University of California
Santa Barbara

Social Efficiency of Natural Resource Management

A dissertation submitted in partial satisfaction
of the requirements for the degree

Doctor of Philosophy
in
Environmental Science and Management

by

Samuel B Collie

Committee in charge:

Professor Christopher Costello, Chair
Professor Andrew Plantinga
Professor Kyle Meng

September 2021

The Dissertation of Samuel B Collie is approved.

Professor Andrew Plantinga

Professor Kyle Meng

Professor Christopher Costello, Committee Chair

June 2021

Curriculum Vitæ
Samuel B Collie

Education

2021	Ph.D. in Environmental Science and Management (Expected), University of California, Santa Barbara.
2015	M.S. in Agricultural and Resource Economics, Colorado State University.
2011	B.S. in Economics, B.A. in Political Science, University of Rhode Island.

Publications

Suter, J. F., Collie, S., Messer, K. D., Duke, J. M., & Michael, H. A. (2019). "Common pool resource management at the extensive and intensive margins: experimental evidence". *Environmental and Resource Economics*, 73(4), 973-993.

Loomis, J., Collie, S., González-Cabán, A., Sánchez, J. J., & Rideout, D. (2019). Wildfire fuel reduction cost analysis: Statistical modeling and user model for fire specialists in California. Gen. Tech. Rep. PSW-GTR-261 (English). Albany, CA: US Department of Agriculture, Forest Service, Pacific Southwest Research Station: 85-95., 261, 85-95.

Appointments

2018-2020	Teaching Assistant, Economics of Environmental Management.
2020	Teaching Assistant, Cost Benefit Analysis.
2019	Teaching Assistant, Statistics for Economics.
2018	Teaching Assistant, Statistics and Data Analysis for Environmental Science and Management
2017	Teaching Assistant, Intermediate Microeconomic Theory.
2013-2015	Graduate Student Researcher, Colorado State University. Assisted Principal Investigators John Loomis and Jordan Suter.

Fellowships and Awards

2020	Teaching Assistant Award, Bren School of Environmental Science and Management.
2015	Sustainable Water Markets Fellowship, Walton Family Foundation.

Abstract

Social Efficiency of Natural Resource Management

by

Samuel B Collie

Natural resources are essential to society, yet inefficiencies in their use are common, leaving potential for greater social benefits. This thesis explores contemporary issues in natural resource management, with an underlying theme of social efficiency. In Chapter 1, I develop a theory of climate change adaptation for dynamically evolving natural resources and apply it to the implications of warming ocean temperatures on global fishery harvests. I find the benefits of adaptation are surprisingly small, if absent of adaptation resource management is otherwise optimal. The findings suggest that exactly predicting and responding to future climate changes may be less important than other barriers to efficiency, such as open-access losses. In Chapter 2, I examine an institutional barrier to efficient use of fresh water resources. In the water-scarce Western United States, surface water resources are allocated hierarchically through a 19th century institution called prior appropriation. I pair daily administrative records of water allocations under prior appropriation with remotely-sensing irrigated crop outcomes to test if water is used efficiently in Colorado's agricultural sector, and find evidence of unrealized gains from trade. Lastly in Chapter 3, I estimate the economic value of reservoir storage, and compare its value to a planned infrastructure development project.

Contents

Chapter 1

Dynamic Adaptation

1.1 Introduction

Great efforts are being made to estimate the impact of climate change and to design climate-robust policies. Adaptation plays a central role because it directly impacts the benefits and costs of climate change and because it confounds empirical estimators of these benefits and costs. We address an unresolved, and pivotal, conflict within this literature. Climate change, defined as a shift in the distribution of weather experienced over time, is unequivocally dynamic. Yet existing theory either abstracts to static decision problems (Carleton et al. 2020; Deryugina & Hsiang 2017; Gammans, Mérel, & Paroissien 2020; Guo & Costello 2013; Hsiang 2016) or illustrates empirical challenges imposed by dynamics (Lemoine 2017, 2018; Mérel & Gammans 2018). We provide a constructive theory of dynamic adaptation and formalize the role, and value of adaptation in a decision environment with a dynamic state variable.

We show that in an optimally-functioning economic system, the value of adapting to a marginal climate change is precisely zero. This result is not driven by the infinitesimal magnitude of the marginal climate change. Rather it is due to an envelope theorem result first introduced by Guo and Costello (2013) and developed further by Hsiang (2016) and Deryugina and Hsiang (2017). That is, if behavior is optimized prior to climate change, the value of a marginal change in behavior is precisely zero. Our first contribution is to show that this result, first shown in a static environment, extends to this much more complex and realistic dynamic decision environment. Because climate change may be abrupt and "non-marginal", we then contribute a theory of adaptation to discrete climate changes in dynamic decision settings. We show the value of adaptation to discrete climate changes is positive, but need not scale monotonically with the size of the change. Lastly, we show that if climate change is uncertain, its overall effect can be decomposed into a direct effect, the value of adaptation and the value of information.

Others have shown that the entirety of climate change effects may be disaggregated into "direct effects" of weather and "belief effects" of actions based upon expected future weather (Carleton et al. 2020; Deryugina & Hsiang 2017; Hsiang 2016). This is observation is true in principal, but static theories based upon it require an implicit assumption that payoffs may be neatly written as a function of weather and beliefs. Yet it is not immediately clear how to write down such a function for any likely empirical target- be it a profit maximizing firm, or human mortality risk. Explicit treatment of the dynamics in play affords the opportunity, and complication, of doing so. Our theory essentially delineates the "direct" and "belief effects" of climate change for a conventional decision problem with a capital stock.

Explicit treatment of dynamics introduces several complications to the analysis of climate change effects. In a dynamic setting, the effects of a change in any entity, including a change in climate, ripple through time via their effects on the trajectory of the capital stock. This in turn affects the choice of actions in all periods, as these depend, in part, upon the level of the capital stock when these decisions are made. Thus, our treatment of a dynamic capital stock is an analogue to a popular view that adaptive potential is lesser in the short run than it is in the long run (Auffhammer 2018; Mérel & Gammans 2018). That is, estimates of short-run effects of climate change do not account for the optimal readjustment of capital they induce. Through explicit treatment of a dynamic capital stock, our theory captures climate-induced capital adjustments in a natural way, and does so without an ad-hoc decision to hold some inputs fixed in the short run.

Supposing the change in actions due to a change in climate are derived, one is forced to consider how much of these changes constitute adaptation. We define "dynamic adaptation" as the change in the *plan* of action due to a change in climate. This is represented mathematically as the change in the policy function, or functions, that prescribes an action in each time period as a function of the contemporaneous level of the capital stock and the weather. Because a plan is a function (or functions), a change in climate evokes changes in actions even when the planner does not adapt. This occurs because climate change affects the realizations of weather as well as the trajectory of the capital stock, and the levels of each dictate the action prescribed by a plan in every period. We show that a plan that is optimized prior to a change in climate "bakes-in" a surprising extent of the optimal response to the climate change. This is perhaps more intuitive if we call a plan by its alternate name of a "feedback control rule". Supposing that the climate changes, but the plan of action does not, the transition of the capital stock between periods will not match what the planner expected to occur when the plan was devised. This in turn results in a different prescribed action in each period, affecting the trajectory of the capital stock, and so on. As a result of this feedback, the actions and capital stock trajectory prescribed by a climatically ill-conceived policy may not differ markedly from those of the climatically optimal policy.

Expectations, and the timing of information, present further nuances. Whereas static theories treat expectations in reduced form, via the "belief" channel, we are forced to reconcile the role of expectations in greater detail. First off, we must decide if the relevant definition of climate effects is foreword-looking, based on what is expected to happen, or backwards-looking, based on what did happen. Since most interest in climate change pertains to future predicted impacts, we adopt the former view. Our definitions of climate climate change effects are forward-looking, or “ex-ante” of future climate changes (from the perspective of the initial time period), based on current beliefs about how the climate will change. The role of expectations (or beliefs), we find, hinges on a planner’s information set. We analytically derive climate effects under several assumptions regarding information, doing so using methods from the economic theory of natural capital.

In our initial model, we suppose the climate changes instantaneously at time zero, and this is known to planners, should they choose to respond. In this context, we imagine an “unresponsive” planner type that acts as though the climate has not changed, whose incorrect expectations of future weather are based on the original climate prior to the change. Second, we imagine a “responsive” planner type with correct expectations of future weather, based on the new climate after the change. For both planner types, we evaluate their expected rewards under the true climate after the change, and attribute their differences as “the value of adaptation”. Although here given a different name, this is equivalent to the “belief effect” hypothesized in static frameworks.

While this treatment of discrete climate changes is equivalent in the limit to the analysis of marginal climate changes studied previously, it is not entirely satisfying because it rests on an unrealistic assumptions that climate change occurs instantaneously and that the change is known with certainty. To address this shortcoming, we consider a model of gradual and uncertain climate changes, previously studied by Kelly, Kolstad, and Mitchell (2005) and Kala (2017). Here, we introduce a “sophisticated” planner type, whom is not informed of the change in climate but attempts to learn it through their experience of weather. Additionally, a “perfect foresight” planner is informed of the climate change, and has the correct expectation of future weather. In this model of uncertain climate change, we show that the belief effect comprises the value of adaptation as well as the value of information. Said differently, the overall effect of climate change includes the value of knowing beforehand how the climate will change.

To summarize our theoretical findings, we extend the analytic framework of climate change effects to explicitly account for dynamic linkages in decision-making. Despite the ripple-effects of climate changes over time on the trajectory of a capital stock, and the effects on expectations of future weather, we conclude that the value of adaptation is likely small. For marginal climate changes, we formalize that the value of adaptation is precisely zero, extending an envelope result previously based on static analysis of climate changes. We then allow the climate to change discretely, and show that although the value of adaptation is non-zero, it is likely small due to the feedback-response of a dynamic plan of action. These results are not because adaptation will not occur, or because the magnitude of adaptation itself is small. Rather, they are because the value of changing a plan of action that was previously optimal are small, and precisely zero on the margin. Finally, we allow for uncertain climate changes and find that their total effect additionally comprises the value of information.

This theory of dynamic adaptation to climate change has applications economy-wide, in growth, household decision making, natural resources, and elsewhere. We apply the theory in an empirical setting of global marine fishery management. Our analysis builds on a recent advance in fisheries science that estimated the effect of warming sea surface temperatures on the productivity of the world's commercial fisheries stocks (Free et al. 2019). Applying our theory using numerical stochastic dynamic programming, we derive distinct policy functions (or harvest control rules in the fisheries management parlance) for each of the 235 global fishery stocks, adopting the empirically derived effects of climate change on each stock. This approach allows us to estimate the effects of uncertain climate changes for each fishery. While climate change in this example is decidedly non-marginal, we find that the values of adaptation and information are orders of magnitude smaller than the direct effects of climate change without adaptation. In other words, we find that maintaining a stationary harvest-control rule in the face of a changing climate "bakes-in" the majority of the optimal response to the change. This accords with our theoretical results and indicates that the majority of costs or benefits of climate change arise due to its direct effects without adaptation or learning.

Our structural approach to estimating climate change effects in the context of global fisheries is advantageous for our purposes, as it allows us to exactly derive the components of our theoretical model. This allows us to side-step empirical challenges raised elsewhere of matching theory to empirically-testable implications (see for example, Gammans et al. (2020); Mérel and Gammans (2018) and Lemoine (2018)). However, since much interest in climate change effects pertains to reduced-form estimation, it is worth discussing how our results inform that literature. A principal concern therein is when the envelope result of zero marginal value of adaptation applies, because when it does the entirety of climate change effects are readily obtained by reduced-form estimates of direct effects (without adaptation). We argue that virtually all economic decision problems contain some dynamic component (i.e. durable capital investment) and that analysis of marginal climate changes is of little relevance given that the climate has already changed substantial and is expected to continue to do so. Even so, we demonstrate, in theory and in the context of global fisheries, that the value of adaptation is small. Therefore, even when the envelope theorem does not hold exactly, it is likely the case that the role of adaptation does not substantially impact the effects of climate change.

1.1.1 Relation To Previous Literature

We build on a literature that considers climate change adaptation from a theoretical standpoint. Beginning chronologically, Guo and Costello (2013) derives an analytic framework for estimating an economic value of adaptation and an empirical application to California timber management under climate change. This paper was the first to formalize the connection between the envelope theorem and the value of adaptation: in situations where agents are optimizing their behavior, the value of adaptation to marginal environmental change is exactly zero. This result implies that the role of adaptation can be safely ignored in estimates of climate change impacts- so long as agents are optimizing behavior and the environmental change is indeed marginal.

These ideas are taken up and refined by Hsiang (2016) and Deryugina and Hsiang (2017), who develop rigorous mathematical treatments of climate and weather with an eye towards reduced-form empirical analysis. Ultimately they rely on the same envelope-theorem logic as Guo and Costello (2013), in which the value of adaptation vanishes when agents are optimizing. Instead of assuming the role of an optimizing planner, Deryugina and Hsiang (2017) motivates optimizing behavior with the First Welfare Theorem. The competitive equilibrium of a perfect economy is optimal, therefore the envelope theorem readily applies. An empirical payday arises from these insights because the econometrician can recover climate change impacts from short-run variations in weather, whenever the assumption of optimizing behavior is met.

Carleton et al. (2020) advances the climate change adaptation literature by explicitly deriving the benefits and costs of adaptation in a revealed preference framework of climate change mortality. Their key insight is that the difference in full effect of climate change with adaptation and the partial effect without adaptation reveals the benefit of undertaken adaptation. Moreover, this value is readily estimable as the difference of two reduced-form estimators of climate change impacts, one that includes the role of adaption and the other that does not. We explore many of the same ideas as Carleton et al. (2020), but extend the analysis to the uncertain dynamic decision environment. In what follows, we use slightly different nomenclature than Carleton et al. (2020). They define "adaptation surplus" as adaptation benefits minus costs. Their revealed preference argument is that the benefit of observed adaptation must outweigh its unobserved cost. In contrast, we define "the value of adaptation" as the maximum attainable benefit of adaptation without regard to its cost. Since we cannot parse out adaptation costs empirically, we omit adaptation costs in our theory. Instead, we note that since it ignores adaptation costs, our value of adaptation must exceed adaptation surplus and therefore our empirical estimates reflect an upper bound for economically realizable benefits of adaptation.

All of the papers mentioned so far feature static decision environments. Yet the climate effects they describe allude to "belief effects" of future weather, and hence an implicit dynamic decision environment. In contrast, two papers by Derek Lemoine (2017; 2018) introduce dynamic decision-making to the literature outlined above. In these papers, Lemoine shows that the envelope-theorem logic does not apply to certain empirical situations where agents make decisions over time. He recognizes that the envelope-theorem only applies to an agent's objective, which may not be the object being studied. As a concrete example, suppose the quantity of interest is the impact of climate change on agriculture. A classic approach is to estimate the effect of short-run weather variation on annual agricultural profits (Deschênes & Greenstone 2007). Lemoine points out that the envelope theorem only applies in this scenario if annual-profits are indeed the target of agricultural producers' optimization; that is so long as their decision problem is truly static. If instead an empiricist instead studies land prices, which capitalize future rents, the envelope theorem indeed applies.

We contribute to the theory of adaptation, outlined above, in several ways. We provide a theoretical treatment of the effects of exogenous climate changes on an infinite-horizon dynamic decision problem with a capital stock. We demonstrate that the envelope theorem applies in such settings and prove that the value of adaptation to a marginal climate change is zero. Applying methods from the theory of natural capital (Arrow, Dasgupta, & Mäler 2003; Dasgupta & Mäler 2000), we provide the first derivation of the effects discrete (i.e. non-marginal) climate changes in dynamic decision environments. The value of adaptation to discrete climate changes is greater than zero and hence an empirical question. Our final theoretical contribution is due to the introduction of climate change uncertainty, in the spirit of Kelly et al. (2005) and Kala (2017). When climate changes are uncertain, we show their overall effect also comprise the value of information.

Our empirical analysis contributes to an established fisheries-management literature sharing our methodology of stochastic dynamic programming, that quantify the benefits of adaptation in ocean fisheries management. Stationary harvest control rules are commonly employed to manage fishery stocks, despite a consensus of evidence that fishery stock-recruitment varies widely, due to sea temperature warming among other factors (Britten, Dowd, & Worm 2016; Maunder & Thorson 2019; Minto, Myers, & Blanchard 2008; Rose 2000; Rothschild 2000; Thorson, Jensen, & Zipkin 2014; Vert-Pre, Amoroso, Jensen, & Hilborn 2013). This evidence points to potential gains from adapting to climatic conditions as they unfold. For example, Vert-Pre et al. (2013) argue that: *"Fisheries management agencies need to recognize that irregular changes in productivity are common and that harvest regulation and management targets may need to be adjusted whenever productivity changes."* While true in principal, studies that estimate the returns to adaptive management strategies find only modest increases in performance as compared to simpler, non-adaptive, strategies (Ludwig & Hilborn 1983; Parma 1990; Smith & Walters 1981; C. Walters & Parma 1996; C. J. Walters 1981). Our theory of dynamic adaptation formalizes this numerical approach to estimating the gains from adaptation and provides a plausible mechanism for the routine finding that the benefits of adaptation are small. We also provide the first global empirical analysis of this kind, thus extending previous results based on case-studies of individual fishery stocks.

1.2 Theory

The theory section is organized as follows. First we lay out a general model of decision-making in the context of a dynamic capital stock and exogenous climate. Our setup is very general and builds constructively on the theories of adaptation in static settings. Climate is viewed as the distribution of possible weather, and climate changes are shifts in the distribution of weather outcomes. Following the setup, we provide two examples of climate change. The first example helps build intuition and is designed to be as simple as possible, while retaining the features relevant to our analysis. In that example, climate change is viewed as a once-and-for-all change in the distribution of weather. In that context, we derive the effects of marginal climate changes, and the effects of a discrete change from climate C^a to climate C^b. Recognizing that climate change is a more dynamic process, we develop a second example where we allow climate to change gradually over time and for these changes to be uncertain. Planners may learn about climate change given their experience of past weather. In this context the full effect of climate change is shown to additionally comprise the value of information.

1.2.1 Model Preliminaries

Time is discrete and infinite, indexed by $t = 1, 2, \ldots, \infty$. A climate C is a vector of sufficient statistics for the distribution of weather: $w_t \sim f(C)$. Weather is independent between time periods implying that if $W^t \equiv (w_1, \ldots, w_t)$ is the history of weather at time t, then $\Pr(w_{t+1} \mid W^t) = \Pr(w_{t+1})$. For convenience, further assume the weather in all climates has the same domain. In other words, the climate does not alter the set of possible weather. Evolution of a scalar capital stock x_t (e.g. a fishery stock's biomass), is influenced by a scalar control h_t (e.g. harvest) and contemporaneous weather w_t. Choice of the control in any period is restricted to the set-valued correspondence $\mathcal{H}$ which may depend on the current capital stock and the the weather; $h_t \in \mathcal{H}(x_t, w_t)$. Single-period rewards in period are $r(h_t, x_t, w_t)$ and the state evolves according to $x_{t+1} = g(h_t, x_t, w_t)$. These modeling choices reflect the argument that climate, as a distribution, affects outcomes solely through realizations of weather and through beliefs of future weather (Deryugina & Hsiang 2017; Hsiang 2016). The timeline in a single period is as follows: the weather shock is realized, the planner chooses the control and subsequently period rewards are realized. Lastly, $\delta \in (0, 1)$ is a discount factor.

Collecting everything so far yields the objective for a planner facing climate C.

$$\max_{\{h_t\}_{t=1}^{\infty}} \mathbb{E}_w \left[\sum_{t=1}^{\infty} \delta^{t-1} r(h_t, x_t, w_t) \right] \tag{1.1}$$

subject to:

$$w_t \sim f(C)$$

$$h_t \in \mathcal{H}(x_t, w_t)$$

$$x_{t+1} = g(h_t, x_t, w_t)$$

given: x_1, w_1, C

We require that for any initial state x_1, w_1 and climate C, there exists a unique solution to 1.1 that can be represented by a continuous policy function. To ensure this condition is met, we exchange variables in objective 1.1 to instead choose the trajectory of the capital stock itself. The result is succinct yet general and has convenient analytic properties. Assumption 1 ensures this exchange of variables is feasible.

