Water Externalities: Tragedy of the Common Canal*

Charles A. Holt, Cathleen A. Johnson, Courtney A. Mallow, and Sean P. Sullivan† ‡

July 14, 2011

Abstract

Laboratory experiments are used to investigate alternative solutions to the allocation problem of a common-pool resource with unidirectional flow. Focus is on the comparative economic efficiency of nonbinding communications, bilateral "Coasian" bargaining, allocation by auction, and allocation by exogenous usage fee. All solutions improve allocative efficiency, but communication and bilateral bargaining are not generally as effective as market allocations. An exogenously imposed optimal fee results in the greatest allocative efficiency, closely followed by an auction allocation that determines the usage fee endogenously.

---

* Post-print version of Charles A. Holt, Cathleen A. Johnson, Courtney A. Mallow, & Sean P. Sullivan, *Water Externalities: Tragedy of the Common Canal*, 78 S. ECON. J. 1142 (2012).

† University of Virginia (Holt, Mallow, and Sullivan) and University of Arizona (Johnson). We wish to thank Claudia Antonacci, Andrew Barr, Rachel Blank, Anna Draganova, Stephanie Lawrence, Ricky Sahu, and Sara St. Hilaire for research assistance. We also received helpful comments from Ben Cohen, Lee Coppock, William Shobe, Nicholas Smith and Ilya Zlatkin. This research was funded in part by the National Science Foundation (SES0098400) and the University of Virginia Bankard Fund.

‡ Corresponding author: Sean P. Sullivan, sps2d@virginia.edu.

I. Introduction

Few economic insights are as celebrated or bemoaned as the tragedy of the commons. Interest stems from the large number of important natural and man-made resources with common-pool characteristics: that is, resources which are *rivalrous*, so that appropriation by any one user depletes the overall stock of resource available to other users, and *non-excludable*, so that it is difficult or impossible for any one user to prevent another from appropriating a portion of the resource. The *tragedy* of common-pool resources is allocative inefficiency. Acting in their own self-interest, users tend to over-exploit such resources, failing to limit appropriation to efficient levels, and failing to direct appropriation to the most efficient users.

An interesting special case of commons problems arises when resource availability follows a unidirectional flow, such that users make appropriation decisions in sequence. Examples include cumulative contributions to air pollution in persistent jet streams, the harvest of migratory fish populations, and the irrigation decisions of farmers located along a shared canal. The distinguishing feature of these common-pool resource problems is that the external cost of resource exploitation is only felt by users "downstream" of the appropriator. Sequential extraction commons problems may be especially painful when the value of exploitation is negatively related to the order of appropriation. An example is when favorable growing conditions mean field irrigation is more productive for downstream farmers, but upstream farmers exploit their positional advantage by over-appropriating the available water resource in irrigating their less productive fields. For example, Ostrom and Gardner (1993) describe the Thambesi water system, where "headlanders" have established first-priority water rights against those downstream. Overuse of water by the Thambesi "headlanders" during the pre-monsoon season results in an aggregate irrigation of fields only one tenth the size of what could have been achieved with a reallocation to downstream farmers (Yoder, 1986).[1] Similarly, Cárdenas et al. (2011) describe the Coello River Watershed in Columbia, where increases in water scarcity and the demand for irrigation by downstream users are highlighting inappropriate water-use habits of upstream users.

The textbook solution to the tragedy of the commons is to imbue the common-pool resource with broad-based property rights. In the absence of prohibitive transaction costs and

[1] Of course, not all watersheds have the property that those with primary water rights have the least productive land. For example, Cifdaloz, Regmi, Anderies, and Rodriguez (2010) describe a watershed in Nepal where irrigation infrastructure was built up over time to include ever more marginal lands, so that upstream users in fact have the most productive land.

liquidity constraints, the creation of fungible appropriation rights causes users to internalize the externality of appropriation; subsequent purchases and sales of appropriation rights then limit appropriation to efficient levels and reallocate appropriation to efficient users.[2] When property rights are difficult to enforce or politically infeasible to assign, however, a host of exogenous direct regulations may provide alternative solutions to the commons problem. Examples include government enforced taxes or quotas on resource appropriation. In addition, the path-breaking work of Elinor Ostrom and her coauthors has uncovered a rich variety of endogenous/institutional solutions to the commons problem—most of which avoid the need for exogenous imposition of property rights or heavy-handed regulation.[3]

In the special case of a common-pool resource with unidirectional flow, many of the same potential solutions remain available. Following Coasian logic, marketable resource appropriation rights may provide users with incentives to trade in such a way that resource allocation is diverted to its highest-value uses (see, e.g., Yoder, 1986). Alternatively, resource allocation might be directed by government intervention in the form of either an auction or exogenous pricing scheme on resource appropriation.[4] In some cases, informal social arrangements may effectively mitigate the "tragedy" without the need for any government intervention. For example, community coordination on investment decisions (such as contributions to a community-shared resource infrastructure) may provide a context for coordination on resource appropriation and for assigning social sanctions to violators of agreed appropriation limits (see, e.g., Ostrom and Gardner, 1993).

This is not to say that all of these potential solutions are equal, or even desirable in every situation involving a common-pool resource. Property rights are of limited value if appropriation cannot be monitored, and will not tend to affect efficient allocation outcomes when the transaction costs of bargaining are prohibitive. Efforts at direct government intervention through use limits or fees may be hampered by misleading information provided to regulators or by

[2] Readers from many disciplines will recognize this as the famous Coase theorem (Coase, 1960).

[3] See e.g. Ostrom et al. (1994) and the special section of the Fall 1993 *Journal of Economic Perspectives* devoted to "Management of the Local Commons."

[4] Practical market solutions to many water allocation problems appear increasingly viable. For example, Alevy et al. (2010) use a field experiment to compare revenue generated by alternative techniques for auctioning water rights, with particular focus on right-to-choose auctions. Murphy et al. (2000, 2009) use laboratory experiments to show that a two-sided uniform price auction reallocates water rights in a highly efficient manner, comparing favorably to a centrally-managed regulatory solution. Cummings, Laury, and Holt (2004), use laboratory experiments to design an auction-based procedure for reducing irrigation during a drought in the state of Georgia.

offsetting activities taken by the users.[5] Government intervention might even be counter-productive if it disrupts extant social norms that dictate relatively efficient appropriation habits (see, e.g., Cárdenas, Stranlund, and Willis, 2000).[6]

Social solutions may be particularly suspect for common-pool resources with unidirectional flow. Absent some *ex ante* community investment effort, the sequential structure of unidirectional externalities may preclude the opportunities for reciprocity and punishment needed to establish and enforce social appropriation norms. Information problems may also make social appropriation agreements difficult to maintain. For example, the geographic structure of large watersheds may compound informational problems if significant distances and heterogeneous farming operations limit opportunities for user interactions (Swallow et al., 2006).

Empirical properties of the general tragedy of the commons are well studied in an extant literature that explores common-pool resource problems through observational study, as well as laboratory and field experiments (see generally, Ostrom et al., 1994; Ostrom, 2006). Results conclusively demonstrate over-appropriation of common-pool resources (e.g. Walker et al., 1990), but also suggest the potential for mitigation through self-governance or internal social coordination (Ostrom et al., 1992; 1993; Cárdenas, 2000). The apparent efficacy of social solutions is variable (see, e.g., Ostrom et al., 1994), and appears to depend in part on the homogeneity of the social group (Cárdenas et al., 2002; Cárdenas, 2003). Endogenously developed appropriation rules can mitigate allocative inefficiency (Ostrom et al., 1992; Walker et al., 2000), and may outperform exogenously imposed rules when the government is a poor monitor of conformance (see Cárdenas et al., 2000).

The particular problem of a common-pool resource with unidirectional flow presents unique challenges, due to the exploitation of positional advantage in appropriation (see, e.g., Budescu et al., 1997; Cárdenas et al., 2008).[7] A recent study by Cardenas et al. (2011) considers whether over-appropriation of the common-pool resource by upstream users may cause downstream users to under-contribute to the upkeep of a common irrigation system, thus creating

[5] In fisheries, for example, limits on the season length result in larger boats. For an irrigation system, limits on pipe size may result in the use of more powerful pumps, etc. There is a saying in Spanish: "el que hace la regla, hace la trampa" (he who makes the rule, makes the trick).

[6] Cárdenas, Stranlund, and Willis (2000) conducted a field experiment in rural villages in Columbia. They find the application of rules and regulations that are imperfectly monitored and outside of informal community institutions tend to increase selfish, individualistic behavior—resulting in resource overuse.

