior user's water demand when σ_j is *decreased* causes the direction of trade to reverse. The junior user now sells water to the senior user, who now faces a substitution parameter larger than that of the junior user. Since the junior user's demand for water is now shifted downward relative to the senior user, gains are larger than in the case where both user's demand curves are identical. The relative difference is much lower than in the "High σ_j" specification, because of the higher initial water allocation to the senior user across all water realizations. Furthermore, gains over the baseline case now *increase* slightly with water availability due to higher volumes of water transfers to the senior user as the junior user receives more water. This is illustrated by the increasing magnitude of transfers from the junior to senior user as water becomes more abundant (right panel of Figure 10).

Figures 11 and 12 represent precent differences in gains to trade between the high and low specifications of σ_j and our baseline for our representative negative and positive bias forecasts, respectively. In Figure 11, the scale of differences for the "High σ_j" over comparable water realizations are decreased when water is relatively scarce relative to when no forecast bias is present (Figure 10). This is because in negative bias (i.e. when the mean of the forecast is lower than the realization mean), users hire less labor than when no bias is present, making the scale of the labor discrepancy smaller when water is more scarce. By parallel reasoning, in the case of positive forecast bias (Figure 12), the premium in gains from trade observed for the "High σ_j" run are much larger when water is relatively scarce. This is because in positive bias the magnitude of the labor discrepancy is much larger when water is scarce because the forecast mean is higher than when no bias is present (Figure 9).

Differences in gains for the "low σ_j" over comparable water realizations are changed negligibly when negative bias is introduced, this is due to relatively small shifts in the degree of the labor discrepancy for the senior user between forecasts, who can guarantee a much larger portion of its water supply than the junior user can when facing uncertainty.

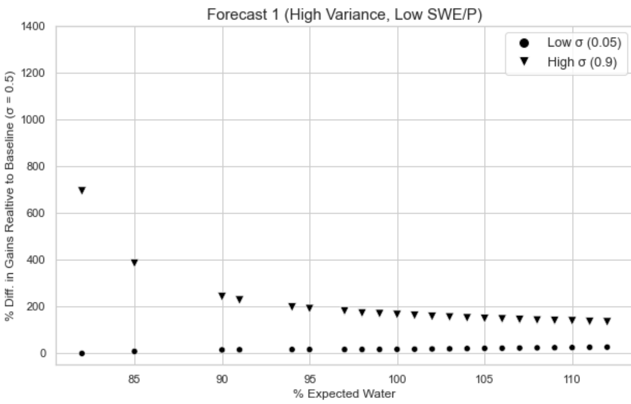

Figure 9. The percent difference in gains to trade between our "low"(0.05) and "high" (0.9) specification for the junior user's input substitution parameter (σ_j), across water realizations in our "High Variance, low SWE/P" forecast.

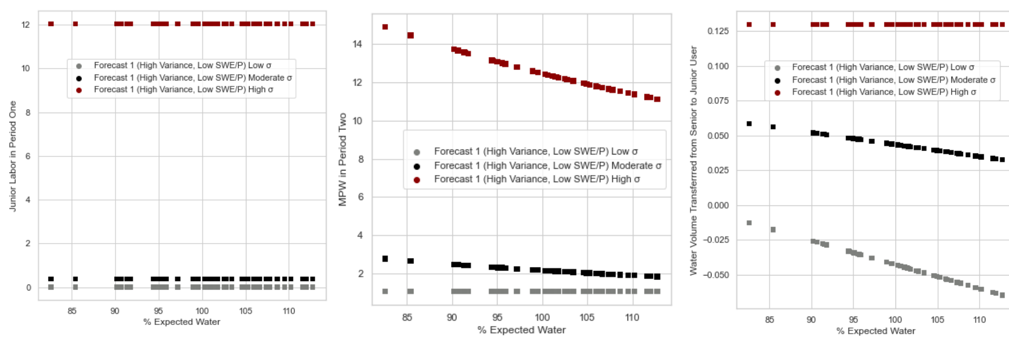

Figure 10. Junior labor investment (left) marginal product of water (middle), and water volume transferred from senior to junior user for "low", "moderate" and "high" values of σ_j Labor investment in the forecast period, marginal product in the realization period and water transfers to the junior user increase as water becomes more substitutable with the labor input.

