# Tragedy of the common canal

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

University of Virginia, University of Arizona

19 February 2010

# Water Externalities: Tragedy of the Common Canal

Charles A. Holt, Cathleen Johnson, Courtney A. Mallow, and Sean P. Sullivan[*]

February 20, 2009

**Abstract**

This paper uses laboratory experiments to investigate the effects of alternative solutions to a common-pool resource with a unidirectional flow. The focus is on the comparative economic efficiency of communications, bilateral "Coasian" bargaining, auctions and price-based allocations. All treatments improve allocative efficiency relative to a baseline environment. Communication and bilateral bargaining are not generally as effective as market allocations. An exogenously imposed, optimal fee results in the greatest efficiency gain, followed by auction allocations that determine the usage fee endogenously.

## I. Introduction

The standard solution to the tragedy of the commons is to assign broad-based property rights, thereby internalizing the externality. Subsequent purchases and sales can then reallocate these rights to efficient producers. When property rights are difficult to enforce or when their assignment is politically infeasible, a host of direct regulations may arise. The path-breaking work of Elinor Ostrom and her coauthors has uncovered a rich variety of institutional solutions, most of which do not involve property right assignments or heavy-handed regulation.[1]

A particularly interesting common-pool resource problem arises when resource availability follows a unidirectional flow, such that usage by upstream producers only imposes externalities on those farther downstream. An example is the situation of farmers aligned along a

---

[*] 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, Nicholas Smith and Ilya Zlatkin. This research was funded in part by the National Science Foundation (SES0098400) and the University of Virginia Bankard Fund.

[1] See Ostrom, E., R. Gardner, and J. K. Walker (1994) and the special section of the Fall 1993 *Journal of Economic Perspectives* devoted to "Management of the Local Commons."

canal, who make sequential use of the scarce water resource to irrigate their fields.[2] Even with well defined property rights, overuse by upstream farmers may prevent water from reaching fertile downstream areas.[3] In such cases, market-based solutions may offer some advantages: e.g., marketable shares of water flow provide farmers with incentives to trade in such a way that water is diverted to its highest-value uses (Yoder, 1986). Social solutions may also exist. For example, opportunities for efficient coordination of canal repairs or joint marketing may put property owners in social settings where they are in a better position to negotiate mutual reductions (Ostrom and Gardner, 1993).

The first laboratory experiments on this topic are reported by Gardner, Ostrom, and Walker (1990), who provide a conceptual framework for evaluating problems associated with overuse of a common-pool resource. In these experiments, subjects simultaneously select the intensity of their harvest from a common-pool resource by allocating "tokens" between two activities. Walker, Gardner, and Ostrom (1990) consider the effects of changing the endowment of tokens. Cardenas, Janssen and Bousquet (2008) introduce a water irrigation experiment: the first attempt to investigate sequential extraction of a common-pool resource, in contrast to simultaneous request.

Experiments that include simultaneous provision and sequential appropriation of water resources suggest that the lack of trust from downstream players towards upstream players curbs cooperation.[4] Social pressures may result in some improvements, especially for small groups of economically or ethnically homogeneous users.[5] In experimental work where the requests for appropriations are made sequentially, but distribution occurs at the conclusion of the decision sequence, individual requests and position sequence are negatively correlated.[6] Late movers request less and early movers request more, regardless of the information conditions and the

---

[2] Other examples of directional flows include pollution that is blown by prevailing winds, and the harvest of migratory fish.

[3] Ostrom and Gardner (1993) report an example from Nepal in which overuse by "headlanders" during the pre-monsoon season results in crop values that are much lower than what could be achieved with a reallocation to downstream rice farmers.

[4] Cardenas, Johnson and Rodrigues (2009) conducted both a "water irrigation" experiment and a "water trust" experiment.

[5] Cardenas (2003) and Cardenas *et al.* (2002) report field experiments in which non-homogeneous groups have more difficulty in dealing with common-pool resource problems.

