Contingent Resource Claims and Coordinating Ex-Ante Investment
DRAFT – WORK IN PROGRESS

Daniel E. Osgood Alexander Pfaff Arthur A. Small, III*

May 23, 2005

Abstract

In light of proposals for mechanisms (insurance or futures contracts, options or forward contracts, and "reliability" contracts) for allocating natural resource inputs with uncertain supplies, we identify and discuss a key feature of uncertainty for actors who make ex-ante investments complementary to the resource input that exists even for those with risk-neutral individual preferences and production characteristics. We then show the value, in coordinating total investment, of differentiated reliabilities arising from a first best ex-ante (pre-rain) contingent claims market. We examine next how other markets and institutions compare to and interact with the proposed first best system, focusing on spot markets and other markets as well as and social processes that differentiate reliabilities.

Key words: risk, uncertainty, renewable resources, water, options, missing markets, climate variation, climate prediction, forecasting, property rights

JEL Codes: Q21, Q25, D83, D23, D02

1. Introduction

How can individuals best make investment choices in the face of an uncertain natural resource supply and insecure resource claims? And how can societies best allocate such resources across the individuals making those choices? This situation has generated scores of responses across the centuries in agriculture, in fisheries and in other settings. They include rules of thumb developed over the centuries. In the state of Ceara in Northeast Brazil, in subsistence rain-fed agriculture, poorly educated farmers invest under uncertainty by tilling their land repeatedly until Saint Joseph's Day. If rains sufficient for the plants have not yet come by then, the farmers cease such investment for the season. This rule of thumb appears to reflect informal learning over time about the ITCZ (inter-tropical convergence zone) that climate scientists see as critical for rains in Northeast Brazil.[1] It reaches its southernmost point around Saint Joseph's Day and its subsequent movement northward lowers the probability of sufficient rains.

*Senior authorship shared equally. All authors are affiliated with Columbia University. Correspondence to aas69@columbia.edu

[1](see XXXX YYYY ZZZZ)

Concerning both individual investment choices and allocation across individuals, if human needs are modest relative to nature's bounty then an open access regime may handle matters smoothly. *De facto* or *de jure* rights would simply attach to units of the resource as they are withdrawn (see Locke's 'rule of capture' [**?**]). Otherwise, when the resource not employed by one user will be employed by another (as in irrigated agriculture if not the rain-fed example above), open access yields undesirable resource depletion or degradation, dissipation of rents and perverse incentives for entry. This latter case seems more likely. The rising demands for irrigation, municipal and industrial water supply from Ceara's water management system, for instance, and more generally rising demands within many sectors around the world that feature competition for scarce natural resource inputs, indicate the increasing value of better answers than open access to the challenge of managing of natural resources under uncertainty.

One leading proposal has been the development of better private property rights regimes, which if feasible might avoid the difficult question that we opened with above. In the case of water, the penalties of uncertainty and variation in water availability have been one of the basic arguments for the original establishment of rights sufficient for spot water markets [**?**, **?**, **?**]. Uncertainty in water rights and water availability is known to impose severe costs to farmers, with information and uncertainty issues substantially impacting investment decisions [**?**, **?**, **?**, **?**, **?**] A related branch of literature identifies reliability and investment impacts on urban water users [**?**, **?**, **?**].

While the secure private resource rights being advocated might in principle generate efficient resource use, we believe that in practice there exist at times insuperable (and at times understandable) hurdles to that approach. Spot markets for water, for instance, are rare and those that exist are typically very thin due to the difficulty of adapting existing institutions [**?**]. Water allocation is typically characterized by ambiguous property rights. While formal water rights per se might well be sufficient for the markets and contracts of interest to arise, the only water rights that actually exist may be those implied through delivery contracts.

The reasons for a lack of functioning markets include private information, investment decisions, property rights, transactions costs, and market thinness, as noted by economists studying markets emerging in the 1990s [**?**, **?**, **?**, **?**, **?**]. In this literature, it was recognized that new markets cannot necessarily be mandated in their most efficient theoretical form, but that they are typically developed through adaptation of existing institutional constraints. For instance, the state may not wish to relinquish some degree of control over the allocation of water. Also, the state is unlikely to make delivery promises that a user can treat as a water right, if that makes the state liable for water, or its value, when nature does not deliver. Thus we develop here an approach that need not involve private absolute rights to the uncertain resource in question.

Specifically, for our starting point we note that common-pool resources are often managed by a public or quasi-public agency set up specifically and solely to handle that task. Within a designated jurisdiction, the resource management agency: assumes residual control rights to the resource *in situ*; organizes the creation and allocation of withdrawal rights; and maintains a system for monitoring and enforcing rules as well as for adjudicating disputes. Sometimes the agency will also handle the physical distribution of the resource itself and/or provide ancillary services such as information about weather or market conditions. The variety of different architectures that deal with common-pool natural resources is wide and fascinating [**?**].

1.1. Public Management with Uncertain Supply

We assume that a common-pool management agency is a local monopolist for a specific common-pool resource. It decides how to allocate permission (rights) to harvest a renewable natural resource when aggregate supplies are uncertain. A canonical example is a water agency, before knowing exactly how much water it will have available to deliver, choosing whether or not to sign delivery contracts with (i.e., commitments to) water users and, if so, how many and for whom. Another example is a fishery management agency setting rules governing harvest rates by fishing vessels while uncertain about aggregate stocks.