[7] As the spatial and temporal dynamics of appropriation decisions become complex, the fine details of resource governance may become increasingly relevant to practical efficacy (Janssen et al., 2010).

a voluntary provision problem. To investigate this potential provision problem, the authors use a two-stage experiment involving a voluntary contributions stage followed by sequential extraction decisions.[8] Starting from about 50% of the full efficiency level, contributions increase to about 60% when participants are allowed to communicate prior to the contributions decision, but the introduction of fines for over-appropriation does not increase average contribution levels. For a similar game, Janssen, Anderies, and Joshi (2011) note that contribution inefficiencies become more severe as the number of users making contribution and sequential extraction decisions increases (from 2 to 5 subjects per group). Because users in these studies have constant and equal values for units of the resource, excessive extraction by upstream users does not imply an allocative inefficiency—rather, efficiency losses are entirely attributable to under-contribution in the provision of the resource.

The present paper focuses more sharply on the *allocative inefficiencies* that arise in simple unidirectional commons problems with a fixed supply of the resource. A simple experiment is explored in which each user has a collection of "fields" with productivities that range from high to low, and where over-appropriation by upstream users prevents downstream users from irrigating their generally more productive fields.

After demonstrating the stark allocative inefficiency that attains under these circumstances, we consider several potential solutions to the particular allocation problem of a common-pool resource with unidirectional flow. While much of the research on commons problems relies on field experiments for their high external validity, our focus on the comparative efficacy of solutions motivates greater attention to internal validity, and thus favors a research design based on a controlled laboratory experiment. Four potential solutions to the allocation problem are considered: (1) non-binding communication, (2) binding bilateral bargaining, (3) resource allocation by auction, and (4) resource allocation by imposition of an optimal appropriation fee.

The remainder of the paper proceeds as follows. Section II describes the general game and specific treatment environments studied in the experiment. Section III explores the results of the experiment, focusing on behavior and observed efficiency in each environment. Section IV discusses the apparent robustness of experimental results. Section V concludes with a discussion of observed themes and potential extensions.

---

[8] This was a framed lab experiment run in the field, with participants recruited from villages in Columbia and Kenya.

II. Procedures

The experiment described in this paper was conducted in 25 independent sessions. Each session involved 6 unique and inexperienced subjects recruited from the undergraduate class of the University of Virginia. A session consisted of compensation-incentivized decision making in 3 rounds of a baseline environment, followed by 3 rounds of one of 5 exogenously assigned treatment environments. The decision environment in every round was a simple sequential extraction commons problem with an institutional solution that varied by treatment.

To model a sequential extraction problem with a static order of resource extraction, the order of decision making was kept the same in every round of every session.[9] The relatively small number of rounds in the experimental design reflects the desire to maintain subject interest, and mitigate hurt feelings in subjects assigned the final extraction decision—i.e. in subjects who are most severely affected by the commons problem in every round of the session. This design lends power to inferences drawn from session averages, but provides insufficient data for careful analysis of individual-subject decisions. All sessions were run with web-based VeconLab software, using the Water Externalities program.[10] The rich terminology of farmers using water from a shared canal to irrigate their fields was used to keep the decision-making context clear to the participants.[11] The same context was used in all treatments.

Participants in this experiment were given the role of farmers located at various points along a shared canal that provided an exogenous and scarce water resource. Each session assigned subjects numbered identities according to the "address" of their farm: e.g. ID 1 moved first, ID 2 moved second, etc. The decision sequence was repeated over several rounds, with subject addresses the same in every round.

Each participant was endowed with 4 fields of randomly determined productivity. The productivity value for a given field corresponded to the cash value of the crops the field yielded

9 This is an appropriate model of real-world sequential extraction commons problems where the order of extraction is fixed by location or property rights. For example, farmers located along a canal will generally make extraction decisions in the same order every growing season.

10 The program is available online at http://veconlab.econ.virginia.edu/admin.php for instructor setup and at http://veconlab.econ.virginia.edu/login.php for participant login. Setup options are flexible in terms of the numbers of participants, the numbers of fields and the ranges of their random productivity draws, the possibility of random changes in the water stock, etc. Instructions for participants are configured automatically to match the selected setup. These instructions are presented to participants prior to the first round and prior to the round following a treatment change.

11 Cf. Janssen, Anderies, and Joshi (2011) who contextualize contribution and sequential extraction decisions in terms of bandwidth and file downloads to present a decision environment familiar to student participants.

in the absence of irrigation. An irrigated field yielded a cash value of triple its productivity value. To the extent that water was available, each participant decided in sequence whether or not to irrigate each of his or her 4 fields in a given round. A total stock of 12 "units" of water was available in each round, and every decision to irrigate a field depleted one unit from the available stock.

Rounds corresponded to different growing seasons such that each round reset the water supply to 12 units and randomly regenerated field productivity values. Stochastic field productivities are meant to capture sources of individual-specific exogenous productivity variation, such as exogenous shocks in local conditions, insect populations, and worker health. Variable productivities also help to promote subject interest throughout several rounds of a similar decision-making environment.

In deciding whether or not to irrigate a field, subjects in all treatments were able to view their own field productivity draws, but not the productivity draws of other users. Subjects were always told the available stock of water remaining at the time of decision, but were provided with the amounts of water appropriated by specific upstream users only in specific treatments, as noted below.

Table 1 displays the ranges of random field productivity values for the upstream producers (IDs 1-3) and the downstream producers (IDs 4-6). This design assigns multiple subjects to each type (rather than having only a single upstream and downstream user) to reflect the bargaining and information complexities inherent in situations with multiple users of the common-pool resource. Productivities were distributed as discrete-uniform random variables: e.g. high-productivity fields were equally likely to have values of $7, $8, $9, $10, or $11, and low-productivity fields were equally likely to have values of $2, $3, $4, $5, or $6. Participants did not know the distributions used to generate field productivity draws.

Table 1. Fields and Ranges of Base Productivity Values (Tripled with Irrigation)

Field Number	Upstream Producers (IDs 1-3)	Downstream Producers (IDs 4-6)
1	$7-$11	$7-$11
2	$2-$6	$7-$11
3	$2-$6	$7-$11
4	$2-$6	$2-$6

The efficient allocation applies the 12 available water units to irrigate the 12 highest productivity fields in a given round. Since the productivity ranges for the two types of fields do not overlap, this corresponds to allocating 1 unit of water to each of the upstream producers (IDs 1-3), and allocating 3 units of water to each of the downstream producers (IDs 4-6). If upstream producers behave selfishly, however, they will each take 4 units of water, leaving no available stock for the downstream producers with the greater number of high-productivity fields.[12]

Since irrigation triples yield values, the net gain from irrigation is twice the yield value of a given field. Notice that an optimal fee is a price for water of $13, which would always deter farmers from irrigating low-productivity fields, but would never deter farmers from irrigating high-productivity fields. The imposition of such a fee theoretically yields 100% efficiency, as compared to the approximate 75% efficiency that results from purely selfish behavior under these parameter values.

A *status quo* baseline environment, as just described, was assigned to the first three rounds of every session. In the final three rounds of a session, one of five different treatments was exogenously assigned: a repetition of the baseline environment ("baseline"), non-binding communication ("chat"), bilateral bargaining with chat ("bargaining"), an auction of water rights ("auction"), or an optimal irrigation fee ("optimal fee").

The repeated baseline treatment was used to gauge potential learning or order effects that may bias repeat measurements in the experimental design. To the extent that any such artifact might be observed, data from the repeat baseline treatment provide a basis for comparison that is corrected for the effects of participant experience.

The chat treatment provided participants 3 minutes to communicate in an online chat room. After the chat period ended, participants made extraction decisions just as in the baseline, except that each subject was able to view the specific water use decisions of upstream farmers. Providing the specific appropriation decisions of upstream users contemplates a solution in which communication is tied to a monitoring capacity: e.g. where communication allows farmers to aggregate individual accounts of the appropriation decisions of their immediate neighbors. As discussed previously, there is a large literature on the effects of communication in common-pool resource dilemmas. The purpose of this treatment was to determine how a controlled amount of

[12] Although not universal, this inverse relationship between productivity and the order of appropriation appears typical of many watersheds. In cases where upstream users have more productive land, sequential extraction may not result in a loss of efficiency.

social interaction and monitoring might enhance efficiency in this simple model of sequential appropriation, in order to provide a basis for comparison with market-based policies to be discussed next.