[6] Budescu et al. (1997) considered information manipulations on the effect of having a sequence of extraction requests from a shared common-pool resource.

randomness of order assignment.

In some cases, informal social arrangements at a local level seem to outperform regulations imposed by a higher government authority.[7] However, the sequential structure of unidirectional externalities can create problems. For example, let sequential locations along a canal be represented by users 1, 2, and 3, so that user 1 acts first, user 2 acts second, etc. User 3 could pay user 1 to reduce water usage, but this private arrangement will be of little benefit to 3 if user 2 exploits the extra water flow when given the move. In a long sequence of usage decisions, "defection" by a single upstream user, if observed or incorrectly inferred, may induce a cascade of downstream defections. Efforts to impose use limits or fees may be hampered by misleading information provided to regulators or by offsetting activities taken by the users.[8]

This paper describes a laboratory experiment designed to evaluate the efficiency gains provided by four potential solutions to a sequential common-pool appropriation problem: communication ("chat"), bilateral bargaining with chat ("bargaining"), an auction of water rights ("auction"), and an optimal irrigation fee ("optimal fee"). A between-subjects design is used to compare the allocative efficiencies of these potential solutions. Section II describes the general game and specific treatment environments in detail. Section III explores results of the experiment, specifically subject behavior and observed efficiency in each environment. Section IV concludes with a discussion of observed themes and potential extensions.

## II. Procedures

Participants in this experiment are given the role of "farmers" located along a common "canal" that flows by each of their farms in sequence. Each session consists of 6 participants, with numbered addresses corresponding to their identities (IDs 1-6). Addresses determine the sequence in which water use decisions are made: ID 1 moves first, ID 2 moves second, etc. Address locations do not change between the rounds.

Each participant is endowed with 4 fields of randomly determined productivities. The productivity value for a given field corresponds to the cash value of the crops that field yields in

---

[7] Cardenas, 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 overuse.

[8] In fisheries, for example, limits on the season 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).

the absence of irrigation. If irrigated, a field yields a cash value of triple its productivity value. Provided water is available, each participant decides whether or not to irrigate each field. A total stock of 12 "units" of water is available in each round; irrigating a field is a binary decision that depletes a "unit" of water from the total available stock. When deciding whether to irrigate each field, the available amount of water is visible to the participant, but the amounts taken by upstream user are not visible (except in specific treatments, as noted below).

Rounds correspond to different growing seasons with renewed water supplies. To reflect local climate variations, productivities are randomly generated from discrete-uniform distributions in each round. All sessions are run with web-based Veconlab software, using the Water Externalities program.[9] Rich terminology (farmers, fields, water) is used to help make the decision-making context clear to the participants. The same context is used in all treatments.

Table 1 displays the ranges of random field-productivity values, which differed between the three upstream producers (IDs 1-3) and the three downstream producers (IDs 4-6). Both productivity ranges and their realizations in each round were private information. Distributed as discrete-uniform random variables, productivities are constrained to be integer amounts: e.g. high-productivity fields are equally likely to have values of $7, $8, $9, $10, or $11, and low-productivity fields are equally likely to have values of $2, $3, $4, $5, or $6.

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

| Field Number | Upstream Producers (IDs 1-3) | Downstream Producers (IDs 4-6) |
|---|---|---|
| 1st | $7-$11 | $7-$11 |
| 2nd | $2-$6 | $7-$11 |
| 3rd | $2-$6 | $7-$11 |
| 4th | $2-$6 | $2-$6 |

The optimal allocation of the 12 water units is to irrigate the 12 high-productivity fields. Since the productivity ranges for the two types of fields do not overlap, this corresponds to

9 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. This program can also be used in a classroom setting to induce discussions of common-pool resource problems.

allocating 1 unit of water to each of the upstream producers (IDs 1-3), and allocating the remainder evenly over the downstream producers (IDs 4-6). If upstream producers behave selfishly, they will each take 4 water units, leaving no available stock for the downstream producers with more high-productivity fields. Since irrigation triples yield values, the net gain from irrigation is twice the yield value. Notice that an optimal fee is a price for water of $13, which would deter farmers from irrigating low productivity fields, but would not deter those with high productivity fields. The imposition of an optimal fee, in theory, would yield maximum earnings (100% efficiency), as compared with the approximate 75% efficiency that would result from purely selfish behavior under these parameter values.

