

Working Paper Series

No. 47

Time and Uncertainty in Resource Dilemmas: Equilibrium Solutions and Experimental Results

Anabela Botelho

Ariel Dinar

Lígia M. Costa Pinto

Amnon Rapoport

September 2012

Núcleo de Investigação em Microeconomia Aplicada

Universidade do Minho



Time and Uncertainty in Resource Dilemmas: Equilibrium Solutions and Experimental Results

Anabela Botelho, Ariel Dinar, Lúgia M. Costa Pinto, Amnon Rapoport*

Abstract

Most common pool resource (CPR) dilemmas share two features: they evolve over time and they are managed under environmental uncertainties. We propose a finite-horizon, stochastic, dynamic model that integrates these two dimensions. A distinguishing feature of our model is that the duration of the game is determined *endogenously* by the players' collective decisions. In the proposed model, if the resource stock level below which the irreversible event occurs is known in advance, then the optimal resource use coincides with a unique symmetric equilibrium that guarantees survival of the resource. As the uncertainty about the threshold level increases, resource use increases if users adopt decision strategies that quickly deplete the resource stock; however, resource use decreases if they adopt path strategies guaranteeing that the unknown threshold level is never exceeded. Our experimental results show that CPR users frequently implement decision strategies that terminate the game immediately. When the uncertainty about the resource level is reduced, users maintain a positive resource level for a longer duration.

*University of Minho and NIMA (Botelho and Pinto); Water Science and Policy Center, Department of Environmental Sciences, University of California, Riverside (Dinar); and A. Gary Anderson School of Business Administration, University of California, Riverside (Rapoport). This paper was prepared while Botelho and Pinto were on sabbatical leave at the Water Science and Policy Center, University of California, Riverside, and University of Arizona, Tucson. Botelho and Pinto thank the Fundação para a Ciência e Tecnologia for sabbatical scholarships SFRH/BSAB/1226/2012 and SFRH/BSAB/1159/2011, respectively.

1. Introduction

Most natural common pool resource (CPR) dilemmas share two features: they evolve over time and they are managed under environmental uncertainties. While each of these two features has been analyzed separately in the experimental literature, no attempt has been made to integrate them. In particular, the analysis of strategic behavior in the face of environmental uncertainty about the size of the CPR has been conducted under the assumption of single-period interaction, and the analysis of strategic behavior in time-dependent settings has ignored environmental uncertainties.

The bulk of the experimental literature (e.g., Ostrom, Gardner, and Walker 1994) has analyzed unrestricted resource-use decisions by placing subjects in the context of repeated time-independent CPR dilemmas characterized only by strategic uncertainty about the planned behavior of others. A major finding is aggregate behavior consistent with the theoretically predicted resource misallocations (Gordon 1954; Hardin 1968). It has been noted, however, that while the time-independent framework may be an adequate representation of CPRs characterized by flows, in which availability of the resource in the future is independent of current requests, it fails to capture the important temporal feature of stock resources, like groundwater systems, fisheries, forests, etc. (Brown 2001). In these instances, decisions concerning resource-use are typically made in the “shadow of the future” in the sense that current requests from the resource pool compete with resource availability in the future.

It is a priori unclear if the presence of such temporal factors would elicit behavior different from that observed in time-independent settings. Because current requests not only affect the future profits of other group members, but also their individual profits as well, CPR users may adopt precautionary strategies that lead them closer to Pareto optimal outcomes. For example, Reinganum and Stokey (1985) show that efficient resource-use decisions are to be expected in time-dependent settings when users can jointly commit to “path” strategies over the entire planning horizon. On the other hand, the consideration that current request decisions affect future request possibilities also raises a dynamic optimization problem that complicates the attainment of Pareto optimal outcomes even in single-agent contexts (Messick and McClelland 1983; Hey, Neugebauer and Sadrieh 2009). Moreover, as shown by Dutta (1995), the standard intuition from infinitely repeated time-independent games, whereby Pareto optimal outcomes can be sustained in equilibrium through threat of credible punishment by patient players, does not necessarily carry over to time-dependent games with stock variables. In these games, players’ payoffs depend not only on current and previous periods’ decisions, but also on state variables that change from period to period. Theoretically, this renders tacit agreements on Pareto optimal paths much more difficult to attain than in purely repeated frameworks.

Seminal experimental investigations placing subjects in time-dependent CPR contexts have been conducted by Herr, Gardner, and Walker (1997) and Mason and Philips (1997). Herr et al. (1997) considered the effect of increasing exploitation costs at a predetermined fixed rate from period to period in a 10-period supergame. They concluded that subjects did not internalize the future increased costs, and that behavior in the time-dependent setting intensified the race for resources relative to time-independent settings. Although current decisions in this study affected future exploitation costs, an effect on the size of the resource stock was not considered. Mason and Philips (1997) considered an infinite time horizon supergame in which subjects were given an initial stock, and request strategies endogenously determined the stock

size (and exploitation costs) thereafter. They concluded that lack of cooperative behavior is exacerbated when time-dependency is included in CPR dilemmas. Osés-Eraso, Udina, and Viladrich-Grau (2008) have modified this game by implementing a more realistic finite horizon supergame and allowing for early extinction of the stock as a function of cumulative group requests. Exogenously manipulating the initial stock size (scarce or abundant), they found that early extinction of the resource occurs irrespective of the initial scarcity condition and costs.

Although Osés-Eraso et al. (2008), and one of the experimental conditions in Mason and Philips (1997), allowed for extinction of the resource stock as a function of the cumulative requests, they excluded the possibility for more complex resource dynamics involving threshold effects and sudden changes in the resource state. The only study that partially addressed this possibility experimentally was Gardner and Walker (1992). They implemented a 20-period supergame in which resource extinction (end of the game) could occur within a given period with an endogenously determined probability modeled as an increasing function of total group requests. In this model, the critical threshold stock level triggering the catastrophic event was known in advance, but subjects were unable to avoid the damage, and the resource was quickly destroyed with a median duration of six periods. However, the model does not capture the effects of the lack of precise information regarding the size and growth of natural resources, which characterizes most real-world commons. The conditions that trigger an irreversible event are often imperfectly known, or are affected by stochastic environmental conditions outside CPR users' control, rendering the critical threshold stock level unknown a priori.

Experimental assessments of the impacts of environmental uncertainty on group requests from a CPR have mainly been conducted by Rapoport and co-authors (Rapoport and Suleiman, 1992; Rapoport, Budescu, Suleiman, and Weg 1992; and Budescu, Rapoport, and Suleiman 1995) in the context of repeated single-trial experiments. In their experiments, subjects could request resources from pools whose parameters were randomly selected from a set of commonly known uniform probability distributions. Using mean-preserving spreads to capture increasing levels of uncertainty regarding the resource size, these experiments have demonstrated that increased uncertainty causes subjects to request more from the shared resource. Along with the observed uncertainty effect, Rapoport and co-authors have provided several explanations for its occurrence. One explanation posits that subjects perceive the central tendency of a probability distribution and its variability to be positively correlated. Thus, as the uncertainty increases, subjects' estimates of the mean value of the resource size increases, prompting them to request more. Another explanation, in line with an "outcome desirability bias" (Gustafsson, Biel, and Garling 1999), posits that subjects overweight the most desired upper bound of the possible resource size. This, again, prompts subjects to request more of the resource as the uncertainty increases in mean-preserving resource distributions. Despite the merits of these explanations, the literature seems to have overlooked that the observed relationship between environmental uncertainty and individual requests pertains to single-trial experiments. Under these circumstances, a significant restraint in individual requests by the group members that can be considered cooperative behavior may not yield the highest collective payoffs as it may constitute resource under-use from an economically efficient perspective. Indeed, increased requests as a response to an increase in environmental uncertainty levels in time-independent settings conforms to Pareto-efficient solutions, and not just Nash behavior. This observation, however, is unlikely to be true in time-dependent settings. A

number of theoretical articles on various resource management problems, including intrusion of saltwater in coastal aquifers (Tsur and Zemel 1995), forest fires (Yin and Newman 1996), and pollution-related events (Clarke and Reed 1994), have generally established that efficient solutions for the dynamic management of resources under uncertain critical threshold levels require more prudence and conservative behavior than those under conditions of certainty. Thus, whether or not the positive relationship between environmental uncertainty and request behavior from a shared resource is likely to be observed in a time-dependent laboratory setting, in which both considerations of individual future payoffs and efficient solutions require more conservative requests, is still an open empirical question.