The bargaining treatment afforded two types of agreements: (1) a contract to pay an upstream user $P for restricting their irrigation to at most Q units, and (2) a contract to accept $P from a downstream user in exchange for restricting one's own irrigation to at most Q units. All contracts were bilateral, but participants could form contracts with any number of upstream and downstream users. For example, ID 1 might agree to restrict irrigation to 3 units in exchange for a payment of $2 from ID 4, and also to restrict irrigation to 2 units in exchange for a payment of $10 from ID 5. In this case, ID 1 would receive a total of $12 and would be limited to use at most 2 units of water. Note that the property rights being exchanged are an upstream user's rights to appropriate water, not the upstream user's right to exclude others from appropriating water—as noted previously, this non-excludability property is a driving element of the commons problem. This treatment was motivated by the Coase theorem, which turns critically on the absence of transaction costs. Although there were no explicit transaction costs in the experiment, time limits and the need to engage in multiple interrelated negotiations could generate substantial indirect transaction costs. Participants were also hampered by not knowing each other's productivity values when negotiating contracts. To facilitate coordinated bargaining, subjects in this treatment were permitted to communicate in a chat room during the bargaining process. The chat time was set at 6 minutes in each round to provide participants sufficient time to negotiate binding bilateral contracts. Public extraction decisions were employed to allow subjects to see that their contracts were being enforced.

The auction treatment imposed a permit requirement for field irrigation. All farmers, regardless of address, had the opportunity to bid for up to 4 permits each. The highest 12 bids were selected, and the price paid for all permits was the highest rejected bid (i.e. the 13th bid), with collected revenue retained rather than returned to the users. This is a multi-unit, uniform-price auction with private values, so it was never optimal for a user to bid above value.[13] Bidding below value at the rejection margin could, however, reduce the price paid for other

[13] This setup is similar to the multi-unit uniform-price auction that was implemented by the Regional Greenhouse Gas Initiative (RGGI) for the sale of allowances for carbon dioxide emissions from electric power generators in 10 northeast states. Laboratory experiments were used to refine recommended auction procedures (Holt, et al., 2007 and Burtraw, et al., 2009).

permits. Therefore, bidding at value was not necessarily an equilibrium strategy, as would be the case in a second-price auction with a single prize. If bids mirror values, an auction would select the high-value users. The resulting allocation would be efficient, and the clearing price would constitute an optimal usage fee. The purpose of the auction treatment was to determine how effectively a market process could approximate an optimal usage fee.

Finally, the optimal fee treatment also imposed a permit requirement for irritation. In contrast to endogenous determination of the fee in the auction treatment, the optimal fee treatment exogenously imposed an optimal per-unit fee of $13 for each water unit used, simulating an appropriately set Pigouvian tax. To keep the decision process simple and comparable to the auction treatment, revenue from the fee was not returned to the participants.[14]

Reported results are based on a total of 25 six-person sessions, run between March and December 2009. Sessions lasted from 35 to 60 minutes, depending on which treatment environment was assigned. Participants received $6 for showing up, and were paid a cash amount equal to 4% of the money they earned in the experiment. Earnings depended on the treatment, but generally ranged from $12 to $30, including the initial $6 payment.

III. Results

In every environment, there exists a unique optimal allocation in which water is used to irrigate the 12 most productive fields. Efficiencies are calculated as a percentage of this optimal allocation. Round-average efficiencies are shown in Figure 1, where each line represents an average over all 5 sessions for a specific environment. The dashed gray line lying below the others tracks the predicted efficiency in the "selfish" outcome where all water is taken by the three upstream farmers. Note that selfish predictions are always around 75% efficiency, with some slight variability due to random productivity draws. The legend labels on the right indicate the treatment applied in rounds 4-6.

[14] We also ran 5 sessions in which the treatment involved an optimal fee, but fee revenues were equally divided among the farmers. This treatment is not reported, since the results are comparable to the optimal fee treatment with no rebate.

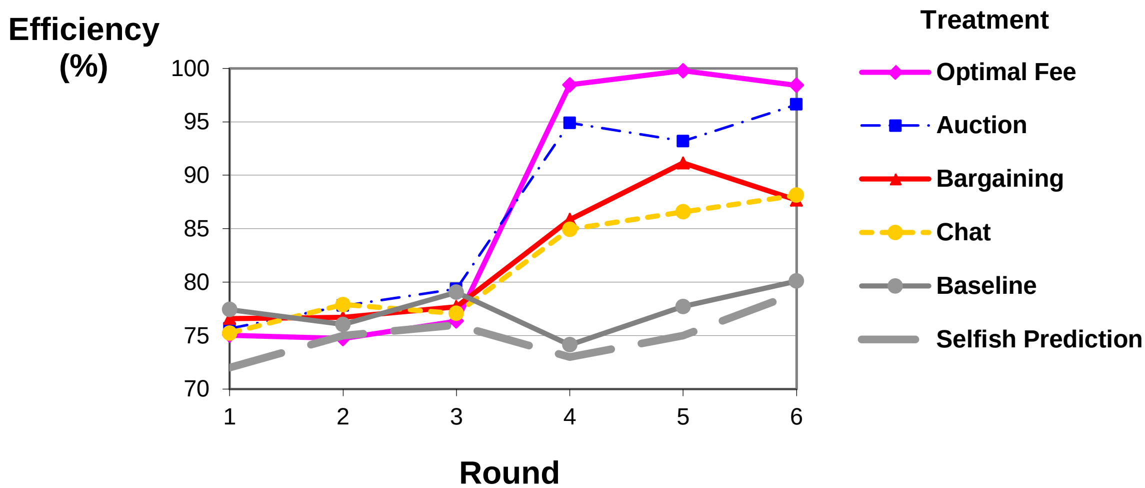

Figure 1. Efficiency Averages by Treatment

As Figure 1 illustrates, the highest allocative efficiencies are observed when an optimal fee is exogenously imposed, though an auction of water use rights is a close second. Bargaining and chat environments are less efficient, and exhibit little difference from each other. Under baseline conditions, average efficiencies are 1 to 3 percentage points higher than the purely selfish predictions, indicating a small amount of altruistic behavior in this experiment. Recall that in the experimental environment, the unidirectional flow of water and static location of subjects means that acts of generosity cannot be reciprocated, nor can acts of selfishness be punished unless learned of and acted upon by an upstream user.

Session-average efficiencies are arrayed in Figure 2 for the baseline environment (rounds 1-3), and in Figure 3 for the treatment environments (rounds 4-6). The order of sessions from left to right in Figure 2 matches that in Figure 3. Notice that there are 5 bars in each treatment cluster, representing average efficiencies for each of the 5 sessions.