Assumption 1. *State transitions are monotonic in the control, implying g has a well-defined inverse.*

$$g^{-1}(x_t, x_{t+1}, w_t) \equiv h_t \tag{1.2}$$

Assumption 1 allows us to recast period-rewards as a function of x_t, x_{t+1} and w_t, leaving the choice of the control h_t implicit.

$$u(x_t, x_{t+1}, w_t) \equiv r(g^{-1}(x_t, x_{t+1}, w_t), x_t, w_t) \tag{1.3}$$

Additionally, define the inverse mapping of the feasibility correspondence as follows:

$$\mathcal{X}(x_t, w_t) \equiv \{g(h_t, x_t, w_t) : h_t \in \mathcal{H}(x_t, w_t)\} \tag{1.4}$$

We are now in position to recast objective 1.1 in terms of a choosing the optimal trajectory of the stock. The state equation now holds implicitly, having been subsumed by the transformations.

$$\max_{\{x_{t+1}\}_{t=1}^{\infty}} \mathbb{E}_w \left[\sum_{t=1}^{\infty} \delta^{t-1} u(x_t, x_{t+1}, w_t) \right] \tag{1.5}$$

subject to:

$$x_{t+1} \in \mathcal{X}(x_t, w_t)$$

$$w_t \sim f(C)$$

$$\text{given: } x_1, w_1, C$$

We restrict attention to the class of problems with a unique solution to 1.5, that is characterizable by a continuously differentiable policy function. To this end, we maintain the ensuing concavity assumption.

Assumption 2. *Period rewards $u(x_t, x_{t+1}, w_t)$ are concave in x_t, x_{t+1} and continuously differentiable in the interior of its domain.*

Assumption 2 yields several well known results.[1] Under the stated assumptions and the model primitives, there exists a unique solution to 1.5, which can be characterized by a continuous policy function (or functions) $x^*_{t+1} = \pi_t(x_t, w_t)$. The model presented thus far is stationary, meaning it has no dependence on calendar time. In this case the optimal plan can be represented by a single policy function $x^*_{t+1} \equiv \pi(x_t, w_t)$. In addition, there is a continuously differentiable functional value equation $V(x_t, w_t)$ equal to the continuation value of the optimal plan beginning in any time period t. The climate is implicit in the optimal policy, as it determines the expectation of future weather. It is useful to make that dependence explicit with $\pi(x_t, w_t \mid C)$. And similarly, it is useful to establish that $V(x_t, w_t \mid C)$ is the value function under climate C, equal to the maximized value of objective 1.5. The full set of assumptions and regularity conditions are provided in Appendix A.

Some notes on the model and its assumptions are warranted. Our aim here is to capture the relevant features of dynamic adaptation as simply as possible. To do so we restrict the class of dynamic decision problems in several ways (i.e. concavity, independence of weather), but purely for the sake of parsimony. These assumptions can be relaxed without damage to our core findings; Key examples are discussed momentarily.

Uniqueness of the optimal plan requires the period reward function, $u(x_t, x_{t+1}, w_t)$, be concave in its first two arguments: the current capital stock and the stock in the next period. This "inter-period complementarity" (Lemoine 2017) entails that a moderate level of the capital stock in adjoining time periods is preferred to an otherwise comparable situation of extremes. This assumption is non-restrictive for our case-study of renewable fishery management, where inter-period complementarity is due to density-dependent

[1] See for example Acemoglu (2009) or Stokey and Lucas Jr. (1989).

growth. Other decision problems where multiple solutions arise are outside the scope of the current paper, but can be handled with care to boundary conditions.[2]

In casting the decision problem as the choice of trajectory of the capital stock, we assume that conditional on the current stock and weather, the stock in the next period is known with certainty. This assumption has convenient analytic properties, but does not cover all empirically relevant cases. For example, Weitzman (2002) argues that in fisheries management, the relevant objective is to choose harvest each prior to the realization of the productivity shock. One could reflect that sentiment by amending our setup slightly, in which case the weather would not appear as a state variable and instead appear as a distribution each period. Although we do not explore that case theoretically, we do implement it in the empirical section and find a small increase in the estimated value of adaptation. Heuristically, the difference between choosing actions before versus after the weather is realized each period amounts to a first versus second order impact of climate on actions (Costello, Polasky, & Solow 2001). However, our theoretical results obtain either way and so here we maintain the assumption that the capital stock in the next period is known with certainty, as to provide the most succinct derivations and definitions of dynamic climate change effects.

A final note on our setup is that the assumption that weather is independent is also made for convenience and is unnecessarily restrictive. It is straightforward to accommodate serially-correlated weather that satisfies a Markov property of any order. However doing so does not provide additional insight to our problem and so we take the simpler independence assumption.

[2]For example Hotelling's non-renewable resource with stock dependent extraction costs (Hotelling 1931), a situation Lemoine (2017) describes as having 'inter-period substitutes'.

1.2.2 Characterizing the Optimal Plan

We establish a few useful tools en route to our main theoretical results. Write the Value function in its recursive Bellman form:

$$V(x_t, w_t \mid C) = \max_{x_{t+1}} \{u(x_t, x_{t+1}, w_t) + \delta \mathbb{E}_w [V(x_{t+1}, w_{t+1} \mid C)]\} \tag{1.6}$$

subject to:

$$x_{t+1} \in \mathcal{X}(x_t, w_t)$$

$$w_{t+1} \sim f(C)$$

Denoting the optimal choice as $x^*_{t+1} \equiv \pi(x_t, w_t \mid C)$, the Euler Equations are:

$$\mathbb{E}_w \left[u_2(x_t, x^*_{t+1}, w_t) + \delta u_1\left(x^*_{t+1}, x^*_{t+2}, w_{t+1}\right) \mid w_t\right] = 0, \;\; t \in \mathbb{N}^1 \tag{1.7}$$

(shown in Appendix B). Note that u_1, u_2 and u_3 respectively denote the partial derivatives of $u(x_t, x_{t+1}, w_t)$ with respect to its first, second and third arguments. The conditional expectation indicates that at time t, weather w_t is known but w_{t+1} is uncertain.

Euler equations (1.7) are a cryptic rendition of the familiar equality of the marginal benefits and costs. The first term is the change in contemporaneous payouts due to a marginal change in the next-period's stock. The second term is the expected discounted change in the next period's payouts due to a marginal change in the next period's stock.

1.2.3 Effects of Marginal Climate Changes

We are now in position to analyze the effects of a marginal climate change. To do so we temporarily assume existence of a smooth-valued function $w(C) = w_t$, which maps the climate to a draw of weather.[3] We use $w(C)$ here purely for analytic convenience, supposing it has a gradient $w'(C)$ describing the change in (stochastic) weather due to a marginal change in climate. The assumption that $w(C)$ exists and is differentiable is needed because we have not yet imposed any structure on climate, nor have we described what it means for that distribution to change. We defer further discussion of this to section 1.3, where we provide a structural representation of climate change under normally distributed weather.

We first derive the effects of a marginal climate change: the derivative of value function at time 1 with respect to C. In other words, suppose the climate changes marginally at time 1 and that the change persists indefinitely. Following the nomenclature of Guo and Costello (2013), we decompose the total change into the "direct effect without adaptation" and the "value of adaptation".

Definition 1. *The* ***full effect with adaptation*** *of a marginal climate change is equal to the sum of the* ***direct effect without adaptation*** *and the* ***value of adaptation****.*

$$\overbrace{\frac{dV(x_1, w_1 \mid C)}{dC}}^{\substack{\text{Full Effect} \\ \text{w/ adaptation}}} = \sum_{t=2}^{\infty} \delta^{t-1} \mathbb{E}_w \Bigg[\overbrace{u_3(x_t^*, x_{t+1}^*, w(C)) w'(C)}^{\substack{\text{Direct effect} \\ \text{w/o adaptation}}} + \dots$$

$$\dots \underbrace{\left(u_2(x_t^*, x_{t+1}^*, w(C)) + \delta u_1(x_{t+1}^*, x_{t+2}^*, w(C)) \right) \frac{dx_{t+1}^*}{dC}}_{\text{Value of adaptation}} \Bigg] \tag{1.8}$$

(Derivation provided in appendix C).

The components of definition 1 are intuitive. The direct effect without adaptation is the first-order effect of climate change on discounted rewards, mediated through climate's effect on weather. The value of adaptation is the change in expected rewards due to the

[3]For details of the implicit assumptions that underlie $w(C)$, refer to Deryugina and Hsiang (2017).

change in the optimal plan, or policy function. As the value of an optimal response, the value of adaptation is weakly greater than zero by definition. Our first result is that for marginal climate changes, the value of adaptation is exactly zero. This is an envelope result.

Theorem 1. *There is zero value of adaptation to a marginal environmental change.*

$$\frac{dV(x_1, w_1 \mid C)}{dC} = \frac{\partial V(x_1, w_1 \mid C)}{\partial C} \tag{1.9}$$

Proof. Euler equations (1.7) hold in all time periods, thus the inner parentheses of 1.8 collapse. Therefore on the margin, the full effect of a marginal climate change is equal to the partial effect without adaptation. □

1.2.4 Discrete Climate Changes

While Theorem 1 is instructive and helps build intuition, the requirement that climate change is "marginal" is restrictive, and may be empirically irrelevant in many applications. Thus, we now turn to discrete environmental changes, supposing the climate changes from C^a to C^b, once-and-for-all at time 1.[4] Deriving the consequences of this change requires a description of what happens when the planner fails to adapt. To that end, we introduce two types of planners, one which does and one which does not respond to the climate change.

Responsive planners maximize expected rewards under the new climate C^b, as prescribed by the optimal policy function $\pi(x_t, w_t \mid C^b)$. Unresponsive planners retain the optimal plan under climate C^a when in actuality climate C^b generates the weather. The expected rewards for both planner types are evaluated ex-ante of weather generated by

[4] It is convenient to assume climate change does not affect period 1 weather, given that it is known to the planner in the first time-period.

the true climate C^b, from the perspective of time 1.

We draw from the theory of natural capital to analytically derive the effect of each planner type on expected rewards. That literature provides, through *allocation mechanisms* and *economic programs*, methods to accommodate sub-optimal behavior in dynamic environments (Arrow et al. 2003; Dasgupta & Mäler 2000) which will be essential for this analysis.

Allocation Mechanisms

Allocation mechanisms provide means for analyzing arbitrary sub-optimal behavior in dynamic decision environments. They are set-valued correspondences that map model givens to a feasible plan. A feasible plan is a recursive sequence of functions that prescribe the optimal stock in the next period, given the current period's stock and weather realization (see Appendix A for details). The climate enters an allocation mechanism directly through the draw of weather in every period. It may also enter indirectly through beliefs. Since feasible plans need not be optimal, there is no requirement that these beliefs be consistent with realized weather; in other words, one could believe the climate is C^a even if weather is consistently being drawn from C^b. Building towards the analysis of discrete climate changes, we define allocation mechanisms for the responsive and unresponsive planner types. The expected value of each planner type's behavior, as determined by their allocation mechanisms, is called the value of their economic programs.

Responsive Planners

A responsive planner optimizes expected rewards given correct climate beliefs. Their allocation mechanism is therefore a mapping from the climate, C, to the optimal policy function given that climate, $\pi(\cdot, \cdot \mid C)$. We denote an allocation mechanism for a responsive planner $\mathcal{R}(C)$, making the dependence on climate explicit.

Since a responsive planner's behavior is optimal, the value of the economic program described by $\mathcal{R}(C)$ is equal to the corresponding value function, as previously derived:

$$V(x, w \mid \mathcal{R}(C)) \equiv V(x, w \mid C) \tag{1.10}$$

It follows that Theorem 1 (zero marginal value of adaptation) applies to a responsive planner.

Unresponsive Planners

An unresponsive planner enacts the optimal policy function assuming the climate is C^a, when in actuality it is climate C^b that generates the weather. By assumption, this policy is feasible.[5] The stock's trajectory then depends upon climate C^a, through the uniformed planner's beliefs (or policy function), and upon climate C^b through realizations of weather.

Denote the allocation mechanism for an unresponsive planner as $\mathcal{U}(C^a, C^b)$, where C^a is the planner's perceived climate and C^b is the true climate that generates weather. $\mathcal{U}(C^a, C^b)$ is most easily defined implicitly through the value of its economic program.

$$V(x_1, w_1 \mid \mathcal{U}(C^a, C^b)) \equiv \mathbb{E}_w \left[\sum_{t=1}^{\infty} \delta^{t-1} u(x_t, x_{t+1}, w_t) \right] \tag{1.11}$$

subject to:

$$x_{t+1} = \pi(x_t, w_t \mid C^a)$$

$$w_{t+1} \sim f(C^b)$$

given: x_1, w_1, C^a, C^b

[5] Climate does not enter the constraint correspondence $\mathcal{X}(x_t, w_t)$, which depends upon the current state and realized contemporaneous weather. Additionally, we assume the range of weather outcomes is equal under all climate scenarios.

The value of the uniformed planner's economic program is the ex-ante value of following their allocation mechanism, from the perspective of time 1, taken over the new climate C^b. Notice that the assumed climate C^a enters through the choice of actions, whereas the true climate C^b enters through the draw of weather in each period.

Effects of Discrete Climate Changes

We can now analyze the effects of discrete climate changes. The direct effect without adaptation due to a discrete climate change is defined as the difference in expected values under climates C^a and C^b holding the policy function constant at $\pi(\cdot, \cdot \mid C^a)$. The value of adaptation to a discrete climate change is defined as the difference in expected values due to a change in plans from policy function $\pi(\cdot, \cdot \mid C^a)$ to $\pi(\cdot, \cdot \mid C^b)$, holding the climate fixed at C^b. These effects can be constructed from the responsive and unresponsive planners' allocation mechanisms.

Definition 2. *The* ***full effect with adaptation*** *of a discrete climate change from* C^a *to* C^b*, once-and-for-all in period* $t = 1$*, is equal to the* ***direct effect without adaptation*** *plus the* ***value of adaptation****.*

$$\begin{aligned}\overbrace{V(x_1, w_1 \mid C^b) - V(x_1, w_1 \mid C^a)}^{\text{Full Effect w/ Adaptation}} &= \overbrace{V(x_1, w_1 \mid \mathcal{U}(C^a, C^b)) - V(x_1, w_1 \mid \mathcal{R}(C^a))}^{\text{Direct Effect w/o Adaptation}} \\ &+ \underbrace{V\left(x_1, w_1 \mid \mathcal{R}(C^b)\right) - V\left(x_1, w_1 \mid \mathcal{U}(C^a, C^b)\right)}_{\text{Value of Adaptation}}\end{aligned} \tag{1.12}$$

The definition is an identity, given definition 1.10. Note that the first and last terms have been added and subtracted.

Our next result establishes an equivalence of the discrete and marginal analyses in the limit as the discrete change in climate approaches zero. Unlike the marginal analysis, the discrete effects are defined in absolute terms, meaning they are not divided by the magnitude of the change in climate. This is because climates C^a and C^b are vectors of sufficient statistics for arbitrary probability distribution functions, making it difficult to objectively define the distance between them. Therefore we enlist a simplifying assumption that climates are scalar to establish equivalence between the discrete and marginal analysis as the magnitude of climate change shrinks.

Theorem 2. *Assuming $C^a, C^b \in \mathbb{R}$, the effects of a discrete climate change from C^a to C^b (definition 2), when divided by the size of the change $C^b - C^a$, are equal to the effects of a marginal climate change (definition 1) in the limit as $C^b \to C^a$.*

1. *Equivalence of the Direct Effects:*

$$\lim_{C^b \to C^a} \left[\frac{V(x_1, w_1 \mid \mathcal{U}(C^a, C^b)) - V(x_1, w_1 \mid \mathcal{R}(C^a))}{C^b - C^a} \right] = \frac{\partial V(x_1, w_1 \mid C^a)}{\partial C^a}$$

2. *Equivalence of the Value of Adaptation:*

$$\lim_{C^b \to C^a} \left[\frac{V\left(x_1, w_1 \mid \mathcal{R}(C^b)\right) - V\left(x_1, w_1 \mid \mathcal{U}(C^a, C^b)\right)}{C^b - C^a} \right] = 0$$

Proof provided in Appendix D.

It can also be shown that the value of adaptation to a discrete environmental change is generally greater than zero. This is because the envelope theorem only holds along the optimal trajectory of the stock. The discrete case involves sub-optimal behavior and thus a reduction in the value function, relative to what would have arisen under optimal behavior. Furthermore, the value of adaptation need not increase monotonically with the size of the environmental change. This would be the case if and only if the value function $V(x, w \mid C)$ were concave in C. This is not generally the case under the assumptions of our model.

So far we have established that the value of adaptation to once-and-for-all climate change is zero for marginal changes and positive for non-marginal changes. But this model is a building block and is contrived by design. We have assumed thus far that climate changes occur in period 1, persist indefinitely, and that this change is known to planners, although some fail to respond. We confront these limitations by introducing a more empirically relevant model of gradual and uncertain climate change. This, we will show, conflates the value of adaptation with the value of information. We disentangle the two and provide a modified definition of the value of adaptation.

1.3 Climate Change Uncertainty

Here we suppose that climate changes gradually period by period and that planners are uncertain of these changes and must thus form expectations about the future. Maintaining analytical traction in this context requires we structure the problem carefully. In particular we restrict attention to a model of climate change uncertainty with self-conjugate priors. To that end we build on the climate change models of Kelly et al. (2005) and Kala (2017).

The weather in all periods is distributed normally with a time varying mean and known constant variance. Period-t Climate, $C_t \equiv (\mu_t, \rho_w)$, is completely characterized by the mean and precision (inverse of variance) of period t weather. We assume that all planner types are informed of mean weather prior to period 1, μ_0, and its precision, ρ_w. Instead of an abrupt once-and-for-all shift, we implement gradual climate change as follows: The mean of climate evolves according to $\mu_{t+1} = \mu_t + d$. Under these assumptions, d completely describes the climate in every period (i.e. $C_0 = (\mu_0, \rho_w)$, $C_1 = (\mu_0 + d, \rho_w)$ etc.) so we may refer to d as the "climate".

We will allow for the possibility that some planner types are aware of the drift parameter d, in which case climate changes along a known trajectory, and other planner types are uncertain of d. Planners have prior beliefs over the drift parameter: $d \sim \mathbb{N}\left(\underline{d}_1, 1/\rho_d\right)$, where $\underline{d}_1$ is the prior mean of d at time 1 and ρ_d is the precision of these beliefs. The prior mean of weather in period t, from the perspective of time 1, is:

$$\mu_t \sim \mathbb{N}(\mu_0 + td, t^2/\rho_d) \tag{1.13}$$

which follows because $\mu_t = \mu_0 + td$ is a linear transformation of d. Note that the prior precision of weather decreases in t (or the prior variance of weather increases in t). This is an intuitive representation of climate change, since the climate is more certain in the near future than it is in the distant future.

Our goal is to derive the value of adaptation to climate change, where climate change is now due to the drift in mean weather d. We do so by introducing several planner types. A "Perfect Foresight" planner, $\mathcal{F}$, is informed of the trajectory of climate, given by the value of d. A "Sophisticated" planner, $\mathcal{S}$, is uncertain of d and behave optimally given this uncertainty. Lastly, an "Unresponsive" planner, $\mathcal{U}$, retains the optimal plan under the assumption of no climate change (i.e. presuming $\mu_t = \mu_0$ for all t). Each planner type has a distinct objective, from which a policy function is derived.