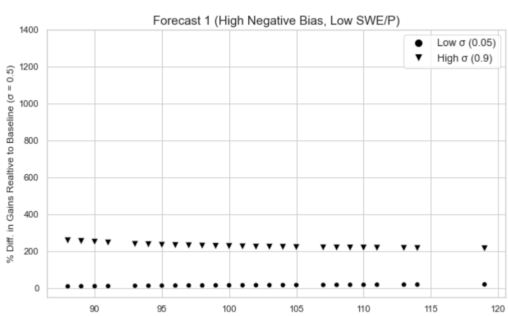

Figure 11. The percent difference in gains to trade between our "low"(0.05) and "high" (0.9) specification for the junior user's input substitution parameter (σ_j), in our "High Negative Bias, Low SWE/P" forecast

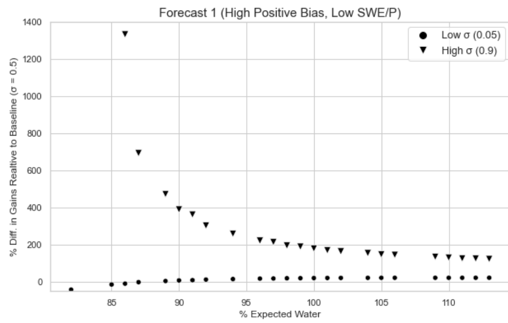

Figure 12. The percent difference in gains to trade between our "low"(0.05) and "high" (0.9) specification for the junior user's input substitution parameter (σ_j), in our "High Positive Bias, Low SWE/P" forecast.

4.4 Station Level Analysis

Results from our station-level analysis are displayed in Figures 13-15. Figure 13 represents boxplots of gains to trade in each SWE/P category across all SNOTEL stations in forecast sample (n=772. Figure 14 shows frequency weighted average returns to trade by SWE/P level. Each boxplot is comprised of profit gains from trade across individual stations within a SWE/P category. For each station we run our producer model once over the trade and rigid scenarios, where the forecast faced by producers in the forecast period represents the station-level median forecast variance over the (1986-2020) sample period. Just as in the previous experiments, we normalize median observed values to control for scale differences between basins.

The least snow fed category (SWE/P = 0.0-0.2) benefits from trade the most. On average, returns from trade are 6.95% higher for the lowest SWE/P category (Figure 13) relative to the highest SWE/P category. The boxplots and histograms in Figure 13 and Figure 14 show that relatively extreme returns to gains are much more common for lower SWE/P basins, meaning that the potential for very high returns is much higher as basins receive less of their streamflow from snowpack.

Interestingly, the premium in gains from trade in the lowest SWE/P category over the others is dampened by the sign of the median bias in each category. Median bias in the lowest SWE/P category is negative while the median in categories spanning SWE/P 0.4-0.8 are positive. Despite absolute bias decreasing in SWE/P category (Table 2), our results from the Bias and Variance experiments (Figures 8a.i and 8b.i,) show that negative bias results in relatively lower returns to trade than positive bias, due to the higher junior marginal product of water in the latter. This effect slightly offsets the impact of relatively more extreme variance and absolute bias levels observed in the lowest SWE/P category relative to the rest.