A classic example is the system of senior and junior water rights used by agriculturalists in California's Central Valley. Other systems are more egalitarian, at least formally, relying on variable restrictions that are set by the resource management agency and apply uniformly to all users. An example would be the fishing periods, or outright fishing bans, used by many fisheries management agencies as imperfect tools to calibrate allowed catch according to the perceived size and vulnerability of the underlying fish stock.

The agency's control over the natural resource's supply is sharply limited. For surface water resources, in the absence of reservoirs the aggregate supply is truly exogenous, i.e. unresponsive to changes in total demand in the short- and the long-run. For ground water and fisheries, the agency may be able to exert some control over the aggregate deliveries in any period. Further, depending on the situation, the resource agency may be able to ease constraints on resource supplies relative to demand through investments in infrastructure (e.g., a water reservoir) or institutions (e.g., a tradable permit system for use rights). However, there are limits on how much infrastructure and institutions can impact supply. Although a reservoir can transfer surplus water across time and institutions may allow transfers of water across users, when the available water is exhausted, a high shadow value for water does not necessarily bring more rain.

Even when there is no control over the natural supply, information about it may exist. Concerning the Northeast Brazil example above, for instance, there exists some skill in forecasting seasonal precipitation in Brazil. Forecasts for the region rank among the most accurate in the world [**?**]. Concerning the water management system specifically, there is also skill in forecasting streamflows, i.e. the inputs into the reservoirs, which is achieved by adding sea surface temperature indices to typical stream flow regressions [**?**]. Within a reservoir management optimization problem, this skill can raise yield conditional on reliability or, conversely, raise the reliability conditional on a given yield.

No institutional design or forecast, however, eliminates the basic feature of uncertainty in natural resource arrivals into the system. Nature's uncertain total delivery to the agency must be translated into uncertainties in the aggregate supplies passed on to users—or at least, to some users.[2] Whether by design or by default, every institutional architecture incorporates a risk-transmission mechanism, from nature to end user.

[2] Theoretically, this is not quite true. It would in principle be possible for the resource management agent to eliminate all uncertainty for users by pursuing a policy based on a concept of "*maxi-min* sustainable yield". The agency could commit to make only a small aggregate quantity of deliveries, such that it knows with certainty that these obligations can be honored in full. To eliminate all user uncertainty, the agency would then need credibly to commit to waste, destroy, or otherwise render valueless any residual quantities of the resource that happened by good fortune to arrive, but for which deliveries could not be assured fully in advance. The agency would operate rather like a firm that committed to pay off bond holders, and then to burn the residual earnings that would otherwise have been paid out to owners of common stocks. This unusual dividends policy would indeed shield equity holders from exposure to risk.

Our contribution in this paper is to highlight the fundamental risk penalties driven by coordination in ex ante investment, suggest a first best framework, contingent resource claims, for coordination of investment and transmission of such aggregate risk across users through contracts that can be honored in every state of nature to improve investments and outcomes. Through analysis of this framework we address three central questions: (1) what is risk-neutral reliability, i.e. in what risk-neutral setting does a "reliability" that has value exist? (2) what can be achieved, in particular concerning the coordination of investment choices, by a market that permits individuals to purchase differentiated levels of individual reliability? and (3) how does such a market relate to other institutions, e.g. what is its role in relation to spot markets, options markets, institutional constraints, interventions to support the vulnerable, and social processes that have generated unequal marginal products of water and reliabilities?

The basic instrument traded in a market for differentiated contingent claims on an uncertain stock of a given natural resource is the right to buy one unit of a natural resource, at a stated price if and only if the aggregate supply of that resource exceeds a stated threshold. That is, users may purchase a particular 'place in line for water' before the total aggregate amount of water is known. In addition, the contingent claims contract is a generalization of common water rights systems. Its implementation is relatively natural in these settings, since it simply involve temporary leases of prior appropriative rights.

Our instrument is distinguished from a price option primarily in its simplicity and that it is possible to exercise the contingent claims contract in every state of nature. When water supplies are completely exhausted, it is not possible for a seller to deliver water at any price and some sellers will not be able to honor water option contracts. To our knowledge, no market exists with regular trading of water options. Instead, existing water price options contracts consist of a few heavily negotiated individual contracts, perhaps illustrating the difficulty of implementing price options and demonstrating a value of proposing alternative ex ante mechanisms. The simplicity of the contingent claims contract not only allows its value to be calculated in a much more straightforward manner than price options, but makes it a useful benchmark for an investigation into the roles of risk, investment coordination, markets, and institutions in natural resources.

To further distinguish our innovations, we proceed to discuss the literature, particularly that of water options markets, before presenting the formal modeling. In particular, we note that the literature on water price options contracts has focused on techniques the valuation of the options as opposed to explicitly specifying the unique role that ex ante contracts can play in more efficient investment and resource allocation.

2. Previous Literatures

There is a mature literature on information, water rights, water reliability and investment, and water markets and an infant literature on water options. However, the water literature has yet to address the fundamental relationship between coordination of water related investments under uncertainty and ex-ante verses ex-post market based exchange mechanisms, particularly in a world of imperfectly defined rights and imperfect information.