A baseline environment as just described is used for the first three rounds of each session. In the final three rounds, one of five different treatments is applied: a repeat of the baseline (no change), communication ("chat"), bilateral bargaining with chat ("bargaining"), an auction of water rights ("auction"), or an optimal irrigation fee ("optimal fee"). Repeating the baseline environment allows the experimenter to obtain a basis for comparison that is corrected for experience.

In the chat treatment, participants are given three minutes to communicate in an online chat room. After the chat period ends, they made decisions in sequence as before, except that each person is able to view the water use decisions of upstream participants by ID. There is, of course, a large literature on the effects of communication in common-pool resource dilemmas. The purpose of this treatment is to determine how a controlled amount of social interaction might enhance efficiency in this sequential setup, in order to provide a basis for comparison with market-based policies to be discussed next.

Like the communication treatment, the bargaining treatment also involves a chat room and public decision making. The chat time is extended to 6 minutes in each round to provide participants sufficient time to negotiate binding bilateral contracts. Two types of contracts are possible: (1) an offer to pay an upstream user to their restrict irrigation to at most Q units in exchange for a payment of $P, and (2) an offer to accept $P from a downstream user in exchange for agreeing to restrict one's own irrigation to at most Q units. Anyone who receives a proposed contract can accept it or not. All agreements are bilateral and binding, but participants can make agreements with any number of upstream and downstream users. For example, ID 1 might agree to restrict irrigation to 3 units in exchange for payment of $2 from ID 4, and the same person (ID

1) might agree 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.

This treatment is motivated by the Coase theorem, which suggests that bargaining under the umbrella of well-defined property rights should result in an efficient allocation even in the presence of externalities, as long as certain assumptions are satisfied (Coase, 1960). The most critical assumption is the absence of transactions costs. Although, there are no explicit bargaining costs in the experiment, time limits and the need to engage in multiple, interrelated negotiations may generate substantial indirect transactions costs. Participants are also hampered in terms of not knowing others' productivity values when negotiating contracts.

In the auction treatment, a permit is required to irrigate a field. All farmers, regardless of address, have the opportunity to bid for as many as 4 permits each. The highest 12 bids are selected, and the price paid for the permit is the highest rejected bid (i.e. the 13th bid). This is a multi-unit, uniform-price auction with private values, so it is never optimal to bid above one's value.[10] Bidding below value at the rejection margin could, however, reduce the price paid for other permits. Therefore, bidding at value is not necessarily an equilibrium strategy, as would be the case in a second-price auction with a single prize. If bids do mirror values, then 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 is to determine how effectively a market process could approximate an optimal fee.

In contrast to endogenous determination of the fee in the auction treatment, the optimal fee treatment simply imposed an exogenous per-unit fee of $13 for each water unit used, simulating a Pigouvian tax. The revenue from the fee is not returned to the participants.[11] No chat was allowed in the auction and optimal fee treatments.

Reported results are based on a total of 25 six-person sessions, run between March and December 2009, using student subjects recruited from the University of Virginia. Session lasted from 35 to 60 minutes, depending on the treatment. Participants received $6 for showing up, and were paid a cash amount equal to 4% of the money they earned in the experiment. Earnings

[10] This setup is similar to the multi-unit uniform-price auction 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).

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

depended on the treatment, but generally ranged from $12 to $30, including the initial $6 payment.