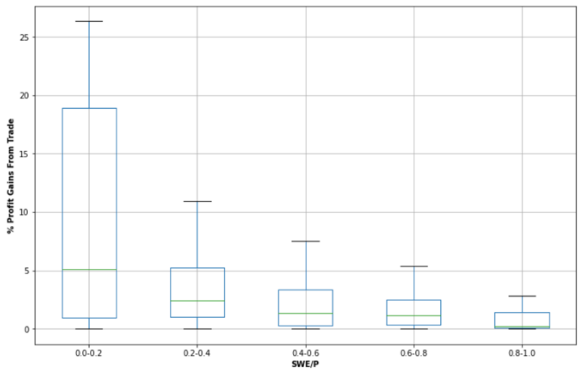

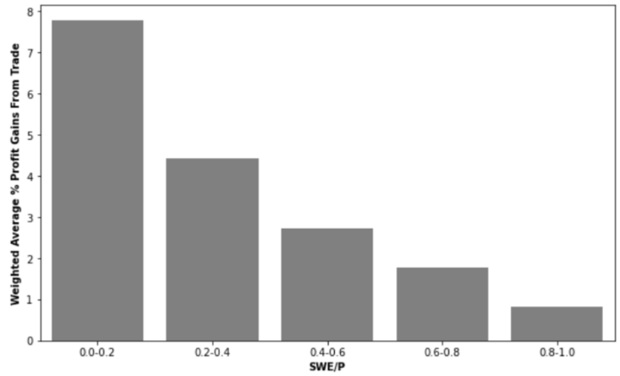

Figure 13. Boxplot showing gains to trade across all SNOTEL stations in forecast sample (n=772). For each station median forecast variance and bias was used to inform the producers' forecast in period 1 of the trade model.	*Figure 14. Bar chart representing frequency weighted percent gains from trade in each SWE/P category.*

Figure 15. Histogram plots illustrating the distribution of percent profit gains to trade from our station-level analysis. Each histogram bin represents the count of stations in which a given bucket of percent profit gains from trade were observed

5 Conclusion

Our findings contribute to the literature on water trade in several important ways. We diverge from much of the previous literature concerning water markets to establish that the importance of trade depends both upon the severity of drought as well as the degree to which season-ahead water forecasts are uncertain. We find that when producers must make irreversible investments before the realization of water availability, water trade helps offset the efficiency losses from sub-optimal investment positions made at seasonal lead times. These offsetting effects, i.e. the value of water trade relative to rigid water allocation, increase with forecast bias and variance. Thus, snow-fed basins benefit less from water trade than rain-fed basins. We develop the theoretical underpinnings for this result and show how decreasing snowpack can alter the magnitude of returns from trade due to differences in forecast precision derived from a hydrological forecasting dataset comprised of stations in the Western US. Additionally, because the value of trade under uncertainty is driven partly by user's ability to "make-up" for a suboptimal pre-season investment due to uncertainty, gains from trade increase with the drought-affected user's ability to substitute between inputs. This further diverges from previous research which has not incorporated considerations for uncertainty in water markets, under such assumptions, theory would support that water trade value decreases with the ability for producers to substitute into inputs other than water.

As climate change is expected to reduce the fraction of precipitation falling as snow, these results suggest that water trade will become a more valuable mechanism for allocating water supplies; as warming shifts basin's streamflow from snow fed to rain fed, season ahead supply forecasts will likely become less accurate. In turn, water reliant economies will face increased supply uncertainty and the water trade as a policy tool may become more critical. Because of climate change's role in reducing total water supply, these findings complement the literature which studies water market value in the context of increased drought severity. In particular (Arellano-Gonzalez et al. 2021) estimates that introducing more flexible markets to agricultural markets in water scarce regions of California primarily governed by prior appropriation rights would result in economic gains between $348 and $362 million using drought projections for the period 2070-2099. Taken together with our findings, these results further stress the economic importance of well-functioning water markets in the Western US, perhaps even in areas that have not historically faced water scarcity or variability issues. Future research should incorporate the effects of supply uncertainty due to SWE/P in large scale analyses of water markets similar to those carried out by Arellano-Gonzalez et al. (2021) and Hagerty (2019).

Our model does not explicitly account for adaptive drought mechanisms other than trade (though in regions that already face a large degree of water scarcity, the margin for further adaptation may be limited). For example, to reduce losses from capital investments associated with perennial crops, farmers may switch to annual plantings (S), plant drought resistant or higher value crops (Howitt et al. 2010), or adopt more efficient irrigation technology (Schuck et al. 2005). Instead, we capture these adaptive margins indirectly via our sensitivity analysis, in which we vary the degree to which the drought impacted junior user can substitute between water and labor in production. Importantly, we find that when substitution capability is increased, gains from trade are increased further, due to the uncertainty in investment that occurs in the pre-growth season. Intuitively, increased flexibility between inputs in productions helps producers better use water markets to offset suboptimal investment decisions, decreasing losses from uncertainty.