To compare the performance of each planner type, we wish to evaluate their allocation mechanisms under the same,"true", climate. But what climate should be used? For the sake of theory, it does not matter, since for any particular climate the Perfect Foresight planner will perform best, the Sophisticated planner second best, and the Unresponsive planner third. Policy analysis, however, merits judicious choice of the "true" climate to analyze, but also motivates forward-looking analysis of climate effects yet to come. Therefore we define uncertain climate change effects by evaluating each planner type under the same distribution of prior climate beliefs as the Sophisticated planner (equation

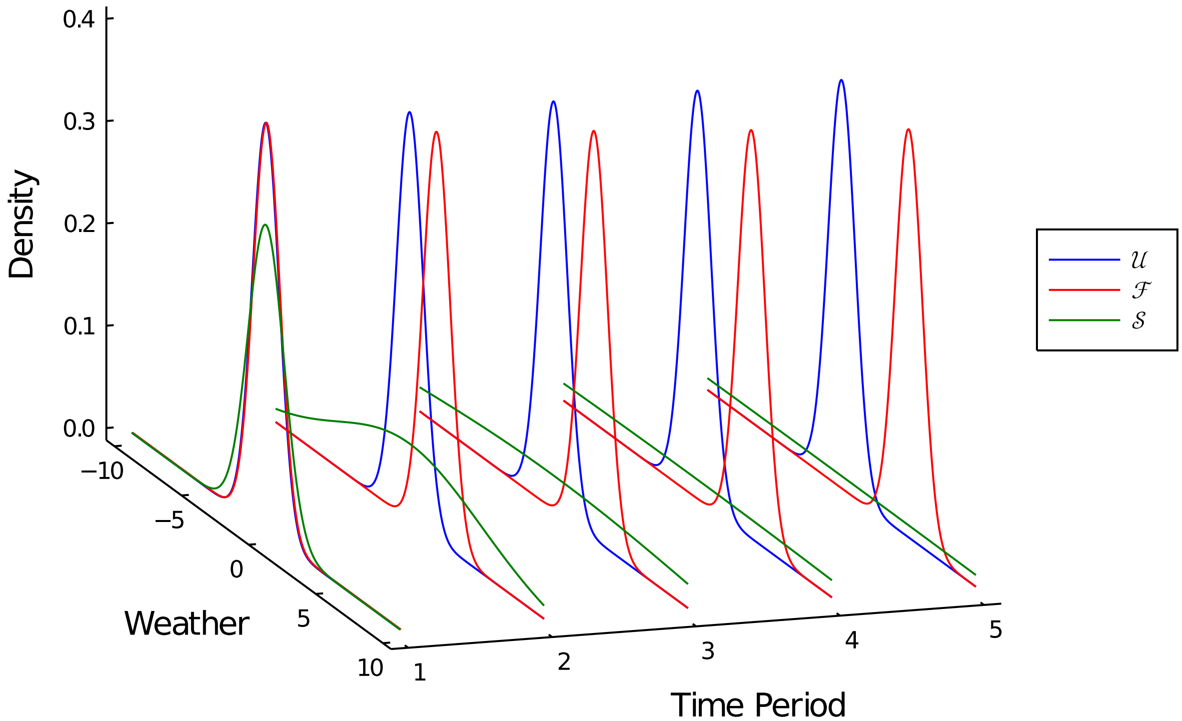

Figure 1.1: Distributions of expected weather for unresponsive ($\mathcal{U}$), perfect foresight ($\mathcal{F}$) and Sophisticated Planners ($\mathcal{S}$).

1.13). As such, our theory reflects the expected effects of future climate changes, given uncertainty over how the climate will actually change.

As in the case without climate uncertainty, we segment the full effect of climate change into the sum of mutually exclusive parts. Without climate change uncertainty, we showed that the full effect of climate change is composed of the value of adaptation and the direct effect without adaptation. With climate change uncertainty, we will show an additional term is needed: the value of information. This arises because the maximum attainable value under climate change is infeasible given planners' information. To deal with this, and close the model, we introduce the Perfect Foresight planner, whom is perfectly informed of the trajectory of climate change (i.e the value of d). Perfect Foresight planners remain uncertain of future weather, but they know the distribution from which future weather will be drawn.

1.3.1 Perfect Foresight Planners

Perfect Foresight planners are perfectly informed of the trajectory of climate change and optimize behavior accordingly. Denote the allocation mechanism for a Perfect Foresight planner as $\mathcal{F}$. The objective for Perfect Foresight Planners is:

$$V_1^{\mathcal{F}}(x_1, w_1 \mid d) \equiv \max_{\{x_{t+1}\}_{t=1}^{\infty}} \left\{ \sum_{t=1}^{\infty} \delta^{t-1} \mathbb{E}_w \left[u(x_t, x_{t+1}, w_t) \mid d \right] \right\} \tag{1.14}$$

$$\text{subject to:}$$

$$x_{t+1} \in \mathcal{X}(x_t, w_t)$$

$$w_t \mid d \sim \mathbb{N}(\mu_0 + td, 1/\rho_w)$$

$$\text{given:}$$

$$x_1, w_1, \mu_0, d, \rho_w \tag{1.15}$$

In all time periods, a Perfect Foresight Planner solves the Bellman equation:

$$V_t^{\mathcal{F}}(x_t, w_t \mid d) \equiv \max_{x_{t+1}} \left\{ u(x_t, x_{t+1}, w_t) + \delta \mathbb{E}_w \left[V_{t+1}^{\mathcal{F}}(x_{t+1}, w_{t+1} \mid d) \mid d \right] \right\} \tag{1.16}$$

$$\text{subject to:}$$

$$x_{t+1} \in \mathcal{X}(x_t, w_t)$$

$$w_{t+1} \mid d \sim \mathbb{N}(\mu_0 + (t+1)d, 1/\rho_w)$$

$$\text{given: } x_t, w_t, \mu_0, d, \rho_w$$

Bellman equation 1.16 gives rise to policy functions $\pi_t^{\mathcal{F}}(x_t, w_t \mid d)$. Like the value function, these are indexed by t to indicate the dependence on calendar time, which arises due to the gradual climate change. The Perfect Foresight Planner makes optimal decisions for any "true" value of the drift parameter, d. But what is the correct value of d? To facilitate comparison with the other planner types, we assume that the true value of d is drawn from the prior beliefs, given in equation 1.13. Thus, the expected value of a Perfect Foresight Planner's economic program, from the perspective of period 1, is:

$$\mathbb{E}_d \left[V_1^{\mathcal{F}} \left(x_1, w_1 \mid d \right) \right] \equiv \mathbb{E}_d \left[\sum_{t=1}^{\infty} \delta^{t-1} \mathbb{E}_w \left[u(x_t, x_{t+1}, w_t) \mid d \right] \right] \tag{1.17}$$

subject to:

$$x_{t+1} = \pi_t^{\mathcal{F}}(x_t, w_t \mid d)$$

$$w_t \mid d \sim \mathbb{N}(\mu_0 + (t+1)d, 1/\rho_w)$$

$$d \sim \mathbb{N}(\underline{d}_1, 1/\rho_d)$$

The value of a Perfect Foresight planner's economic program represents the best possible outcome if the trajectory of climate is known, evaluated ex-ante of the prior distribution of climate change scenarios. It is derived by supposing the trajectory of climate (given by drift parameter d) is unknown ex-ante, but subsequently all actions are based as if d is known with certainty. This represents the best case scenario for uncertain climate changes, and is in practice infeasible, so it provides an upper bound on possible outcomes.

1.3.2 Sophisticated Planners

A Sophisticated planner takes the best feasible course of action given climate change uncertainty and update beliefs about climate in a standard Bayesian fashion as they experience weather. Their economic program is identical to a Perfect Foresight planner's, save for the maximization being taken outside of climate expectations. Denote the allocation mechanism for sophisticated planners as $\mathcal{S}$.

$$V_1^{\mathcal{S}}(x_1, w_1) \equiv \max_{\{x_{t+1}\}_{t=1}^{\infty}} \left\{ \mathbb{E}_d \left[\sum_{t=1}^{\infty} \delta^{t-1} \mathbb{E}_w \left[u(x_t, x_{t+1}, w_t) \mid d \right] \right] \right\} \tag{1.18}$$

subject to:

$$x_{t+1} \in \mathcal{X}(x_t, w_t)$$

$$w_t \mid d \sim \mathbb{N}(\mu_0 + td, 1/\rho_w)$$

$$d \sim \mathbb{N}(\underline{d}_1, 1/\rho_d)$$

given:

$$x_1, w_1, \mu_0, \underline{d}_1, \rho_w, \rho_d$$

Solving a Sophisticated planner's objective requires deriving their updated beliefs over climate as they experience weather. Thus in every period we must derive the posterior distribution of d given the history of weather W^t. Because this problem has been constructed to have self-conjugate priors, there is a closed-form solution for the posterior in all periods. Specifically, $d \mid W^t \sim \mathbb{N}(\overline{d}_t, 1/\overline{\rho}_{d,t})$ with (proof in appendix E):

$$\overline{d}_t = \frac{\rho_d \underline{d}_1 + \rho_w \sum_{j=1}^{t} j(w_j - \mu_0)}{\overline{\rho}_{d,t}} \tag{1.19}$$

$$\overline{\rho}_{d,t} = \rho_d + \left(\frac{t(t+1)(2t+1)}{6} \right) \rho_w \tag{1.20}$$

The posterior distribution can also be written recursively (see appendix F). The recursive derivation is advantageous because it allows the posterior to written as a function of two state variables: the time period t and the prior mean $\underline{d}_t$. Specifically, the posterior mean of d at time t, $\overline{d}_t$, is given by:

$$\overline{d}_t = \frac{t(w_t - \mu_0)\rho_w + \underline{d}_t\left(\rho_d + \frac{t(t-1)(2t-1)}{6}\rho_w\right)}{\rho_d + \left(\frac{t(2t+1)(t+1)}{6}\right)\rho_w} \tag{1.21}$$

This allows the Sophisticated Planner's problem be written recursively in a Bellman equation. In any time period, a Sophisticated Planner solves:

$$V_t^{\mathcal{S}}(x_t, w_t, \underline{d}_t) = \max_{x_{t+1}} \left\{u(x_t, x_{t+1}, w_t) + \delta\mathbb{E}_d\left[\mathbb{E}_w\left[V_{t+1}^{\mathcal{S}}(x_{t+1}, w_{t+1}, \overline{d}_t) \mid d\right]\right]\right\} \tag{1.22}$$

subject to:

$$x_{t+1} \in \mathcal{X}(x_t, w_t)$$

$$\overline{d}_t = \frac{t(w_t - \mu_0)\rho_w + \underline{d}_t\left(\rho_d + \frac{t(t-1)(2t-1)}{6}\rho_w\right)}{\rho_d + \left(\frac{t(2t+1)(t+1)}{6}\right)\rho_w}$$

$$w_{t+1} \mid \overline{d}_t \sim \mathbb{N}(\mu_0 + (t+1)\overline{d}_t, 1/\rho_w)$$

given:

$$x_t, w_t, \mu_0, \underline{d}_1, \rho_w, \rho_d$$

The policy functions for a sophisticated planners are denoted $\pi_t^{\mathcal{S}}(x_t, w_t, \underline{d}_t)$, which depends upon the prior mean of d in addition to the stock and weather state variables. Unlike the other planner types, it is unnecessary to evaluate the value of a Sophisticated planner's economic program over the prior distribution of climate, given that this uncertainty is built-in to their objective.

1.3.3 Unresponsive Planners

The Unresponsive planner incorrectly assumes that the climate in every period is equal to μ_0. Denote the allocation mechanism for unresponsive planners as $\mathcal{U}$. The objective for the Unresponsive planner is:

$$V_1^{\mathcal{U}}(x_1, w_1) \equiv \max_{\{x_{t+1}\}_{t=1}^{\infty}} \left\{ \sum_{t=1}^{\infty} \delta^{t-1} \mathbb{E}_w \left[u(x_t, x_{t+1}, w_t) \right] \right\} \tag{1.23}$$

subject to:

$$x_{t+1} \in \mathcal{X}(x_t, w_t)$$

$$w_t \sim \mathbb{N}(\mu_0, 1/\rho_w)$$

given:

$$x_1, w_1, \mu_0, \rho_w$$

In any period t, the Unresponsive planner solves the Bellman equation:

$$V_t^{\mathcal{U}}(x_t, w_t) \equiv \max_{x_{t+1}} \left\{ u(x_t, x_{t+1}, w_t) + \delta \mathbb{E}_w \left[V_{t+1}^{\mathcal{U}}(x_{t+1}, w_{t+1}) \right] \right\} \tag{1.24}$$

subject to:

$$x_{t+1} \in \mathcal{X}(x_t, w_t)$$

$$w_{t+1} \sim \mathbb{N}(\mu_0, 1/\rho_w)$$

given:

$$x_t, w_t, \mu_0, \rho_w$$

The Unresponsive planner's policy functions are denoted $\pi_t^{\mathcal{U}}(x_t, w_t)$.[6] The expected value of the Unresponsive planner's economic program is taken over the same expectation as

[6]The Unresponsive planner's objective is time-autonomous, so technically their value and policy functions need not be indexed by t. We leave the time index here in preparation for the empirical analysis, where it is necessitated by a non-infinite time-horizon.

the other planner types: The prior distribution of climate.

$$\mathbb{E}_d\left[V_1^{\mathcal{U}}(x_1, w_1)\right] \equiv \mathbb{E}_d\left[\sum_{t=1}^{\infty} \delta^{t-1} \mathbb{E}_w\left[u(x_t, x_{t+1}, w_t) \mid d\right]\right] \tag{1.25}$$

$$\text{Subject to: } x_{t+1} = \pi_t^{\mathcal{U}}(x_t, w_t)$$

$$w_t \mid d \sim \mathbb{N}(\mu_0 + (t+1)d, 1/\rho_w)$$

$$d \sim \mathbb{N}(\underline{d}_1, 1/\rho_d)$$

The Unresponsive planner's actions suppose the climate in every period is equal to μ_0, while the value of their economic program is evaluated over weather drawn from the true sequence of climate $\mu_t = \mu_0 + td$, where, as above, we assume d is drawn from the prior distribution given in equation 1.13. The Uniformed planner provides a lower bound on the reasonable value of a program under climate change, as they simply fail to respond to the change that is happening before their eyes. Their actions are incompatible with their incentives, because there is no cost of adaptation.

1.3.4 Defining Effects of Continuous and Uncertain Climate Change

At last we may define the effects of continuous and uncertain climate changes. First we must define the *ex-ante* full-effect of climate change.

Definition 3. *The* ***ex-ante full effect*** *of an uncertain climate change is the difference between the value of the Perfect Foresight planner's economic program, and the value of a baseline scenario where the climate change does not occur.*

$$\textit{Ex-Ante Full Effect} = \mathbb{E}_d\left[V_1^{\mathcal{F}}(x_1, w_1 \mid d)\right] - V_1^{\mathcal{U}}(x_1, w_1) \tag{1.26}$$

The baseline value is given by the unresponsive planner's value function. It is the optimized value when the climate does not change and is not integrated over the prior distribution of climate.

As in the previous section, the full effect of the change is segmented into parts. Here, the optimized value with climate change is infeasible given that climate changes are uncertain. This gives rise to the value of information.

Definition 4. *The* ***ex-ante full effect*** *of an uncertain climate change is equal to the sum of the* ***value of information****, the* ***value of adaptation****, and a* ***direct effect*** *without information or adaptation.*

$$\textit{Ex-Ante Full Effect} = \underbrace{\left(\mathbb{E}_d\left[V_1^{\mathcal{F}}\right] - V_1^{\mathcal{S}}\right)}_{\textit{Value of Information}} + \underbrace{\left(V_1^{\mathcal{S}} - \mathbb{E}_d\left[V_1^{\mathcal{U}}\right]\right)}_{\textit{Value of Adaptation}} + \underbrace{\left(\mathbb{E}_d\left[V_1^{\mathcal{U}}\right] - V_1^{\mathcal{U}}\right)}_{\textit{Direct Effect}}$$

(Dependence of each term on x_1 and w_1 are omitted for brevity.)

The value of information is weakly positive because for any given value of d, the Perfect Foresight planner's expectation of future weather is more accurate than the Sophisticated planner's. If information were not valuable, the Perfect Foresight planner would elect not to use it. The value of adaptation is weakly positive by a similar argument: The Sophisticated planner will out-perform the Unresponsive planner if their posteriors are informative, and will do just as well if their posteriors are not informative. There is no theoretical sign on the direct effect because climate change can either confer positive or negative consequences. The magnitude of all three Ex-Ante climate effects, in addition to the sign of the direct effect, depend upon the nature of the objective function and climate's impact on it, as well as the prior belief of climate change.

1.4 Adaptation to Climate Change in Global Fisheries

While our theoretical results can lend insight into a range of dynamic economic decision problems, here we focus on a particular example to illustrate and extend our main contribution. We employ the model developed above to estimate the effects of climate change, and the role of adaptation, in global marine fisheries. This is an ideal test-bed for our theory of dynamic adaption for several reasons. First, fisheries management is a dynamic problem in the truest sense because fish populations expand and contract over time due to fishing pressure and other naturally-occurring factors. Second, that climate change affects fish populations is widely acknowledged; warming of the world's oceans has already had a measurable impact on marine fisheries productivity. Third, commercial fisheries are often managed with the clearly defined objectives, making it straightforward to map our model to the real-world optimization problem being addressed. Finally, global data on the status, effects of climate, and policy actions in the world's commercial fisheries have recently been compiled by fishery scientists.

1.4.1 Estimation

The starting point of our analysis is a recent study, Free et al. (2019), that derives the effect of warming sea-surface temperatures on fisheries productivity over the period from 1930 to 2010. Free et al. estimates temperature-dependent biological growth models for 235 individual fisheries, representing about 33% of global fish catch, finding that on average productivity declined by 4.1% between 1930 and 2010. We use these estimated growth models as the basis for 235 individual structural models for each of the Perfect Foresight, Sophisticated and Unresponsive planner types. Our ultimate goal is to estimate the effects of uncertain, gradual climate change (definition 4) for each fishery, disentangling the values of adaptation, information, and direct effects.

Free et al. (2019) specify a temperature dependent Pella-Tomlinson production function. Translating their notation to match ours, let x_{it} be biomass and h_{it} be the harvest of fishery stock i at time t. Weather w_{it} is stock i's annual mean sea-surface temperature, the basis of the temperature-dependent productivity. The stock evolves between periods as:

$$g(h_{it}, x_{it}, w_{it}) = x_{it} + \frac{r_i x_{it}}{p} \left(1 - \left(\frac{x_{it}}{k_i}\right)^p\right) \exp(\theta_i w_{it}) - h_{it}$$

Although the preferred specification of Free et al. (2019) includes an auto-regressive productivity shock, or "process error", we omit that source of uncertainty in our numerical analysis, because it is inconsequential to our focus on climate change effects and adds an extraneous state variable to the dynamic programming exercise. The variables and parameters of the fisheries production model are further described in Table 1.1.

We adopt an objective that is consistent with the way most fisheries are actually managed: To maximize the sum of discounted harvests over time. Save for the discount rate being applied, this objective is consistent with the principle of maximum sustained yield (MSY), a mainstay of fisheries policy and management.

To maintain consistency with the data on fishery productivity and climate change, we focus on the historical record between 1930 and 2010. First, we estimate climate effects *ex-ante*, from the perspective of 1930, exactly as described in the theory section, by integrating over climate uncertainty given by the prior belief of temperature drift d. Second, we estimate climate effects *ex-post*, using the historical record of ocean temperatures that actually occurred. We also explore the consequences of the timing of information by deriving policies where harvest is chosen after the weather is revealed in each period

Table 1.1: Variables & Parameters

Category	***Variable***	***Description***
Control Variables	h_t	harvest
State Variables	x_t	stock biomass
	w_t	sea-surface temp. (SST)
	t	year
Parameters	r	intrinsic rate of growth
	k	carrying capacity
	θ	temperature-dependence
	p	Pella-Tomlinson shape
	δ	discount factor
Climate Variables	μ_0	mean temperature in 1930
	ρ_w	precision of weather
Priors	$\underline{d}_1$	prior mean of mean temperature drift d
	ρ_d	precision of prior

and where harvest is chosen beforehand.