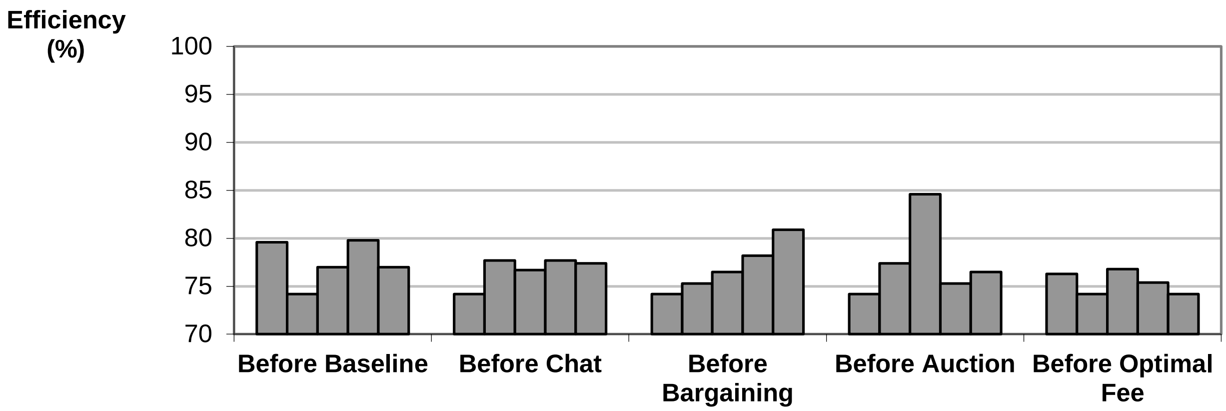

Figure 2. Baseline Efficiencies (Rounds 1-3) by Session and Treatment

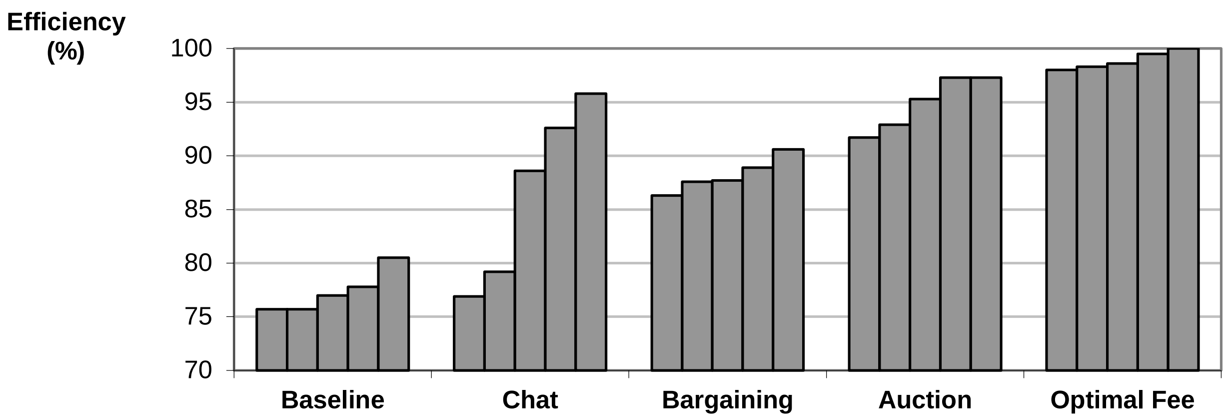

Figure 3. Treatment Efficiencies (Rounds 4-6) by Session and Treatment

A quick glance at these data affords several qualitative observations. Despite some variability, efficiencies in the baseline environment look basically homogeneous across experimental sessions. Other than in the bargaining sessions, there is no apparent correlation between session-average efficiencies in baseline rounds 1-3 and treatment rounds 4-6. Efficiencies do vary considerably both within and across treatments. For example, the minimum efficiency in the optimal fee treatment is greater than the maximum efficiency in any other treatment environment. While it is tempting to declare that observed efficiencies admit a monotone ordering by treatment, the variability of efficiencies in the chat treatment suggests the need for a more nuanced analysis.

Reading too much into patterns based on limited numbers of observations is always dangerous, and the prudent question is whether observed relationships can be explained as more than chance variation. The remainder of this section discusses what inferences these data provide about the relative merits of the considered solutions to the problem of a common-pool resource with unidirectional flow.

Result 1: *Average efficiency in the baseline environment is slightly greater than would be expected under purely selfish behavior*.

The selfish prediction for this experiment has the first three farmers consuming four units of water each, leaving no residual irrigation for the downstream farmers who have more productive fields. Across the baseline environment (rounds 1-3) this allocation corresponds to a predicted average "selfish efficiency" of 74.3%. Casual inspection of Figures 1-3 suggests that observed efficiencies in the baseline environment are slightly greater than the selfish prediction. This conclusion is supported by statistical inference, as we are firmly able to reject the claim that the average baseline efficiency per session equals the selfish prediction at any reasonable level of significance.[15] Of course, rejection of equality does not imply a large inequality, and a 95% confidence interval places the average baseline efficiency only between 75.8 and 77.9%.[16] Thus, while we are confident that average baseline environment efficiency exceeds the selfish prediction, the difference is evidently small.

Result 2: *All non-baseline treatment environments provide efficiency gains over the baseline environment.*

Every potential solution to the common-pool resource problem studied in this paper is meant to improve upon the efficiency of the *status quo* baseline environment. To test the merits of each solution, we explore within-session efficiency gains between baseline and treatment rounds by treatment. This amounts to calculating the difference between average baseline and treatment efficiencies for each session, and then comparing the average differences to zero by treatment type.[17]

---

[15] A sign test comparing session-average efficiency in the baseline environment to the selfish prediction results in a *p-value* of less than 0.01.

[16] Interval constructed by the usual *t* test inversion.

[17] Within-session comparisons exploit pairing of baseline and treatment environments within each session to help

In testing the one-sided alternative that average efficiency is greater under the treatment environment than it is under the baseline environment, we find compelling evidence that each non-baseline treatment does in fact improve upon the average baseline efficiency. One-sided exact *p*-values for each treatment are provided in Table 2; these correspond to a one-sided application of Wilcoxon's signed-rank test.

Table 2: P-values from One-Sided Tests that Average Efficiency Gain Exceeds Zero[18]

Treatment	*p*-value
Baseline	0.59380
Chat	0.03125
Bargaining	0.03125
Auction	0.03125
Optimal Fee	0.03125

Individual tests conform to *a priori* expectations, indicating strong evidence of efficiency gains under every non-baseline treatment. In contrast, the repeated baseline treatment provides no evidence of an efficiency gain, which helps to mitigate concerns that repeated play or sequence effects may be driving experimental results.

When performing many simultaneous tests, there is always a concern that some rejections may result from random variation alone.[19] Thus, when attempting to draw inferences from the combined results of many individual tests, it is sometimes prudent to check whether conclusions differ under stronger rejection rules than simple per-test rejection criteria. A common technique is to use a test which controls the familywise error rate, defined as the probability of *even a single false rejection* among k simultaneous tests.[20] For these data, a joint test of all non-baseline treatment environments leads to the same conclusion—that all non-baseline treatments lead to efficiency gains over the baseline—at the familywise 0.1 level.[21]

mitigate the consequences of unobserved heterogeneity.

[18] Equality of p-values across non-baseline treatments is an artifact of the discrete null distribution of the Wilcoxon signed-rank test.

[19] For example, consider running 20 statistically independent tests at the 0.05 level, and suppose all null hypotheses are in fact true. Since the probability of a false rejection is 5% in each individual test, one false rejection is expected of the 20 tests performed. In fact, the probability of at least one false rejection is $1 - (0.95)^{20}$, or 64%.

[20] Note that this is a very conservative test that is appropriate when a false rejection of the null can have serious consequences, e.g. administering a drug when it actually has no beneficial effect.

[21] The reported rejection corresponds to a sequential Bonferroni-type test described by Hochberg (1988).

Having determined that all examined solutions increase efficiency over the baseline, the next logical question is how much of an improvement each solution affords. To address this question, Figure 4 provides 95% confidence intervals for average efficiency gains over the baseline measure under each treatment environment.[22]

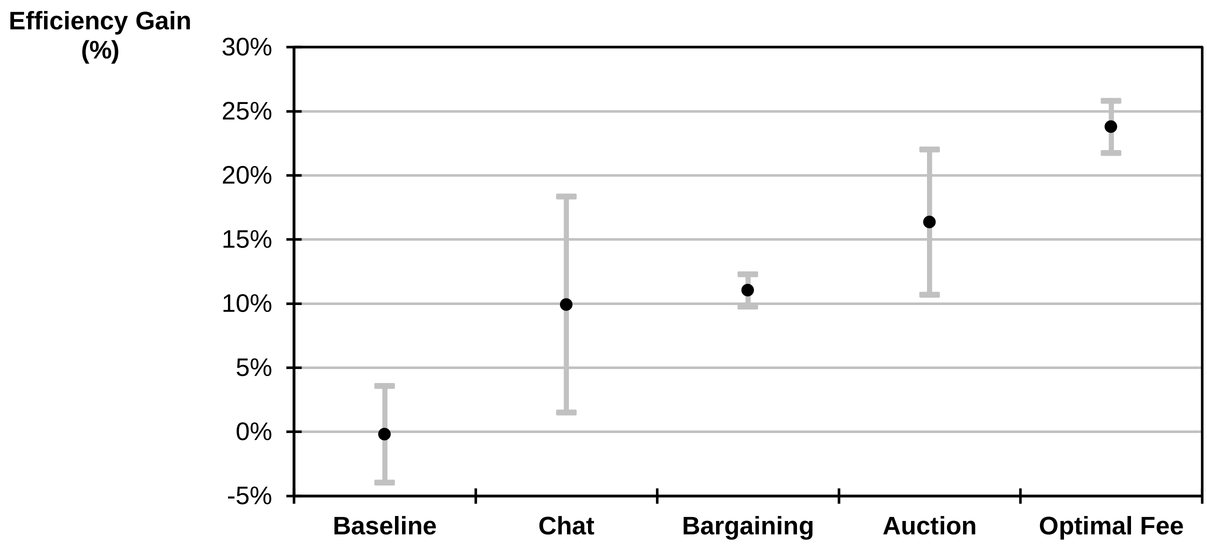

Figure 4. 95% Confidence Intervals for Average Efficiency Gain by Treatment

Confidence intervals for the average efficiency gain vary considerably by solution. The 95% confidence interval for the optimal fee solution is relatively tight, and includes as much as a 25% average efficiency gain over the baseline environment—corresponding to approximately 100% efficiency under this solution. In contrast, outcomes under the chat treatment are sufficiently variable that a 95% confidence interval contains nearly zero efficiency gains, as well as gains of almost 20%. Clearly, we cannot use these data to speak with much precision about the average efficiency gain resulting from non-binding communication. This is not terribly surprising, as the efficacy of social solutions to commons problems are known to be variable (see, e.g., Ostrom et al., 1994). Chat logs from the experiment reveal that some groups managed to effectively establish loose behavioral norms or some degree of social responsibility, reinforced by a mixture of pleas and positive reactions to generosity. In constrast, other groups failed to establish any such covenant, and chat instead devolved into a series of complaints and frustrations.