The present paper addresses this question by developing and experimentally testing a dynamic stochastic game-theoretic model integrating the effects of environmental uncertainty in time-dependent CPR dilemmas. The paper is organized as follows. Section 2 presents the model and solves it for general theoretical benchmarks. Section 3 outlines the experimental design and presents the theoretical predictions that are later used as benchmarks for the analysis of the experimental data. Section 4 reports the results of the experiment, and Section 5 concludes.

2. Dynamic-stochastic CPR game

We develop a dynamic-stochastic game-theoretic model of the appropriation of a CPR by non-cooperative players under conditions of environmental uncertainty. The game involves a fixed set of n players who play a stage game Γ_t in each period t , where an upper bound T to the length of the game is common knowledge, and earnings accumulate through the course of play. However, in contrast to purely dynamic time-dependent games with no uncertainty in the evolution of the game environment (Dutta 1995), the particular game to be faced by the players at time t in the present setup is randomly selected from a commonly known finite set of games, thereby falling in the category of stochastic games (Shapley 1953; Sobel 1971). In addition, in order to capture the effects of environmental uncertainty on the CPR dilemma, players do not know which game has been selected when the game at time t is to be played.

Each of the stage games Γ_t that make up the dynamic game is drawn from Suleiman and Rapoport (1988). In each stage game Γ_t , a group of n players decide simultaneously and anonymously on how much to request from a shared resource (CPR) whose precise size is unknown. However, it is commonly known that the resource size, denoted by S_t , is uniformly distributed on the $[\alpha, \beta]$ closed interval. Each of the n individuals may request anything between 0 and β from the shared resource, and *after* the requests are made, the precise size of the resource is publicly announced, corresponding to the random realization s_t of S_t . Thus, the precise value of s_t corresponds to the particular stage game Γ_t randomly selected by “nature” at time t . Furthermore, if the sum of group requests is smaller than or equal to s_t , then each individual keeps his or her own requests. On the other hand, if the sum of group requests exceeds the size s_t of the resource, then each individual’s payoff is zero.

Assuming linear utility functions for all players, and letting r_{jt} stand for the request made by player j on trial t , the expected payoff to the player in stage game Γ_t is given by

$$\pi_{jt} = \begin{cases} r_{jt} & \text{if } \sum_{j=1}^n r_{jt} \leq \alpha \\ r_{jt} \times \text{Prob}(\sum_{j=1}^n r_{jt} \leq s_t) & \text{if } \alpha < \sum_{j=1}^n r_{jt} \leq \beta \\ 0 & \text{if } \sum_{j=1}^n r_{jt} > \beta \end{cases} \quad (1)$$

where $\sum_{j=1}^n r_{jt}$ is the sum of group requests in stage game Γ_t , and $\text{Prob}(\sum_{j=1}^n r_{jt} \leq s_t) = (\beta - \sum_{j=1}^n r_{jt})/(\beta - \alpha)$.

We introduce structural time dependence in this model through the definition of transition probabilities governing how the game proceeds from period t to period $t+1$, which we condition on the actual game Γ_t played in period t as randomly selected by “nature,” and on the actions chosen by the players in period t . Specifically, if the aggregate requests are smaller than or equal to the resource size at time t , then the game continues to period $t+1$. If the aggregate requests are infeasible, in the sense that they exceed the resource size at time t , then the game terminates. Formally, the continuation probability from period t to period $t+1$ is given by:

$$p_t = \begin{cases} 1 & \text{if } \sum_{j=1}^n r_{jt} \leq \alpha \\ \frac{\beta - \sum_{j=1}^n r_{jt}}{\beta - \alpha} & \text{if } \alpha < \sum_{j=1}^n r_{jt} \leq s_t \\ 0 & \text{if } \sum_{j=1}^n r_{jt} > s_t \end{cases} \quad (2)$$

In other words, if the group request is below a minimum pre-determined quantity α , the game continues to a subsequent period with certainty, implying an economically unchanged resource size between the periods. If the group request exceeds the randomly determined resource size, then the resource is degraded and the game is terminated. If the group request belongs to this interval of quantities, the game continues to a subsequent period with a positive probability corresponding to the *ex-ante* probability that the group request does not exceed the resource size. Thus, while players may request resources over a predetermined and commonly known time horizon, a distinguishing feature of this model is that the precise duration of the game is determined *endogenously* by the players whose collective decisions determine the probability of an irreversible environmental event. In addition, the model captures those circumstances characterized by both pool-size uncertainty and regeneration-rate uncertainty, features that are present in many real-world commons (Hine and Gifford 1996). In particular, when group requests exceed α but the resource is not depleted, the inter-temporal effect of group requests may be interpreted as captured by a stochastic regeneration rate g_t applied to end of period remaining stock, $s_t - \sum_{j=1}^n r_{jt}$. The parameter g_t determines the stock available in the subsequent period, $s_{t+1} = (s_t - \sum_{j=1}^n r_{jt})g_t$, where g_t is uniformly distributed with limits endogenously determined by group requests and the stochastic resource size, $[\alpha/(s_t - \sum_{j=1}^n r_{jt}), \beta/(s_t - \sum_{j=1}^n r_{jt})]$.

Because the CPR game is composed of interdependent stochastic dynamic programming problems, it may be solved by dynamic programming/Bellman’s equation method, and because it is symmetric we focus on symmetric outcomes as benchmarks for data analysis. In particular, we solve the game for three types of outcomes: the social optimum (joint payoff maximization) outcome, the subgame perfect outcome, and a conservative outcome guaranteeing survival of the resource over the entire time horizon.