## III. Results

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

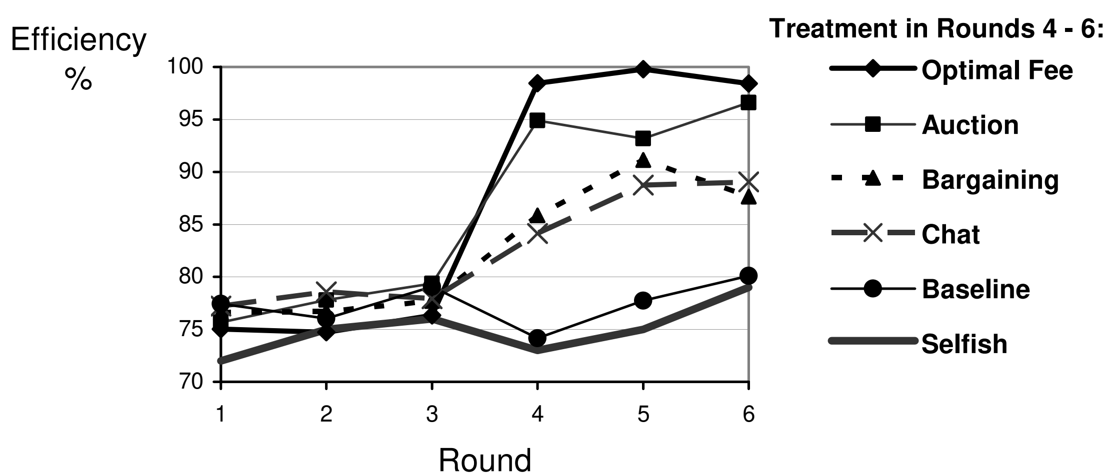


**Figure 1**. Baseline Efficiencies (Rounds 1-3) by Session and Treatment

The highest efficiencies are observed when an optimal fee is exogenously imposed, followed by an auction. Bargaining and chat are less efficient, and exhibit little difference. Under baseline conditions, average efficiencies are 1 to 3 percentage points higher than the purely selfish predictions, indicating a small amount of altruistic behavior. Recall that in this environment, the unidirectional flow of water and static location of participants means that acts of generosity cannot be reciprocated.

Average efficiencies, disaggregated by session, 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, each representing the average efficiency in a session-environment.