2.1. Water Options

Because of the importance of investment decisions and uncertainty, and information in water use, futures markets and other contingent trading systems may be particularly useful as tools for managing risk in water management. The literature on contingent exchanges of water is in its early stages of development and is focused on price options contracts [?]. Following years of unsuccessful attempts between agricultural and urban water users to develop options contracts in California, research was done to propose the structure of an efficient water options contract[?]. The eventual successful options contract established between these parties in California was more exotic and complex than options contracts found in finance. This contract motivated the development of new methodology for pricing and valuing options [?]. The traditional price option studied can only be approximately implemented in the water management situation. Because rainfall is a fixed quantity that is inelastic to water price, it is possible that an option contract cannot be honored at any price. The literature on options has focused on the technical aspects of proposing existing instruments from finance or pricing a the complex contract that has been negotiated as opposed to relating the contribution of the option mechanism to insurance or spot markets, or studying their efficiency role as a vehicle for the coordination of investments, or proposing instruments that are consistently implementable.

2.2. Contingent Claims & Seniority

2.2.1. Finance

Intricate examples of contracts denominated by seniority arise in corporate finance. Many central questions in applied corporate finance concern how to structure the allocation of a firm's earnings between claimants of differential seniority. Common examples involve the division of financial claims between debt and equity. More involved structures may be created to finance large infrastructure projects, with complicated schedules for allocating the "cash waterfall" generated by project operations. Naturally, claimants further back in the queue face greater risk, and may demand correspondingly higher returns on their invested capital. Much research among academics and finance practitioners concerns how to quantify the risks associated with claims of differential seniority, and the corresponding returns that security holders will demand under various market conditions.

The translation of this research into the realm of natural resources is not straightforward, however. Simply put, the cash generated from one asset or firm is fungible with the cash generated by any other. To a very strong approximation, the only difference between various financial claims of equivalent maturities concerns the risks that affect the quantity of cash delivered. In this case, it makes sense to posit, and even to compute, an economy-wide "price of risk" in the manner of [?].

This fundamental equivalence between all forms of cash cannot be assumed when the analysis concerns a non-cash natural resource. First, the user community is small and local and thus any "price of risk" for uncertain deliveries of the resource will likewise be local. This is not necessarily the case for all natural resources, as oil and other commodities are traded on international markets, but water trading is constrained as is assumed in the consideration only of local rains within our model. One implication is that aggregate risk can not be mitigated simply by pooling of uncorrelated risks across localities. Second, the natural resource is presumably an input to a production process, in which it serves as an imperfect substitute, or as a complement, to other factors of production. A crucial implication is that investments in other inputs to production, such as tilling the land, change the marginal value of water. To first order, there is no such analogous investment for cash[3]. Our focus is precisely such ex-ante investments.

2.3. Other Water Literatures

< this is pasting in Dan's "tangential but maybe useful later": [?] Water Constraints and Environmental Impacts of Agricultural Growth, water reliability and its relationship to food security. [?] ex ante dairy technology adoption uncertainty [?] phd Drived Willingess-To-Pay for Water: Effects of Probabilistic Rationing and Price [?] Land costs a lot because of water rights [?] Hedonics of global warming irrigated ag [?] land rights and investment incentives in ghana, transition between systems [?] electricity pricing [?] electricity >

3. General Model of Investment and Resource Demand

There is a continuum of resource users ("farmers"), indexed by $t \in [0, 1]$, each with measure zero. Each farmer chooses an investment level k. Investment at level k yields payoff $\lambda(k)$ if and only if the farmer receives delivery of one unit of the resource ("water"). Otherwise, the investment yields a zero payoff. We assume that conditional investment payoffs exhibit decreasing returns to scale. To simplify the analysis, we assume that the investment function λ is twice continuously differentiable on $k > 0$, with $\lambda'(k) > 0$ and $\lambda''(k) < 0$ for all $k > 0$. We normalize returns so that $\lambda(0) = 0$. To rule out degenerate cases, it is assumed that $\lambda'(0_+) > 1$. (Otherwise, positive levels of investment would never be individually rational, even for farmers whose water deliveries were guaranteed with certainty.) In this first-pass analysis, it will be assumed that the menu λ of investment options is the same for all farmers.

It is also assumed that agents share common and accurate beliefs about the probability distribution of the random variable X, representing aggregate water supplies.[4] These beliefs are represented by a cumulative density function $\Phi(.)$ defined on $[0, 1]$: $\Phi(q) := \Pr[X \leq q]$. It shall be assumed that Φ is differentiable (i.e., the underlying distribution contains no atoms), with associated probability density function $\phi(q) = \Phi'(q)$.

3.1. Strict seniority, no trading

The analytically simplest case is an institution in which each unit of water is awarded according to strict seniority. In this setting, a farmer at position q receives one unit of water if and only if $X \geq q$. Let $\pi_q(k|X)$ denote this farmer's realized returns, conditioned on investment at level k and a realization X of aggregate water supplies:

$$\pi_q(k|X) = \left\{ \begin{array}{cl} -k + \lambda(k) & \text{if } X \leq q; \\ -k & \text{if } X > q. \end{array} \right\} \tag{3.1}$$

[3] *(if the point is heterogeneity across people, are there not portfolio differences across investors that affect the marginal value of a given risky cash claim?)*

[4] This assumption can be relaxed somewhat in some versions of the model. For example, when resources are assigned using a strict seniority system, with no possibility for trading, it will only be necessary to assume that each resource user has accurate belief's about his or her *own* probability of receiving a delivery.