We first construct the symmetric subgame perfect Nash equilibrium outcome, in which players are assumed to adopt “decision rule strategies” (Reinganum and Stokey 1985) since they cannot credibly commit to future requests. In this context, each player j

independently seeks to maximize the value of the resource at any time t by choice of request strategy, taking the decision rule strategies of all the other players exploiting the resource as given. Assuming no discounting of future payoffs, the value of the resource for player j at time t , $V_{jt}(r_{jt})$, satisfies the Bellman equation:

$$V_{jt}(r_{jt}) = \pi_{jt} + p_t \times V_{jt+1}(r_{jt+1}), \quad t = 1, 2, \dots, T - 1. \quad (3)$$

The transversality condition for this maximization problem is that the value of the resource after the final period T is zero (meaning that players leave behind no resources, or if they do those extra resources do not contribute anything to maximized value):

$$V_{jT+1}(r_{jT+1}) = 0. \quad (4)$$

The recursive equation defining player's j value function at final time T is therefore:

$$V_{jT}(r_{jT}) = \pi_{jT}. \quad (5)$$

Maximizing $V_{jT}(r_{jT})$ in the quadratic region of (1) with respect to r_{jT} , and invoking symmetry to write the sum of requests by all the n players excluding player j as $r_{n \setminus jT} = (n - 1)r_{jT}$, yields the subgame perfect request at time T :

$$r_{jT}^* = \beta / (n + 1). \quad (6)$$

The value function at time T is then given by (using (6) in (5)):

$$V_{jT}(r_{jT}^*) = [\beta / (n + 1)]^2 / (\beta - \alpha). \quad (7)$$

Similarly, the value function at time $T-1$ is given by:

$$V_{jT-1}(r_{jT-1}) = \pi_{jT-1} + p_{T-1} \times V_{jT}(r_{jT}^*). \quad (8)$$

Maximizing $V_{jT-1}(r_{jT-1})$ with respect to r_{jT-1} , and assuming that $r_{n \setminus jT-1} = (n - 1)r_{jT-1}$, yields the subgame perfect request at time $T-1$:

$$r_{jT-1}^* = [\beta / (n + 1)][1 - \beta / (n + 1)^2(\beta - \alpha)]. \quad (9)$$

The value function at time $T-1$ is then given by (using (9) and (7) in (8)):

$$V_{jT-1}(r_{jT-1}^*) = [(\beta / (n + 1))^2 / (\beta - \alpha)][1 + n\beta / (n + 1)^2(\beta - \alpha)]^2. \quad (10)$$

Letting $a = n\beta / (n + 1)^2(\beta - \alpha)$, the value function at time $T-1$ can be re-written as:

$$V_{jT-1}(r_{jT-1}^*) = V_{jT}(r_{jT}^*)[1 + a]^2, \quad (11)$$

and the subgame perfect request at time $T-1$ can be written as:

$$r_{jT-1}^* = r_{jT}^*[1 - a/n]. \quad (12)$$

By mathematical induction, one can show that the equilibrium value function at any time t is given by:

$$V_{jt}(r_{jt}^*) = V_{jT}(r_{jT}^*)[1 + a\gamma_t]^2 = [(\beta / (n + 1))^2 / (\beta - \alpha)][1 + a\gamma_t]^2, \quad (13)$$

and the subgame perfect request at time t is given by

$$r_{jt}^* = r_{jT}^*[1 - (a/n)\gamma_t] = [\beta / (n + 1)][1 - (a/n)\gamma_t], \quad (14)$$

where the recursive factor γ_t is given by:

$$\gamma_t = (1 + a\gamma_{t+1})^2. \quad (15)$$

One starts solving the recursion by noting that (7) can also be written as:

$$[\beta/(n+1)]^2/(\beta - \alpha) = r_{jT}^*[\beta/(n+1)(\beta - \alpha)], \quad (16)$$

and that using (14) and (15), the equation above can also be written as:

$$[\beta/(n+1)]^2/(\beta - \alpha) = [\beta/(n+1)][1 - (a/n)(1 + a\gamma_{t+1})^2][\beta/(n+1)(\beta - \alpha)]. \quad (17)$$

Solving (17) with respect to γ_{T+1} yields:

$$\gamma_{T+1} = -1/a. \quad (18)$$

The value of γ_{T+1} in (18) can then be substituted into (15) to get γ_T , and working backward from there to the first period at time $t=1$.

Note, however, that the solution in (14) does not constitute the subgame perfect equilibrium request in all cases. At any time t , if $r_{jt} + r_{n \setminus jt} \leq \alpha$, any vector of requests $\mathbf{r}_t^* = (r_{1t}, r_{2t}, \dots, r_{nt})$, whose elements satisfy the condition $\sum_{j=1}^n r_{jt}^* = \alpha$, $r_{jt} \geq 0$, is also an equilibrium solution whereby the group ensures a continuation probability equal to one. Assuming a symmetrical solution, the result following this condition can be written as:

$$r_{jt}^* = \alpha/n. \quad (19)$$

The subgame perfect request at any time t in all cases is then given by:

$$r_{jt}^* = \max(\alpha/n, [\beta/(n+1)][1 - (a/n)\gamma_t]). \quad (20)$$

Moving next from the equilibrium solution to joint payoff maximization, the social optimum path can be constructed by applying dynamic programming to (3) under the assumption that only a single agent is in charge of the resource. The social optimum request by player j at any time t in all cases is then given by:

$$r_{jt}^* = \max(\alpha/n, [\beta/(2n)][1 - (b)\lambda_t]), \quad (21)$$

where $b = \beta/4(\beta - \alpha)$, $\lambda_t = (1 + b\lambda_{t+1})^2$, and $\lambda_{T+1} = -1/b$.

Comparison of equations (20) and (21) reveals that the subgame perfect path involves higher requests than the social optimum path as long as $(\beta - \alpha) > \alpha/n$, rendering the decision rule strategies Pareto deficient for levels of uncertainty beyond a relatively small threshold level. In turn, the social optimum path only guarantees the survival of the resource over the entire time horizon for moderate levels of uncertainty, that is, as long as $\beta < 2\alpha$. Beyond this level of uncertainty, survival of the resource over the entire time horizon could only be attained by the adoption of “path” strategies requiring each player to commit to a “conservative” request equal to α/n at each stage of the game.

3. Experimental design and theoretical predictions

A. Procedures, parameters and treatments

We designed a simple experiment operationalizing the game described by equations (1) and (2) with groups composed of six ($n=6$) subjects and a time horizon of ten periods ($T=10$). Each subject participated in thirty repetitions of the same dynamic game. Prior to the first game, each subject was randomly and anonymously assigned to

a fixed group for the duration of a session. We implemented two mean-preserving uncertainty conditions in a between-subject design. In one of the uncertainty conditions (hereinafter, “high” uncertainty condition), the commonly known resource size was uniformly distributed on the [150, 850] closed interval, for an uncertainty range of 700 and an expected value of 500. Subjects were provided with written instructions informing them that they could, individually, request from 0 up to 850 tokens, and that the precise value of the resource (called “random draw”) in any period was to be randomly extracted (and announced) after all group members made their requests. They were also informed that if the sum of group requests were larger than the randomly determined resource size in any period, then their individual payoffs in that period would be zero, and the game would be terminated; otherwise, their individual payoffs in that period would equal their individual requests, and the game would continue to a subsequent period unless the game had reached the final period. Specifically, subjects were informed that the game would be terminated either if the sum of the group requests exceeded the value of the resource or after 10 periods, whichever came first. In addition to a \$5 participation fee, at the end of the session subjects were paid for the tokens accumulated in four (randomly determined for each subject) out of the thirty repetitions, in which each token was worth 2 cents. This procedure was implemented to prevent wealth effects. The exact same procedures were used to implement a second uncertainty condition (hereinafter, “low” uncertainty condition) in which the commonly known resource size was uniformly distributed on the [270, 730] closed interval, for a smaller uncertainty range of 460 and the same expected value of 500. In each condition, the sessions lasted for about two hours.

The experiment was implemented using the *z-Tree* (Fischbacher 2007) software. No communication between the subjects was allowed. All experimental sessions were conducted at the Behavioral Research Lab of the School of Business Administration at the University of California, Riverside, which is a standard computerized laboratory with subjects’ stations placed in separate “cubicles” ensuring privacy. Subjects were recruited from the pool of UCR students registered to participate in research studies through the web-based subject recruitment system available at UCR, ensuring that no subject had participated in a similar experiment before. A total of 114 subjects participated in this experiment, 60 of them in the high uncertainty condition (10 different groups) and 54 of them in the low uncertainty condition (9 different groups).