While the sensitivity analysis we include focuses on substitution opportunities producers may have between water and other inputs, in some cases agricultural producers may also adapt to uncertainty in surface water availability by substituting into groundwater pumping (Ward 2014), which is well known to contribute to negative externalities faced by other aquifer users (Lawell, 2016, Pfeiffer and Lin 2012, Brozovic et al. 2010). An interesting avenue for further research would be to investigate the impact of water trade as a means of reducing the cost of these inefficiencies. Water storage volume and timing presents a natural alternative to ameliorating effects of water uncertainty (Josué Medellín-Azuara et al 2008, Smith and Edwards 2021), though this method of adaption does not conserve water, it merely allows for previous storage to become available in periods of drought, which helps smooth use over time. Increased storage capacity without the capability to re-allocate water from low to high value production is thus likely not as effective as water trade.

Importantly, while we have modelled transfers between users as costless, we have ignored transaction and transport costs associated with such trades, which can be prohibitively costly, especially for small volumes (Womble and Hanneman 2020a). However, recent research has also shed light on legal changes that show potential for reducing transaction costs significantly Womble and Hanneman 2020b). For example, in Colorado, laws limiting trade volume to historical use result in substantial transaction costs due to the cost of surveys required to establish such use. Womble and Hanneman (2020b) find that adjusting this requirement lowers transaction costs associated with trades significantly with limited effects on third parties An interesting topic for further research could be to develop policy recommendations based on similar criteria in other states, or between states.

References

Arellano-Gonzalez, Jesus, Amir AghaKouchak, Morgan C Levy, Yue Qin, Jennifer Burney, Steven J Davis, and Frances C Moore. 2021. "The Adaptive Benefits of Agricultural Water Markets in California." Environmental Research Letters 16 (4). IOP Publishing: 44036. doi:10.1088/1748-9326/abde5b.

Arellano-Gonzalez, Jesus, and Frances C. Moore. 2020. "Intertemporal Arbitrage of Water and Long-Term Agricultural Investments: Drought, Groundwater Banking, and Perennial Cropping Decisions in California." American Journal of Agricultural Economics 102 (5). Wiley: 1368–82. doi:10.1111/ajae.12123.

Berghuijs, W. R., R. A. Woods, and M. Hrachowitz. 2014. "A Precipitation Shift from Snow towards Rain Leads to a Decrease in Streamflow." Nature Climate Change 4 (7). Springer Science and Business Media LLC: 583–86. doi:10.1038/nclimate2246.

Brewer, Jedidiah, Robert Glennon, Alan Ker, and Gary Libecap. 2007. "2006 Presidential Address Water Markets in the West: Prices, Trading, and Contractual Forms." Economic Inquiry 46 (2). Wiley: 91–112. doi:10.1111/j.1465-7295.2007.00072.x.

Brozović, Nicholas, David L. Sunding, and David Zilberman. 2010. "On the Spatial Nature of the Groundwater Pumping Externality." Resource and Energy Economics 32 (2). Elsevier BV: 154–64. doi:10.1016/j.reseneeco.2009.11.010.

Burness, H, and James Quirk. 1979. "Economic Aspects of Appropriative Water Rights." Journal of Environmental Economics and Management 7 (4). Elsevier BV: 372–88. doi:10.1016/0095-0696(80)90028-5.

Cai, Ximing, Claudia Ringler, and Jiing-Yun You. 2008. "Substitution between Water and Other Agricultural Inputs: Implications for Water Conservation in a River Basin Context." Ecological Economics 66 (1). Elsevier BV: 38–50. doi:10.1016/j.ecolecon.2008.02.010.

Calatrava, Javier, and Alberto Garrido. 2005. "Spot Water Markets and Risk in Water Supply." Agricultural Economics 33 (2). Wiley: 131–43. doi:10.1111/j.1574-0864.2005.00402.x.

Dettinger, Michael D., and Daniel R. Cayan. 1995. "Large-Scale Atmospheric Forcing of Recent Trends toward Early Snowmelt Runoff in California." Journal of Climate 8 (3). American Meteorological Society: 606–23. doi:10.1175/1520-0442(1995)008<0606:lsafor>2.0.co;2.

Dixit, Avinash. 2009. "Governance Institutions and Economic Activity." American Economic Review 99 (1). American Economic Association: 5–24. doi:10.1257/aer.99.1.5.