Our analysis requires we estimate several climate parameters beyond those estimated by Free et al. (2019). These are: μ_0, the average temperature in period zero (in this case the year 1929), and ρ_w, the constant precision of weather. Using Free et al.'s data, we estimate these parameters by estimating a linear time-trend in sea-surface temperature for each fishery stock. We take μ_0 as the models' predicted temperature in 1930, and ρ_w as the inverse of the residual variance. We also use the estimated trend in sea-surface temperature in our ex-post analysis of climate effects, where it serves as the basis for the Perfect Foresight planner's policy.

We also must supply prior climate beliefs, $\underline{d}_1$ and ρ_d; A subjective matter. The choice of priors is particularly consequential for the *ex-ante* climate change effects, because these effects are integrated across the prior distribution of climate beliefs. Priors are less consequential to *ex-post* climate effects, where they only affect the sophisticated planner. Evaluated *ex-post*, the prior affects the magnitudes of the value of information and the value of adaptation relative to one another, but does not affect their sum, nor does it affect the direct effect. In other words, the closer prior climate beliefs are to what turns out to happen, the larger is the value of adaptation and the smaller is the value of information. In the end, we set $\underline{d}_1 = 0.0$ and $\rho_d = 10000.0$, consistent with the view that in 1930, the mean ocean temperature was not thought to be changing, and that in 100 years, the mean temperature was expected to change by no more than plus or minus 2.57 degrees with 99% confidence.

For each fishery we solve the full dynamic program over a finite horizon (1930-2010), obtaining policy and value functions in each period. Model specification directly mirrors the theory of uncertain climate changes and is omitted here to avoid repetition. Details are provided in Appendix G. This procedure gives rise to an expected present value of the fishery, from the perspective of 1930, for each planner type. We then compare these values across planner types to estimate the overall effect of climate change on that fishery, and to disaggregate it into the value of information, value of adaptation, and the direct effect of climate change, as shown in Definition 4. The numerical analysis is implemented in the Julia programming language and is available on GitHub. Despite the relative complexity of our model, the program runs quickly on a standard desktop computer.

1.4.2 Results

Figure 1.2 shows the results for all 235 fisheries. Blue dots are the direct effect of climate change, as a % of the baseline net present value (NPV) without climate change, orange dots are the value of adaptation, and green dots are the value of information. By displaying these as a percent of baseline fishery value, we are able to convey that fisheries facing no climate effects (those near $\theta = 0$ on the horizontal axis) should enjoy a very small value of adaptation, information, and direct effect. For fisheries to the right ($\theta > 0$), warming sea surface temperature improves their productivity over time. For these fisheries, all three effects are expected to be positive. Fisheries to the left ($\theta < 0$), are negatively affected by warming. For these fisheries, the direct effect is negative if the ocean does in fact warm. Figure 1.2 displays results for four modelling scenarios. The first (scenario A), matches our theory exactly. That is, harvest is chosen *ex-post* of the weather realization each period, and the effects of climate change are evaluated *ex-ante* (from the perspective of 1930). Scenario B entails the same policies, however climate effects are evaluated *ex-post* of the actual historic record of temperatures. In scenarios C and D, harvest is chosen *ex-ante* of the weather realization each period, and climate change effects are evaluated *ex-ante* and *ex-post* correspondingly. Aggregate results across all fishery stocks are provided in Table 1.2. These are calculated by summing across all fishery stocks (and thereby weighting by biomass).

The principal finding of our empirical analysis is that the estimated total effect of climate change is dominated by the direct effect and the values of adaptation and the values of information are comparatively small. This is reflected in Table 1.2, where we see that in the *ex-post* scenarios B and D, sea-surface temperature warming over the 1930-2010 period resulted in a 0.75 - 0.6 % decrease in the average net present value of these fishery stocks (this compares to the 4.1% average decrease in average MSY reported by Free et al. (2019)). The values of adaptation and information meanwhile are two orders of magnitude smaller. Although theory predicts the values of adaptation and information are weakly greater than zero, this only holds *ex-ante* of the actual weather realizations. In *ex-ante* scenarios A and C, these values are indeed positive, and close to zero. However for any particular sequence of weather, adaptation and information may not be beneficial, reflected in negative values in the *ex-post* scenarios B and D.

Our numerical results also illuminate several nuances within our theory of dynamic

Table 1.2: Aggregate Climate Effects By Modelling Scenario

Scenario	Direct Effect	Value of Adaptation	Value of Information	Total Effect
A	0.01683 %	0.00019 %	0.00046 %	0.01747 %
B	-0.75564 %	0.00109 %	-0.00106 %	-0.75561 %
C	-0.01818 %	0.01187 %	0.03787 %	0.03156 %
D	-0.59377 %	-0.01286 %	-0.00629 %	-0.61292 %

Modelling scenarios correspond to Figure 1.2. Table values are aggregated over all fishery stocks, and are percentages relative to the baseline scenario without climate change.

adaptation. Recall that our theory predicts that *ex-ante*, the values of adaptation and information are both weakly positive and the direct effect takes either sign. In the *ex-ante* scenarios (A and C), we do in fact estimate positive values of adaptation and information. As for the direct effect, the numerical results do not reveal a clear relationship with the temperature dependence parameter θ. This is because these effects are evaluated over the prior distribution of climate, where we chose a prior mean of zero mean temperature drift. If we instead had a prior belief of warming temperatures, the expected trend would emerge. This is reflected in the *ex-post* scenarios (B and D), where the historic record of sea-surface temperature did in fact entail warming temperatures for most stocks. As a result, the direct effect is positively correlated with θ.

Our findings regarding the timing of information are intuitive. In scenarios A and B, harvest is chosen each period before the realization of weather. In scenarios C and D, harvest is chosen beforehand. In the later case, climate (the distribution of weather) has a greater influence, reflected in the greater magnitudes of the values of adaptation and the values of information between scenario C versus A, and scenario D versus B. Although we did not explore this distinction theoretically, it appears to be true that policies with actions chosen *ex-ante* of weather shocks will enjoy greater values of adaptation.

1.5 Conclusion

Climate change has already impacted most economic sectors, and these effects are expected to increase in the coming decades. The consequences of these changes will depend on adaptive and learning strategies made dynamically as they occur. The purpose of this paper has been to show that despite the complexity of these dynamic decision problems, it is possible and relatively straightforward, to analytically sub-classify constituent parts of the net effects of climate changes. These include a direct effect, the value of adaptation, and when climate changes are uncertain, the value of information. In so doing, we essentially derive the "direct-effects" of weather and "belief-effects" of future weather, based on the model primitives of a standard economic decision problem. In that way, our theory gives credence to existing static analyses that assume existence of a reduced-form payout function wherein the climate and beliefs enter directly. We further demonstrated that insights gleaned from these previous reduced-form static analyses continue to hold in a fully-fledged structural dynamic setting.

Foremost is that there is zero value of adaptation to a marginal climate change, even in a dynamic decision environment. This is due to a dynamic envelope theorem that was first recognized by Euler. When an agent chooses an optimal policy in a dynamic decision environment, a marginal change to that policy must not change expected rewards, otherwise it would not have been optimal in the first place. It follows that the value of the optimal response, or adaptation, to an exogenous marginal climate change is precisely zero. Our theory of marginal climate changes showed this envelope result remains true when the exogenous marginal change in question is a shift in the climate that generates future weather.

Enlisting tools from the theory of natural capital, we examined the effects of discrete climate changes. The value of adaptation to a discrete climate change is weakly greater than zero, but its magnitude depends on the curvature of the value function. Our final theoretical contribution is a derivation of uncertain climate change effects using what Kala (2017) describes as the "workhorse" model of dynamic learning. A novel result is that the full effect of climate change includes the value of information.

A caveat to our analysis of uncertain climate change is that analytical traction requires that we make strong structural assumptions. Specifically, it assumes the mean of weather increases linearly and that the variance of weather is constant over time. As put succinctly by Kelly et al. (2005): *"A different structural model may give quantitatively different results on learning."*. If one wished to avoid such structural assumptions, it would be straightforward to do so at the cost of entangling the value of adaption and the value of information. That is, it would be straightforward to derive an upper-bound for the the value of adaption without specifying a model of learning, by simply deriving the value of Perfect Foresight and Unresponsive planner types. Although the perfect-foresight policy is unfeasible *ex-ante*, this proposition is less egregious then one might think, since the combined additional values of information and adaptation are minuscule anyhow.

Our empirical analysis illustrates the theory of uncertain climate change effects in the context of global fisheries management over the period of 1930-2010. For each of 235 individual fish stocks, we derive the value of three different planner types and construct the values of information, adaptation and a direct-effect. These results depend on structural assumptions and a subjective prior belief of future climate changes. Still, it appears to be generally true that a stationary policy function (or harvest control rule) "bakes-in" a considerable degree of climate adaptation. These results indicate that in the context of an optimized dynamic setting, most of the optimal response to climate change occurs without explicitly adapting, and that failing to capture adaptation costs and benefits in empirical estimates of climate change costs and benefits is hardly damaging.

References

Acemoglu, D. (2009). *Introduction to modern economic growth.* Princeton and Oxford: Princeton University Press.

Arrow, K. J., Dasgupta, P., & Mäler, K.-G. (2003). Evaluating projects and assessing sustainable development in imperfect economies. *Environmental and Resource Economics*, *26*(4), 647–685.

Auffhammer, M. (2018). Quantifying economic damages from climate change. *Journal of Economic Perspectives*, *32*(4), 33–52.

Britten, G. L., Dowd, M., & Worm, B. (2016). Changing recruitment capacity in global fish stocks. *Proceedings of the National Academy of Sciences*, *113*(1), 134–139.

Carleton, T. A., Jina, A., Delgado, M. T., Greenstone, M., Houser, T., Hsiang, S. M., ... others (2020). Valuing the global mortality consequences of climate change accounting for adaptation costs and benefits. *National Bureau of Economic Research*.

Costello, C., Polasky, S., & Solow, A. (2001). Renewable resource management with environmental prediction. *Canadian Journal of Economics/Revue canadienne d'économique*, *34*(1), 196–211.

Dasgupta, P., & Mäler, K.-G. (2000). Net national product, wealth, and social well-being. *Environment and development economics*, *5*(1), 69–93.

Deryugina, T., & Hsiang, S. (2017). The marginal product of climate. *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.

Free, C. M., Thorson, J. T., Pinsky, M. L., Oken, K. L., Wiedenmann, J., & Jensen, O. P. (2019). Impacts of historical warming on marine fisheries production. *Science*, *363*(6430), 979–983.

Gammans, M., Mérel, P., & Paroissien, E. (2020). Reckoning climate change damages along an envelope. *Agricultural & Applied Economics Association Annual Meeting*.

Guo, C., & Costello, C. (2013). The value of adaption: Climate change and timberland management. *Journal of Environmental Economics and Management*, *65*(3), 452–468.

Hotelling, H. (1931). The economics of exhaustible resources. *Journal of political Economy*, *39*(2), 137–175.

Hsiang, S. (2016). Climate econometrics. *Annual Review of Resource Economics*, *8*, 43–75.

Kala, N. (2017). Learning, adaptation, and climate uncertainty: Evidence from indian agriculture. *MIT Center for Energy and Environmental Policy Research Working Paper*, *23*.

Kelly, D. L., Kolstad, C. D., & Mitchell, G. T. (2005). Adjustment costs from environmental change. *Journal of Environmental Economics and Management*, *50*(3),

468–495.
Lemoine, D. (2017). Expect above average temperatures: Identifying the economic impacts of climate change. *National Bureau of Economic Research*.
Lemoine, D. (2018). Estimating the consequences of climate change from variation in weather. *National Bureau of Economic Research*.
Ludwig, D., & Hilborn, R. (1983). Adaptive probing strategies for age-structured fish stocks. *Canadian journal of fisheries and aquatic sciences*, *40*(5), 559–569.
Maunder, M. N., & Thorson, J. T. (2019). Modeling temporal variation in recruitment in fisheries stock assessment: A review of theory and practice. *Fisheries Research*, *217*, 71–86.
Mérel, P., & Gammans, M. (2018). Climate econometrics: Can the panel approach account for long-run adaptation?
Minto, C., Myers, R. A., & Blanchard, W. (2008). Survival variability and population density in fish populations. *Nature*, *452*(7185), 344–347.
Parma, A. M. (1990). Optimal harvesting of fish populations with non-stationary stock-recruitment relationships. *Natural Resource Modeling*, *4*(1), 39–76.
Rose, K. A. (2000). Why are quantitative relationships between environmental quality and fish populations so elusive? *Ecological Applications*, *10*(2), 367–385.
Rothschild, B. J. (2000). Fish stocks and recruitment: The past thirty years. *ICES Journal of Marine Science*, *57*(2), 191–201.
Smith, A., & Walters, C. J. (1981). Adaptive management of stock–recruitment systems. *Canadian Journal of Fisheries and Aquatic Sciences*, *38*(6), 690–703.
Stokey, N. L., & Lucas Jr., R. E. (1989). *Recursive methods in economic dynamics*. Harvard University Press.
Thorson, J. T., Jensen, O. P., & Zipkin, E. F. (2014). How variable is recruitment for exploited marine fishes? a hierarchical model for testing life history theory. *Canadian Journal of Fisheries and Aquatic Sciences*, *71*(7), 973–983.
Vert-Pre, K. A., Amoroso, R. O., Jensen, O. P., & Hilborn, R. (2013). Frequency and intensity of productivity regime shifts in marine fish stocks. *Proceedings of the National Academy of Sciences*, *110*(5), 1779–1784.
Walters, C., & Parma, A. M. (1996). Fixed exploitation rate strategies for coping with effects of climate change. *Canadian Journal of Fisheries and Aquatic Sciences*, *53*(1), 148–158.
Walters, C. J. (1981). Optimum escapements in the face of alternative recruitment hypotheses. *Canadian Journal of Fisheries and Aquatic Sciences*, *38*(6), 678–689.
Weitzman, M. L. (2002). Landing fees vs harvest quotas with uncertain fish stocks. *Journal of environmental economics and management*, *43*(2), 325–338.

1.6 Appendix

A Complete Set of Assumptions for Decision Problem 1.5

A plan is a trajectory of the capital stock, starting at x_1 and thereafter given by a recursive sequence of functions $\{x_{t+1} = \tilde{\mathbf{x}}[W^t], t \in \mathbb{N}^0\}$, where $W^t \equiv (w_1, \ldots, w_t)$ is the history of weather up to and including the current time period. A plan prescribes the capital stock in the next period, given the history of weather up to and including that period (as a convention $\tilde{\mathbf{x}}[W^{t-1}] = x_0$ (Acemoglu 2009)). A feasible plan satisfies the requirement that:

$$\tilde{\mathbf{x}}[W^t] \in \mathcal{X}(\tilde{\mathbf{x}}[W^{t-1}], w_t)$$

We maintain the minimal assumptions such that a unique feasible plan optimizes objective equation 1.5, and that it can be characterized by a continuously differentiable policy function $x^*_{t+1} = \pi(x_t, w_t)$.

A. The constraint correspondence $\mathcal{X} : \Omega_x \times \Omega_w \rightrightarrows \Omega_x$ is non-empty valued for all $x \in \Omega_x$ and $w \in \Omega_w$. Moreover, $\mathcal{X}$ is a convex set.

B. Period rewards $u(x_t, x_{t+1}, w_t)$ are concave in x_t, x_{t+1} and continuously differentiable in the interior of its domain.

C. Weather shocks are independent. That is, $\Pr(w_{t+1} \mid W^t) = \Pr(w_{t+1})$

D. (Transversality Condition) For all possible x_t, w_t and $\tilde{\mathbf{x}} \in \mathcal{F}(x, w)$ the limit of discounted expected utility $\lim_{T\to\infty} \mathbb{E}_w \left[\sum_{t=1}^{T} \delta^t u(\tilde{\mathbf{x}}[W^{t-1}], \tilde{\mathbf{x}}[W^t], w_t)\right]$ exists and is finite.

Under these assumptions, a unique and continuous optimal policy function $x^*_{t+1} = \pi(x_t, w_t)$ exits, as does the continuous value function $V(x_t, w_t)$. See for example Acemoglu (2009).

Back to Section 1.2.1.

B Derivation of Euler Equations

The Bellman Equation is:

$$V(x_t, w_t \mid C) = \max_{x_{t+1}} \{u(x_t, x_{t+1}, w_t) + \delta \mathbb{E}_w [V(x_{t+1}, w_{t+1} \mid C)]\} \tag{1.27}$$

$$\text{subject to:}$$
$$x_{t+1} \in \mathcal{X}(x_t, w_t)$$
$$w_{t+1} \sim f(C)$$

The first-order condition is:

$$\frac{d}{dx_{t+1}} = \frac{\partial u(x_t, x_{t+1}, w_t)}{\partial x_{t+1}} + \delta \frac{\mathbb{E}_w\left[V(x_{t+1}, w_{t+1} \mid C)\right]}{dx_{t+1}} = 0$$

The envelope theorem gives:

$$\frac{dV(x_t, w_t \mid C)}{dx_t} = \frac{\partial V(x_t, w_t)}{\partial x_t} = \frac{\partial u(x_t, x_{t+1}, w_t)}{\partial x_t}$$

Inserting this envelope result into the first-order condition yields (denoting $x^*_{t+1} \equiv \pi(x_t, w_t \mid C)$):

$$u_2(x_t, x^*_{t+1}, w_t) + \delta \mathbb{E}_w \left[u_1\left(x^*_{t+1}, x^*_{t+2}, w_{t+1}\right)\right] = 0, \;\; t \in \mathbb{N}^1$$

Which is equivalent to:

$$\mathbb{E}_w\left[u_2(x_t, x^*_{t+1}, w_t) + \delta u_1\left(x^*_{t+1}, x^*_{t+2}, w_{t+1}\right) \mid w_t\right] = 0, \;\; t \in \mathbb{N}^1$$

as stated in the text.

Back to Section 1.2.2.

C Derivation of Marginal Climate Change Effects

First note that it takes two time-periods for the climate C to permeate actions. That is:

$$\begin{aligned}
&x_1, w_1: \text{ given} \\
&x^*_2 = \pi(x_1, w_1), w_2 = w(C) && \implies dx^*_2/dC = 0 \\
&x^*_{t+1} = \pi(x^*_t, w(C)) && \implies dx^*_{t+1}/dC = \frac{\partial \pi(x^*_t, w(C))}{\partial x}\frac{dx^*_t}{dC} + \frac{\pi(x^*_t, w(C))}{\partial w} w'(C)
\end{aligned}$$

Where the last line holds for all $t \geq 2$. Now write the the value function an infinite sequence beginning at time 1:

$$V(x_1, w_1 \mid C) = u(x_1, x^*_2, w_1) + \sum_{t=2}^{\infty} \delta^t \mathbb{E}_w\left[u(x^*_t, x^*_{t+1}, w(C))\right]$$

$$\text{subject to:}$$

$$x^*_2 = \pi(x_1, w_1)$$

$$x^*_{t+1} = \pi(x^*_t, w(C)), \; t \geq 2$$

The derivative with respect to C is then:

$$\frac{dV(x_1, w_1 \mid C)}{dC} = \sum_{t=2}^{\infty} \delta^t \mathbb{E}_w \Bigg[u_3(x_t^*, x_{t+1}^*, w(C))w'(C) + \ldots$$

$$\ldots \big(u_2(x_t^*, x_{t+1}^*, w(C)) + \delta u_1(x_{t+1}^*, x_{t+2}^*, w(C)) \big) \frac{dx_{t+1}^*}{dC} \Bigg]$$

as is stated in the text. Back to section 1.2.3

D Equivalence of discrete and marginal climate change effects

The following lemma establishes equivalence of the responsive and unresponsive planners' allocation mechanisms when $C^a = C^b$.

Lemma 1. *If $C^a = C^b = C$, the responsive and unresponsive allocation mechanisms are identical and the values of their economic programs are equal to the optimal value function $V(\cdot, \cdot \mid C)$.*

$$C^a, C^b = C \implies V(x, w \mid \mathcal{U}(C^a, C^b)) = V(x, w \mid \mathcal{R}(C)) = V(x, w \mid C) \tag{1.28}$$

This holds by definition, since the allocation mechanism $\mathcal{U}(C^a, C^b)$ is optimal when the original environment C^a is equal to the new environment C^b.