B. Theoretical predictions

We present the theoretical predictions that are used as benchmarks for the analysis of the data from the two treatments. The top panel in Table 1 shows the dynamic programming paths for the high uncertainty condition. If players follow a conservative path (not shown in the Table), the symmetric individual request (r) is 25 tokens in each period of the game, for a total group request (R) of 150 each period of the game. The overall payoff across the 10 periods of the game (II) for each player is, therefore, $25 \times 10 = 250$ tokens. Next, consider the social optimum (SO) path displayed in Table 1 (top panel). In the last period of the game, when only a single period remains to the end of the game, the optimal solution is an individual request of 71. The probability of receiving this request (p) is about 0.61, yielding an expected payoff of $II = 43$ tokens. If only two periods remain, the optimal solution is an individual request of 49. The probability of receiving this request is about 0.79. The individual’s expected payoff across these two periods is, therefore, $49 \times 0.79 + 0.79 \times 43 = 73$ tokens. Working backwards in this fashion, the individual’s expected payoff across the 10 periods of the game from following the social optimum path is equal to 274 tokens, the value of II

shown in Table 1 when ten periods are remaining (i.e., at the beginning of the game). This corresponds to the maximum symmetric expected payoffs that subjects may achieve in this game. Comparing the expected payoffs from following a conservative path to the social optimum path yields an efficiency index of $(250/274) \times 100 = 91\%$ for the conservative path. This means that subjects are expected to achieve 91 percent of the maximum expected payoffs that may be achieved in this game if they follow a conservative path. The subgame perfect equilibrium (SPNE) path shown in Table 1 is constructed in the same manner, considering the predicted (symmetric) Nash equilibrium requests by each player. In this case, the individual's expected payoff across the ten periods of the game from following this path is 31 tokens, yielding a meager efficiency index of only about 11%.

The bottom panel in Table 1 shows the dynamic programming paths for the low-uncertainty condition. If players follow a conservative path in this condition, the symmetric individual request is 45 tokens in each period of the game, for a total group request of 270 each period of the game. The overall payoff across the 10 periods of the game for each player is $45 \times 10 = 450$ tokens. The individual's expected payoff across the 10 periods of the game from following the social optimum path is 453 tokens, and is 44 tokens from following the subgame perfect equilibrium path. In this case, the subgame perfect path yields an efficiency index of about 10 percent.

Comparison of the upper and lower parts of Table 1 shows that, as might be expected, payoffs increase under all three benchmarks as the uncertainty about the size of the resource decreases. Table 1 shows that each of these two uncertainty conditions yields different predictions concerning players' requests from the shared resource, with the social optimum path entailing substantially lower requests than the respective subgame perfect path. Therefore, the subgame perfect paths are Pareto deficient in both of the uncertainty conditions implemented in the laboratory. Importantly, the efficiency index of the subgame perfect path is maintained approximately equal in both uncertainty conditions, so that the incentives for any cooperative behavior do not differ much between the two conditions. Moreover, any increase in the requests that might be observed in response to an increase in the uncertainty levels cannot simultaneously make part of a competitive (subgame perfect) and a cooperative (conservative or social optimum) path. Indeed, as shown in Table 1, while the increased uncertainty in the high uncertainty condition elicits higher requests than in the low condition, if subjects follow the subgame perfect path, it overall elicits significantly lower requests than in the low-uncertainty condition, if subjects follow a conservative path or the social optimum path.

4. Experimental results

Our analysis of the experimental data focuses on the effects of environmental uncertainty on resource-use decisions at the group level. We organize the analysis of group behavior by examining in order: (A) behavior in the high uncertainty condition, (B) behavior in the low uncertainty condition, and (C) comparison of the behavior across the two uncertainty conditions.¹ In each case, the main results are presented in the form of summary observations.

¹ Because subjects participated in 30 repetitions/series of the same dynamic game, we first investigated whether play of the games changed as subjects gained experience. The general finding in purely repeated CPR games is that behavior is consistent with efficient outcomes in the first rounds of play, and

A. High-uncertainty condition

Tables 2 and 3 summarize the main results of the high-uncertainty condition. Table 2 presents the length of the games played by each group. Pooling across all 10 groups, the median length of the games is one period, with three of the 10 groups registering a median length of two periods. Clearly, none of the groups adopted a conservative path, and depletion of the resource stock occurred rather quickly.

Figure 1 depicts the probability of resource destruction as implied by the social optimum and equilibrium paths (broken lines) along with the observed proportions. Using the predictions in the top panel of Table 1, the probability of resource destruction prior to period eight is 0 percent at the social optimum path, increasing to 8 percent $((1-0.92)\times 100)$ prior to period nine and to 27% $((1-0.92\times 0.79)\times 100)$ prior to period 10. In the context of dynamic games, theory tells us that we should observe an immediate depletion of the resource stock if groups are unable or unwilling to make commitments about future extraction rates (Reinganum and Stokey 1985), corresponding to the assumption that behavior is guided by decision rule strategies underlying the predicted SPNE path. In fact, as implied by the numbers in Table 1, the probability of resource destruction prior to period two is 79 percent $((1-0.21)\times 100)$ at the equilibrium path, increasing to 96 percent $((1-0.21\times 0.21)\times 100)$ prior to period three, and reaching about 100 percent in subsequent periods. With such high destruction probabilities, the chance of observing games lasting for more than a single period in the data is quite small if subjects do not deviate from the predicted equilibrium path. As Figure 1 reveals, despite the variability of group behavior, the rates of resource destruction are quite above those predicted by the social optimum path, and closer to the equilibrium path.

Table 3 shows that group requests terminating the game immediately, which accounted for about 55 percent of the data, average 688 tokens. This mean compares closely to the equilibrium prediction of 702 tokens. As expected, first-period requests are negatively associated with the length of the game. The average first-period group requests for games longer than one period, which account for 45 percent of the data, is 487 tokens. These requests are in between the efficient and the equilibrium values.

Using Wald tests while adjusting standard errors for clustering at the group level, the signed mean differences between the observed first-period group requests (R^{Obs}) and the equilibrium (R^{SPNE}) and efficient (R^{SO}) values were tested for statistical significance. The results are summarized in Table 4. They show that group requests in games terminated in the first period ($\text{Length}=1$) are not significantly different from the predicted equilibrium value. The mean estimated difference between R^{Obs} and R^{SPNE} is about -14 tokens; it is not significantly different from zero at conventional significance levels using the Wald statistics. The mean differences between first-period group requests and predicted equilibrium and efficient values are both significantly different from zero for games lasting for more than one period ($\text{Length}>1$). Mean group requests are 214 tokens below the equilibrium value, and 337 tokens above the efficient value. In order to evaluate whether the observed deviations from equilibrium requests are larger or smaller than the observed deviations from efficient requests, the delta method

approaches the equilibrium prediction in the last rounds. Under this pattern of behavior, we would expect to observe longer games in the first series of play, and shorter games as the series approach the end. The figure in the Appendix plots the maximum number of periods played by each group in each of the games, where the title in each of the panels identifies the uncertainty condition. In each case, the figure suggests that there is no systematic association between the length of the games and order of play. This impression was confirmed by several statistical analyses (available from the authors). Therefore, we pool the data across the games for the statistical analysis of the data in both uncertainty conditions.