[22] Confidence intervals are constructed by inversion of Wilcoxon's signed-rank test. Note that these confidence intervals are analogous to two-sided tests, while the hypotheses tested above are one-sided.

Result 3: *Average efficiencies differ between most treatment environments.*

Given that all the solutions considered in this paper appear to increase efficiency to some degree, the next step is to decide whether anything can be said about which ones work better than others in this simple model of sequential extraction. To address this question, we rely on between-session variation in comparing average efficiencies across the various treatment environments. At the most fundamental level, the question is whether we can be certain of any difference between treatments in the first place. Casual inspection of Figure 4 strongly suggests we can, and formal statistical tests agree: we reject the possibility that efficiency gains are equal across treatments at every reasonable level of significance.[23]

Of course, the interesting question is not whether the treatment effects of the various solutions differ, but how they differ. To address this point, we conduct a multiple comparisons test of all pairwise contrasts between treatments using the Wilcoxon-Mann-Whitney test.[24] Table 3 summarizes inferences gained from each comparison: reported *p*-values are exact.

Table 3: *P*-values from Two-Sided Tests of Common Location.[25]

Comparison	*p*-value
Bargaining vs Chat	1.00000
Chat vs Auction	0.09524
Bargaining vs Auction	0.00794
Auction vs Optimal Fee	0.00794
Chat vs Optimal Fee	0.00794
Bargaining vs Optimal Fee	0.00794

[23] Kruskal-Wallis tests for equality of location yield asymptotic *p*-values of less than 0.005 whether or not the repeated baseline treatment is included in the comparison.

[24] The intuition behind this test is easily illustrated for the special case where all 5 observations under one treatment are lower than all 5 observations under another treatment. Under the null hypothesis that differences between treatments are simple due to random noise, then of the "10 take 5" = 252 ways of permuting these numbers, only 2 of these (all 5 greater under one treatment and all 5 less under one treatment) are as or more extreme that what was observed. Under the null hypothesis, the probability of this outcome is 2/252 = 0.00794, as shown by the bottom 4 rows of Table 3.

[25] The equality of several p-values is an artifact of the discrete null distribution of the Wilcoxon-Mann-Whitney rank-sum test.

Specific conclusions drawn from this family of tests are provided in the next three results. As noted previously, the more statistical tests one conducts, the more false rejections one can be expected to produce. This is never a problem on a per-comparison basis, but can sometimes muddy conclusions drawn from looking comprehensively at the results of a family of tests. For completeness, we comment on how conclusions differ under the stronger requirement of controlling the familywise error rate, where appropriate.

Result 4: *The optimal fee treatment yields higher average efficiency than any other treatment environment.*

In terms of simply increasing efficiency over the *status quo,* the optimal fee solution is a clear winner. This conclusion stands whether or not one chooses to take the more conservative approach of controlling the familywise error rate.[26] The observation of nearly 100% efficiency in this treatment is, of course, consistent with economic theory. Since the price of irrigation is fixed at a level that causes all farmers to internalize the social opportunity cost of water appropriation, even a small dose of individual rationality should be sufficient to affect a socially optimal allocation.

Unfortunately, the practicality of this solution to the common-pool resource problem appears limited. There is no reason to expect that an optimal fee would be obvious in a typical policy-making setting, particularly when users have incentives to selectively report valuations and lobby for lower fees. Because a fee-based solution could fail quite miserably if the fee were set at the wrong price, difficulty in determining the proper fee may translate into substantially lower practical efficacy in many settings.

Result 5: *There is weak evidence that the auction treatment yields higher average efficiency than either the chat or bargaining environments.*

Because an auction uses a market mechanism to "discover" the market-clearing fee, it is

[26] The optimal fee treatment is concluded to provide higher average efficiency than any other non-baseline treatment when using the Hochberg (1988) algorithm to control the familywise error rate at the 0.025 level.

not surprising that it should closely follow the optimal fee treatment in terms of average efficiency. Efficiency of the auction treatment clearly surpasses that of the bargaining environment, and superiority of the auction over the chat environment is also evident, albeit with a less impressive *p*-value. We draw the same conclusion when controlling familywise error rates at the 0.2 level, but fail to reject that average efficiency is the same under chat and auction treatments at lower levels of the familywise error rate.[27]

Result 6: *There is no evidence that bilateral bargaining results in greater average efficiency than simply allowing participants to communicate in a non-binding way.*

Because externalities are fundamentally problems of property rights, the Coase theorem argues that private bargaining in the context of well-defined property rights should result in socially optimal allocations. By contrast, allowing farmers to engage in non-binding communication without the ability to make and enforce contracts provides no theoretic argument for an efficiency gain over selfish behavior. While we would have expected the bargaining treatment environment to exhibit greater average efficiency than the chat treatment, the data fail to support this claim.

One possible explanation for the dismal performance of private bargaining is the potentially serious obstacle of transaction costs, which are assumed away in the Coase theorem. Although property rights are well-defined and there are no explicit transaction costs in this treatment, a downstream farmer has to make multiple contracts with upstream farmers in order to ensure water availability. With no centralized coordinator, the difficulty of forming an appropriate menu of contracts can represent a substantial implicit transaction cost.[28] There is

[27] The auction treatment is concluded to provide higher average efficiency than either the bargaining or chat treatments when using the Hochberg (1988) algorithm to control the familywise error rate at the 0.2 level. At lower levels, there is not sufficient evidence to statistically distinguish the auction and chat environments. Although 0.2 is higher than contemporary standards of "statistical significance" as applied to individual hypothesis tests, it is reasonable among tests controlling the familywise error rate. Intuitively, rejection at this level corresponds to allowing for no more than a 20% chance of experiencing even a single false rejection among all six comparisons conducted in Table 3.

[28] A strikingly similar result is found in a network formation experiment. Connecting to the network is a contribution to the public good. If each player (node) connected to its nearest neighbor(s), players would enjoy higher earnings. If all players did not connect, the players that had connected would suffer a loss. The coordination problem created enough of a barrier that no "chain networks" could form in the laboratory (Deck and Johnson, 2001).

also a free-riding problem, since various farmers may benefit from contracts to which they are not a party: in the words of an ID 6 participant during the chat phase of bargaining, “I would sign with you player 1 [ID 1], but the water doesn’t seem to get to me anyway.”

IV. Robustness Tests

The experimental design’s reliance on 3 rounds of replication for every decision environment is relatively short compared to other experiments in the extant literature. The small number of rounds is motivated by several considerations. First, there is a high time-cost to running many replications of chat and bargaining treatments due to the required chat periods. Second, the relative simplicity of the decision problem suggests that behavior should stabilize rapidly in these environments. Third, focus on session-averages in data analysis means that additional rounds of play do not increase effective sample sizes. Fourth, fairness concerns counsel against assigning subjects to last-mover (i.e. persistent loser) status across too many rounds of the experiment.

These considerations aside, it is possible that treatment effects could be diminished or augmented when averaged over a longer sequence of rounds (which would require either higher total payoffs or lower per-round incentives). As noted in Section III, results from the repeated baseline treatment reveal no evidence that behavior tends to change over 6 rounds of the baseline environment. To provide additional confidence, however, 5 additional sessions of the experiment were conducted in September 2010. These sessions were exact replications of chat-treatment sessions described in Section III, except that the treatment environment was replicated 10 times instead of only 3. Of all the decision environments considered in this study, it was thought that the social reinforcement mechanisms of cheap-talk and publicly visible extraction decisions in the Chat environment had the greatest chance of yielding behavior that evolved over time (either toward cooperative or fully-selfish outcomes).