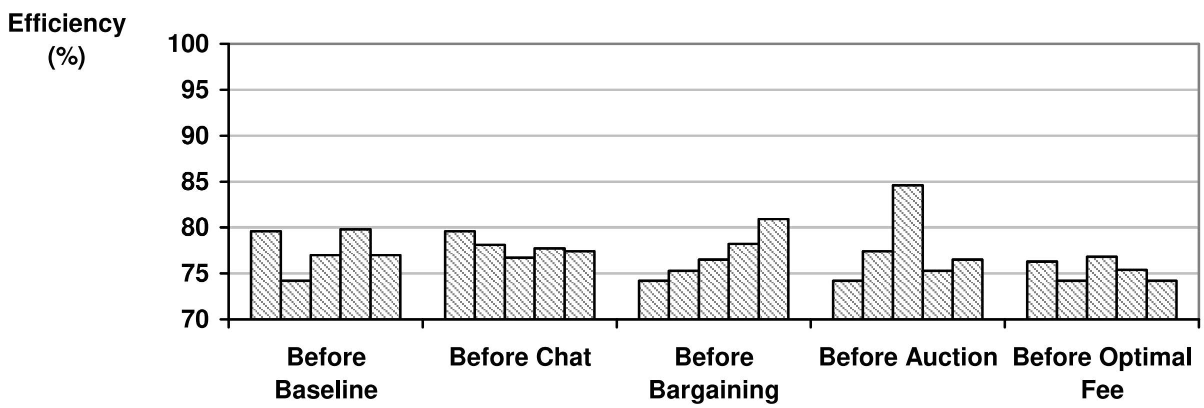


**Figure 2**. Baseline Efficiencies (Rounds 1-3) by Session and Treatment

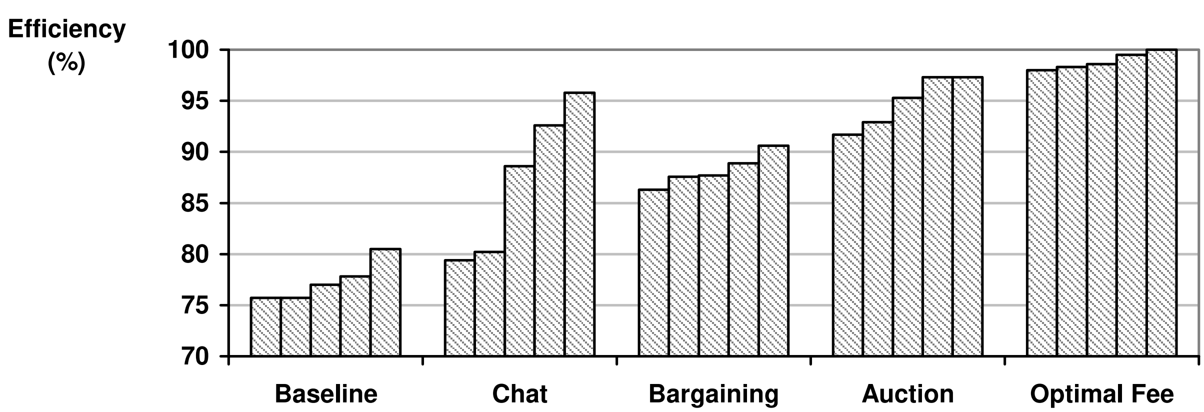


**Figure 3.** Treatment Efficiencies (Rounds 4-6) by Session and Treatment

A quick glance at these data affords several qualitative observations. Despite some variability, efficiencies in the baseline environment look basically homogeneous across experimental sessions. There is no correlation between baseline efficiencies and treatment efficiencies, except in the bargaining sessions, where efficiencies are increasing from left to right in both figures. Efficiencies 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 more 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 a common-pool resource problem with sequential extraction.

**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 an average "selfish efficiency" of 74.33%. Casual inspection of Figures 1-3 suggests that observed efficiencies 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 (pooled over all 25 sessions) equals the selfish prediction at any reasonable level of significance.[12]

Of course, rejection of equality does not imply a large inequality, and a 10% confidence interval places the average baseline efficiency only between 76.1 and 78.0%.[13] 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

[12] Student's *t*-test provides a p-values of less than 0.0001.

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

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 environments. This amounts to calculating the difference between baseline and treatment efficiencies for each session, and then comparing the average differences to zero by treatment type.[14]

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

| 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, finding strong evidence of efficiency gains under every non-baseline treatment. For its part, the repeated baseline treatment shows 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.[15] 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 rejections. A common technique is to use a test which controls of the family-wise error rate, defined as the probability of *even a single*

[14] Within-session comparisons exploit pairing of baseline and treatment environments within each session to help mitigate the consequences of potentially unobserved heterogeneity.

[15] 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, we can expect one false rejection out of the 20 tests performed. In fact, the probability of at least one false rejection is $1 - (0.95)^{20}$, or 64%.

*false rejection* among k simultaneous tests.[16] 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 family-wise 0.1 level.[17]

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 illustrates 95% confidence intervals for average efficiency gains under each treatment environment.[18]

