Games and Economic Behavior 104 (2017) 760–774

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Dilemmas, coordination and defection: How uncertain tipping points induce common pool resource destruction

Alexander Maas [a,*], Christopher Goemans [b], Dale Manning [b], Stephan Kroll [b], Thomas Brown [c]

[a] Department of Economics and Rural Sociology, University of Idaho, 875 Perimeter Dr., Moscow, ID 83844, United States
[b] Department of Agricultural and Resource Economics, Colorado State University, 1200 Center Ave. Mall, Fort Collins, CO 80523, United States
[c] United States Forest Service, United States

ARTICLE INFO

Article history:
Received 8 June 2016
Available online 4 July 2017

JEL classification:
C7
C9
D7
D8
Q2

Keywords:
Common pool resource
Uncertainty
Experimental economics
Tragedy of the Commons

ABSTRACT

Many common pool resources (CPRs) have tipping points—stock levels below which the resource is permanently damaged or destroyed—but the specific levels at which these thresholds are crossed are rarely known with certainty. We model a CPR in which uncertainty simultaneously creates a Prisoner's Dilemma and a Coordination Game. This model highlights a novel mechanism through which uncertainty incentivizes the overuse of a CPR. In the model, two Nash Equilibria exist, both of which lead to a Tragedy of the Commons, but one is an inferior solution because it leads to assured resource destruction. We use a single-period laboratory experiment to investigate the effects of uncertain tipping points on constituents' resource extraction decisions. Experimental results suggest that uncertainty reduces coordination in this type of CPR setting and increases the likelihood of resource destruction. We also find that tax and fine policies reduce consumption rates and prevent resource destruction.

1. Introduction

Allocating scarce natural resources in the absence of clearly defined property rights can be a contentious and complicated process, particularly when those resources are depletable or destructible. Consequently, common pool resources (CPRs) are a well-documented source of both political turmoil (Adams et al., 2003) and economic inefficiency (Gordon, 1954; Hardin, 1968). While traditional conceptualizations of CPRs create inefficiency by incentivizing marginal overextraction and rent dissipation, there is another feature that complicates the incentive structure of resource consumers. Specifically, many natural resources also have tipping points—stock levels at which the resource is permanently damaged or destroyed—that lead to severe personal or communal losses when they are exceeded. The specific locations of these thresholds are rarely known with certainty. For example, fishermen and scientists may agree that sufficiently depleting fish stocks will cause a fishery to collapse, but the exact level at which the population crashes is largely unknown (Botsford et al., 1997). This paper extends the existing literature on CPRs by developing a theoretical model of resource extraction when individuals share a CPR with an uncertain threshold. Natural experiments with exogenously determined threshold uncertainty shocks

* Corresponding author.
E-mail address: alexander.s.maas@gmail.com (A. Maas).

http://dx.doi.org/10.1016/j.geb.2017.06.009

and policies rarely exist. Because of this, and consistent with previous investigations (Ostrom, 2006), we use a laboratory experiment to examine the impact of uncertainty on consumption decisions.

A theoretical model highlights the existence of two Nash Equilibria in these settings: one in which marginal overconsumption dissipates rents and threatens resource stocks, and one in which individuals behave as if the resource will be destroyed with certainty because they anticipate others doing the same. Thus, individuals make one decision if they believe the resource will likely survive and a different decision if they believe it will be destroyed. For example, fishermen may harvest at levels that keep a fish stock healthy, even if the stock is below the level of maximum economic yield. This strategy reflects some amount of rent dissipation or risk of resource collapse since the stock is marginally overfished. However, as uncertainty—in both the ecosystem and the consumption rates of others—increases, fishermen may be more likely to catch as much as possible because they anticipate that fishery collapse is inevitable.

For clarity, we refer to the first Nash Equilibrium (NE) as the *partially defect NE*—where *defect* refers to a deviation from the socially optimal solution and *partially* refers to the fact that the resource may survive. The second NE will be referred to as the *fully defect NE*, because individuals do not choose the socially optimal outcome nor do they coordinate on the *partially defect NE*, and thus, end up with a lower expected payoff because the resource is destroyed with certainty. Note that both NE are inferior to the socially optimal levels of resource extraction, which will be referred to as the *SOR solution*. The presence of both NE captures an important feature of many CPR systems that has not been given significant attention in the literature.

CPR settings have traditionally been characterized as having two main features: 1) the payoff to each individual acting in her self-interest is higher than the payoff for acting in the interest of the group, and 2) all individuals receive a lower payoff when everyone acts in their own self-interest (Dawes, 1980). The model used in this experiment is consistent with Dawes's definition, but it distinguishes between two underlying causes of overextraction: 1) the classic Prisoner's Dilemma in which marginal overextraction occurs because individuals do not account for the external costs imposed on others (*partially defect NE*), and 2) a deliberate exhaustion in which a coordination problem can lead to a suboptimal NE and resource collapse (*fully defect NE*).