Then

$$E\pi_q(k) := E_X\left[\pi_q(k)\right] = -k + \lambda(k) \cdot \Pr[X \geq q]. \tag{3.2}$$

A risk-neutral profit-maximizing farmer taking his queue-location q as given will choose an investment level to maximize $E\pi_q(k)$. This simple problem has first order condition

$$\frac{dE\pi_q}{dk} = -1 + \lambda'(k) \cdot \Pr[X \geq q] = 0 \tag{3.3}$$

implying

$$\lambda'(k) = 1/\Pr[X \geq q]. \tag{3.4}$$

Since λ' is strictly decreasing, with $\lambda'(0_+) > 1$, investment will be positive for at least some (low) positions $q > 0$.

Case: $\Pr[X = 1] = 0$**.** In this case, production is fundamentally constrained by resource availability: the marginal farmer at the end of the allocation queue (position $q = 1$) has no chance of getting any of the resource. Then there will exist a critical location $Q < 1$ at which investment drops to zero, determined by the relation $\lambda'(0_+) = 1/\Pr[X \geq Q]$. For $q < Q$, equation (3.4) defines an optimal investment level $k^*(q)$. Since λ' is strictly decreasing and continuously differentiable, an expression for $k^*(q)$ can be derived explicitly:

$$k^*(q) = \left(\lambda'\right)^{-1}\left(\frac{1}{1 - \Phi(q)}\right). \tag{3.5}$$

It is easy to show that k^* is continuously differentiable and strictly decreasing on $(0, Q)$.[5] The complete optimal investment function is the (continous) extension of k^* defined by letting $k^* \equiv 0$ on $Q \leq q \leq 1$.

Case: $\Pr[X = 1] > 0$**.** In this case, the marginal farmer at postion $q = 1$ has at least some chance of getting water. If moreover $\lambda'(0_+) > 1/\Pr[X = 1]$, then positive investment will be optimal for all q, according to the investment function (3.5). Otherwise, there will exist a $Q \leq 1$, defined as above, with $k^*(q) = 0$ for all $Q \leq q \leq 1$.

3.1.1. Comments on strict seniority with no trading

Note that

In other words, in this setting, a system of resource allocation based on strict seniority "works", very well. It induces all potential users of the resource to make efficient investment decisions, conditioned on their respective positions in the allocation queue. Moreover, given this distribution of investment, strict seniority becomes optimal ex post: resources are allocated exactly to the highest-value consumers — those at the front of the queue, who made the heaviest investments in complementary resource-specific assets.

A strict seniority system has other features to recommend it. It does not require that farmers engage in heroic computational tasks, in order to choose a correct level of investment: each farmer need only compare his or her own marginal productivity of capital with his or her own probability of receiving water. Efficient investment does not require that each farmer form

[5] Since the probability distribution over water supply is continuous, we have that $\Pr[X \geq q] = \Pr[X > q] = 1 - \Phi(q)$.

expectations about the choices of other farmers, based on assumptions about other farmers' productivity, beliefs, etc. A strict seniority system can actually serve as a useful co-ordination device, keeping farmers away from making mistakes based on mis-matched beliefs about the choices made by their colleagues.

Indeed, most of the

It must be cautioned, of course, that the assumptions underpinning this result are not trivial. First, it is assumed that a well-defined queue has been established — that each successive unit of the resource is allocated to a pre-identified claimant, through a system of differentiated contingent resource claims. It is further assumed that all potential users of the resource share a *common* investment opportunity set $\lambda(k)$. It is also assumed implicitly that there are no productivity shocks — once investment decisions are taken, productivity $\lambda(k)$ is known with certainty for all resource users, conditioned only on the uncertain delivery of the resource itself. These and other stylized assumptions compell an advisory that application to real-world settings should be taken with care.

Nonetheless, the analysis suggests lessons about the potential benefits of strict seniority that should hold true even under more general assumptions. Indeed, many of the notorious inefficiencies of seniority systems arise in fact because in these institutions *seniority is not strict enough.* In many such systems — as, for example, with agricultural rights to water withdrawls in much of California's Central Valley — senior claimants do not hold claims to a specified quantity of the resource. Rather, they are in many cases allowed to withdraw as many units as they wish, paying a fixed price per unit. It is typically this flexibility in the seniority system that leads to the notorious waste of water on low-value applications — not the seniority queuing system *per se*.[6]

Our analysis suggests that, in such cases, many of the anticipated economic efficiencies that are supposed to attend the creation of a resource market may be acheivable in substantial part simply through a clarification of *non-tradable* seniority rights. Given the strong opposition from some users to the creation of full trading in resource rights, it may be advisable for policy advisors to focus first on this step, the redefinition of seniority rights in terms of quantity, rather than in terms of price. This first step can enhance efficiency through its salutory effect on ex ante investment behavior, without requiring the dramatic institutional overhaul implied by the movement to a full market.