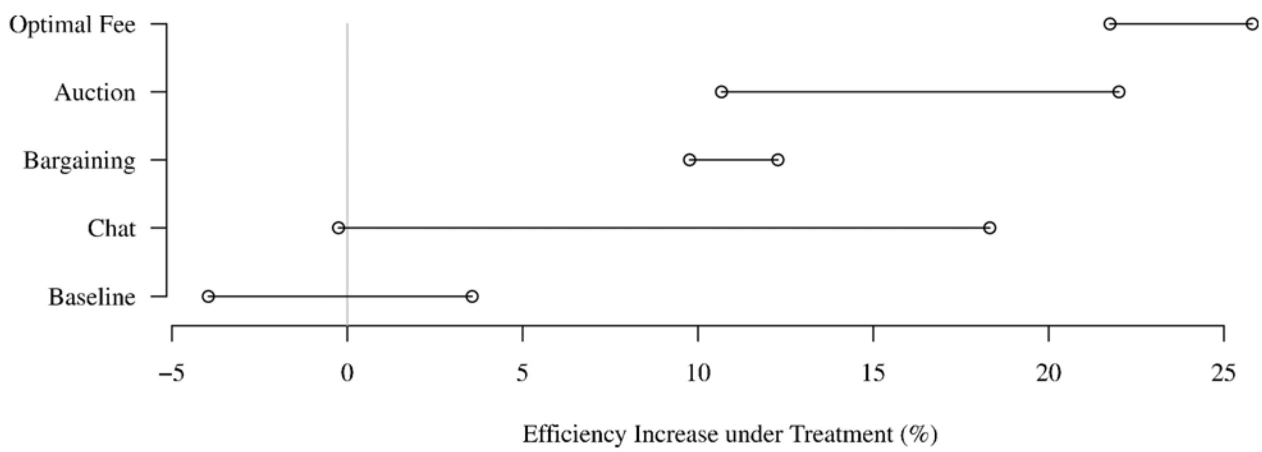


**Figure 4**. 95% Confidence Intervals for Average Efficiency Gain by Treatment

Confidence intervals for the average efficiency gain under each solution vary considerably in both breadth and location. It is interesting that a 95% confidence interval on the optimal fee solution includes as much as a 25% average efficiency gain over the baseline environment—corresponding to approximately 100% efficiency under this solution. Outcomes under the chat treatment are sufficiently variable that a 95% confidence interval contains both zero efficiency gains, and gains of nearly 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: chat logs reveal that some groups manage to establish loose

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

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

[18] 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.

behavioral norms or some degree of social responsibility, while other groups fail to establish such norms, and chat devolves into a series of complaints and frustrations.

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

Given that all the solutions considered in this paper do increase allocative efficiency to varying degrees, the next logical inquiry is whether we can say anything about which ones work better than others. To address this question, we rely on between-session variation in comparing average efficiencies across our 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.[19]

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 pair-wise contrasts between treatments using the Wilcoxon-Mann-Whitney test.[20] Table 3 summarizes inferences gained from each comparison: reported p-values are exact.

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

| 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 |

[19] 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.

[20] The intuition behind this test is easily illustrated. For example, if all 5 observations under one treatment are lower than all 5 observations under another, 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 extreme or more extreme that what was observed. Under the null, the chances of this are 2/252 = 0.00794, as shown by the bottom 4 rows of Table 3.

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 it 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 family-wise 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 family-wise error rates.[21] 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 the water, even a small dose of individual rationality should suffice to affect a socially optimal allocation.

Unfortunately, the practicality of this solution to the common-pool resource problem is limited. There is no reason to expect an optimal fee would be obvious in a typical policy-making setting, particularly when users have incentives to report valuations selectively and to 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 a great decrease in practical efficacy in many settings.

**Result 5:** *The Auction treatment yields higher average efficiency than either the Chat or Bargaining environments.*

Because an auction solution uses a market mechanism to "discover" the optimal fee, it is not surprising that it should closely follow the optimal fee treatment in terms of average

---

[21] The Optimal Fee treatment is concluded to provide higher average efficiency than any other non-baseline treatment when using the Hochberg algorithm to control the family-wise error rate at the 0.025 level.

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 family-wise 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 family-wise error rate.[22]

**Result 6:** *There is little evidence that bilateral bargaining results in a 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 theoretical 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 transactions costs, which are assumed away in the Coase theorem. Although property rights are well defined and there are no explicit transactions costs in this treatment, a downstream farmer has to make multiple contracts with upstream farmers in order to insure water availability. With no centralized coordinator, the difficulty of forming an appropriate menu of contracts can represent a substantial implicit transactions cost.[23] There is also a free-riding problem, since various farmers may benefit from contracts to which they are

---

[22] The auction treatment environment is concluded to provide higher average efficiency than either the bargaining or chat treatments when using the Hochberg (1988) algorithm to control the family-wise 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 family-wise 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.