Garen, David. 1992. "Improved Techniques in Regression-Based Streamflow Volume Forecasting." Journal of Water Resources Planning and Management 118 (6). American Society of Civil Engineers (ASCE): 654–70. doi:10.1061/(asce)0733-9496(1992)118:6(654).

Goemans, Christopher, and James Pritchett. 2014. "Western Water Markets: Effectiveness and Efficiency." In Water Markets for the 21st Century, 305–30. Springer Netherlands. doi:10.1007/978-94-017-9081-9_16.

Gohar, Abdelaziz A., and Frank A. Ward. 2010. "Gains from Expanded Irrigation Water Trading in Egypt: An Integrated Basin Approach." Ecological Economics 69 (12). Elsevier BV: 2535–48. doi:10.1016/j.ecolecon.2010.07.030.

Grafton, R. Quentin, Gary Libecap, Samuel McGlennon, Clay Landry, and Bob O'Brien. 2011. "An Integrated Assessment of Water Markets: A Cross-Country Comparison." Review of Environmental Economics and Policy 5 (2). University of Chicago Press: 219–39. doi:10.1093/reep/rer002.

Graveline, N. 2016. "Economic Calibrated Models for Water Allocation in Agricultural Production: A Review." Environmental Modelling &Amp; Software 81 (July). Elsevier BV: 12–25. doi:10.1016/j.envsoft.2016.03.004.

Greif, Avner, Paul Milgrom, and Barry R. Weingast. 1994. "Coordination, Commitment, and Enforcement: The Case of the Merchant Guild." Journal of Political Economy 102 (4). University of Chicago Press: 745–76. doi:10.1086/261953.

Hagerty, N. 2019. In Liquid Constrained in California: Estimating the Potential Gains from Water Markets Work. Pap., Univ. Calif. Berkeley.

Howitt, R, D Macewan, J Medellín-Azuara, and J Lund. 2010.

Johnson, Ronald N., Micha Gisser, and Michael Werner. 1981. "The Definition of a Surface Water Right and Transferability." The Journal of Law and Economics 24 (2). University of Chicago Press: 273–88. doi:10.1086/466984.

Just, Richard E. 1975. "Risk Aversion under Profit Maximization." American Journal of Agricultural Economics 57 (2). Wiley: 347–52. doi:10.2307/1238513.

Lehner, Flavio, Andrew W. Wood, Dagmar Llewellyn, Douglas B. Blatchford, Angus G. Goodbody, and Florian Pappenberger. 2017. "Mitigating the Impacts of Climate Nonstationarity on Seasonal Streamflow Predictability in the U.S. Southwest." Geophysical Research Letters 44 (24). American Geophysical Union (AGU). doi:10.1002/2017gl076043.

Leonard, B., and G. D. Libecap. 2016. "Economic Analysis of Property Rights: First Possession of Water in the American West." Working Paper Series, National Bureau of Economic Research.

Li, Hongxia, and Yongqiang Zhang. 2017. "Regionalising Rainfall-Runoff Modelling for Predicting Daily Runoff: Comparing Gridded Spatial Proximity and Gridded Integrated Similarity Approaches against Their Lumped Counterparts." Journal of Hydrology 550 (July). Elsevier BV: 279–93. doi:10.1016/j.jhydrol.2017.05.015.

Li, Man, Wenchao Xu, and Mark W. Rosegrant. 2017. "Irrigation, Risk Aversion, and Water Right Priority under Water Supply Uncertainty." Water Resources Research 53 (9). American Geophysical Union (AGU): 7885–7903. doi:10.1002/2016wr019779.

Lin Lawell, C.-Y. Cynthia. 2016. “The Management of Groundwater: Irrigation Efficiency, Policy, Institutions, and Externalities.” Annual Review of Resource Economics 8 (1). Annual Reviews: 247–59. doi:10.1146/annurev-resource-100815-095425.

Livneh, Ben, and Andrew M. Badger. 2020. “Drought Less Predictable under Declining Future Snowpack.” Nature Climate Change 10 (5). Springer Science and Business Media LLC: 452–58. doi:10.1038/s41558-020-0754-8.