(Oehlert 1992) was used to calculate the standard error and 95 percent confidence interval of the absolute value of the ratio of the estimated difference between observed requests and the respective equilibrium and efficient requests (Δ). The confidence interval for the ratio is [0.4; 0.9]. It indicates that the observed deviations from equilibrium requests are smaller than the observed deviations from efficient requests.

To complement the analysis of group behavior, we computed the per-period mean square deviation (MSD) of requests from predicted requests (either SPNE or SO) for each group in each of the 30 played games. For each group separately, Table 5 indicates the number of games in which the MSD from the SPNE path is smaller than the MSD from the SO path. Also reported in the table are the binomial probabilities associated with the observed number of games under the null hypothesis that it is equally likely for either of the two predicted paths to result in the smaller MSD in any given game. The results show that the SPNE path is the best predictor of behavior for eight of the 10 groups, and that for two of the groups we cannot reject the hypothesis that both paths are equally likely at a significance level of 5 percent. Defining success as an observation in which the SPNE path is the best predictor of group behavior, the probability of observing 8 or more groups following the SPNE path is 0.003 under the null hypothesis that the three events (SPNE, SO, or both) are equally likely. For any one-tailed significance level lower than 5 percent, we reject the null hypothesis in favor of the alternative hypothesis that the SPNE path is the overall best predictor of behavior for the groups in this uncertainty condition.

Taken together, these findings are summarized in the following observation.

Observation 1. Groups in the high-uncertainty condition adopt decision strategies that quickly deplete the resource stock. Group requests are uniformly closer to the SPNE path than to the SO path.

B. Low-uncertainty condition

Tables 6 and 7 summarize the main results in the low-uncertainty condition. The median length of the games in the low-uncertainty treatment across all groups is two periods, with four of the nine groups (44 percent) registering a median length of one period. It is twice as large as the same median in the high-uncertainty condition. Again, none of the groups adopted a conservative path, with depletion of the resource stock occurring rather quickly.

Figure 2 depicts the probability of resource destruction as implied by the SO and SPNE paths, along with the observed proportions. Although the destruction probability curves are more dispersed than in Figure 1, they are closer to the SPNE than the SO path.

Table 7 shows that group requests terminating the game immediately, which accounted for about 46% of the data, average 569. This mean request compares closely to the equilibrium prediction of 588 tokens. The mean first-period group requests for games longer than one period is 412 tokens. These requests are in between the efficient and the equilibrium values.

Table 8 addresses the issue of whether first-period requests are significantly different from the SPNE and SO paths. In games terminating in the first period, the mean difference between R^{Obs} and R^{SPNE} is -19 tokens; the null hypothesis of zero difference could not be rejected by the Wald test at conventional significance levels. The mean differences between first-period group requests and predicted equilibrium and efficient values are both significantly different from zero for games lasting for more

than one period ($\text{Length} > 1$). Mean group requests are 177 tokens below the equilibrium value, and 142 tokens above the efficient value. Although the absolute value of the ratio of these two differences is greater than 1, suggesting that the distance between observed requests and the SPNE prediction is larger than the distance between observed requests and the SO prediction, the computed confidence interval for the ratio is [0.7; 1.8]. Thus, at the 5 percent significance level, we cannot reject the hypothesis that the observed deviations from the SPNE path are equal to the observed deviations from the SO path.

Table 9 addresses the issue of whether groups' behavior is better described by the SPNE or the SO path. The results show that the SPNE path is the best predictor of behavior for four of the nine groups. The SO path is the best predictor of behavior for two other groups, and the results of the testing procedure are inconclusive for the remaining three groups. Defining success as an observation in which the SPNE path is the best predictor of group behavior, the probability of observing four or more groups following the SPNE path is 0.35 under the null hypothesis that the three events (SPNE, SO, or both) are equally likely. Therefore, we do not reject the null hypothesis at conventional significance levels.

We summarize these findings in the following observation.

Observation 2. Groups in the low-uncertainty condition tend to adopt decision strategies that quickly deplete the resource stock. However, the SPNE path is not uniformly the best predictor of group requests, with some groups adopting behavior closer to the SO path and other groups adopting behavior falling in between these two polar cases.

C. Comparing uncertainty conditions

We would expect higher uncertainty about the size of the resource to elicit *higher* group requests, if groups adopt decision strategies, but to elicit *lower* group requests, if groups adopt path strategies leading to perfectly efficient outcomes. As seen above (Tables 3 and 7), and consistent with the adoption of decision strategies, group requests are *higher* in the high-uncertainty condition than in the low-uncertainty condition. Given that the same differences in requests generate different probabilities of resource destruction across different manipulations of uncertainty ranges, a general assessment of the effects of increased uncertainty is better accomplished by analyzing the implied differences in destruction probabilities rather than by analyzing the differences in requests observed across the different manipulations of uncertainty ranges.

Panels A and B in Table 10 report the estimated effects of the higher-uncertainty level on the implied probabilities of resource destruction by first-period requests. For completeness, also reported in Table 10 (panel C) is the estimated effect implied by all non-first-period requests. Given that the dependent variable is naturally bounded between 0 and 1, the estimation of treatment effects uses the specification developed by Papke and Wooldridge (1996) for fractional-dependent variables. In addition, because the conditional expectation function in the specification used is nonlinear (so as to generate predictions naturally bounded between 0 and 1), the estimated parameter value associated with the treatment variable does not directly measure the treatment effect on the mean value of the dependent variable. Thus, to aid in interpretation, the coefficient estimates reported in Table 10 are the marginal effects of a discrete change in explanatory variable HIGH taking the unit value for the high-uncertainty condition and the zero value for the low-uncertainty condition.

Table 10 shows that the implied probabilities of destruction induced by the higher request in the high-uncertainty condition are significantly higher than the probabilities of destruction observed in the low-uncertainty condition. Considering only the subset of games terminated in the first period, the probabilities of destruction are, on average, 12 percentage points higher in the high than the low-uncertainty condition. Moreover, the width of the 95 percent confidence interval indicates that we cannot reject the hypothesis that the difference in destruction probabilities between the two treatments is 10 percentage points, corresponding to the predicted difference generated by the SPNE paths. This result is not particularly surprising, given that group requests in both treatments are consistent with the respective SPNE paths for this subset of the data.

Considering only the first-period requests in the subset of games lasting for more than one period, the difference in destruction probabilities between the treatments is 17 percentage points, significantly higher than the predicted difference by the SPNE paths. Because first-period requests in games terminated after the first period are lower than the predicted SPNE values in *both* treatments, it could still be the case that the difference in the implied destruction probabilities remained at about the 10 percentage points, generated by the respective SPNE paths. Clearly, this is not the case, suggesting that reducing uncertainty levels positively impacts resource conservation beyond what would be predicted by common inability of the groups to commit to future extraction rates (i.e., by behavior consistent with the SPNE path).

This observation is further corroborated by the results in Panel C, considering only the subset of all non-first requests in both treatments. Had groups approximated their requests to the SPNE paths after the first period of the game, the mean differences in destruction probabilities between the treatments would be 8.5 percentage points, since the difference declines systematically as the game evolves. Consistent with this pattern of behavior, we observe lower differences in destruction probabilities between the treatments in subsequent periods. However, as indicated by the 95 percent confidence interval, the difference is again significantly higher than would be predicted by groups approximating their respective SPNE paths.