Recall that C^a and C^b are scalars by assumption. First consider the direct effect without adaptation:

$$\begin{aligned}
&\lim_{C^b \to C^a} \left[\frac{V(x_1, w_1 \mid \mathcal{U}(C^a, C^b)) - V(x_1, w_1 \mid \mathcal{R}(C^a))}{C^b - C^a} \right] \\
&= \lim_{C^b \to C^a} \left[\frac{V(x_1, w_1 \mid \mathcal{U}(C^a, C^b)) - V(x_1, w_1 \mid \mathcal{U}(C^a, C^a))}{C^b - C^a} \right] \\
&= \left. \frac{dV(x_1, w_1 \mid \mathcal{U}(C^a, C^b))}{dC^b} \right|_{C^b = C^a} \\
&= \frac{dV(x_1, w_1 \mid \mathcal{R}(C^a))}{dC^a} \\
&= \frac{\partial V(x_1, w_1 \mid \mathcal{R}(C^a))}{\partial C^a}
\end{aligned}$$

The first equality holds by lemma 1 and the second by the definition of a gradient. The third equality applies lemma 1 once again and the final equality applies the envelope theorem.

As for the value of adaptation, again using lemma 1, we have:

$$\begin{aligned}
&\lim_{C^b \to C^a} \left[\frac{V\left(x_1, w_1 \mid \mathcal{R}(C^b)\right) - V\left(x_1, w_1 \mid \mathcal{U}(C^a, C^b)\right)}{C^b - C^a} \right] \\
&= \lim_{C^b \to C^a} \left[\frac{V\left(x_1, w_1 \mid \mathcal{U}(C^b, C^b)\right) - V\left(x_1, w_1 \mid \mathcal{U}(C^a, C^b)\right)}{C^b - C^a} \right] \\
&= \left. \frac{dV(x_1, w_1 \mid \mathcal{U}(C^a, C^b))}{dC^a} \right|_{C^a = C^b} \\
&= 0
\end{aligned}$$

Where the first equality follows from lemma 1, the second by the definition of a derivative and the third by the envelope theorem. The envelope theorem applies here because climate C^a influences the unresponsive planner's economic program solely through their actions. Since these actions are optimal when $C^a = C^b$, there is zero value of a marginal change.

E Posterior Likelihood of Climate

Here we derive the posterior distribution of climate given the previous history of weather. In particular interest to the Sophisticated planners is $P(d \mid W^t)$, as it informs their Bellman equation (1.22).

Let $w_t \mid d \sim \mathbb{N}(\mu_0 + td, 1/\rho_w)$ with $w_s \perp w_t$ for all $s \neq t$, and let $d \sim \mathcal{N}\left(\underline{d}_1, 1/\rho_d\right)$. Then $d \mid W^t \sim \mathbb{N}(\overline{d}_t, 1/\overline{\rho}_{d,t})$ where:

$$\begin{aligned}
\overline{d}_t &= \frac{\rho_d \underline{d}_1 + \rho_w \sum_{j=1}^t j(w_j - \mu_0)}{\overline{\rho}_{d,t}} \\
\overline{\rho}_{d,t} &= \rho_d + \left(\frac{t(t+1)(2t+1)}{6}\right)\rho_w
\end{aligned}$$

Proof. Bayes Theorem gives:

$$\begin{aligned}
P(d \mid W^t) &\propto P(W^t \mid d)P(d) \\
&= P(d) \textstyle\prod_{j=1}^t P(w_j \mid d)
\end{aligned}$$

where the second line follows by independence of weather. Plugging in and expanding the normal probability density functions gives:

$$P\left(d \mid W^t\right) \propto \exp\left(\frac{-1}{2}\left(\rho_d(d - \underline{d}_1)^2 + \rho_w \textstyle\sum_{j=1}^t (w_j - \mu_0 - jd)^2\right)\right)$$

Next completing the squares, consigning terms excluding d to the proportionality constant and reorganizing gives:

$$P\left(d \mid W^t\right) \propto \exp\left(\frac{-1}{2}\left(d^2\left(\rho_d + \rho_w \sum_{j=1}^t j^2\right) - 2d\left(\rho_d \underline{d}_1 + \rho_w \sum_{j=1}^t j(w_j - \mu_0)\right)\right)\right)$$

Noting that $\sum_{j=1}^t j^2 = t(t+1)(2t+1)/6$, define:

$$\begin{aligned}
\overline{\rho}_{d,t} &\equiv \rho_d + \left(\frac{t(t+1)(2t+1)}{6}\right)\rho_w \\
\overline{d}_t &\equiv \frac{\rho_d \underline{d}_1 + \rho_w \sum_{j=1}^t j(w_j - \mu_0)}{\overline{\rho}_{d,t}}
\end{aligned}$$

Plugging these into the above, multiplying everything by exp $\left(-\overline{\rho}_{d,t}(\overline{d}_t)^2/2\right)$, simplifying, and adding the proportionality constant, yields:

$$P\left(d \mid W^t\right) = \frac{\overline{\rho}_{d,t}}{\sqrt{2\pi}} \exp\left(\frac{-1}{2}\left(\overline{\rho}_{d,t}(d-\overline{d}_t)\right)^2\right)$$

which establishes that $d \mid W^t \sim \mathbb{N}(\overline{d}_t, 1/\overline{\rho}_{d,t})$. □

Back to section 1.3.2

F Recursive Derivation of the Posterior Likelihood of Climate

It is well know that the posterior likelihood of climate can be derived recursively. This formulation is advantageous numerically, since the posterior in any period can be derived from two variables: the prior mean $\underline{d}_t$ and the period index t.

For our purposes it suffices to show the recursion works for a given time period t. Let the prior in period t equal the posterior in period $t-1$ (i.e. $\underline{d}_t = \overline{d}_{t-1}$ and $\underline{\rho}_{d,t} = \overline{\rho}_{d,t-1}$). Using 1.19 and 1.20:

$$\underline{d}_t = \frac{\rho_d \underline{d}_1 + \rho_w \sum_{j=1}^{t-1} j(w_j - \mu_0)}{\underline{\rho}_{d,t}}$$

$$\underline{\rho}_{d,t} = \rho_d + \left(\frac{t(t-1)(2t-1)}{6}\right)\rho_w$$

Let $w_t \mid d \sim N(\mu_0 + td, 1/\rho_w)$, with μ_0, t, and ρ_w given. Suppose $d \sim N\left(\underline{d}_t, 1/\underline{\rho}_{d,t}\right)$. Then $d \mid w_t \sim N\left(\overline{d}_t, 1/\overline{\rho}_{d,t}\right)$, where:

$$\overline{d}_t = \frac{t(w_t-\mu_0)\rho_w + \underline{d}_t \underline{\rho}_{d,t}}{\overline{\rho}_{d,t}}$$

$$\overline{\rho}_{d,t} = t^2\rho_w + \underline{\rho}_{d,t}$$

Which establishes that t and $\underline{d}_t$ are sufficient to describe the Posterior likelihood over d (the precision is a function of t and parameters). Furthermore, the recursively-derived posterior is identical to the non-recursive derivation (equations 1.19 and 1.20).

Proof.

$$P(d \mid w_t) \propto P(w_t \mid d)P(d)$$

$$\propto \exp\left(\frac{\rho_w(w_t-\mu_0-td)^2 + \underline{\rho}_{d,t}(d-\underline{d}_t)^2}{-2}\right)$$

$$\propto \exp\left(\frac{d^2(\rho_w t^2 + \underline{\rho}_{d,t}) - 2d(\rho_w t(w_t - \mu_0) + \underline{\rho}_{d,t}\underline{d}_t)}{-2}\right)$$

Define:

$$\begin{aligned}
\overline{d}_t &\equiv \frac{t(w_t - \mu_0)\rho_w + \underline{d}_t\underline{\rho}_{d,t}}{\overline{\rho}_{d,t}} \\
\overline{\rho}_{d,t} &\equiv t^2\rho_w + \underline{\rho}_{d,t}
\end{aligned}$$

then using the same technique as the non-recursive derivation:

$$P(d \mid w_t) = \frac{\overline{\rho}_{d,t}}{\sqrt{2\pi}} \exp\left(\frac{\overline{\rho}_{d,t}(d - \overline{d}_t)^2}{-2}\right)$$

which establishes that $d \mid w_t \sim N\left(\overline{d}_t, 1/\overline{\rho}_{d,t}\right)$.

To write the moments of the posterior as a function solely of $\underline{d}_t$, t and parameters, simply plug in for $\underline{\rho}_{d,t}$ and simplify. That is:

$$\begin{aligned}
\overline{\rho}_{d,t} &= t^2\rho_w + \underline{\rho}_{d,t} \\
&= t^2\rho_w + \rho_d + \left(\frac{t(t-1)(2t-1)}{6}\right)\rho_w \\
&= \rho_d + \left(\frac{t(2t+1)(t+1)}{6}\right)\rho_w
\end{aligned}$$

and so:

$$\overline{d}_t = \frac{t(w_t - \mu_0)\rho_w + \underline{d}_t\left(\rho_d + \frac{t(t-1)(2t-1)}{6}\rho_w\right)}{\rho_d + \left(\frac{t(2t+1)(t+1)}{6}\right)\rho_w}$$

Straightforward algebraic manipulations establish equivalence between this and the non-recursive derivation, by plugging in $\underline{d}_t = \overline{d}_{t-1}$. □

Back to section 1.3.2

G Empirical Bellman Equations

Bellman Equations for the empirical methodology follow directly from the theory of uncertain climate change effects. We must simply insert the dynamic biomass equation provided by Free et al.'s analysis, as well as the objective of maximizing the discounted stream of harvests. A minor nuance is that the time horizon is now finite, spanning the years 1930-2010.

Perfect Foresight Planners

$$V_t^{\mathcal{F}}(x_t, w_t \mid d) = \max_{h_t} \left\{ h_t + \delta \mathbb{E}_w \left[V_{t+1}^{\mathcal{F}}(x_{t+1}, w_{t+1} \mid d) \mid d \right] \right\}$$
$$\text{subject to:}$$
$$0 \leq h_t \leq x_t + r(x_t, w_t)$$
$$x_{t+1} = x_t + r(x_t, w_t) - h_t$$
$$w_{t+1} \mid d \sim \mathbb{N}(\mu_0 + (t+1)d, 1/\rho_w)$$
$$\text{given:}$$
$$x_t, w_t, d, \rho_w$$

Denote the optimal policy function for Perfect Foresight planners as $h_t^* \equiv \pi_t^{\mathcal{F}}(x_t, w_t \mid d)$.

Sophisticated Planners

$$V_t^{\mathcal{S}}(x_t, w_t, \underline{d}_t) = \max_{h_t} \left\{ h_t + \delta \mathbb{E}_d \left[\mathbb{E}_w \left[V_{t+1}^{\mathcal{S}}(x_{t+1}, w_{t+1}, \overline{d}_t) \mid d \right] \right] \right\}$$
$$\text{subject to:}$$
$$0 \leq h_t \leq x_t + g(x_t, w_t)$$
$$x_{t+1} = x_t + g(x_t, w_t) - h_t$$
$$\overline{d}_t = \frac{t(w_t - \mu_0)\rho_w + \underline{d}_t \left(\rho_d + \frac{t(t-1)(2t-1)}{6} \rho_w \right)}{\rho_d + \left(\frac{t(2t+1)(t+1)}{6} \right) \rho_w}$$
$$w_{t+1} \mid \overline{d}_t \sim \mathbb{N}(\mu_0 + (t+1)\overline{d}_t, 1/\rho_w)$$
$$d \sim \mathbb{N}(\underline{d}_1, 1/\rho_d)$$
$$\text{given:}$$
$$x_t, w_t, \mu_0, \underline{d}_1, \rho_w, \rho_d$$

Unresponsive Planners

$$V_t^{\mathcal{U}}(x_t, w_t) = \max_{h_t} \left\{ h_t + \delta \mathbb{E}_w \left[V_{t+1}^{\mathcal{U}}(x_{t+1}, w_{t+1}) \mid d = 0 \right] \right\}$$
$$\text{subject to:}$$
$$0 \leq h_t \leq x_t + r(x_t, w_t)$$
$$x_{t+1} = x_t + r(x_t, w_t) - h_t$$
$$w_{t+1} \sim \mathbb{N}(\mu_0, 1/\rho_w)$$
$$\text{given:}$$
$$x_t, w_t, \mu_0, \rho_w$$

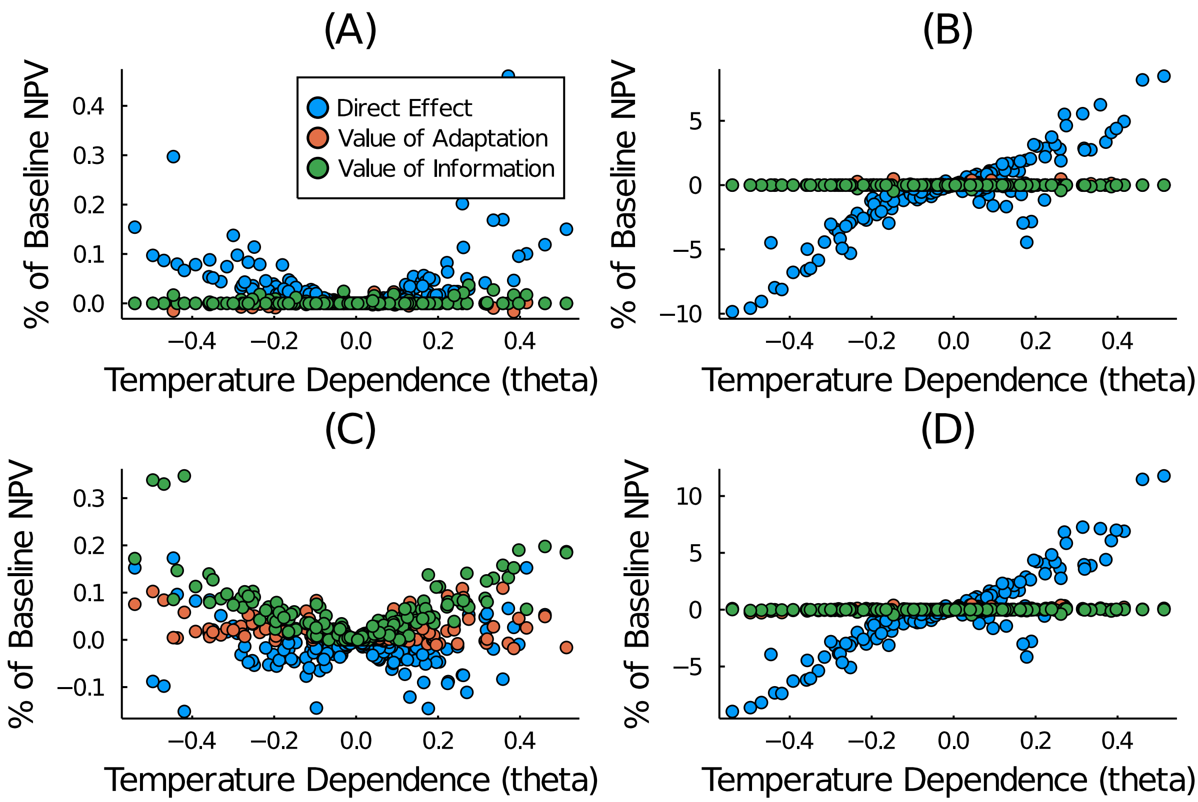

Figure 1.2: Climate Effects By Fishery Stock. The horizontal axis is the estimated effect of sea-surface temperature on productivity. Negative values indicate a negative effect of warming on productivity, and vice versa. The vertical axis is scaled by the fishery stocks' baseline net present value (NPV) to aid comparison between stocks. (A) Matches the theory section: climate effects evaluated *ex-ante* and harvest chosen *ex-post* of the weather realization each period. (B) Harvest chosen *ex-post* of the weather, simulated using the actual historic weather from 1930-2010. (C) Climate effects evaluated *ex-ante* and harvest is chosen *ex-ante* of the weather realization each period. (D) Harvest is chosen *ex-ante* of the weather, simulated using the actual historic weather from 1930-2010.

Chapter 2

Irrigation By Appropriation

2.1 Introduction

The prior appropriation doctrine's seniority-based rights system is considered an inefficient means for allocating water resources. As evidence, Burness and Quirk (1979) provides an eloquent theoretical depiction of prior appropriation and its failure to allocate a stochastic water supply optimally. The inefficiency arises because water allocations are fixed ex-post of their realization, meaning water users cannot adjust consumption as to equate their marginal profits. While Burness and Quirk (1979) appears frequently in critiques of the prior appropriation doctrine, it also proves the Coasian argument that a perfectly competitive market for appropriative water rights is efficient: *"Most of the allocative problems associated with the appropriative doctrine would be eliminated if water rights could be freely transferred or sold"* (Burness & Quirk 1979). As currently practiced in the Western United States, efficiency of prior appropriation is ambiguous because water rights are traded, but doing so incurs significant transaction costs (Brewer, Glennon, Ker, & Libecap 2006; Hagerty 2019b; Howe & Goemans 2003). The region studied in this paper, the state of Colorado, has a particularly active water market (Womble & Hanemann 2020), comprising permanent sales of water rights, single-year leases and informal short-term leases (Brown 2006; Colorado Water Conservation Board 2020). Whether prior appropriation is inefficient in practice is therefore an open empirical question, and one of grave importance. With drought, increasing demands, and new emphasis on environmental water uses straining water supplies in the Western United States, there is no room for waste.

In this paper I empirically test a necessary condition for allocative efficiency, known as the "independence property" (Fowlie & Perloff 2013), in the context of irrigated agriculture and appropriative water rights in the State of Colorado. Essentially a restatement of the Coase theorem, the independence property holds if water consumption is independent of the initial allocation of water among firms, conditional on the aggregate allocation of water to all firms. I show that the classic difference-in-difference research design estimates the independence property precisely; Intuitively, due to its identification based on relative differences in, not absolute magnitudes of, water allocations. The independence property therefore allows a simple, reduced-form, test of allocative efficiency under prior appropriation in situ, that does not require direct observation of firms' marginal profits nor does it require estimating them structurally. Specifically, I identify the causal effect of individual water allocations, prescribed by prior appropriation's seniority rule, on a proxy for irrigated agricultural yields (the remotely-sensed "greenness-index" NDVI), while statistically controlling for the aggregate allocation of water. To confront the threat to identification posed by endogeneity between water demand and allocations, I exploit the stochastic, idiosyncratic timing of water curtailments for identifying variation in water allocations. Although curtailments are expected and planned for, their exact timing are unknowable beforehand. To capture this high-frequency identifying variation, I compare changes in the progress of irrigated crops at a roughly weekly time increment to changes in the priority status of appropriative water rights. I control for the crop being grown, the type of irrigation used, and the weather. In a robustness check, I also allow for locality-based heterogeneity in physical water access (and spatially correlated unobservables) by comparing directly-adjacent water users. Under the null hypothesis that the independence property holds, relative differences in water allocations elicit zero effect on relative differences in NDVI. Across all specifications, I find their effect is significantly positive, meaning that relatively greater water allocations lead to relatively greener crops. By rejecting the null hypothesis of independence, I find that the prior appropriation doctrine, as currently practiced in Colorado, leaves efficiency gains unrealized.