Lund, Jay, Josue Medellin-Azuara, John Durand, and Kathleen Stone. 2018. “Lessons from California’s 2012–2016 Drought.” Journal of Water Resources Planning and Management 144 (10). American Society of Civil Engineers (ASCE). doi:10.1061/(asce)wr.1943-5452.0000984.

Macdonnell, Lawrence. 2014. “Prior Appropriation: A Reassessment.” SSRN Electronic Journal 18. Elsevier BV: 228. doi:10.2139/ssrn.2691098.

Marques, Guilherme F., Jay R. Lund, and Richard E. Howitt. 2005. “Modeling Irrigated Agricultural Production and Water Use Decisions under Water Supply Uncertainty.” Water Resources Research 41 (8). American Geophysical Union (AGU). doi:10.1029/2005wr004048.

Medellín-Azuara, J., R.E. Howitt, and J.J. Harou. 2012. “Predicting Farmer Responses to Water Pricing, Rationing and Subsidies Assuming Profit Maximizing Investment in Irrigation Technology.” Agricultural Water Management 108 (May). Elsevier BV: 73–82. doi:10.1016/j.agwat.2011.12.017.

Medellín-Azuara, Josué, Julien J. Harou, Marcelo A. Olivares, Kaveh Madani, Jay R. Lund, Richard E. Howitt, Stacy K. Tanaka, Marion W. Jenkins, and Tingju Zhu. 2007. “Adaptability and Adaptations of California’s Water Supply System to Dry Climate Warming.” Climatic Change 87 (S1). Springer Science and Business Media LLC: 75–90. doi:10.1007/s10584-007-9355-z.

Miller, E. 2008. “An Assessment of CES and Cobb-Douglas Production Functions, Working Paper # 2008-05.” Washington D.C: Congressional Budget Office (CBO.

Modi, Parthkumar A., Eric E. Small, Joseph Kasprzyk, and Ben Livneh. 2022. “Investigating the Role of Snow Water Equivalent on Streamflow Predictability during Drought.” Journal of Hydrometeorology 23 (10). American Meteorological Society: 1607–25. doi:10.1175/jhm-d-21-0229.1.

Munoz Oscar Javier Rojas, Rojas. 2021. “The Role of Clouds and Surface Processes on Local Temperature Variability : Analysis of Multi-Variables Observations.” Ocean, Atmosphere: English.

Pagano, Thomas, David Garen, and Soroosh Sorooshian. 2004. “Evaluation of Official Western U.S. Seasonal Water Supply Outlooks, 1922–2002.” Journal of Hydrometeorology 5 (5). American Meteorological Society: 896–909. doi:10.1175/1525-7541(2004)005<0896:eoowus>2.0.co;2.

Page, Rebecca, and Lisa Dilling. 2019. “The Critical Role of Communities of Practice and Peer Learning in Scaling Hydroclimatic Information Adoption.” Weather, Climate, and Society 11 (4). American Meteorological Society: 851–62. doi:10.1175/wcas-d-18-0130.1.

Pfeiffer, Lisa, and C.-Y. Cynthia Lin. 2012. “Groundwater Pumping and Spatial Externalities in Agriculture.” Journal of Environmental Economics and Management 64 (1). Elsevier BV: 16–30. doi:10.1016/j.jeem.2012.03.003.

Rafey, Will. 2023. “Droughts, Deluges, and (River) Diversions: Valuing Market-Based Water Reallocation.” American Economic Review 113 (2). American Economic Association: 430–71. doi:10.1257/aer.20201434.

Rigby, D., F. Alcon, and M. Burton. 2010. “Supply Uncertainty and the Economic Value of Irrigation Water.” European Review of Agricultural Economics 37 (1). Oxford University Press (OUP): 97–117. doi:10.1093/erae/jbq001.

Schuck, Eric C., W. Marshall Frasier, Robert S. Webb, Lindsey J. Ellingson, and Wendy J. Umberger. 2005. “Adoption of More Technically Efficient Irrigation Systems as a Drought Response.” International Journal of Water Resources Development 21 (4). Informa UK Limited: 651–62. doi:10.1080/07900620500363321.