Coupled with those summarized in Observations 1 and 2 above, these findings indicate not only that treatment effects cannot solely be attributed to Nash behavior, but also that it is groups' behavior in the low-uncertainty condition that explains the differential treatment effect with respect to equilibrium predictions.

These findings are summarized in the following observation.

Observation 3. Compared with the high-uncertainty condition, the low-uncertainty condition elicits lower requests from the shared resource. Moreover, it also induces a qualitative change in groups' behavior in the sense that it positively impacts resource conservation beyond what would be predicted by groups adopting decision rule strategies under both conditions.

5. Conclusion

The dynamic stochastic game-theoretic model proposed in this paper focuses on the effects of environmental uncertainty in time-dependent CPR dilemmas. While CPR users may extract resources over a predetermined and commonly known time horizon, a distinguishing feature of our model is that the duration of the game is determined

endogenously by the players whose collective decisions determine the probability of an irreversible environmental outcome. The abrupt intrusion of salt water in coastal aquifers once the groundwater table declines below an unknown threshold level is an example of such an event. In the present model, if the resource stock level below which the irreversible outcome occurs is known in advance, then the optimal resource use coincides with a unique symmetric equilibrium use guaranteeing survival of the resource over the finite horizon. As the uncertainty about an otherwise equally expected threshold level increases, resource use increases if users adopt decision strategies that quickly deplete the resource stock. Resource use decreases if users adopt path strategies guaranteeing that the unknown threshold level is never exceeded over the entire horizon.

In an experiment that manipulates the common uncertainty about the threshold resource level, we find that CPR users implement decision strategies that terminate the game immediately. Notwithstanding, reducing the uncertainty about the resource level induces a *qualitative* change in behavior with users more frequently maintaining a positive resource level for a longer duration. If replicated and extended, these results have potentially important theoretical and policy implications. At the theoretical level, they suggest decision strategies that CPR users may use when they may not make credible commitments. At the policy level, these results provide evidence that the reduction of environmental uncertainty by creating and disseminating better scientific information may play a major role in long-range planning to elicit synergy between the economic and ecological systems that jointly govern the dynamic management of shared natural resources. Estimated as the difference between the high- and the low-uncertainty outcomes, the value of this information is an indicator to the policy maker about how much to invest in acquiring and disseminating information to the user that reduces uncertainty about the CPR.

References

- Brown, G. M. (2001) Renewable natural resource management and use without markets. *Journal of Economic Literature*, 38(4), 875-914.
- Budescu, D. V., Rapoport, A. and Suleiman, R. (1995) Common pool resource dilemmas under uncertainty: Qualitative tests of equilibrium solutions. *Games and Economic Behavior*, 10, 171–201.
- Clarke, H. R. and Reed W. J. (1994) Consumption/pollution tradeoffs in an environment vulnerable to pollution-related catastrophic collapse. *Journal of Economic Dynamics and Control*, 18, 991– 1010.
- Dutta, P. (1995) Collusion, discounting and dynamic games. *Journal of Economic Theory*, 66, 289-306.
- Fischbacher, U. (2007) Z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2), 171–178.
- Gardner, R. and Walker, J. M. (1992) Probabilistic destruction of common-pool resources: Experimental evidence. *The Economic Journal*, 102(414), 1149-1161.
- Gordon, H. S. (1954) The economic theory of a common property resource: The fishery. *Journal of Political Economy*, 62, 124–142.
- Gustafsson, M., Biel, A. and Garling, T. (1999) Overharvesting of resources of unknown size. *Acta Psychologica*, 103, 47-64.
- Hardin, G. (1968) The tragedy of the commons. *Science*, 162, 1243-1248.
- Hey, J. D., Neugebauer, T. and Sadrieh, A. (2009) An experimental analysis of optimal renewable resource management: The fishery. *Environmental and Resource Economics*, 44, 263–285.
- Herr, A., Gardner, R. and Walker, J.M. (1997) An experimental study of time-independent and time-dependent externalities in the commons. *Games and Economic Behavior*, 19(1), 77–96.
- Hine, D. W. and Gifford, R. (1996) Individual restraint and group efficiency in commons dilemmas: The effects of two types of environmental uncertainty. *Journal of Applied Social Psychology*, 26(11), 993-1009.
- Mason, C. F. and Phillips, O. R. (1997) Mitigating the tragedy of the commons through cooperation: An experimental evaluation. *Journal of Environmental Economics and Management*, 34(2), 148–172.
- Messick, D. M. and McClelland, C. L. (1983) Social traps and temporal traps. *Personality and Social Psychology Bulletin*, 9(1), 105-110.
- Oehlert, G. W. (1992) A note on the delta method. *The American Statistician*, 46(1), 27-29.
- Osés-Eraso, N., Udina, F. and Viladrich-Grau, M. (2008) Environmental versus human-Induced scarcity in the commons: Do they trigger the same response? *Environmental and Resource Economics*, 40, 529–550.
- Ostrom, E., Gardner, R. and Walker, J. (1994) *Rules, Games, & Common-pool Resources*. Ann Arbor: University of Michigan Press.

- Papke, L. E. and Wooldridge, J. M. (1996) Econometric methods for fractional response variables with an application to 401(K) plan participation rates. *Journal of Applied Econometrics*, 11, 619-632.
- Rapoport, A. and Suleiman, R. (1992) Equilibrium solutions for resource dilemma. *Group Decision and Negotiation*, 1, 269–294.
- Rapoport, A., Budescu, D. V., Suleiman, R. and Weg, E. (1992) Social dilemmas with uniformly distributed resources. In *Social Dilemmas: Theoretical Issues and Research Findings*, W. Liebrand, D. M. Messick, and H. Wilke (eds). Oxford: Pergamon Press, pp. 43–57.
- Reinganum, J. F. and Stokey, N. L. (1985) Oligopoly extraction of a common property natural resource: The importance of the period of commitment in dynamic games. *International Economic Review*, 26(1), 161-173.
- Shapley, L. (1953) Stochastic Games. *Proceedings of the National Academy of Science*, 39, 1095-1100.
- Sobel, M. J. (1971) Noncooperative stochastic games. *The Annals of Mathematical Statistics*, 42(6), 1930-1935.
- Suleiman, R. and Rapoport, A. (1988) Environmental and social uncertainty in single-trial resource dilemmas. *Acta Psychologica*, 68, 99-112.
- Tsur, Y. and Zemel, A. (1995) Uncertainty and irreversibility in groundwater resource management. *Journal of Environmental Economics and Management*, 29, 149–161.
- Yin, R. and Newman, D. H. (1996) The effect of catastrophic risk on forest investment decisions. *Journal of Environmental Economics and Management*, 31, 186– 197.