This paper contributes new causal evidence to a literature that analyzes the efficiency of water use under prior appropriation. Perhaps the most compelling existing evidence of inefficiency is due to analyses of water right transaction prices, that document wide price dispersion, especially between agricultural, municipal and industrial water uses (Brewer et al. 2006; Hagerty 2019b). Water markets are thinly-traded and opaquely operated, making it difficult to obtain quality economic data and establish causal evidence of inefficiency, although this itself is suggestive evidence of inefficiency. Contemporary research increasingly employs remote-sensing data to overcome these empirical challenges. While water-rights transactions are made privately, administrative records of water rights ownership are public-domain. In combination with remotely-sensed agricultural data, researchers now have powerful means for analyzing agricultural water use. Curiously, research in this vein documents several adaptations to water scarcity risk that may mute allocative inefficiency under prior appropriation. It has been shown that crop-choice is used as an adaptive strategy to water scarcity risk (Hagerty 2019a; Ji & Cobourn 2018; Manning, Goemans, & Maas 2017), meaning that agricultural producers make cropping decisions in expectation of water curtailments and thus mediate the effect of curtailments on crop yields. Ji and Cobourn (2018) showed that irrigation organizations buffer water scarcity risk under prior appropriation by holding portfolios of appropriative water rights, with differing in seniority, and allocating water to their members proportionally. Likewise Mukherjee and Schwabe (2014) showed the value of having water-right portfolios across multiple water sources. This collection of evidence, that appropriative water rights are traded, that producers self-select the water intensity of their operations via cropping decisions and that water rights are commonly pooled and shared proportionally, all represent departures from the stark textbook model of prior appropriation that Burness and Quirk (1979) used to demonstrate its inefficiency. Taken together, they indicate that the inefficiency of prior appropriation in practice is nuanced by adaptive behaviors explicitly taken to mollify it.

Several recent papers address also address prior appropriation's inefficiency but do so using different methodologies than I do in this paper. Smith (2021) leverages a natural experimental setting to compare the relative merits of prior appropriation and an alternative system of proportional sharing, and concludes a higher marginal value of water emerges with proportional sharing. Arellano-Gonzalez et al. (2021) takes a structural approach to estimate the costs of water curtailments in California agriculture under simulations of prior appropriation's seniority rule and a market environment, finding that costs are significantly higher in the former. While my reduced-form methodology has the benefits of causal identification and relatively-mild identifying assumptions, Arellano-Gonzalez et al.'s simulation exercise allows them to estimate the magnitude of the cost of inefficient water use- a key advantage compared to my approach where this is not feasible.

Legal institutional scholorship advances several mechanisms underlying prior appropriation's persisting inefficiency. Colorado became the first State to adopt prior appropriation as its water allocation institution with the adoption of its Constitution in 1876. This legislation cemented existing appropriative claims made for irrigated agriculture and ushered in an era of agricultural development. Already in need of water to support Colorado's growing population, a 1891 ruling by Colorado's Supreme Court allowed agricultural water rights to be detached from the land they irrigated, subject to a clause that they not injure third parties. Although this development seemingly promoted water-markets, by encompassing alternative water uses and allowing trade, it also created lasting perverse incentives.

The ability to sell an appropriative claim gave a strong incentive to develop and maintain water rights for the prospect of future sales. As a result, prior appropriation overly-incentivized the development of irrigation infrastructure. In 1911, Coman documented that hopeful Colorado irrigators had made claims far in excess of available supplies. *"The early settlers were accustomed to put in claims for a water supply far in excess of their needs, and the sum total of these claims was often in excess of the maximum output of the stream. The important water resources of the state are thus pledged to two and three times their utmost capacity, and the land now being developed can only be provided with water by the construction of mammoth reservoir."* (Coman 1911).

The beneficial use standard emerged as a solution to the problem of speculative claims. Claims to water out-sized actual supplies in part because they were made long in advance of the infrastructure that would be required to exercise them. Under the beneficial use standard, these claims were deemed as "conditional" and only perfected once water was physically diverted and put to "beneficial" use. While solving a temporary problem, the beneficial use clause created several others that have remained since. No longer was prior-appropriation an autonomous institution without need for regulatory oversight. Colorado's water court was now tasked with tabulating water rights and checking that the water was actually being used under "reasonable" irrigation practices. More subtly, the nature of property rights for water changed from the absolute ownership enjoyed by the minors, to a usufruct right for certain agricultural practices.

These century-old developments in Western water law, here documented in Colorado but replicated throughout the West, left lasting perverse incentives that hinder the efficiency of water allocations today (Libecap 2011). Prominent among these are the beneficial-use requirement and the no-injury clause of water right transfers. The beneficial use clause is now synonymous with *"use-it or lose-it"*. Perfecting water rights through use introduces a perverse incentive to consume more water than would be otherwise be optimal in effort to secure future allocations (MacDonnell, Howe, Rice, et al. 1990; Tarlock 2001). One example is the low adoption of water-efficient irrigation systems. Water saved through efficiency upgrades, called "salvage-water", typically cannot be repurposed or sold. This is thought to be the reason why the original method of flood irrigation, now outdated by more efficient sprinkler systems, is still widely practiced (Huffaker, Whittlesey, & Hamilton 2000).

The no-injury clause of water right trades further barriers to gains from trade. When a crop is irrigated, a portion of the irrigation water applied percolates through the soil into the water table. These "return-flows" intertwine appropriative water rights, as one irrigator's return flow becomes the source of another's water right. The no-injury standard addresses the return-flow problem by limiting the quantity of water eligible for trade to "consumptive" use (the amount diverted less return flows). An analysis of the right's historic average use typically serves as the basis for this calculation (MacDonnell 2015). The return flow problem and its workaround, the no-injury standard, are widely seen as a barrier to the reallocation of water through trade (Thompson Jr 1993). They increase the transaction costs of water right trades, because the legal onus falls to the trading parties to prove that no-injury will occur as the result of a water right transaction. Much of the expense of water right trades is due to sorting out the complex hydrology of return-flows in effort to prove no-injury. A more nuanced issue is due to resulting emphasis on historic average use. Basing a water right's claim to future allocations upon its historic usage reinforces the perverse incentive to *use-it-or-lose-it* and distances prior appropriation from the theoretically optimal rights system of state-contingent claims (Berck & Lipow 1994).

The institutional shortcomings of prior appropriation described here provide potential mechanisms behind this paper's findings. Although appropriative rights can be traded in principle, the cost of doing so is preventative. Short-term water leases are dissuaded by the incentive to maximize historic average use, for the prospect of selling water rights in the future. These obstacles to efficient water re-allocation potentially explain my finding that irrigation decisions in Colorado are made on the basis of individual water allocations, despite that water can be traded in principle.

In summary, this paper contributes a reduced-form test of allocative efficiency under prior appropriation. High frequency remote-sensing data are paired with daily water allocations to provide a natural experimental setting with precise statistical controls for potentially confounding endogeniety between water demand and allocations. In a robustness check to the main specification, a spatial-first-difference specification compares outcomes between water users that are directly adjacent to one another, growing the same crop and using the same type of irrigation system. The analysis intentionally favors the null hypothesis of independence, but uniformly rejects it. Its focus is within the agricultural sector, meaning it disregards the price dispersion that is known to exist between agricultural and urban water uses. Furthermore, it controls for the water scarcity adaptations previously documented within the agricultural sector. The analysis is *ex-post* of cropping decisions and occurs within Colorado's institutional setting wherein irrigation organizations hold portfolios of water rights. These strategies are known to be used to mitigate water scarcity risk and "round-the-corners" off of the starkest application of prior appropriation's seniority rule. In spite of this narrow focus and precise controls, I uniformly reject the null hypothesis of independence between individual allocations and water consumption. This evidence suggests that frictions and perverse-incentives inhibit gains from water re-allocation under the prior appropriation doctrine as practiced in the State of Colorado.

2.1.1 Call Mechanism

This section details the mechanics of prior appropriation's seniority rule as to provide an understanding of the theoretical test for allocative efficiency and the data-generating-process underlying the daily water allocations data used in the empirical analysis. Prior appropriation's priority rule comes into effect when demand for water exceeds supply. In this situation, the owner of a senior water right may "call water" from diversions being made upstream from them, in exercise of water rights junior to theirs. The set of all water rights located upstream of a call are know as the "call reach", and among it diversions are curtailed in order of increasing seniority until the call for water is met. The mechanics of a water call are best explained by means of example, provided in Figure 2.1.

The status of the priority system is constantly in flux, due to changes in supply and demand for water. Calls may be lifted if they are no longer needed, or superseded by new calls for additional water. In Colorado, the priority system is administered on a daily basis by water commissioners. They determine the critical cutoff in seniority within active call reaches such that the calls for water are met exactly. "Ditch-riders" check that active diversions are in compliance to the priority rule, a task now made easier by digital meters installed at most major points of diversion.

The priority system in actuality is much more complex than is depicted in Figure 2.1. In practice water users are not lined up along a single source of water. Rather, they access water from a vast network of tributaries, streams and rivers. Therefore a call reach is small if the right calling for water is located in the headwaters of a water basin, and vice versa.

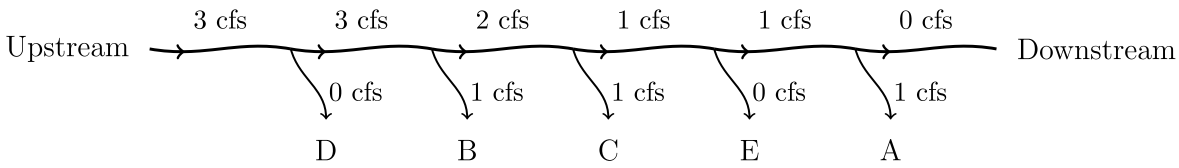

Figure 2.1: *Illustration of a water call.* Five water rights, A-E, are located along a river reach. Seniority is ordered from A, the most senior to E, the most junior. Each right claims 1 cubic feet per second (cfs), for a total of 5 cfs of claimed, "paper", water on the river reach. Only 3 cfs of river flow are available, prompting right A to exercise their priority and call for water. The water supply is sufficient for rights B & C to continue diverting. Rights D & E are out of priority and must not divert the 1 cfs that is physically available to them but legally entitled to right A.

Water diversions are principally made by collectives of individual water users called municipal ditch companies in Colorado. These ditch companies arose when individual land owners pooled their financial resources to develop canal and reservoir infrastructure for the purpose of irrigation. When these infrastructure projects were completed, the ditch companies perfected associated water rights. As a result, ditch companies (the unit of analysis in this paper) typically hold a portfolio of individual water rights across the spectrum of seniority. These rights also include alternate sources of water, including the "direct-flow" irrigation rights studied in this paper, as well as reservoir storage rights, groundwater and trans-basin diversions. Within a ditch company, the most common arrangement is for available water supplies to be allocated proportionally among ownership shares.

2.2 Theoretical Illustration of the Independence Property

Here I demonstrate the independence property of allocative efficiency within a stylized theoretical representation of my empirical setting, based heavily on Burness and Quirk

(1979)'s seminal model of prior appropriation. Although the independence property is widely understood, it perhaps not immediately obvious that it applies to the *intra-seasonal* shocks to agricultural water allocations that provide my identifying variation. To aid intuition, I abstract the complex decision problem of irrigation planning to a simple profit function with two inputs: water consumption x_i and planted area c_i. In the short run (i.e. within a growing season), crop choice is a fixed input chosen *ex-ante* of the realization of water allocations, whereas water consumption is a variable input chosen *ex-post*. I show that the independence property obtains within this sequential decision problem wherein payoffs depend on past actions made in advance of water allocations. I then map the theory to a reduced-form estimable model and causal test of the independence property.

Suppose firms $i = 1, \ldots, N$ share an aggregate stochastic water supply $x \sim f$. Let $x \geq 0$ and F(x) be its cumulative density, with $F(0) = 0$ and $\lim_{x\to\infty} F(x) = 0$. Firms are endowed with appropriative water rights, with seniority ordered from 1 (the most senior) to N (the most junior). A water right, $\overline{a}_i$, specifies the maximum allocation of water to firm i, under the condition that the aggregate supply, x, is sufficient to first meet the demands of the $i-1$ rights senior to i. Expressing this formally, denote the sum of rights senior to i as $A_{i-1} \equiv \sum_{j=1}^{i-1} \overline{a}_j$. Then firm i's water allocation, conditional on x, is:

$$a_i(x) \equiv \begin{cases} 0 & \text{if } A_{i-1} \geq x \\ x - A_{i-1} & \text{if } A_i \geq x \geq A_{i-1} \\ \overline{a}_i & \text{if } x \geq A_i \end{cases}$$

Let profits be given by a strictly concave function $\pi_i(x_i, c_i)$, with $\frac{\partial^2 \pi_i}{\partial x_i \partial c_i} \geq 0$. To avoid extraneous corner solutions, let $\lim_{0\leftarrow x_i} \frac{\partial \pi_i(x_i,c_i)}{\partial x_i} = \infty$ and $\lim_{0\leftarrow c_i} \frac{\partial \pi_i(x_i,c_i)}{\partial c_i} = \infty$, so that ideally all firms plant some crops and consume some water.

I explore two cases in turn. In the first, strict adherence to prior appropriation, water consumption x_i must not exceed firm i's stochastic allocation of water $a_i(x)$. In the second, prior appropriation with trade, firms may buy or sell water with respective a cost or proceed given by $p(a_i(x) - x_i)$. I assume firms are price-takers, so that the price of water, p, is set by the market clearing condition $\sum_i x_i \leq x$. In both cases, the cropping decision c_i is made in advance of the realization of a_i, and constrained by the total land available: $c_i \leq \bar{c}_i$. I derive the optimal choices of x_i and c_i sequentially by backwards recursion.

In the case of strict adherence to prior appropriation, I show that a firm's input decisions depends on their water rights through several channels. If their water allocation is constraining, their water consumption will equal their allocation $a_i(x)$. Water consumption may also depend on the firm's cropping decision, a choice that is based upon their expected allocation of water. In contrary, in the case of prior appropriation with trade, I show that a firm's water consumption is independent of their water allocation, and that their crop decision solely depends on the expected price of water.

2.2.1 Strict Adherence to Prior Appropriation

Firm i's objective under strict adherence to prior appropriation's seniority rule is:

$$\max_{0 \leq c_i \leq \bar{c}_i} \left\{ \mathbb{E}_x \left[\max_{0 \leq x_i \leq a_i(x)} \pi_i(x_i, c_i) \right] \right\} \tag{2.1}$$

Solving the inner problem first, denote the optimal water input $x_i^T(c_i; a_i(x))$. It depends on the endogenous crop choice c_i as well as the exogenous water allocation $a_i(x)$. Plug in x_i^A and expand the expectation to re-write the objective as:

$$\max_{0 \leq c_i \leq \bar{c}_i} \Big\{ \pi_i(0, c_i) F(A_{i-1})$$

$$+ \int_{A_{i-1}}^{A_i} \pi_i(x_i^A(c_i; x - A_{i-1}), c_i) f(x) d(x)$$
$$+ \pi_i(x_i^A(c_i; \overline{a}_i), c_i)(1 - F(A_i)) \Big\}$$

The optimal crop choice is then implicit in the first-order condition. It clearly will depend on the size of the firm's water allocation $\overline{a}_i$, as well as its priority A_{i-1}.

2.2.2 Prior Appropriation With Trade

Here firm i's objective is:

$$\max_{0 \le c_i \le \overline{c}_i} \left\{ \mathbb{E}_x \left[\max_{x_i} \pi_i(x_i, c_i) + p(a_i(x) - x_i) \right] \right\} \tag{2.2}$$

Notice that the water input x_i is no longer constrained by the allocation $a_i(x)$. Again solving the inner problem first (that has an interior solution by assumption), the first order condition is:

$$\frac{\partial \pi_i(x_i, c_i)}{\partial x_i} - p = 0$$

The implicit solution gives firm i's demand for water $x_i^T(c_i; p)$, and the market clearing condition is: $\sum_i x_i^T(c_i; p) = x$. It is unnecessary to expand the expectation in this case, so we simply plug into the objective to obtain:

$$\max_{0 \le c_i \le \overline{c}_i} \left\{ \mathbb{E}_x \left[\pi_i(x_i^T(c_i; p), c_i) + p(a_i(x) - x_i^T(c_i; p)) \right] \right\}$$

At an interior solution, the optimal crop choice c_i^T is implicit in the first order condition:

$$\mathbb{E}_x \left[\left(\frac{\partial \pi_i(x_i^T(c_i; p), c_i)}{\partial x_i} - p \right) \frac{\partial x_i^T}{\partial c_i} + \frac{\partial \pi_i(x_i^T(c_i; p), c_i)}{\partial c_i} \right] = 0$$

and is otherwise equal to $\overline{c}_i$. Either way, c_i^T only depends on the expected price of water, and is entirely independent of the hierarchy of appropriative water rights.

2.2.3 Summary of Theoretical Predictions

The key insight of this illustration is that if there is a competitive water market, firms' water inputs and returns to crop production are independent of their water rights, as well as their water allocations. Rather, only the aggregate supply of water affects firms' input decisions, via the price of water. Although I do not show it here, these results are straightforward to extend to more detailed models of agricultural production, such as with multiple crops or with multiple irrigation applications within a growing season.

2.3 Empirical Strategy

This section details the empirical strategy used to test for allocative efficiency context of Colorado's irrigated agriculture. Motivated by the theoretical prediction, the empirical specification must tease apart the effects of aggregate water supplies and allocations to individual firms. This is necessary because the aggregate supply of water affects firms' profits regardless of allocative efficiency. All else equal, a greater water supply decreases the marginal value of water to all firms if water is allocated optimality, and decreases the marginal value of water to some firms if water is allocated by prior appropriation. I control for the aggregate water supply using the classic difference-in-difference model specification. I will show momentarily that differencing across individuals at a particular point in time purges the effect of the aggregate supply of water.

The second difference, between time periods, purges stationary unobservable characteristics. Fortuitously, this includes the ownership of water rights, that would otherwise pose a threat to identification. For instance, if the highest quality farmland was developed for irrigation first, water rights seniority might be correlated with unobservable farmland characteristics. Causal identification requires that I purge the endogenous variation in water allocations due to water-rights ownership, retaining only the plausibly exogenous variation due to the stochastic timing of water curtailments. Since water rights are fixed in the short run, unit fixed effects produce this desired result.

A remaining complication is that neither water consumption itself, nor marginal profits, are empirically observable. Instead, I use the "greenness-index" NDVI, a workhorse metric of vegetation biomass, as a proxy for crop yields. This necessitates several identifying assumptions. Namely, I assume that crop yields are concave in water inputs and that (unobserved) input costs are not concave. Moreover, I assume that crop yields and NDVI are correlated. Both assumptions are supported by agronomic studies. Trout and DeJonge (2017) conduct field trials of maize yield responses to irrigation in northeast Colorado, finding evidence of decreasing returns to water inputs. de Lara, Longchamps, and Khosla (2019) perform a similar experiment, and compare maize yields to high-resolution NDVI images and find positive correlation between yields and NDVI. Kayad et al. (2016); Payero, Neale, and Wright (2004); Wang, Rich, Price, and Kettle (2005) also find evidence of correlation between NDVI and ground-truthed crop yields.

I specify the following estimating equation. The outcome variable, $NDVI_{idcs}$ is the average "greenness" for unit i, captured on calendar day d, planted with crop c and with irrigation system s. The variable of interest, $Alloc_{id}$, is i's allocation of surface water on day d, measured in cubic feet per second (cfs) per acre. The unobserved price of water on day d is denoted $p(x_d)$. Additionally, $Precip_{id}$ and $Temp_{id}$ are controls for weather. Lastly, $\beta_i, \beta_d, \beta_c, \beta_s$ denote fixed effects respectively for individuals, calendar days, crops and irrigation systems.

$$NDVI_{idcs} = \beta_0 + \beta_1 Alloc_{id} + \beta_2 Precip_{id} + \beta_3 Temp_{id} + \beta_4 p(x_d) + \beta_i + \beta_d + \beta_c + \beta_s + \varepsilon_{idcs} \quad \text{(S.1)}$$

Under the identifying assumptions, if the independence property holds, the coefficient β_1 is equal to zero, because conditional on the price of water $p(x_d)$, profits are independent of initial allocations a_{id}. In this specification, I am controlling for $p(x_d)$, so that I can estimate the effect of individual allocations, $Alloc_{id}$, in isolation. Note that the effect of $p(x_d)$ gets picked up by the fixed effect β_d, so that specification S.1 can be estimated by a classic fixed-effect regression.