Results for these additional sessions are illustrated in Figure 5. As illustrated by the black average line in Figure 5, efficiency varied considerably by round, but evinced no substantial trend that would indicate design bias from relying on results of only 3 rounds per environment. Statistical inference corroborates this qualitative conclusion.[29]

[29] Comparison of average efficiency in rounds 4-8 and 9-13 (the first and second half of the chat treatment) fail to reject equality of average efficiency at every interesting level of significance. Regression of allocative efficiency on round number suggests a loss of less than 0.5% efficiency per round; this trend is only significantly different from

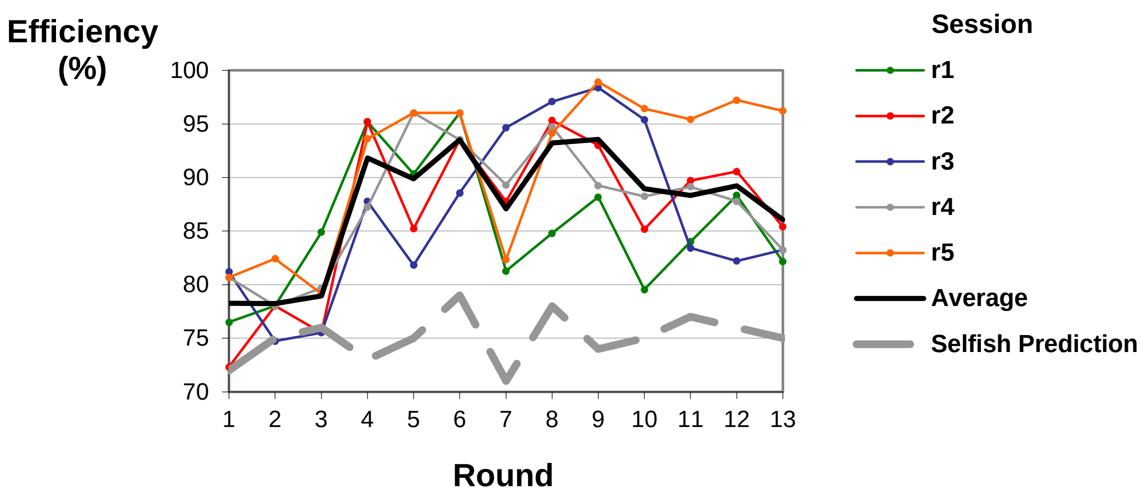

Figure 5. Average Efficiency by Round for Longer Chat Sessions (Round 4-13)

V. Conclusion

This paper is inspired by the rich array of commons problems studied by Elinor Ostrom and her collaborators. One lesson of these field studies is that there is not necessarily a "tragedy" in all common-pool resource environments. The laboratory experiment reported here is intended to compare various solutions to the practical allocation problems that may arise when a common-pool resource has a unidirectional flow. Results confirm the existence of an allocation problem at a level consistent with selfish behavior. This problem is mitigated by treatments designed to mirror certain social institutions and government interventions.

In the setting we investigate, there exists a unique optimal fee, imposition of which causes full internalization of all usage externalities. Unsurprisingly, experimental results show that the exogenous imposition of this fee yields nearly 100% efficiency. A solution relying on a uniform price auction for water permits is not as efficient, but the difference appears to be relatively small. The advantage of an auction solution is that the usage fee is endogenously discovered, which is of great practical importance when the optimal fee is not generally known *a priori*.

When property rights are well-defined and contracts are binding, the Coase Theorem suggests that private bargaining should result in optimal allocations, at least in the absence of significant transaction costs. The bargaining treatment of the experiment implements binding bilateral contracts without explicit negotiation costs. However, because property rights do not

zero at the 0.1 level, and is insignificantly different from zero if final round data are omitted from the regression.

support exclusion and contracts are constrained to be bilateral, participants may have to arrange sequences of contracts in order to ensure water flow to the fertile downstream fields; this source of complexity may represent an implicit transaction cost. Specifically, there is a free-rider problem in the sense that participants located between two parties to a contract may take the water that the upstream party agrees not to use. In this setting, we observe that bargaining has no more effect than a somewhat mild social-pressure treatment that permits participants to talk to each other in a chat room and observe others' decisions (the bargaining treatment also permitted a chat phase and social observation).

An extension of this study may be to revise the bargaining environment to allow users to see which other pairs of users have formed bilateral agreements; the content of the agreement (i.e. how many units of water a user has agreed not to exploit) may also be displayed. This deviation could mitigate some of the coordination and free-riding problems associated with the current bargaining environment, but its relevance depends on the ability of users to cheaply monitor others' private agreements in practical situations. In cases where monitoring is costly or impossible (i.e. if parties agree to sign nondisclosure agreements), the present bargaining environment may be a more appropriate model.

A related extension of the bargaining environment allows bilateral agreements to stipulate direct transfers of water units. For example, instead of contracting to reduce an upstream user's total water usage, a downstream user might simply "buy" a unit of water from the upstream user—circumventing both contractual complexity and free-riding problems. While this alternative specification seems likely to achieve greater average efficiency, its practicality depends on the enforceability of direct transfers that bypass intermediate users.

It is well known from Ostrom's original studies (and a large subsequent literature on voluntary contributions with punishments) that direct punishment opportunities can often solve a commons problem. While sequential extraction forecloses punishment in our model, the capacity for ex-post rewards may achieve similar results if upstream users expect to be sufficiently compensated by downstream users for self-limiting irrigation decisions. Finally, an alternative extension of our experiment would be to determine whether there is also a political solution in which participants vote on irrigation restrictions or usage fees, with fee revenues being distributed to participants in some manner.

References

Budescu, D.V., W.T. Au, and X.P. Chen. (1997) "Effects of Protocol of Play and Social Orientation on Behavior in Sequential Resource Dilemmas," *Organizational Behavior and Human Decision Processes*, 69(3): 179-193.

Burtraw, D., J. Goeree, C. Holt, E. Myers, K. Palmer, and W. Shobe. (2009) "Collusion in Auctions for Emissions Permits: An Experimental Analysis," *Journal of Policy Analysis & Management*, 28(4): 672–691.

Cárdenas, J.C. (2000) "How Do Groups Solve Local Commons Dilemmas? Lessons from Experimental Economics in the Field," *Environment, Development and Sustainability*, 2: 305-322.

Cárdenas, J.C. (2003) "Real Wealth and Experimental Cooperation: Evidence from Field Experiments," *Journal of Development Economics*, 70(2): 263-289.

Cárdenas, J.C., M. Janssen, and F. Bousquet. (2008) "Dynamics of Rules and Resources: Three New Field Experiments on Water, Forests and Fisheries," *Handbook on Experimental Economics and the Environment,* ed. J. A. List and M. Price, Edward Elgar Publishing.

Cárdenas, J.C., L.A. Rodríguez, and N. Johnson. (2011) "Collective Action for Watershed Management: Field Experiments in Colombia and Kenya," *Environment and Development Economics,* 16(3): 1-29.

Cárdenas, J.C., J.K. Stranlund, and C.E. Willis. (2002) "Economic Inequality and Burden-sharing in the Provision of Local Environmental Quality," *Ecological* Economics, 40: 379-395.

Cárdenas, J.C., J. Stranlund, and C. Willis. (2000) "Local Environmental Control and Institutional Crowding Out," *World Development*, 28(10): 1719-1733.

Cifdaloz, O., A. Regmi, J.M. Anderies, and A.A. Rodriguez. (2010) "Understanding Robustness in Small-Scale Social-Ecological Systems: the Pumpa Irrigation System in Nepal. *Ecology and Society*, 15(3): 39.

Cummings, R., S. Laury, and C.A. Holt. (2004) "Using Laboratory Experiments for Policy Making: An Example from the Georgia Irrigation Reduction Auction," *Journal of Policy Analysis and Management,* 3(2): 341-363.

Coase, R. (1960) "The Problem of Social Cost," *Journal of Law and Economics,* 3: 1-44.

Deck, C. and C. Johnson. (2004) "Link Bidding in a Laboratory Experiment," *Review of Economic Design*, 8(4): 359-72.

Gardner, R., E. Ostrom, and J.M. Walker (1990) “The Nature of Common-Pool Resource Problems,” *Rationality and Society*, 2(3): 335-358.