Moreover, a "quantity redefinition" of property rights is in any case a necessary precondition to the development of a full trading regime.

3.2. Spot Market Only

Another possible institutional innovation would allow resource claimants to buy, sell or transfer their withdrawl rights. (The initial allocation of withdrawl rights could be carried out using an auction, a seniority queue, or by any sensible alternative system.) We shall assume that trading in this market is frictionless. Since every unit of the resource is fungible in use, the Law of One Price would quickly assert itself: once total resource supply X is realized, there would emerge a well-defined equilibrium spot price $S(X)$ for each unit of the resource.

A resource user with productivity λ will purchase water (or, equivalently, retain his own allocation for his own use) if and only if $\lambda > S$. In this case, he earns quasi-rents $\lambda - S$.

[6] A full analysis of this "prices vs. quantities" problem awaits further study.

Otherwise, he forgoes the resource (selling off his own allocation, if he has one), and earns no quasi-rents. His realized profits $\pi(k|X)$ from an ex ante investment at level k are

$$\pi(k|X) = \begin{Bmatrix} -k + \lambda(k) - S(X) & \text{if } \lambda \geq S(X); \\ -k & \text{if } \lambda < S(X). \end{Bmatrix} \tag{3.6}$$

His expected earnings $E\pi(k) := E_X[\pi(k)]$ are given by

$$E\pi(k) = -k + \Pr[\lambda \geq S(X)] \cdot \{\lambda(k) - E_X[S(X)|\lambda \geq S(X)]\}. \tag{3.7}$$

Given beliefs about the probability distribution of $S(X)$, our risk-neutral profit-maximizing farmer will choose an investment level k designed to maximize the expression on the left. To say more about the resulting choice requires imposing some assumptions about the process by which farmers form their expectations about future prices, at the prior time when they make their investment decisions.

A farmer familiar with the concept of an endogenously determined rational expectations equilibrium will recognize that the future demand for water will be determined by the investment decisions of other farmers, each of whom will also follow the rules leading to a profit function π as displayed in equation (3.6). Their collective investment choices will yield a *productivity ordering function* $\Lambda(q)$. Without loss of generality, this ordering can be assumed montone decreasing. (For convenience, we for the moment restrict attention to cases in which farmers' collective investment choices yield a *strictly* decreasing productivity ordering.) Our particular farmer, by investing at level k, will in effect be selecting his place $q(k)$ in the allocation queue, defined by $q(k) := \Lambda^{-1}(\lambda(k))$, where $\Lambda^{-1}(.)$ denotes the inverse function of Λ. Then equation (3.7) becomes

$$E\pi(k) = -k + \int_{q(k)}^{1} [\lambda(k) - \Lambda(x)] \cdot \phi(x)dx \quad . \tag{3.8}$$

Since $\Lambda(q(k)) = \lambda(k)$, the first order condition for optimal ex ante investment takes a simple form reminiscent of equation (3.4):

$$\lambda'(k) = 1/\Pr[X \geq q(k)]. \tag{3.9}$$

3.2.1. Comments on the Spot Market Model

IA rational expectations Nash equilibrium in this setting is a pattern of investment $k(q)$ inducing a productivity ordering $\Lambda(q) = \lambda(k(q))$ in which all agents satisfy optimality conditions including equation (3.9). It is possible to show that there are, under a wide range of assumptions, well-defined equilibria of this kind.

It bears noting, however, that acheivement of this outcome requires stiff assumptions about the abilities of resource users to compute equilibria, and to co-ordinate on a jointly agreeable optimum.

In future work we will examine the use of forward markets in contingent resource claims, as a device for reducing these computational and coordination challenges.

4. Modeling Reliability

4.1. Ex Ante Reliability For Coordination in Investing

We propose an ex-ante market in which users can effectively choose their levels of reliability, conditional on what nature delivers. In our formal model of a market for differentiated contingent claims on an uncertain stock of a given natural resource, the basic instrument traded in this market is the right to buy one unit of a natural resource, at a stated price[7], if and only if the aggregate supply of that resource exceeds a stated threshold . That is, users who are willing and able to pay will purchase a particular 'place in the water line', i.e. the right to the Nth unit of water if N or more units are available due to the rains. Such an institution takes whatever aggregate signal of reliability exists, such as from historical climatology or a skillful seasonal forecast, and translates it into differented individual signals. Put another way, after the market has operated, one's individual reliability is determined by the aggregate reliability plus one's place in line. We believe that this can play an important coordination role in guiding private ex-ante investment decisions.

To fix ideas, consider a complete 'atomization' of water use. Each unit of water, say enough for a parcel of land with one fruit tree, is treated separately in contracts. Each actual sizable water user, then, is represented as many distinct 'unit users'. We will consider whether each such unit acts independently of the other units for that water user. If they do, as in our default model, this eliminates bundling of water units that would make risk aversion important for water-related decisions. Thus, in our core model we are not considering risk-averse water users purchasing lotteries for W units of water, such that either W or zero arrives. In that setting, the lottery's price would rise with its probability of W arriving and risk-aversion would affect one's valuation of any given lottery.