A remaining concern is that the independence property only holds exactly if all water users actually do share the same supply of water, and if there is no difference in transportation costs. This could be addressed by separating water rights into hydrologically-connected groups, such as by water sub-division. However these groups would be somewhat arbitrary, and the assumption that water-rights share a common supply would seem most reasonable if the groups were as small as possible. For this reason, the second empirical specification pairs individual water users with their nearest-neighbor and estimates the effect of relative differences in allocations among these pairs. The surface water rights I analyze are gravity-fed within interconnected ditch systems. Therefore differences in physical water access and transmission costs can be assumed to be negligible within the close proximity of each spatial-pair.

Supposing that ditches i and j are spatially-paired, the data are transformed as:

$$\widehat{NDVI}^{i,j}_{dcs} = NDVI_{idcs} - NDVI_{jdcs}$$

The right-hand side variables are transformed equivalently. The fixed-effect specification (equation S.2) includes a compound fixed-effect $(\beta_i - \beta_j)$ for the spatial pair. The remaining effects (date, crop and irrigation systems) are subsumed by the transformation.

$$\widehat{NDVI}^{i,j}_{d,c,s} = \beta_1 \widehat{Alloc}^{i,j}_{d,c,s} + \beta_2 \widehat{Precip}^{i,j}_{d,c,s} + \beta_3 \widehat{Temp}^{i,j}_{d,c,s} + (\boldsymbol{\beta}_i - \boldsymbol{\beta}_j) + \widehat{\varepsilon}^{i,j}_{d,c,s} \tag{S.2}$$

The final specification, again first-differences the data chronologically beginning with the first observation of each growing season. The twice-differeneced transformation is denoted as:

$$\Delta\widehat{NDVI}^{i,j}_{dcs} = NDVI_{idcs} - NDVI_{jdcs} - (NDVI_{i,d-1,cs} - NDVI_{j,d-1,cs})$$

where the first observation for each year, each spatial pair, each crop and irrigation system is dropped. The notation $d-1$ is shorthand for the previous observation, as NDVI is observed roughly once-a-week, described further in the data section. The preferred specification is thus:

$$\Delta\widehat{NDVI}^{i,j}_{d,c,s} = \beta_1\Delta\widehat{Alloc}^{i,j}_{d,c,s} + \beta_2\Delta\widehat{Precip}^{i,j}_{d,c,s} + \beta_3\Delta\widehat{Temp}^{i,j}_{d,c,s} + \Delta\widehat{\varepsilon}^{i,j}_{d,c,s} \tag{S.3}$$

Across all specifications, the independence property predicts that $\beta_1 = 0$ if there exists a perfectly competitive water market. To reiterate, this is because with a market, only the aggregate supply of water affects firms' profits, via its affect on water prices and their expectation. By including fixed-effects for each individual day when NDVI is observed, I am controlling for the effects of aggregate water supplies, leaving only relative differences in individual water allocations.

The baseline specification, S.1, is a standard fixed-effect regression. It has the appeal of being a standard model, and with it I am able to use the entirety of the available data. The spatial first difference specifications, S.2 and S.3, test the robustness of the baseline specification to the possibility that differences in physical access violate the independence property. Furthermore, the spatial first-difference absorbs confounding spatially-correlated unobservables, such as farmland quality (Druckenmiller & Hsiang 2018). In these specifications, I compare only neighboring water users, growing the same crop and with the same irrigation technology. In order to do so, I am forced to drop many observations where this comparison is not possible. Finally, the third specification, S.3, has a strong intuitive appeal. Here I also apply a first-difference over time, so that the resulting data are changes in NDVI (or covariates) from one period to the next.

A final point about the empirical methodology in general is that the independence property is merely a necessary condition for allocative efficiency. I showed that in theory, if there is a market for water, the independence property obtains. By the first welfare theorem, the equilibrium of such a market is Pareto optimal. Because the independence property is a necessary, but not sufficient condition for optimality, one of two things can happen empirically. If I reject the null hypothesis of independence, I can conclude that water is not allocated efficiently, because a necessary condition is not satisfied. However if I fail to reject independence, I cannot conclusively say that water is allocated efficiently, since independence is merely a necessary condition.

2.4 Data

2.4.1 Administrative Water Rights Data

Water rights data are sourced from Colorado's Division of Water Resources (DWR). Several datasets are combined to create a daily account of surface water allocations throughout the State of Colorado. In accordance with the prior appropriation doctrine, individual water rights are defined by a point of diversion, a seniority and a quantity of water measured in cubic feet per second. This water rights data is tabulated by Colorado's Division of Water Resources in a publically-available dataset called 'Net Amounts'. To account for possible changes in water rights over time, archived versions of the Net Amounts dataset were obtained for each year between 2012 to 2018. Archived water rights data prior to 2012 are contained in scanned PDF documents and thus inaccessible for large-scale analysis. Water rights data years of 2008-2011 are infilled with the data from 2012. Minimal change in water rights over the 2012 to 2018 period suggests infilling these data years is not damaging.

Colorado's Division of Water Resources maintains a separate geospatial dataset called "Irrigated Lands" that links each point of diversion its service area: the spatial extent of irrigated farmland where the water may be applied. These geospatial layers are updated every five years. For each year in the 2008-2018 study period, I use the closest available year of the irrigated lands data to define the geographic boundaries of each ditch's service area. Although each service area contains a collection of individual parcels of land, the ownership of ditch shares is not public and hence it is infeasible to segment the data at a level finer than the individual ditches. For this reason, ditches and their associated service areas serve as the primary unit of analysis. To account for the varying size of these ditch companies, the water rights allocations from the 'Net Amounts' data are divided by the number of acres in the ditch's service area. The resulting metric water allocations used in the analysis has units of cubic feet per second per acre.

The datasets mentioned thus far delineate the hierarchy of water rights but do not indicate the actual allocations of water. Allocations, as explained in the previous second, depend on the administration of the priority system in response to calls for water. A third dataset, called "Administrative Calls", contains this time-series. Attributes in this dataset include the date of the call, its duration and the diversion location where it was placed. Critically, the Administrative Calls data also includes the marginal water right that is on the cusp of curtailment (water right "C" in figure 2.1). That is, it gives the critical priority that separates water rights within a call reach that are and are not eligible to divert water.

The Administrative Calls dataset does not contain the list of water rights within each call reach, and thus it alone does not contain the requisite information to assemble a daily account of water allocations. Until recently, it has been impossible to discern this critical piece of information. It can now be done using the United States Geological Survey's National Hydrography Dataset Plus dataset. It contains the location of every natural watercourse and the direction it flows. This hydrography data allows one to select a location on a watercourse and trace it upstream to identify all of the tributaries that feed water to that particular location. This makes it possible to identify the reach of every call recorded in the Administrative Calls dataset. Colorado's Division of Water Resources hosts an online tool for performing this analysis named the 'Structure Call Analysis Tool'. Using this tool, it is possible to identify the priority status of each water right in Colorado at a daily timestep. This information is collected and combined with the Net Amounts and Irrigated Lands datasets to derive daily water allocations across Colorado.

2.4.2 Remote Sensing Data

To assess the effect of water allocations it is necessary to obtain a metric of agricultural outcomes. Unfortunately, Colorado's crop statistics are collected by county, whose large spatial extents obfuscate variation in water allocations between individual ditches. For this reason, the remotely-sensed "greenness index", NDVI, is used as the outcome data. Remote sensing provides the advantage of an outcome measured precisely over space (at 30m resolution) and in time (roughly weekly).

The NDVI data used in the analysis is derived from georeferenced surface reflectence images captured by the remote sensing satellites Landsat 5, 7 and 8.[1] This data is calibrated for factors such as the angle between the satellite and the Earth's surface. These

[1] Landsat 6 failed to reach orbit during its deployment. Landsat 7 images are missing significant portions of data due to a mechanical failure.

calibrations transform raw measurements taken by the satellites' sensors to reflect actual land-surface reflectence. The calibration is performed by the Earth Resources Observation Science center and are referred to as "Level-2" data. Data from Landsat satellites 5, 7 and 8 were combined to cover the 2008 to 2018 study period. In combination, they provide roughly one observations per week for each unit of analysis. Data are collected for the Colorado's growing season of May through September. Missing data due to cloud cover and Landsat 7's broken scan line corrector are discarded, yielding an average of 16 observations per year for each observational unit.

The NDVI data are further categorized by crop planted and by irrigation system. Crop type is sourced from the U.S. Department of Agriculture's (USDA's) Cropland data layer. Each pixel of NDVI is matched to the Cropland data layer individually. Irrigation technology is sourced from Colorado DWR Irrigated Lands dataset of parcel boundaries. The NDVI data is aggregated by averaging among individual ditches, crop, irrigation technology and date of capture, yielding 1,336,360 distinct observations.

2.4.3 Weather Data

Weather data is sourced from the PRISM Climate Group at Oregon State University. The specific weather attributes are daily precipitation and mean daily temperature. The weather data has a coarser spatial resolution than the Landsat and Cropland data. For this reason, weather data is collected at the individual ditch level and is not further refined by crop type and irrigation technology.

2.5 Results and Discussion

The data are used to estimate specifications S.1, S.2 and S.3. Across all specifications, individual water allocations have a statistically significant effect on crops' greenness (table

2.1). The spatial-first difference procedure did not produce qualitatively different results than the base panel-model using the full dataset (specification S.1).

Table 2.1: Coefficient Estimates

	Dependent Variable: NDVI		
Model Specification:	S.1	S.2	S.2
Alloc. (cfs/acre)	0.01804***	0.01725***	0.01436**
	(1.31e3)	(5.03e3)	(7.23e3)
Precip. (mm)	0.00079***	0.00029	0.00047
	(7.89e5)	(4.90e4)	(2.88e4)
Temp. (C)	-0.00776***	-0.01033***	-0.00218
	(1.50e4)	(2.81e3)	(1.92e3)
Crop & Irrig. F.E.	True	False	False
Entity F.E.	True	True	False
Date F.E.	True	False	False
Entities	4203	465	465
Observations	1336360	202977	194609

Significance Levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.
Standard errors, in parentheses, are clustered entity. Entities are individual ditches in specification S.1, and spatial pairs of ditches in specifications S.2 and S.3

These results evidence a failure of the independence property, a necessary condition for allocative efficiency. They indicate that daily irrigation decisions within Colorado agriculture are made on the basis of individual water allocations, meaning there is not a perfectly competitive water-market. Although it is widely believed water-markets do not function optimally, these results are meaningful because they supply a reduced-form test of water-market efficiency with a causal interpretation.

Several caveats to these findings are warranted. While my approach has the advantage of being a causal test for allocative efficiency, it is unfeasible to recover the overall magnitude of the deadweight loss due to this inefficiency without making structural assumptions. In spite of this limitation, the results are worrying. Recall the my estimation procedure intentionally takes a narrow focus within Colorado's irrigated agriculture sector. Therefore it excludes potential gains from trade between alternative water uses. Moreover, the analysis is net of previously-demonstrated water-scarcity adaptations including crop choice and membership of irrigation organizations. This implies these adaptations, made ex-ante of water supply realizations, are insufficient for equating ex-post marginal values of water across uses.

A second caveat is that I cannot distinguish the mechanism driving the results. In the introduction, I posited several institutional weaknesses of prior appropriation as potential barriers to allocative efficiency. These were the beneficial use standard and the no-injury standard of water trades. Both weaken property rights under prior appropriation, distancing the institution from the ideal system of property rights that theoretically give rise to efficiency. Although my results indicate gains from trade are not realized, they do not explain why.

A final limitation is due to data availability. Observation of firms' water inputs, marginal profits, or water market spot-prices would allow for a more straightforward test for efficiency, that would not rely on the few identifying assumptions I was forced to make. In spite of these limitations, in my view, my results represent the best causal test of prior appropriation's allocative efficiency possible given currently available data. In conclusion, the evidence provided here suggests that water use in Colorado's agriculture is in accordance with Burness and Quirk's (1979) theoretical prediction of prior appropriations' inefficiency.

References

Arellano-Gonzalez, J., AghaKouchak, A., Levy, M. C., Qin, Y., Burney, J., Davis, S. J., & Moore, F. C. (2021). The adaptive benefits of agricultural water markets in california. *Environmental Research Letters*, *16*(4), 044036.

Berck, P., & Lipow, J. (1994). Real and ideal water rights: The prospects for water-rights reform in israel, gaza, and the west bank. *Resource and Energy Economics*, *16*, 287-301.

Brewer, J., Glennon, R., Ker, A., & Libecap, G. (2006). Transferring water in the american west: 1987-2005. *U. Mich. JL Reform*, *40*, 1021.

Brown, T. C. (2006). Trends in water market activity and price in the western united states. *Water Resources Research*, *42*(9).

Burness, H. S., & Quirk, J. P. (1979). Appropriative water rights and the efficient allocation of resources. *The American Economic Review*, *69*(1), 25–37.

Colorado Water Conservation Board. (2020). *Alternative transfer methods in colorado; status update, framework for continued support, and recommendations for cwcb action* (Tech. Rep. No. DQ1 PMAA 2020*0166). Colorado Department of Natural Resources.

Coman, K. (1911). Some unsettled problems of irrigation. *American Economic Review*, *1*(1), 1–19.

de Lara, A., Longchamps, L., & Khosla, R. (2019). Soil water content and high-resolution imagery for precision irrigation: Maize yield. *Agronomy*, *9*(4), 174.

Druckenmiller, H., & Hsiang, S. (2018). *Accounting for unobservable heterogeneity in cross section using spatial first differences* (Tech. Rep.). National Bureau of Economic Research.

Fowlie, M., & Perloff, J. M. (2013). Distributing pollution rights in cap-and-trade programs: are outcomes independent of allocation? *Review of Economics and Statistics*, *95*(5), 1640–1652.

Hagerty, N. (2019a). Adapting to water scarcity: Evidence from irrigated agriculture. *Working Paper*.

Hagerty, N. (2019b). Liquid constrained in california: Estimating the potential gains from water markets. *Working Paper*.

Howe, C. W., & Goemans, C. (2003). Water transfers and their impacts: Lessons from three colorado water markets. *JAWRA Journal of the American Water Resources Association*, *39*(5), 1055–1065.

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.

Kayad, A. G., Al-Gaadi, K. A., Tola, E., Madugundu, R., Zeyada, A. M., & Kalaitzidis, C. (2016). Assessing the spatial variability of alfalfa yield using satellite imagery and ground-based data. *PLoS One*, *11*(6), e0157166.

Libecap, G. D. (2011). Institutional path dependence in climate adaptation: Coman's" some unsettled problems of irrigation". *American Economic Review*, *101*(1), 64–80.

MacDonnell, L. J. (2015). Prior appropriation: a reassessment. *University of Denver Water Law Review*, *18*(2), 228-311.

MacDonnell, L. J., Howe, C. W., Rice, T. A., et al. (1990). Transfers of water use in colorado.

Manning, D. T., Goemans, C., & Maas, A. (2017). Producer Responses to Surface Water Availability and Implications for Climate Change Adaptation. *Land Economics*, *93*(4), 631–653.

Mukherjee, M., & Schwabe, K. (2014). Irrigated agricultural adaptation to water and climate variability: the economic value of a water portfolio. *American Journal of Agricultural Economics*, *97*(3), 809–832.

Payero, J., Neale, C., & Wright, J. (2004). Comparison of eleven vegetation indices for estimating plant height of alfalfa and grass. *Applied Engineering in Agriculture*, *20*(3), 385.

Smith, S. M. (2021). The relative economic merits of alternative water right systems. *Journal of Environmental Economics and Management*, *105*, 102389.

Tarlock, A. D. (2001). The future of prior appropriation in the new west. *Natural Resources Journal*, *41*(4), 769–793.

Thompson Jr, B. H. (1993). Institutional perspectives on water policy and markets. *Calif. L. Rev.*, *81*, 671.

Trout, T. J., & DeJonge, K. C. (2017). Water productivity of maize in the us high plains. *Irrigation Science*, *35*(3), 251.

Wang, J., Rich, P. M., Price, K. P., & Kettle, W. D. (2005). Relations between ndvi, grassland production, and crop yield in the central great plains. *Geocarto International*, *20*(3), 5–11.

Womble, P., & Hanemann, W. M. (2020). Water markets, water courts, and transaction costs in colorado. *Water Resources Research*, *56*(4), e2019WR025507.

Chapter 3

The Storage Premium

3.1 Introduction

Infrastructure development, such as canals, dams and reservoirs, is a stalwart remedy for water-scarcity problems. Reservoir storage is highly valuable as it allows for smoothing of consumption for water supplies that are highly cyclical and stochastic. Yet empirical analyses have found that reservoirs' benefits accrue locally (Duflo & Pande 2007), and may not pass the benefit-cost test generally to the tax base that supports them (Fisher, Fullerton, Hatch, & Reinelt 1995; Howe 1987). This evidence, based on analysis of existing reservoirs, motivates cost-benefit analysis of planned future taxpayer-funded reservoir projects. However, doing so is not straightforward, as it entails dynamic optimization of reservoir operations and the underlying hydrology of the reservoir system (Fisher & Rubio 1997). Furthermore, such calculations depend upon expectations of future supply and demand conditions, quantities that are difficult to predict beforehand. An estimate of the market capitalization of reservoir storage capacity therefore provides valuable information, as it captures the market's expected future value of storage, net of the complex dynamics that underpin it in theory.

I estimate the market value of reservoir storage using a proprietary dataset of water transactions prices in the South Platte river basin in Northeastern Colorado, curated by the consulting firm WestWater Research LLC. This water market is among the most active in the United States (Brewer, Glennon, Ker, & Libecap 2008), and is highly competitive due to its abundance of market participants. Ongoing urbanization within the region has led to the gradual sale of water rights from their original agricultural use to municipal water suppliers to meet new housing demand. The particular asset being transferred are ownership-shares of mutual ditch companies, entities that own portfolios of water rights, including in some cases reservoir storage rights, and allocate water to their shareholders pro-rata. The data permit two distinct estimates of the value of reservoir storage. First, I estimate a hedonic model of ditch company share prices, retrieving an estimate of the market capitalization of the shares' partial ownership of reservoir infrastructure and associated water rights. This model estimates a value of $2,597 per acre-foot of expected annual water storage volume (95% C.I. = [$-217.6, $5412]). This estimate pertains to the volume of water *stored*, and the corresponding volume of consumable stored-water is certainly less, due to system losses and the historic consumptive use of previous agricultural practices. The second model estimates the value of storage in terms of consumable yield, meaning the volume of water the buyer expects to receive annually. Here, I estimate a value of $9,500 (95% C.I. = [$4441, $14,600]), equal to a 185% premium compared to direct-flow consumable yield that cannot be stored. The difference in these estimates captures the cumulative effects of storage and transmission losses, as well as court-determined consumptive use of irrigation water based on historical irrigation practices.

Both estimates, of the value of water stored, and of the value of consumable water that is stored, are useful policy metrics for the assessment of planned, but yet unbuilt, reservoir projects. A major and contentious example is the Northern Integrated Supply

Project (NISP). This project involves the construction of two new reservoirs for the supply of predicted housing demand in Colorado's Front Range, and is estimated to cost a minimum of $500 million. These reservoirs, if completed, will expand storage capacity by 215,600 acre-feet (170,000 acre-feet in the proposed Glade Reservoir and 45,600 acre-feet in the proposed Galeton reservoir), to be filled with a junior water right with a priority date of May 2, 1980. In total, the project is expected to yield 40,000 acre-feet of consumable stored water per year. Using my estimated value of reservoir storage, a back of the envelope calculation values this reliable yield at $380 million (95% C.I. = [$178m, $582m]). Applying the estimated market capitalization of reservoir infrastructure, I estimate that the value of the Glade Reservoir's 170,000 acre-feet capacity, in tandem with its junior water right, is $170 million (95% C.I. [-$14.2m, $354m]). This compares to the estimated cost of $124 million for building the Glade reservoir in isolation. The overall NISP project involves a complex suite of planned infrastructure development, whose overall valuation is outside of the scope of this paper. Although these back of the envelope calculations should be viewed cautiously, they fail to reject a null hypothesis that the NISP project's benefits are equal to their cost. These estimates pertain to the market value of water alone, and do not capture social costs and benefits including recreation, wildlife habitat and the adverse-affects of "buy-and-dry" water transfers on agricultural communities that the NISP project aims to displace. They are also point estimates derived from supply and demand conditions in the market for water over the period 2014-2019. By substantially increasing the supply of stored water in the region, the NISP project is apt to decrease the market equilibrium value of storage water, holding all else equal.