Experimental studies have found that uncertainty increases the instances of collective action failures in many contexts (Wit et al., 2004; McBride, 2010; Anderies et al., 2013; Aflaki, 2013), including the standard rent dissipation problem (Walker et al., 1990), probabilistic degradation or destruction (Blanco et al., 2015; Walker and Gardner, 1992), nonpoint source pollution (Poe et al., 2004), varying degrees of externalities (Suter et al., 2012), and unknown thresholds or tipping points (Budescu et al., 1990). While many of these studies conclude that increased uncertainty leads to overconsumption (Biel and Gärling, 1995; Gärling et al., 1998; Rapoport et al., 1992), they do not capture a key feature of many CPR settings—the existence of two distinct NE: one in which the resource survives and one in which it is destroyed. Notable exceptions are Barrett and Dannenberg (2012) and Dannenberg et al. (2015), who investigate constituents' contribution decisions in the context of public good provision under threshold uncertainty. Their model is particularly relevant to our setup because it creates a similar incentive structure in which both an inferior NE and superior NE exist, but neither equilibrium is located at the socially optimal solution. They conclude that public good contributions decrease with uncertainty, which increases the likelihood of ending up at the inferior NE. While their modeling framework is related to ours, previous research suggests that both the framing of a social dilemma as a CPR or Public Good (or in some other form) and the framing of the actual decision (e.g., as contribution, taking, etc.) may have implications for participants' decisions (see, for example, Van Dijk and Wilke, 2000; Apesteguia and Maier-Rigaud, 2006; Cherry et al., 2013a, 2013b).

To better understand the underlying social dilemma associated with real-world renewable CPRs, we investigate not only the effect of uncertainty on how much individuals choose to extract, but also its impact on which NE individuals choose: the *partially defect NE* or the *fully defect NE*. To accomplish this task, we develop a CPR model featuring both NE and generate data from a laboratory experiment to answer the following questions:

1) Does threshold uncertainty incentivize overextraction of a CPR relative to the *SOR solution*?
2) If uncertainty incentivizes overextraction, is it driven by marginal overuse, an increase in the probability that individuals choose to *fully defect*, or both?
3) How effective are tax and fine policies at improving efficiency in a CPR setting with threshold uncertainty?

Our experiment systematically manipulates the magnitude of threshold uncertainty in the context of a CPR and tests how these changes affect individual and group consumption decisions. As part of the experiment, individuals are randomly assigned heterogeneous payoff and damage functions, placed in four-person groups, and tasked with requesting tokens from a group account. If the sum of the group's consumption is greater than the number of tokens available, each player suffers a substantial loss. Their consumption decisions are observed under three levels of uncertainty (certainty, low uncertainty, and high uncertainty) and three policy settings (no regulation, excise taxes, and a fine).

Our results are consistent with previous findings that increased threshold uncertainty leads to increased consumption, but our contribution is to demonstrate that this phenomenon can occur because some individuals give up on preserving the resource entirely (i.e., choose the *fully defect NE*). We find that, *ceteris paribus*, threshold uncertainty has little effect on the request rates of individuals who act as if the resource will not be destroyed with certainty; however, it significantly increases the probability that individuals choose the *fully defect NE* and destroy the resource with certainty. This finding suggests that a key reason for overextraction in the presence of uncertainty may not be the omission of external costs in the individual's extraction decision, but rather a discrete shift in which each resource user defects from the better equilibrium to minimize the possible damage to herself caused by the myopic decisions of others. Our results suggest that accurately identifying resource tipping points can significantly reduce the chances that individuals deliberately exhaust a resource, and that when such scientific knowledge is unavailable, it may be valuable to introduce policies that discourage defecting.

We also consider the effectiveness of an excise tax and a lump-sum fine. While both significantly increase economic efficiency in the experiment, the tax proves to be more effective. We attribute this to the underlying mechanisms through which the policies affect behavior. Both the tax and fine increase the expected marginal cost of resource extraction; however, only the tax can eliminate the *fully defect NE*. With the fine in place individuals who believe resource destruction is certain still have the incentive to ignore the impact of their consumption on others and extract as much as possible.

This paper is composed of five sections. The next section provides context and outlines our theoretical model, Section 3 details the experiment, Section 4 reports our results, and Section 5 concludes.