Our focus is instead the fact that in agriculture, and more generally in water resource management, decisions must be made both ex-ante, i.e. before nature's annual supply is known, and ex-post, after the rains. In agriculture, before the amount of water to be available is known, agents make ex-ante choices investment decisions about crop, fallow, and irrigation practice that affect the entire season to come. After the rains and reservoir inflows are revealed by nature, ex-post optimization is constrained by the ex-ante choices. Such 'putty/clay technology' has long been known to have many implications for behavior in agriculture and other sectors.[8] We suggest that it is the key to understanding the value of "reliability" in water delivery for farmers. Consequently, it is also a key determinant of the gains from market transactions that publicly provide ex-ante information about potential users' resource demands.

As our leading example, consider the ex-ante decision of how much land to till at the start of the growing season, before the rains. In any model, tilled area will be a function of the expected revenues, and thus both the crop price and the crop yield for any given level of water input, as well as the expected costs, and thus fertilizer prices and the cost of hired or own labor. Optimal ex-ante tillage should also be a function of the probability that water arrives. Recall that we refer here to one unit of water, which in our model is sufficient for the 'atomized' unit of land.

Ex-ante knowledge of this probability of the unit of water arriving is our first definition of

[7] *(we had some discussion of setting up with two prices in order to permit a para in which we consider different possibilites, but then making clear the case that we focus on)*

[8] *[so much so that I wonder if McCarl etc. don't mention it for climate forecast use and value that we should cite although probably not water-institution implications of the same idea]*

reliability. Its value derives from the ex-ante nature of the tillage investment. Without that, the value of this knowledge is zero. When the ex-ante probability estimate is high one invests with less trepidation and if it is low enough then one will not invest. Climate forecasts that are used to update prior default probabilities of water arrival, i.e. that change estimated reliability, will also affect investments. The rain-fed farmers who use the Saint Joseph's Day tillage investment rule are essentially forecasting then investing based on estimated reliability, using a time-tested rule for when to issue a 'bad season' forecast and, on that basis, update their default reliability prior.

The relationship between ordering and resource driven returns on investment are fundamental to the problem of efficient resource allocation. For a given ex post quantity of water, the efficient allocation is to provide the first increment of water to the incremental use with the highest return for that water, the next increment of water to the use with the next highest return, and so on until the water is exhausted. With investment given, the essence of efficiency is that the incremental uses that do get water have higher returns than those that do not get water. Since each incremental use either obtains water or not, any system that results in the correct ordering without imposing losses is efficient. The key difference between efficient ordering systems are the transfers and the distribution of rents. Since changes in probability distributions describing the resource are mapped to marginal uses, not users at the beginning or end of the queue, most impacts reveal themselves through the marginal consumer, who may or may not get the resource.

Agents must make investment choices before the ex post resource is known. These investment choices determine the returns to water. In the certain and efficient world, the users who invest are the ones who get water, and those who do not invest are those who do not get water. Innefficiency arises when investment is inconsistant with the aggregate amount available, and the ordering of who gets the resource. If farmers cannot coordinate in their individual investment decisions, the level of investment when individuals are aggregated may not be appropriate for the quantity of the common resource. The requirement of an efficient mechanism is that it provides the signals necessary to prevent a 'train wreck' due to uncoordinated individual action. When investments must be made ex ante under uncertainty, the unique role of a contingent claims mechanism (or comparable options markets) is to efficiently provide this coordination. Although it can provide some of the benefits from reduced uncertainty that other mechanisms (insurance, spot markets, etc.) the feature that distinguishes it is that it is an investment coordination mechanism. Using these observations, one can compare the relative advantages and disadvantages for packages of market and non market institutions, ordering mechanisms, and communication mechanisms. We proceed by formally evaluating a particular mechanism the contingent claims market.

4.2. Contingent Claims As Ex Ante Reliability You Invest In

We propose an ex-ante market in which users can effectively choose their levels of reliability, conditional on what nature delivers. In our formal model of a market for differentiated contingent claims on an uncertain stock of a given natural resource, the basic instrument traded in this market is the right to buy one unit of a natural resource, at a stated price[9], if and only if the aggregate supply of that resource exceeds a stated threshold. That is, users who are willing and able to pay will purchase a particular 'place in the water line', i.e. the right to the Nth unit

[9] *(we had some discussion of setting up with two prices in order to permit a para in which we consider different possibilities, but then making clear the case that we focus on)*

of water if N or more units are available due to the rains. Such an institution takes whatever aggregate signal of reliability exists, such as from historical climatology or a skillful seasonal forecast, and translates it into differentiated individual signals. Put another way, after the market has operated, one's individual reliability is determined by the aggregate reliability plus one's place in line. We believe that this can play an important coordination role in guiding private ex-ante investment decisions.

To examine this claim, we start by modeling the socially optimal level of investment and then what would happen concerning investment of this sort given solely a post-rain spot water market. One natural prior is that there will be socially inefficient investment arising from the spot market since without a contingent claims market everybody has the same average perceived reliability, making it harder for common expectations to coordinate who does and who does not invest. Errors can occur in either direction. Perhaps surprisingly, the ideal spot market involving actors with prescient expectations actually comes close to socially optimal investment, getting it right on average. That is worth understanding, as this coordination has not been a focus in the water literature. However, a number of alternative models of this unsupported coordination process would do much worse.