[23] 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 didn't 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).

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. Conclusion

This paper was 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 common-pool resource environments. The laboratory results reported here are intended to compare various solutions to a commons problem in the context of a sequential structure that arises naturally in settings with a unidirectional flow. Results suggest that the commons problem of inefficient resource allocations can be mitigated by the introduction of proper social institutions or government intervention.

In the setting we investigate, there exists a unique optimal fee, the imposition of which causes full internalization of all usage externalities. Unsurprisingly, experimental results show exogenous imposition of this fee corresponds to nearly 100%, efficiency. A solution relying on a uniform price auction for water permits is not as efficient, but the difference appears relatively small. The advantage of an auction approach 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 transactions costs. The bargaining treatment of the experiment implements binding bilateral contracts without explicit negotiation costs. Because contracts are constrained to be bilateral, however, 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 transactions cost. Moreover, there is a free-riding problem in the sense that participants located between two parties to a contract may take the water that the upstream person agrees not to use. In this setting, we find 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 interesting extension may be to revise the bargaining environment to allow for direct, bilateral trade 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 practical relevance is unclear.

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. 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. Tung 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.

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

Cardenas, J. C., M. Janssen, 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.

Cardenas, J. C., N. Johnson, L.A. Rodriguez (2009) "Vertical Collective Action: Addressing Vertical Asymmetries in Watershed Management," working paper.

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

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

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

Deck, C. and C. Johnson, "Link Bidding in a Laboratory Experiment" *Review of Economic Design*, 8 (4), April 2004, pp. 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.

Hardin, G. (1968) “The Tragedy of the Commons,” *Science*, 162: 1243-1248.

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: New York State Energy Research and Development Authority.

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. K. Walker (1994) *Rules, Games, and Common-Pool Resources*. Ann Arbor: University of Michigan Press.

Ostrom, E., and J. K. Walker (1991) "Communication in a Commons: Cooperation without External Enforcement," in *Laboratory Research in Political Economy*, ed. by T. Palfrey. Ann Arbor: University of Michigan Press, 289-322.

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.

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

**Appendix: Data on Average Efficiencies by Session**

| Session Name | Avg. Efficiency: Rounds 1-3 | Avg. Efficiency: Rounds 4-6 | Treatment Rounds 4-6 |
|---|---|---|---|
| wex21 | 74.290 | 76.912 | Chat |
| wex22 | 77.726 | 79.209 | Chat |
| wep11 | 77.723 | 92.595 | Chat |
| wep30 | 77.423 | 95.752 | Chat |
| wex2 | 76.678 | 88.616 | Chat |
| wex16 | 80.852 | 90.604 | Bargaining |
| wex17 | 76.450 | 87.743 | Bargaining |
| wex18 | 78.248 | 88.867 | Bargaining |
| wex19 | 74.190 | 86.293 | Bargaining |
| wex20 | 75.289 | 87.575 | Bargaining |
| wep10 | 74.190 | 91.718 | Auction |
| wep17 | 76.504 | 97.349 | Auction |
| wep18 | 84.597 | 95.267 | Auction |
| wex1 | 75.289 | 97.301 | Auction |
| wex3 | 77.492 | 92.852 | Auction |
| wex6 | 76.279 | 98.018 | Optimal Fee |
| wex7 | 75.408 | 99.456 | Optimal Fee |
| wex8 | 76.797 | 98.582 | Optimal Fee |
| wex9 | 74.190 | 00.000 | Optimal Fee |
| wex10 | 74.190 | 98.342 | Optimal Fee |
| wex11 | 79.635 | 75.670 | Baseline |
| wex12 | 76.971 | 76.996 | Baseline |
| wex13 | 79.809 | 77.754 | Baseline |
| wex14 | 74.190 | 75.670 | Baseline |
| wex15 | 76.975 | 80.535 | Baseline |