Hochberg, Y. (1988) “A Sharper Bonferroni Procedure for Multiple Tests of Significance,” *Biometrika*, 75(4): 800-802.

Holt, C., W. Shobe, D.Burtraw, K. Palmer, and J. Goeree. (2007) “Auction Design for Selling CO2 Emission Allowances under the Regional Greenhouse Gas Initiative,” RGGI Reports. Albany, NY: New York State Energy Research and Development Authority.

Janssen, M.A., R. Holahan, A. Lee, and E. Ostrom. (2010) “Lab Experiments for the Study of Social-Ecological Systems,” *Science*, 328(5978): 613-617.

Janssen, M.A., J.M. Anderies, and S.R. Joshi. (2011) "Coordination and Cooperation in Asymmetric Commons Dilemmas," Forthcoming in *Experimental Economics*.

Ostrom, E., and R. Gardner. (1993) “Coping with Asymmetries in the Commons: Self-Governing Irrigation Systems Can Work,” *Journal of Economic Perspectives,* 7(4): 93-112.

Ostrom, E., R. Gardner, and J.M. Walker (1994) *Rules, Games, and Common-Pool Resources*. Ann Arbor, MI: University of Michigan Press.

Ostrom, E., J. Walker, and R. Gardner. (1992) “Covenants With and Without a Sword: Self-Governance is Possible,” *American Political Science Review,* 86(2): 404-417.

Ostrom, E. (2006) “The Value-Added of Laboratory Experiments for the Study of Institutions and Common-Pool Resources,” *Journal of Economic Behavior & Organization*, 61(2): 149-163.

Swallow, B., N. Johnson, R. Meinzen-Dick and A. Knox. (2006) ‘The Challenges of Inclusive Cross-scale Collective Action in Watersheds,’ *Water International*, 31(3): 361-376.

Walker, J.M., R. Gardner, and E. Ostrom. (1990) “Rent Dissipation in a Limited-Access Common-Pool Resource: Experimental Evidence,” *Journal of Environmental Economics and Management*, 19: 203-211.

Walker J.M., R. Gardner, A. Herr, and E. Ostrom. (2000) “Collective Choice in the Commons: Experimental Results on Proposed Allocation Rules and Votes,” *Economic Journal*, 110(460): 212-234.

Yoder, R.D. (1986) “The Performance of Farmer-Managed Irrigation Systems in the Hills of Nepal,” *Ph.D. Dissertation*, Cornell University.

Appendix A: Data on Average Efficiencies by Session

Session ID	Avg. Efficiency: Rounds 1-3	Avg. Efficiency: Rounds 4-6	Treatment Rounds 4-6
c1	74.190	76.912	Chat
c2	77.726	79.209	Chat
c3	77.723	92.595	Chat
c4	77.423	95.752	Chat
c5	76.678	88.616	Chat
b1	80.852	90.604	Bargaining
b2	76.450	87.743	Bargaining
b3	78.248	88.867	Bargaining
b4	74.190	86.293	Bargaining
b5	75.289	87.575	Bargaining
a1	74.190	91.718	Auction
a2	76.504	97.349	Auction
a3	84.597	95.267	Auction
a4	75.289	97.301	Auction
a5	77.492	92.852	Auction
f1	76.279	98.018	Optimal Fee
f2	75.408	99.456	Optimal Fee
f3	76.797	98.582	Optimal Fee
f4	74.190	100.000	Optimal Fee
f5	74.190	98.342	Optimal Fee
l1	79.635	75.670	Baseline
l2	76.971	76.996	Baseline
l3	79.809	77.754	Baseline
l4	74.190	75.670	Baseline
l5	76.975	80.535	Baseline

Appendix B: Data on Average Efficiencies by Session for Longer Chat Sessions

Session ID	Avg. Efficiency: Rounds 1-3	Avg. Efficiency: Rounds 4-8	Avg. Efficiency: Rounds 9-13	Treatment Rounds 4-13
r1	79.807	89.530	84.440	Chat
r2	75.287	91.412	88.772	Chat
r3	77.150	89.976	88.536	Chat
r4	79.463	92.164	87.530	Chat
r5	80.757	92.432	96.842	Chat

Appendix C: Experiment Instructions

Baseline Treatment Instructions:

Page 1 of 3

- **Role:** In this game, you have the role of a farmer with a number of distinct fields that can be planted.
- **Productivity:** For each field, you will be given a number that determines the cash value of the crops from that field.
- **Example:** If you receive numbers 7, 5, and 3, then the first field that you plant yields a value of $7, the second a value of $5 and the third a value of $3.
- **Irrigation:** Yield values are multiplied by a factor of 3 if irrigation is available, e.g. $7, $5, and $3 become $21, $15, and $9.
- **Canal:** The watershed is a canal that flows by each person's farm, one at a time.
- **Position:** Each farm will have an address, 1, 2, ...6, with lower numbered addresses being closer to the source of the water on your canal, so the water flows by Canal Street #1 first, by #2 second, etc.
- **Irrigation Decision:** Water is measured in units, with each unit sufficient to irrigate one field.

Page 2 of 3

- **Your Fields:** There are **6 farmers** located on your canal, and each of you has will have **4 fields**.
- **Water Use Decision:** If water is available at your location, you will be able to pump any number of units, up to 4, to irrigate some or all of your fields.
- **Water Supply:** The total amount of water available is **12 units**, which is not enough to irrigate all fields for all of those located along your canal.
- **Public Information:** Your irrigation decision WILL NOT be observed by the other farmers in your group; they will only see how much water remains when it is their time to make an irrigation decision.
- **Water Use:** If you use fewer than 4 units of water, either because it is not available or because you decide to reduce irrigation, then only the fields for which you use water units will have **tripled** yield values.

Page 3 of 3

- **Rounds:** There will be a series of **6 rounds**, which correspond to growing seasons, with a renewed water stock in each round. New irrigation decisions may be made in each round.
- **Water Flow:** The available supply of water will be the same (12 units) at the start of each new round.
- **Productivities:** The productivities for each of your fields may change from one round to the next, due to random variations in local conditions.
- **Note:** You will be located on a canal with the **same** 5 other people in all rounds, and your address will not change.

Summary Page

- You are located on a canal with 5 other people.
- Each person on your canal has an address (#1, #2, etc.) that determines the order in which you can make your irrigation decisions.
- Each person has 4 fields, with productivities that are **tripled** if irrigated.
- For example, if one of your fields has a productivity of $4, it will provide you with earnings of $4 if it is not irrigated, and $12 if it is irrigated
- Your earnings for a round are calculated by adding up the productivites of your fields, after any of these productivities have been tripled as a result of irrigation.
- For each round, you will be told the available water stock for your canal. When it is your turn, you decide how much of the remaining water stock to use for irrigation.
- Everyone begins with an initial cash balance of $0.00. Earnings will be added to this amount. The program will keep track of your cumulative earnings.
- **Special Earnings Announcement:** Your cash earnings will be **4%** of your total earnings at the end of the experiment.

For all non-baseline treatments, only pages containing new or modified instructions are reported. The first page of instructions is the same in every treatment, and so is always omitted.

Chat Treatment Instructions:

Page 2 of 3

- **Your Fields:** There are **6 farmers** located on your canal, and each of you has will have **4 fields**.
- **Water Use Decision:** If water is available at your location, you will be able to pump any number of units, up to 4, to irrigate some or all of your fields.
- **Water Supply:** The total amount of water available is **12 units**, which is not enough to irrigate all fields for all of those located along your canal.
- **Public Information:** Your irrigation decision WILL be observed by all of the other farmers in your group.
- **Water Use:** If you use fewer than 4 units of water, either because it is not available or because you decide to reduce irrigation, then only the fields for which you use water units will have **tripled** yield values.

Page 3 of 3

- **Rounds:** There will be a series of **6 rounds**, which correspond to growing seasons, with a renewed water stock in each round. New irrigation decisions may be made in each round.
- **Communications:** There will be a **chat room** open at the beginning of each round for 3 minutes, which allows you to send and receive messages. Bidders will be identified by their ID numbers, and you are free to discuss any aspect of the irrigation decisions. You should avoid inappropriate language, threats, or attempts to arrange side payments.
- **Water Flow:** The available supply of water will be the same (12 units) at the start of each new round.
- **Productivities:** The productivities for each of your fields may change from one round to the next, due to random variations in local conditions.
- **Note:** You will be located on a canal with the **same** 5 other people in all rounds, and your address will not change.