Table 1 – Dynamic programming paths for high- and low-uncertainty conditions

Time Remaining	Social optimum (SO) path				Subgame perfect (SPNE) path				
	R	r	p	Π	R	r	p	Π	
<i>A. High uncertainty condition: $n=6$, $\alpha=150$, $\beta=850$, Expected Value=500, Range=700</i>									
1	425	71	0.61	43	729	121	0.17	21	
2	296	49	0.79	73	711	118	0.20	28	
3	206	34	0.92	99	705	117	0.21	30	
4	150	25	1.00	124	703	117	0.21	31	
5	150	25	1.00	149	702	117	0.21	31	
6	150	25	1.00	174	702	117	0.21	31	
7	150	25	1.00	199	702	117	0.21	31	
8	150	25	1.00	224	702	117	0.21	31	
9	150	25	1.00	249	702	117	0.21	31	
10	150	25	1.00	274	702	117	0.21	31	
<i>Efficiency Index (%)</i>				<i>100</i>					<i>11</i>
<i>B. Low-uncertainty condition: $n=6$, $\alpha=270$, $\beta=730$, Expected Value=500, Range=460</i>									
1	365	61	0.79	48	626	104	0.23	24	
2	270	45	1.00	93	605	101	0.27	34	
3	270	45	1.00	138	597	99	0.29	39	
4	270	45	1.00	183	593	99	0.30	41	
5	270	45	1.00	228	591	98	0.30	42	
6	270	45	1.00	273	589	98	0.31	43	
7	270	45	1.00	318	589	98	0.31	43	
8	270	45	1.00	363	589	98	0.31	43	
9	270	45	1.00	408	588	98	0.31	44	
10	270	45	1.00	453	588	98	0.31	44	
<i>Efficiency Index (%)</i>				<i>100</i>					<i>10</i>

Note: R is total group request; r is individual (symmetric) request; p is the probability of receiving the request and continuing the game; Π is individual expected payoff from conforming to the paths described. Adoption of a conservative strategy yields an efficiency index of 91% ($25 \times 10 = 250/274$) in the high uncertainty condition, and an efficiency index of 99% ($45 \times 10 = 450/453$) in the low-uncertainty condition.

Table 2 – Number of games played by group (Gi) and length of game: High-uncertainty condition

Length	G1	G2	G3	G4	G5	G6	G7	G8	G9	G10	Total
1	11	20	19	19	12	10	20	20	16	17	164
2	6	6	5	6	6	7	9	6	5	5	61
3	5	3	4	2	7	6	1	3	5	4	40
4	3	1	2	2	2	5	0	1	2	1	19
5	1	0	0	0	0	0	0	0	0	0	1
6	2	0	0	0	1	0	0	0	0	2	5
7	0	0	0	0	1	0	0	0	2	1	4
8	2	0	0	1	1	1	0	0	0	0	5
9	0	0	0	0	0	0	0	0	0	0	0
10	0	0	0	0	0	1	0	0	0	0	1
<i>Median</i>	2	1	1	1	2	2	1	1	1	1	1
<i>Mean</i>	3	2	2	2	2	3	1	2	2	2	2
<i>SD</i>	2	1	1	1	2	2	1	1	2	2	2

Table 3 – Per period mean group requests by period and length of game: High-uncertainty condition

Length of game	Period within the game									
	1	2	3	4	5	6	7	8	9	10
1	688 (219)									
2	549 (139)	589 (174)								
3	435 (103)	442 (75)	495 (160)							
4	459 (102)	468 (76)	421 (89)	432 (77)						
5	375 (0)	375 (0)	400 (0)	325 (0)	325 (0)					
6	427 (53)	398 (73)	366 (46)	387 (76)	357 (52)	564 (362)				
7	454 (58)	417 (83)	349 (66)	378 (33)	384 (54)	366 (29)	395 (23)			
8	393 (81)	369 (107)	353 (74)	369 (46)	403 (82)	398 (84)	410 (68)	458 (132)		
10	395 (0)	385 (0)	320 (0)	337 (0)	335 (0)	309 (0)	343 (0)	282 (0)	335 (0)	277 (0)

Note: Standard deviation is in parentheses. Games of length 5 and 10 were only observed once.

Table 4 – Statistical analysis of differences between observed and predicted values

Length Games	Variable	Coefficient	Wald z Statistics	p-value	Lower 95% CI	Upper 95% CI
Length=1	$R^{Obs} - R^{SPNE}$	-13.757	-0.39	0.699	-83.461	55.947
	$R^{Obs} - R^{SO}$	537.854	15.12	0.000	468.150	607.558
	$\Delta = \frac{R^{Obs} - R^{SPNE}}{R^{Obs} - R^{SO}}$	0.026	0.38	0.706	-0.107	0.158
Length>1	$R^{Obs} - R^{SPNE}$	-214.397	-8.37	0.000	-264.601	-164.194
	$R^{Obs} - R^{SO}$	337.213	13.16	0.000	287.010	387.417
	$\Delta = \frac{R^{Obs} - R^{SPNE}}{R^{Obs} - R^{SO}}$	0.636	5.12	0.000	0.392	0.879

Table 5 – Number of games with smallest MSD from the SPNE path

Group	Nº Games	Proportion	Hypotheses	p-value	Decision
1	17	0.567	H0: p=0.5 H1: p>0.5	0.292	Do not Rej. H0
2	29	0.967	H0: p=0.5 H1: p>0.5	<0.001	Rej. H0
3	27	0.900	H0: p=0.5 H1: p>0.5	<0.001	Rej. H0
4	28	0.933	H0: p=0.5 H1: p>0.5	<0.001	Rej. H0
5	27	0.900	H0: p=0.5 H1: p>0.5	<0.001	Rej. H0
6	17	0.567	H0: p=0.5 H1: p>0.5	0.292	Do not Rej. H0
7	30	1.000	H0: p=0.5 H1: p>0.5	<0.001	Rej. H0
8	29	0.967	H0: p=0.5 H1: p>0.5	<0.001	Rej. H0
9	20	0.667	H0: p=0.5 H1: p>0.5	0.05	Rej. H0
10	21	0.700	H0: p=0.5 H1: p>0.5	0.02	Rej. H0

Note: The per period mean square deviation (MSD) of requests from predicted requests for each group in each game $k=1, \dots, 30$, is computed as $\sum_{t_k=1}^{l_k} (R_{t_k} - R_{t_k}^*)^2 / l_k$, where R_{t_k} is the observed group request in period t of game k , $R_{t_k}^*$ is the respective prediction (either at the SPNE or SO path), and l_k is the length of game k . For each group, we define a “success” as a game in which the MSD from the SPNE path is smaller than the MSD from the SO path. Let s_g represent the number of successes for each group. Under the null hypothesis that it is equally likely for either of the two predicted paths to result in the smaller MSD in any given game (H0: p=0.5), the probability of observing at least s_g successes (H1: p>0.5) in 30 games is given by $\left(\frac{1}{2}\right)^{30} \times \sum_{i=s_g}^{30} \binom{30}{i}$, and the probability of observing at most s_g successes (H1: p<0.5) in 30 games is given by $\left(\frac{1}{2}\right)^{30} \times \sum_{i=0}^{s_g} \binom{30}{i}$.

Table 6 – Number of games played by group (Gi) and length of games: Low-uncertainty condition

Length	G1	G2	G3	G4	G5	G6	G7	G8	G9	Total
1	17	13	10	15	8	17	13	19	11	123
2	4	7	8	6	8	8	8	6	5	60
3	4	3	4	0	1	1	5	1	4	23
4	3	1	1	0	2	0	3	1	2	13
5	0	1	1	4	5	3	0	1	0	15
6	1	2	1	1	1	1	1	0	2	10
7	0	1	3	0	0	0	0	0	4	8
8	0	0	0	0	1	0	0	0	0	1
9	0	0	0	4	4	0	0	1	1	10
10	1	2	2	0	0	0	0	1	1	7
<i>Median</i>	1	2	2	2	2	1	2	1	2	2
<i>Mean</i>	2	3	3	3	4	2	2	2	3	3
<i>SD</i>	2	3	3	3	3	1	1	2	3	2

Table 7 – Mean group requests by period and length of game: Low-uncertainty condition

Length of game	Period within the game									
	1	2	3	4	5	6	7	8	9	10
1	569 (186)									
2	430 (101)	506 (166)								
3	464 (112)	425 (73)	455 (89)							
4	394 (93)	398 (105)	404 (89)	447 (83)						
5	375 (114)	381 (82)	384 (86)	397 (69)	463 (102)					
6	372 (102)	386 (64)	390 (49)	398 (82)	387 (73)	397 (62)				
7	358 (65)	338 (60)	401 (74)	405 (94)	334 (56)	351 (59)	461 (85)			
8	270 (0)	420 (0)	320 (0)	500 (0)	420 (0)	370 (0)	470 (0)	620 (0)		
9	351 (64)	373 (74)	371 (56)	371 (83)	342 (57)	359 (73)	343 (54)	329 (38)	497 (105)	
10	416 (67)	369 (37)	350 (54)	433 (59)	353 (34)	405 (70)	341 (55)	389 (52)	376 (51)	449 (141)

Note: Standard deviation is in parentheses. Games of length 5 and 10 were only observed once.