We then examine what the ex-ante contingent claims market would achieve. The immediate difference is that instead of having a common perceived reliability, those with high places in the water line will have higher reliabilities than those lower in the line. In short, for a given yield, those with higher reliabilities will gain more from investing and will purchase the high places in the water line. This serves as a credible signal to coordinate investments across individuals.[WE HAVE NOT REALLY DONE THE DETAIL DISCUSSION ON THIS, SO WILL NEED REDOING]

To understand what exactly this market is achieving, see that for each yield there is a threshold probability of water, declining with yield, that makes one indifferent to tilling the unit of land. This generates, for any given total rainfall, an 'investment indifferent curve'. Now imagine that as above everybody starts with the same prior perceived reliability. The gains from trading in contingent claims market are the gains from exchanging probabilities. They are investments at higher yields and investments that would not otherwise have happened. [REVISIT THIS WITH RESULT; *does a horizontal line in the Art Graphic, as I think is implied here and rules out backward-sloping diagonal, mean no gains at all ...??*]

(maybe the Art Graphic story is out, maybe the trade to higher lambdas is done in the Spot too? Maybe the point is that the water is not actually used until the rain actually comes, at which time it can be reallocated in the spot market no matter what was done before. That would imply that shifting the probabilities across the yields is not a gain relative to Spot market ... SEE TEXT IN EARLIER Vs)

... How is the surplus created by these trades or shifts in probabilities distributed? Depends on which case arises, scarce or not scarce water in the spot market. Also depends on who has the initial rights to sell the claims or the water. ...

4.3. Other Markets & Policy

Spot Market with Contingent Claims: Given our contingent claims market, would the spot market have any value at all? If water claims have already been ordered along the lines of the marginal productivity of water (i.e., the exogenous yield parameter), then if nothing changes other than the realization of the rain there will not be any reason for trades in the spot market. However, almost surely there are some shocks over time. Two that come to mind are a changed exchange rate, which matters for exported crops, and a labor shock such as illness, perhaps dengue in Northeast Brazil. Thus, since shocks are likely, we would expect that even after a contingent claims market has functioned a spot market for water would create value.

Water Futures & Insurance (our market improves futures function): While our proposed ex-ante contingent claims market does improve investments, as noted above it can not eliminate the uncertainty inherent in nature's delivery of rain. Thus claims holders are still subject to risk. Insurance could be provided through water futures which delivered money if rainfall is below a stated quantity. This could compensate for lost agricultural income in the low-rain states while, as does any insurance premium, lowering agricultural profits in the states in which water is delivered. The connection to our proposal is that water users are much better informed insurance or futures customers after the contingent claims market has been implemented. Any individual knows exactly the quantity of rain which will lead their claim (their 'place in line') to receive water and thus can purchase futures contracts for the personalized trigger quantity.

... not clear how thick the market is, though, with each person using different Qs ...

... insurers could hedge risk by using knowledge of rainfall correlations across basins within the country or anywhere they operate ...

... with risk neutral people we need to have a reason why they want insurance ... don't do the modeling but instead cite literatures ...

Water Futures & Expectations (future market helps Spot out, lowering value of ours?): [*My thought was that we could briefly discuss why some other ex ante mechanisms perform part of the role of the contingent claims but why they dont fully function. I dont have any fully developed thoughts on this, though: Dan*]

Land Markets & Political Economy *(Historical Wealth Distribution & Why Initial Yield and Probability Distributions Are Not Exogenous or Independent)*: [*Dan's version of this kind of statement was that we point out that lots can be happening, e.g. lots of room for playing with distribution, as long as the ordering and the marginals are the same, since those dictate the outcomes in the aggregate.*]

Various institutions in society may effectively differentiate both yields and probabilities, and also correlate their values across people, meaning that greater probabilities of water, i.e. higher reliabilities, are held by those with greater yields. This may be quite efficient in our model's terms. Put another way, if other processes (including credit access) generate differentiated probabilities so that who should invest is clear, what is implied by our proposed new market may not be much different from current reality without any such market. Those who have initially high wealth may have purchased the land with high yields, or perhaps made investments to raise yields, and they may also be able to lobby for higher reliabilities. Our results then suggest that there may be efficient components to what might often be described as a set of outcomes resulting purely from inequity.

...Point: while this sort of story could apply in the form of the rich owning the good soil, it is more dramatically relevant to Ceara, at least, when the rich also invest in higher yield crops like fruit ...

... for the modeling, note in addition that there may be elements of the probability, e.g. being upstream and hard to monitor, that just go with the territory, and those locations may not have been where the rich wanted to go to live ... perhaps more generally, we could find an analogous land market result to our contingent claims market result, along the lines that one should allocate land (like water) to those who will invest (may need to have credit limitations on the poor for this) ...

... for the policy, while we may not have much to say about it, we could then note that a major issue for water or land will be compensation for those moving out of production to something else ... makes a big difference if water or land are sold for their opportunity cost to the lower yield individuals or instead the surplus is split ...

If considering policy interventions to assist the more vulnerable, we might note based on our model that giving water rights could work better than giving the "cash equivalent". Given the surplus achieved by the water or land transfers, the marginal value of their water to them may not be enough to buy back the water on the spot market. Thus giving cash would take away the option to keep the one lifestyle they have known, which some may prefer even if others wish to sell and move.