Siirila-Woodburn, Erica, Cynthia Gerlein-Safdi, Juliane Mueller, James Dennedy-Frank, and Robinson Negron-Juarez. 2021. AI Predicted Shifts in Watershed Hydrodynamics Driven by Extreme Weather and Fire. Office of Scientific and Technical Information (OSTI). doi:10.2172/1769660.

Smith, Steven M, and Eric C Edwards. 2021. “Water Storage and Agricultural Resilience to Drought: Historical Evidence of the Capacity and Institutional Limits in the United States.” Environmental Research Letters 16 (12). IOP Publishing: 124020. doi:10.1088/1748-9326/ac358a.

Stewart, Iris T., Daniel R. Cayan, and Michael D. Dettinger. 2004. “Changes in Snowmelt Runoff Timing in Western North America under a ‘Business as Usual’ Climate Change Scenario.” Climatic Change 62 (1–3). Springer Science and Business Media LLC: 217–32. doi:10.1023/b:clim.0000013702.22656.e8.

United States Census Bureau. 2024. Geographic Levels. Accessed January 2024: https://www.census.gov/programs-surveys/economic-census/guidance-geographies/levels.html#par_textimage_34

Vaux, H. J., and Richard E. Howitt. 1984. “Managing Water Scarcity: An Evaluation of Interregional Transfers.” Water Resources Research 20 (7). American Geophysical Union (AGU): 785–92. doi:10.1029/wr020i007p00785.

Ward, Frank A., Alex S. Mayer, Luis A. Garnica, Nolan T. Townsend, and David S. Gutzler. 2019. “The Economics of Aquifer Protection Plans under Climate Water Stress: New Insights from Hydroeconomic Modeling.” Journal of Hydrology 576 (September). Elsevier BV: 667–84. doi:10.1016/j.jhydrol.2019.06.081.

Weston, J, and Fred. 1950. “Profit as the Payment for the Function of Uncertainty-Bearing.” Journal of Business of the University of Chicago 22 (2). University of Chicago Press: 106. doi:10.1086/232895.

Williams, A. Park, Edward R. Cook, Jason E. Smerdon, Benjamin I. Cook, John T. Abatzoglou, Kasey Bolles, Seung H. Baek, Andrew M. Badger, and Ben Livneh. 2020. “Large Contribution from Anthropogenic Warming to an Emerging North American Megadrought.” Science 368 (6488). American Association for the Advancement of Science (AAAS): 314–18. doi:10.1126/science.aaz9600.

Womble, Philip, and W. Michael Hanemann. 2020. “Legal Change and Water Market Transaction Costs in Colorado.” Water Resources Research 56 (4). American Geophysical Union (AGU). doi:10.1029/2019wr025508.

Womble, Philip, and W. Michael Hanemann. 2020. “Water Markets, Water Courts, and Transaction Costs in Colorado.” Water Resources Research 56 (4). American Geophysical Union (AGU). doi:10.1029/2019wr025507.

Xiao, Dacheng, Yuning Shi, Susan L. Brantley, Brandon Forsythe, Roman DiBiase, Kenneth Davis, and Li Li. 2019. “Streamflow Generation from Catchments of Contrasting Lithologies: The Role of Soil Properties, Topography, and Catchment Size.” Water Resources Research 55 (11). American Geophysical Union (AGU): 9234–57. doi:10.1029/2018wr023736.

Zellmer, Sandra, and A Amos. 2021. “United States: The Emergence of Environmental Considerations.” In The Evolution of the Law and Politics of Water, 205–23. St. Paul, MN, USA: Springer Netherlands. doi:10.1007/978-1-4020-9867-3_13.

Zhang, Wenxia, Kalli Furtado, Peili Wu, Tianjun Zhou, Robin Chadwick, Charline Marzin, John Rostron, and David Sexton. 2021. “Increasing Precipitation Variability on Daily-to-Multiyear Time Scales in a Warmer World.” Science Advances 7 (31). American Association for the Advancement of Science (AAAS). doi:10.1126/sciadv.abf8021.

Zuo, Alec, Robert Brooks, Sarah Wheeler, Edwyna Harris, and Henning Bjornlund. 2014. “Understanding Irrigator Bidding Behavior in Australian Water Markets in Response to Uncertainty.” Water 6 (11). MDPI AG: 3457–77. doi:10.3390/w6113457.