Bargaining Treatment Instructions:

Page 2 of 4

- **Your Fields:** There are **6 farmers** located on your canal, and each of you has will have **4 fields**.
- **Water Use Decision:** If water is available at your location, you will be able to pump any number of units, up to 4, to irrigate some or all of your fields.
- **Water Supply:** The total amount of water available is **12 units**, which is not enough to irrigate all fields for all of those located along your canal.
- **Public Information:** Your irrigation decision WILL be observed by all of the other farmers in your group.
- **Water Use:** If you use fewer than 4 units of water, either because it is not available or because you decide to reduce irrigation, then only the fields for which you use water units will have **tripled** yield values.

Page 3 of 4

- **Rounds:** There will be a series of **6 rounds**, which correspond to growing seasons, with a renewed water stock in each round. New irrigation decisions may be made in each round.
- **Communications:** There will be a **chat room** open at the beginning of each round for 6 minutes, which allows you to send and receive messages. Bidders will be identified by their ID numbers, and you are free to discuss any aspect of the irrigation decisions. You should avoid inappropriate language, threats, or attempts to arrange side payments.
- **Negotiations:** During the chat phase you will be able to negotiate irrigation limits with others in your group. You may either agree to pay someone else a specified amount to reduce their irrigation to a specified limit, or you may agree to reduce your irrigation to a specified limit if someone else pays you a specified amount. You may make payments to more than one person, and you may receive payments from more than one person.

- **Contracts:** Contracts are not final unless both parties enter **exact matching payment/receipt amounts**, and the **exact same irrigation limits**.
- **Example:** If one person proposes to reduce irrigation to at most Q water units in exchange for $1, and another offers to pay the first person $2 for that reduction to Q units, then there is room for agreement, but this would not be a binding contract since the dollar amounts do not match.
- **Enforcement:** A binding contract will constrain irrigation decisions (for that round only) and enforce agreed payments when earnings are calculated for that round.
- **Multiple Contracts:** You may receive payments from one or more people (located downstream) to reduce your irrigation, and you may agree to pay one or more people (located upstream) to reduce their irrigation.
- **Water Flow:** The available supply of water will be the same (12 units) at the start of each new round.
- **Productivities:** The productivities for each of your fields may change from one round to the next, due to random variations in local conditions.
- **Note:** You will be located on a canal with the **same** 5 other people in all rounds, and your address will not change.

Page 4 of 4

- **Note:** Contracts are not final unless both parties enter **exact matching payment/receipt amounts**, and the **exact same irrigation limits**.
- For example, if the prices do not match, there will be no deal.
- A deal is indicated by the phrase "Contract Signed" on your screen.
- No changes in offers will be recorded after the time expires.
- You have to update your screen with the submit button to view offers from others as they come in, to see contract approvals, and to check the clock.

Two-person Negotiation Example, You Are Upstream

Negotiations	Other ID	Irrigation Limit	Payment Amount	Status
Position 1	**(you)**	*	*	*
with Id 2 (position 2)	I will restrict irrigation to [0 units] in exchange for [select] from ID **2**.			

Two-person Negotiation Example, You Are Downstream

Negotiations	Other ID	Irrigation Limit	Payment Amount	Status
with Id 1 (position 1)	I propose that ID 1 restrict irrigation to [0 units] in exchange for [select] from me.			
Position 2	**(you)**	*	*	*

Continue

Summary Page

- You are located on a canal with 5 other people.
- Each person on your canal has an address (#1, #2, etc.) that determines the order in which you can make your irrigation decisions.
- Each person has 4 fields, with productivities that are **tripled** if irrigated.
- For example, if one of your fields has a productivity of $4, it will provide you with earnings of $4 if it is not irrigated, and $12 if it is irrigated
- Your earnings for a round are calculated by adding up the productivities of your fields, after any of these productivities have been tripled as a result of irrigation.
- For each round, you will be told the available water stock for your canal. When it is your turn, you decide how much of the remaining water stock to use for irrigation.
- **Note:** During the chat period, you will also be free to negotiate maximum irrigation amounts with other people in your group. Contracts are not final unless both parties enter the **same** payment/receipt amounts and the **same** irrigation limits.
- Everyone begins with an initial cash balance of $0.00. Earnings will be added to this amount. The program will keep track of your cumulative earnings.
- **Special Earnings Announcement:** Your cash earnings will be **4%** of your total earnings at the end of the experiment.

Auction Treatment Instructions:

Page 2 of 3

- **Your Fields:** There are **6 farmers** located on your canal, and each of you has will have **4 fields**.
- **Water Permit Auction:** In order to use water, you must purchase "permits" in an auction, one permit for each field you wish to irrigate.
- **Water Supply:** The total amount of water available is **12 units** (12 permits), which is not enough to irrigate all fields for all of those located along your canal.
- **Water Use:** If you purchase fewer than 4 permits, either because you decide not to bid or because some of your bids are too low to be accepted, then only the fields for which you purchase permits and use water units will have **tripled** yield values.
- **Auction Procedure:** All bids received from the 6 people in your group will be collected and ranked, and the highest 12 bids will be accepted, with ties broken by a random device. Each accepted bid allows you to irrigate an additional field.
- **Auction Clearing Price:** The price you pay when your bid is accepted is **not** your bid; it is the highest of the rejected bids (the bid of rank 13).
- **Example:** For simplicity, suppose that only 1 permit is being sold, and the bids are $1, $2, and $3. Then the high bid of $3 will win, but that bidder only has to pay $2, which was the highest rejected bid. **Note:** What you pay if you win is not what you bid.

Summary Page

- You are located on a canal with 5 other people.
- Each person on your canal has an address (#1, #2, etc.) that corresponds to the direction of water flow, from 1 to 2 ....
- If you purchase permits in an auction, then there will be sufficient flow for you to irrigate a number of fields equal to the number of permits purchased.
- Each person has 4 fields, with productivities that are **tripled** if irrigated.
- For example, if one of your fields has a productivity of $4, it will provide you with earnings of $4 if it is not irrigated, and $12 if it is irrigated
- Your earnings for a round are calculated by adding up the productivites of your fields (after any of these productivities have been tripled as a result of irrigation) and then subtracting the prices you paid for permits purchased at auction.

- For each round, you will be told the available water stock for your canal, which equals the number of permits to be sold at auction. You then submit bids to purchase permits for your fields.
- All bids for your group will be collected and ranked, and the highest bids will be accepted. The price you pay when your bid is accepted is **not** your bid; it is the highest of the rejected bids.
- Everyone begins with an initial cash balance of $0.00. Earnings will be added to this amount. The program will keep track of your cumulative earnings.
- **Special Earnings Announcement:** Your cash earnings will be **4%** of your total earnings at the end of the experiment.

Optimal Fee Treatment Instructions:

Page 2 of 3

- **Your Fields:** There are **6 farmers** located on your canal, and each of you has will have **4 fields**.
- **Water Use Decision:** If water is available at your location, you will be able to pump any number of units, up to 4, to irrigate some or all of your fields.
- **Water Supply:** The total amount of water available is **12 units**, which is not enough to irrigate all fields for all of those located along your canal.
- **Public Information:** Your irrigation decision WILL NOT be observed by the other farmers in your group; they will only see how much water remains when it is their time to make an irrigation decision.
- **Water Use:** If you use fewer than 4 units of water, either because it is not available or because you decide to reduce irrigation, then only the fields for which you use water units will have **tripled** yield values.
- **Water Use Fee:** An irrigation fee of **$13** must be paid for each unit of water that you use, this fee is deducted from your earnings for the round.

Summary Page

- You are located on a canal with 5 other people.
- Each person on your canal has an address (#1, #2, etc.) that determines the order in which you can make your irrigation decisions.
- Each person has 4 fields, with productivities that are **tripled** if irrigated.
- For example, if one of your fields has a productivity of $4, it will provide you with earnings of $4 if it is not irrigated, and $12 if it is irrigated
- Your earnings for a round are calculated by adding up the productivites of your fields, after any of these productivities have been tripled as a result of irrigation.
- In addition, you have to pay **$13** for each unit of water that you use.
- For each round, you will be told the available water stock for your canal. When it is your turn, you decide how much of the remaining water stock to use for irrigation.
- Earnings for this part will be added to earnings from the previous part (losses will be subtracted). The program will keep track of your cumulative earnings.
- **Special Earnings Announcement:** Your cash earnings will be **4%** of your total earnings at the end of the experiment.