... in terms of interventions, recall Dan's note that environmental reliability demands are represented as an agent in the trading (generalizing from how such demands have been represented by the United States Bureau of Reclamation in spot markets in the San Joaquin Delta Mendota Water Authority... also Dan noted that we should consider the flexibility of a transparent system, versus an ossified historical power structure ...

... a particularly policy relevant set of shocks to consider (noting that Dan's flexibility point raises the shocks idea, e.g. a technological one), for instance in thinking about transition/compensation issues, are the likely future changes towards more water for the port and the new urban centers ...

5. Conclusion/Extensions

We have proposed a market in contingent resource claims for an uncertain resource. This instrument more natural to impliment than the price options contract. Its simplicity allows its use as a framework for understanding the relative roles of ex ante markets, ex post markets, insurance, and existing institutions in investment coordination and efficiency of resource use.

We have found that a first best equilibruim exists for a contingent resource claim market.

Using the framework to compare alternate mechanisms, we find that the feature that distinguishes this instrument (and the similar, but less implimentable price option) from spot markets and insurance contracts is that it provides a mechanism ensuring coordination of ex ante investments. Although the instrument may provide additional benefits under risk aversion, uncertainty costs from risk for individual producers, or curvature of production functions, these assumptions are not neccessary for this basic finding.

Results are robust to specifications in which agents are homogenous and heterogenous. The primary impact of assuming heterogeneity is that agents accrue some scarcity rents if they are different, while the water rights holders capture rents from homogenous agents.

Implimentability:our ordered places in line, which make all claims contingent on the aggregate rainfall, can not leave the state on the hook; in contrast, though, are Manu's reliability-denominated delivery contracts; as Art and I noted early on, they will fail on aggregate if the knowledge of the distribution is imperfect (though this is not a problem for us, even if we advertise "estimated reliabilities" for each place in the water line in order to be useful for the users);

we can point out also it is not clear what should trigger non-delivery compensation with these (if you did not get water, was 90% reliability true?), which raises the fundamental point that this alternative approach to getting out of liability is not nearly as straightforward as getting a place in line; evidence of getting-in-line systems would support this; our contingent claims market brings a nice ability to pay for reliability to such approaches. As mentioned earlier, options markets have the same implimentability problems.

Multi-Season (probably in same paper -> Q:get done before conference?) (*long-term contracts as efficient (duh?)*) (*also just the ability to invest to choose your lamdba from a set of possible investments/crops*): The second definition of reliability is similar but addresses multi-season investments. Even for an atomized water user independent of other unit users, there may well be interdependencies across time. On the single land unit, the value from investing this year may be zero if no water arrives next year. Thus, the estimated ex-ante probability of having water in the relevant future years is also a type of reliability that will influence investment. For example, in fruit, which has fixed costs including early years gone by without any output from new trees, when investing in a tree one must think about the probability of losing output or even the tree if water does not arrive for a year or perhaps for two successive years. Industry investments with significant capital costs, and financing costs that may rise with under-repayment of loans, also require consideration of the stream of outputs over time that are water dependent.

Local versus Export Markets: Up until this point, the financial gains achieved from a marginal unit of water, which we have been calling yield but certainly have implicitly included the output price, have been seen as independent of the realized rainfall. However, in local markets, it is possible that the output price is lower in high-rain, i.e. high-supply years than in low-rain years. In the case of Ceara, this is well known for crops such as xxx, yyy and zzz. r-unit-user interactions when local output price falls endogenously with others' outputs (but fruit exported); also, Art noted that such interactions complicate the task of extracting probability information from posted prices for claims of varying seniority ... this lowers lambda as Q rises for local products, but not for exported products, which happen to be those which require more investment. That creates another efficiency reason to allocate water and land to those who can invest.*{NOTE: last point fits the multi-crop/investment level model, so sticking this here}*

Forecast Value (different paper IF our market affects the forecasts' value; else maybe a note and worth less to IRI, though maybe we can say our market helps allocate risk better and thus adoption of forecasts is more likely for reservoir releases): Given the increasing skill of coming-season streamflow and precipitation forecasting, markets at allow forecasts to be effectively communicated and utilized. have value. [we noted that if our market has already ordered people, assuming that they were responding to the default prior aggregate rainfall probability, then a forecast should not cause any reordering, rather only feeling more or less happy about the L,K investments that you made – this of course leaves room for value in investments from early forecast] [along these same lines, the forecast surely has value for ex-ante investing but surely agricultural economists have shown that, and the forecast value and the ordered contingent claims market value may be completely independent innovations]

An interesting feature of the system is that prices for differentiated claims might be used to estimate the market's consensus beliefs about the probability that claims of differentiated seniority will be honored — i.e., the market's consensus beliefs about the probability distribution governing X. An set of dissemination and policy questions concerns cases where the market's implied distribution differs from expert opinion. [note that this is quite different from informing a water manager about anything concerning his 'subjects' other than their beliefs, something that early on Dan had raised, so if we want to do that we need to think about that in addition i think]

Intertemporal with Storage: DEFINITELY DIFFERENT PAPER (have Art's first thoughts in another file)