Table 8 – Statistical analysis of differences between observed and predicted values

Length Games	Variable	Coefficient	Wald z Statistics	p-value	Lower 95% CI	Upper 95% CI
Length=1	$R^{Obs} - R^{SPNE}$	-18.977	-0.60	0.545	-80.464	42.510
	$R^{Obs} - R^{SO}$	299.415	9.54	0.000	237.923	360.902
	$\Delta = \frac{R^{Obs} - R^{SPNE}}{R^{Obs} - R^{SO}}$	0.063	0.57	0.569	-0.155	0.282
Length>1	$R^{Obs} - R^{SPNE}$	-176.841	-9.50	0.000	-213.315	-140.367
	$R^{Obs} - R^{SO}$	141.551	7.61	0.000	105.077	178.025
	$\Delta = \frac{R^{Obs} - R^{SPNE}}{R^{Obs} - R^{SO}}$	1.249	4.22	0.000	0.670	1.829

Table 9 – Number of games with smallest MSD from the SPNE path

Group	N° Games	Proportion	Hypotheses	p-value	Decision
1	21	0.700	H0: p=0.5 H1: p>0.5	0.02	Rej. H0
2	17	0.567	H0: p=0.5 H1: p>0.5	0.292	Do not Rej. H0
3	11	0.367	H0: p=0.5 H1: p<0.5	0.100	Do not Rej. H0
4	8	0.267	H0: p=0.5 H1: p<0.5	0.008	Rej. H0
5	4	0.133	H0: p=0.5 H1: p<0.5	<0.001	Rej. H0
6	23	0.767	H0: p=0.5 H1: p>0.5	0.003	Rej. H0
7	30	1.000	H0: p=0.5 H1: p>0.5	<0.001	Rej. H0
8	27	0.900	H0: p=0.5 H1: p>0.5	<0.001	Rej. H0
9	15	0.500	H0: p=0.5 H1: p>0.5	0.572	Do not Rej. H0

Note: The per period mean square deviation (MSD) of requests from predicted requests for each group in each game $k=1, \dots, 30$, is computed as $\sum_{t_k=1}^{l_k} (R_{t_k} - R_{t_k}^*)^2 / l_k$, where R_{t_k} is the observed group request in period t of game k , $R_{t_k}^*$ is the respective prediction (either at the SPNE or SO path), and l_k is the length of game k . For each group, we define a “success” as a game in which the MSD from the SPNE path is smaller than the MSD from the SO path. Let s_g represent the number of successes for each group. Under the null hypothesis that it is equally likely for either of the two predicted paths to result in the smaller MSD in any given game (H0: p=0.5), the probability of observing at least s_g successes (H1: p>0.5) in 30 games is given by $\left(\frac{1}{2}\right)^{30} \times \sum_{i=s_g}^{30} \binom{30}{i}$, and the probability of observing at most s_g successes (H1: p<0.5) in 30 games is given by $\left(\frac{1}{2}\right)^{30} \times \sum_{i=0}^{s_g} \binom{30}{i}$.

Table 10 – Maximum likelihood estimates of treatment effects on destruction probabilities

Variable	Coefficient	Wald z Statistics	p-value	Lower 95% CI	Upper 95% CI
<i>A. Length of Games=1</i>					
HIGH	0.1232	4.39	0.000	0.0681	0.1782
<i>B. Length of Games>1 – First Period Requests</i>					
HIGH	0.1713	6.88	0.000	0.1225	0.2201
<i>C. Length of Games>1 – Non-First Period Requests</i>					
HIGH	0.1338	8.73	0.000	0.1038	0.1638

Figure 1 – Probability of destruction: Predicted (SO, SPNE) and observed values under high-uncertainty condition


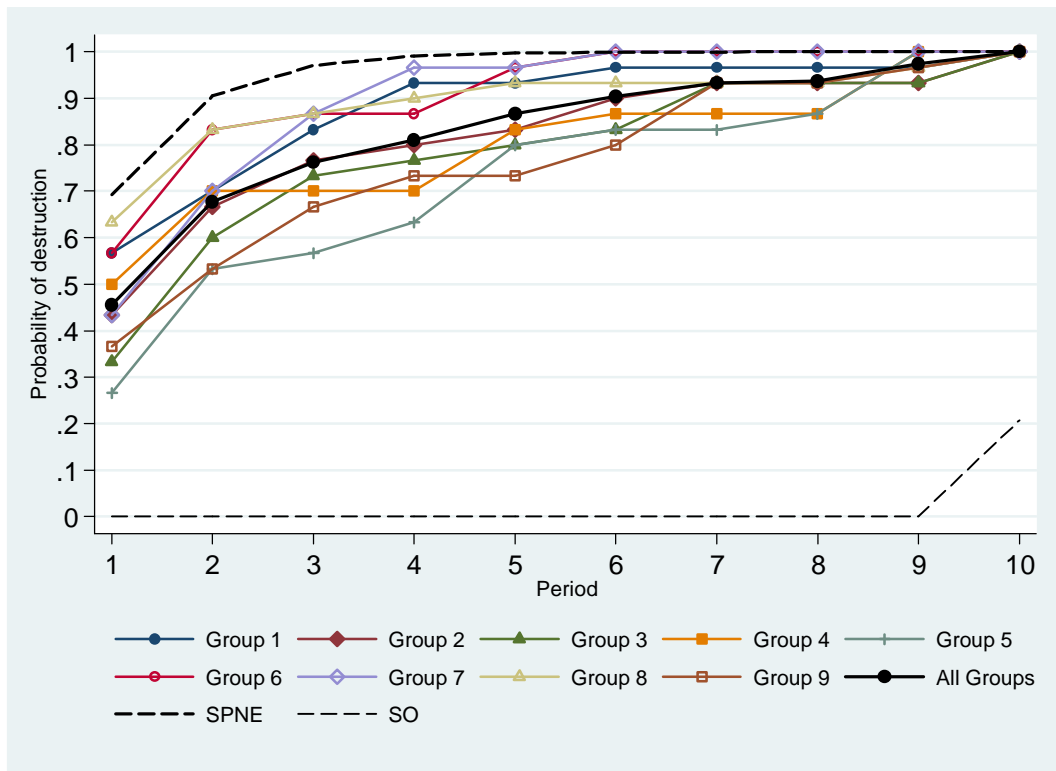


Figure 2 – Probability of destruction: Predicted (SO, SPNE) and observed values under low-uncertainty condition



Appendix - Length of games by each group in each uncertainty condition