3.1.1 Relation to Previous Literature

The economic value of water has been an object of empirical interest for decades, as missing or incomplete markets for water obscure the resource's shadow value (Young & Loomis 2014). As spot markets for water are thinly traded, the value of water is typically derived from analysis of water-rights transactions, or through hedonic analysis of irrigated farmland value to which irrigation water rights are tied.[1] Yet as noted by Mukherjee and Schwabe (2015), there are few estimates of water value differentiated by the quality and source of water supplies. An exception is Payne, Smith, and Landry (2014), that uses a previous version of the water transactions dataset studied here and estimated a storage premium of roughly $1,700 per acre-foot of water delivered annually. The value of stored water derived in the current paper essentially replicates Payne et al. (2014), with the difference in estimates ($1,700 versus $6,166) capturing a dramatic increase in value over the intervening years. My second estimate, the value of reservoir storage capacity, is due to a hedonic analysis of ditch company share prices. This methodology was also used in Goodman and Howe (1997), and is here extended to include ditches' reservoir storage capacity.

3.2 Data and Methods

The primary data source is a subset of Waterlitix, a proprietary dataset of water-rights transactions data curated by the consulting firm WestWater Research LLC. The data pertain to transactions of ditch-company shares in the South Platte River Basin, over the period 2014-2019. Key data attributes include the ditch company name, the price per share, consumable yield per share and the asset class (i.e. direct flow rights, or reservoir

[1] See, for example, Crouter (1987); Faux and Perry (1999); Hartman and Anderson (1962); Mukherjee and Schwabe (2015); Petrie and Taylor (2007); Selby (1945)

storage). Because ditch companies are highly heterogeneous in the water rights they own and the number of outstanding shares, the data include a transformed value of price per unit of consumable water, allowing for equivalent comparisons among the transactions. Historic consumptive use (HCU), is determined by Colorado's Water Court, in effort not to expand water use when water rights are transferred to a new type and place of use. Historic consumptive use is critically important to the value of ditch company shares, as it effectively determines how much water one should expect to receive per share owned.

I collected additional data attributes from public archives of ditch company's Articles of Incorporation and water-court proceedings. These attributes include the number of outstanding shares for each ditch company, and the collection of water rights they own. Water rights are classified as either "direct-flow", permitting diversion of natural streamflow for immediate use, or as "storage" rights, that permit a volume of water to be diverted and stored in a reservoir for later use.[2] Both classes of water rights are appropriative, entailing that their date of first appropriation determines their priority to water supplies.

The final piece of data I collected pertains to the water rights' priorities. All else equal, a senior water right is more valuable than a junior right, as the senior right receives water with greater certainty than the junior. However, the priority date is simply an ordinal ranking, and a water-right's actual supply of water additionally depends on its geographic location and its priority relative rights that share the same water source. To control for this complex effect of priority, I performed a reliability analysis for each individual water right. Specifically, I measured reliability as the percentage of time the water right was in priority within growing season months of April to October, over the 20 year window preceding the study window (1994-2014). I weighted each water right's decreed volume

[2] Transmountain water rights, such as those of the Colorado Big-Thompson (CB-T) project, are excluded from the analysis.

(measured in cubic feet per second for direct-flow rights, and in acre-feet for storage rights) by this reliability metric, in effort to control for the influence of priority.

To facilitate better understanding of the data, I describe the data attributes used in the statistical analysis, and the methodology used to derived them, in turn.

Share Price

The primary price data is the price per share of a municipal ditch company. These ditch companies were formed in the 1800's, when collections of neighboring farmers pooled there financial resources to develop irrigation infrastructure to secure irrigation water supplies (Goodman & Howe 1997). Over time, ditch companies augmented their water supplies by expanding canal infrastructure and constructing reservoirs. As they did so, ditch companies perfected additional water rights with sequentially junior priority dates. As a result, ditch companies now own portfolios of water rights, differentiated by asset class (direct flow versus storage), and with varying seniority. These water supplies are apportioned to ditch company share holders proportionally. Therefore, a ditch company's share prices captures the market value of partial ownership of the company's overall portfolio of water rights. The share price data are supplied by WestWater Research.

Yield Per Share

Colorado water law dictates that when water rights are transferred to a new type and place of use, the new use must not exceed the historical water consumption of the original water use. For water being sold out of the agricultural sector (the typical scenario), an assessment of historical agricultural practices is made to determine the historical average water consumption per year. Consumptive use refers to the volume of water that is fully consumed, meaning it is made unavailable for subsequent water uses. Factors such as canal seepage and return flows of applied irrigation are not considered consumptive, as these water losses return to the water basin and are claimed by downstream water rights. As a result, the consumptive use determination is typically a small fraction of historical diversions, and varies widely depending on the infrastructure historically used for irrigation. The yield per share data supplied by WestWater Research are sourced from archived water court decrees, among other sources.

Unit Price

Unit price is equal to the price per unit of consumable yield, where the later quantity is due to the historic consumptive use determination described above. The unit price is equal to a ditch company's share price, divided by the yield per share, so that the resulting units are dollars per unit of yield ($/ AF). The result is readily comparable across ditch companies, as their varying number of outstanding shares are subsumed by the transformation.

Asset Class

A ditch company's asset class is either direct-flow, storage, or both direct flow and storage; determined by the ditch company's portfolio of water rights. In the South Platte basin, almost all water rights are in priority in the springtime, when supplies are ample due to snowmelt runoff. Only the most senior rights supply water through the late summer, when water supplies dwindle. Storage rights are therefore valuable, as they allow springtime water supplies to be consumed year-round.

Water Rights Reliability

The appropriation water rights studied here are differentiated by their appropriation, or priority, date. Senior water rights have the first priority to water supplies and are therefore highly valuable. However, the mapping from a water rights priority date to its reliability is non-linear, meaning the priority date alone is not a sufficient statistic for the quality, or reliability, of the water right (Payne et al. 2014). To create a better metric of a water right's quality, I analyzed the priority status of each water right in the 20 year period preceding the study window. This analysis was performed using water rights data retrieved from Colorado's Decision Support System (2021). Specifically, I calculated the percent of time each right was in priority between 1994 and 2014, within the growing season months of April through October. Winter months were omitted, because water available then is not used for irrigation, and thus it is not transferable as per the historical consumptive use standard. Because ditch companies own portfolios of water rights, it is necessary to aggregate water right descriptors, such as reliability. Therefore, reliability itself does not appear in the regression analysis, but rather is aggregated by asset class as described below.

Expected CFS Per Share

To quantify and aggregate a ditch company's direct-flow water rights, I calculated a weighted sum of the volume of each water right (measured in CFS), multiplied by the right's reliability and divided by the number of outstanding shares in the ditch company. This calculation puts greater emphasis on the volume of senior water rights that have greater reliability, but does not penalize a ditch company if it also owns less reliable junior rights. The resulting statistic "expected CFS per share", captures the average CFS per share of a ditch company, over the growing season months of April through October.

Expected AF Per Share

I aggregated storage rights in an equivalent fashion as the direct-flow rights, except that storage rights' volumes are measured in units of AF. These volumes are equal to the expected volume of water, per share, *stored* in a reservoir, and the corresponding yield of stored water is less, due to transmission and storage losses (equivalently, CFS per share captures the volume of water that may be diverted, not accounting for consumptive use). Expected CFS and AF per share serve as the primary metrics of direct flow and storage rights in the hedonic analysis of ditch company share prices.

3.2.1 Descriptive Statistics

Selected summary statistics of the water transactions data and collected water rights data are provided in table 3.1. The raw data reveal a premium for stored water, despite that on average storage rights are junior to senior rights and are therefore less reliable. It is unfortunately impossible to provide meaningful summary statistics of ditch-company share prices and their determinants without revealing the underlying proprietary data.

Table 3.1: Summary Statistics of Transactions Data by Asset Class

	Unit Price[1]	Avg. Priority Year		Avg. Reliability	
Asset Class		DF	S	DF	S
Direct Flow (DF)	$10,726	1874	N/A	0.859	N/A
Storage (S)	$14,386	N/A	1893	N/A	0.633
Both (DF & S)	$14,752	1877	1907	0.801	0.588

[1] Unit price is the price per acre-foot of historic average consumptive use.
[2] Direct Flow units are cubic feet per second (CFS), storage units are acre-feet (AF).

3.3 Estimation

Statistical analysis of ditch-company share prices and unit prices of consumable water provide greater insight to the determinants of their values. The first model I estimate is a hedonic price function of the determinants of ditch company share prices. First developed by Rosen (1974), the hedonic price methodology is used to estimate the contributions of various measurable attributes to the price of a differentiated good. In this setting, I analyze the contributions of reliable yield per share, direct flow rights and reservoir storage rights to ditch company share prices. Because the dataset includes a relatively few observations (N = 93), it was necessary to aggregate the water rights attributes into the metrics of expected CFS per share and expected AF per share, as described previously.

$$Share\ Price = \alpha + \beta_1 Yield + \beta_2 CFS + \beta_3 AF + \boldsymbol{\gamma} District + \boldsymbol{\delta} Year + \varepsilon \tag{3.1}$$

The estimating equation for the hedonic analysis of ditch company share prices given by equation 3.1. The variables "Yield", "CFS" and "AF" refer to the yield per share and expected CFS and AF per share, as described previously. Inclusion of yield per share controls for the historic consumptive use determination, that could otherwise bias the estimates of the variables of interest. "District" and "Year" are shorthand for fixed effects by administrative water district and year respectively. These control for the influence of location and time trends in prices that are not the focus of this study. The corresponding estimation results are provided in table 3.2.

Table 3.2: Hedonic Analysis of Ditch Company Share Prices

Dependent Variable: Ditch Company Share Price			
	Coefficient	Std. Error	P-Value
Intercept	-79,041.11	24390.053	0.001
Yield Per Share	5,460.89	1500.327	0.000
Weighted CFS Per Share	257,615.19	89952.835	0.004
Weighted AF Per Share	2,597.16	1436.142	0.071

Number of Observations = 93, $R^2 = 0.752$. Standard errors are clustered by ditch company (N = 42). Water division and year fixed effects are also estimated but omitted here for brevity.

The hedonic estimation results produce an estimated value of \$257,615 per expected CFS of direct-flow diversions, and \$2,597 per expected acre foot of stored water, both received annually in perpetuity. The value of reservoir storage is estimated noisily (standard error = \$1,436), perhaps due to the simplifying assumptions that were necessary to produce this estimate. For example, the reliability of storage rights was taken over the entire growing season (April-October), although reliable spring flows could be ample to fill a reservoir. Considerations like the previous example speak to the limitations of the reduced-form approach of valuing reservoir storage infrastructure. In principal, analysis of reservoir storage volumes, captured by diversion records, could inform a structural model, but this is outside the scope of the current paper. Despite its limitations, the hedonic model produces the first reduced-form estimate of the market capitalization of reservoir infrastructure that I am aware of.

My second empirical specification, analysis of the unit price of reliable *consumable* water, circumvents inherent complications of the hedonic approach. Here, share prices are divided by the water court-determined yield per share, an estimate that factors in the complex hydrology of historic reservoir operations. Following Payne et al. (2014)'s methodology, the estimating equation is very simple. The unit price of water is simply regressed on the asset class of the water rights being transferred, as well as the non-parametric controls discussed previously.

$$Unit\ Price = \alpha + \boldsymbol{\beta} Asset\ Class + \boldsymbol{\gamma} District + \boldsymbol{\delta}\, Year + \varepsilon \tag{3.2}$$

Estimation results are provided in table 3.3. The results reveal wide price dispersion between direct flow and storage water. Note that in contrast to the previous hedonic estimations, here all three asset classes have units of dollars per acre foot of yield (explaining the difference in magnitude of the value of direct flow). The coefficients for the direct flow and storage and storage variables are interpreted as additional value of these asset classes, relative to the baseline category of direct flow. The net value of storage water is then $\$3333 + \$6167 = \$9,500$ (95% C.I. = [$4441, $14,600]).

Table 3.3: Water Prices by Asset Class

Dependent Variable: Unit Price ($/AF Average Annual HCU)	Coefficient	Std. Error	P-Value
Intercept (Direct Flow)	3333.330	0.002	0.000
Direct Flow & Storage	5796.933	6507.589	0.373
Storage	6166.509	2581.050	0.017

[1] Number of observations = 93, R^2 = 0.293. Standard errors are clustered by ditch company (N = 42). Water division and year fixed effects are also estimated but omitted here for brevity.

Taken together, these results indicate that reservoir storage is highly valuable in the South Platte river basin. I presented two methods for valuing reservoir storage, the first capturing the market capitalization of reservoir infrastructure and associated water rights, and the second capturing the value of reliable water yield supplied through reservoir storage. Both methods have useful policy implications for the valuation of planned, but yet unbuilt reservoir development, as detailed in the next section.

3.3.1 Valuation of the Northern Integrated Supply Project

The Northern Integrated Supply Project (NISP) is a proposed suite of reservoir infrastructure development in the South Platte river basin, that has been in the federal permitting process since 2004. The principle components of the project are the proposed Glade and Galeton reservoirs, with respective capacities of 170,000 and 40,000 acre-feet. To put these volumes in perspective, the existing Horsetooth reservoir, a component of the Colorado Big Thompson project that serves much of the region's municipal water supply, has a capacity of 156,735 acre feet. The NISP project aims to provide additional water supplies to the consortium of 15 municipalities and ditch companies that are backing, and funding, the project. Using a water right with a May 2, 1980 priority, the NISP project will capture excess springtime runoff that otherwise flow to Nebraska, unused in the State of Colorado. Crucially, NISP aims to supply water for municipal growth in the region without the adverse affect of "buy-and-dry" transfers on the region's agricultural economy that would otherwise be necessary. Details of the NISP project were sourced from the US Army Corps of Engineer's Final Environmental Impact Statement (2018).

Overall, the NISP project entails a complex set of plans, including reservoir and pipeline construction, enhancing streamflows through the town of Fort Collins, a new recreation site at the Glade reservoir, and re-routing a highway. A complete cost-benefit analysis of the full project, including environmental impacts and amenity values is outside the scope of this paper. Instead, I estimate the market capitalization of the projects expected water yield (40,000 AF annually) and of the proposed Glade reservoir in isolation. Applying the point estimates of table 3.3, a back of the envelope calculation values the project's 40,000 AF yield at $\$9,500 \times 40,000 = \$380,000,000$ (95% C.I. = [$178m, $582m]). Compared to the estimated minimum project cost of $500m, a t-test fails to reject the null-hypothesis that the estimated market value and project costs are equal (P-value = 0.245). This is immediately apparent, given that $500m falls within the 95% confidence interval.

As for the Glade reservoir, I analyzed the reliability of its May 2, 1980 water right in the same fashion detailed earlier. During growing season months in the period of 1996-2014, this right was in priority only 38.5% of the time. Multiplied by the reservoir's proposed 170,000 AF capacity, the expected volume of stored water is 65,402 AF. A back of the envelope calculation, using point estimates from table 3.2, values the Glade reservoir at $170 million (95% C.I. [-$14.2m, $354m]). The wide confidence interval reflects the noisy estimate of this model specification, and again a t-test fails to reject a null hypothesis of equality to the estimated cost of $123,974,000 for building the reservoir in isolation (P-value = 0.625).

In combination, using these back-of-the-envelope calculations, I fail to reject the null hypothesis that the project benefits are equal to their cost. Therefore, I cannot rule out that the project breaks even, nor can I say affirmatively that the project is worth undertaking. However, I have cautioned that these estimates do not capture social amenities or disamenities of the project. Proponents of NISP argue that if the project is not completed, growing municipal water demand will have to be met by "buy-and-dry" water transfers. Since I cannot reject the null hypothesis that the value of NISP is equal to its cost based on the value of water it supplies alone, social values are likely to be the deciding factor for whether the project is worth undertaking. Moreover, this analysis has revealed that the value of storage water has increased rapidly over the past decade, and will likely continue to do so. My estimates are based on market data from the 2014-2019 period, and the value of storage may have already increased to the point that the benefits of the NISP project outweigh their cost.

3.4 Conclusion

In this paper, I have advanced current knowledge of the implicit value of water resources through estimation of the value of reservoir storage in the South Platte river basin. Motivated by previous research, I estimated the value of reservoir storage in two ways. First, I estimated the market capitalization of reservoir infrastructure and associated water rights using a hedonic analysis of ditch-company share prices. Second, I estimated the value of reliable yield of water transferred in these trades, differentiated by the asset class of the water supplies. The principle finding is that storage is highly valuable, and receives a substantial premium ($6,166 per AF yield), compared to direct flow water supplies. I then applied these estimates to a contentious proposed reservoir development project: The Northern Integrated Supply Project (NISP). Concentrating on the economic value of the project's expected water supply and reservoir construction costs, I fail to reject the null hypothesis that the project's value is equal to its cost. These estimates suggest that the net value of social benefits and costs of NISP, not analyzed here, likely determine if the project's net benefits exceed its costs.

References

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.

Colorado's Decision Support System. (2021). *Colorado division of water resources.* Retrieved from `https://https://cdss.colorado.gov/`

Crouter, J. P. (1987). Hedonic estimation applied to a water rights market. *Land Economics*, *63*(3), 259–271.

Duflo, E., & Pande, R. (2007). Dams. *The Quarterly Journal of Economics*, *122*(2), 601–646.

Faux, J., & Perry, G. M. (1999). Estimating irrigation water value using hedonic price analysis: a case study in malheur county, oregon. *Land economics*, 440–452.

Fisher, A., Fullerton, D., Hatch, N., & Reinelt, P. (1995). Alternatives for managing drought: a comparative cost analysis. *Journal of Environmental Economics and Management*, *29*(3), 304–320.

Fisher, A., & Rubio, S. J. (1997). Adjusting to climate change: Implications of increased variability and asymmetric adjustment costs for investment in water reserves. *Journal of Environmental Economics and Management*, *34*(3), 207–227.

Goodman, D. J., & Howe, C. W. (1997). Determinants of ditch company share prices in the south platte river basin. *American Journal of Agricultural Economics*, *79*(3), 946–951.

Hartman, L. M., & Anderson, R. L. (1962). Estimating the value of irrigation water from farm sales data in northeastern colorado. *Journal of Farm Economics*, *44*(1), 207–213.

Howe, C. W. (1987). Project benefits and costs from national and regional viewpoints: Methodological issues and case study of the colorado—big thompson project. *Natural Resources Journal*, *27*(1), 5–20.

Mukherjee, M., & Schwabe, K. (2015). Irrigated agricultural adaptation to water and climate variability: the economic value of a water portfolio. *American Journal of Agricultural Economics*, *97*(3), 809–832.

Payne, M. T., Smith, M. G., & Landry, C. J. (2014). Price determination and efficiency in the market for south platte basin ditch company shares. *JAWRA Journal of the American Water Resources Association*, *50*(6), 1488–1500.

Petrie, R. A., & Taylor, L. O. (2007). Estimating the value of water use permits: A hedonic approach applied to farmland in the southeastern united states. *Land Economics*, *83*(3), 302–318.

Rosen, S. (1974). Hedonic prices and implicit markets: product differentiation in pure competition. *Journal of political economy*, *82*(1), 34–55.

Selby, H. (1945). Factors affecting value of land and water in irrigated land. *The Journal of Land & Public Utility Economics*, *21*(3), 250–258.

US Army Corps of Engineers, O. D. (2018). Northern integrated supply project: Final environmental impact statement. , *1*.

Young, R. A., & Loomis, J. B. (2014). *Determining the economic value of water: concepts and methods*. Routledge.