2. CPR model and predictions

Our model is an extension of work first proposed in the 1980s; Messick et al. (1988) developed a simplified setup to investigate the inefficiencies created by social and threshold uncertainty in a CPR setting. This model was further developed in Budescu et al. (1990, 1995), and Aflaki (2013). While their experiments provide insights into the nature of uncertainty in a CPR, their setup eliminates the motivation of individuals to request a positive amount if they expect a threshold to be exceeded (*fully defect NE*), because the payoff for requesting zero and the payoff for requesting as much as possible are equivalent. Thus, their model creates a Coordination Game in which once a solution is reached, no single player has an incentive to marginally extract more. More recently, Botelho et al. (2014) used a similar setup in a dynamic setting and found that increased levels of uncertainty may lead to quicker depletion of a resource stock, but players may also adopt strategy paths that guarantee the threshold will not be exceeded.

Our setup uses a single-period model but allows for choices similar to those of the dynamic game. Using a single-period game allows for more observations, and can be a tool when time and money constraints are binding. Moreover, a single-period model reduces complexity for experiment participants and makes analysis more tractable—many dynamic games have an infinite (or very large) number of NE.

In our experiment, constituents make a single extraction decision. Two symmetric NE exist, one in which individuals *partially defect* from the *SOR solution* and one in which they *fully defect* and destroy the resource with certainty. The decision to *partially defect* from the *SOR solution* is akin to the defect strategy in a Prisoner's Dilemma, whereas the choice between *partially* and *fully defecting* is one of coordination. An in-depth discussion of our model and its solutions is presented below.

2.1. Basic setup

We use a single-period, non-cooperative, n-person model similar to the one developed by Budescu et al. (1990, 1995). We build upon their model in three significant ways in order to capture key features of many CPR settings: we 1) introduce a lump-sum damage amount to reflect resource collapse, 2) incorporate decreasing marginal returns to individual payoff functions, and 3) utilize a two-step optimization framework to analyze individual behavior. These modifications create significant changes in the incentive structure of resource constituents and provide insight into the possible motivations underlying overconsumption decisions in a CPR setting.

Each individual j $(j = 1, 2, \ldots, n)$ requests a specific number of tokens and receives a payoff based on the amount they request as well as the actions of others. Individual j receives $b_j(r_j)$ dollars for requesting r_j tokens. For the general theoretical model, we assume that $b_j(r_j)$ is twice differentiable and concave. If the total group request exceeds the threshold of available tokens, an individual's total payoff is reduced by D_j. The total number of tokens requested by the group is denoted $R = \sum_{j=1}^{n} r_j$, while the total number of tokens requested by other individuals is $R_{-j} = R - r_j$. The threshold under which the group must stay in order to avoid incurring the lump-sum damage is a random variable $(\tilde{X})$, uniformly distributed between a lower (α) and upper (β) bound.

Each individual's total payoff function is therefore:

$$\pi_j(r_j; R_{-j}) = \begin{cases} b_j(r_j) & \text{if } R \leq X \\ b_j(r_j) - D_j & \text{if } R > X \end{cases} \tag{1}$$

where X, absent the tilde, denotes a particular realization of the threshold.

2.2. Individual's decision problem

We assume that each of the n individuals maximizes her total expected payoff given uncertainty about the threshold and assuming the actions of others are exogenous. Note, however, that because of the piecewise nature of the uniform distribution function, the expected payoff function is kinked across three segments, or

$$E\big[\pi_j(r_j;R_{-j})\big]=\begin{cases} b_j(r_j) & \text{if } r_j\le \alpha-R_{-j} \\ b_j(r_j)-\frac{D_j(R_{-j}+r_j-\alpha)}{\beta-\alpha} & \text{if } \alpha-R_{-j}<r_j<\beta-R_{-j} \\ b_j(r_j)-D_j & \text{if } r_j\ge \beta-R_{-j} \end{cases} \tag{2}$$

Similar to Hanemann (1984) we take advantage of this discontinuity and model the individual's choice of r_j utilizing a two-stage maximization framework. In the first stage, we consider the optimal request of each individual, conditional on being constrained to one of three subsets of possible requests:

$$\begin{aligned} P_{k=1}&=\{r_j\in\mathbb{R}^+ : r_j\le\alpha-R_{-j}\} \\ P_{k=2}&=\{r_j\in\mathbb{R}^+ : \alpha<r_j+R_{-j}<\beta\} \\ P_{k=3}&=\{r_j\in\mathbb{R}^+ : r_j\ge\beta-R_{-j}\} \end{aligned} \tag{3}$$

For expositional purposes we refer to $P_{k=1}$ as the "safe zone," $P_{k=2}$ as the "uncertain zone," and $P_{k=3}$ as the "destruction zone."

In the second stage, we consider each individual's "discrete choice" of which "zone" is optimal for them to select. This two-stage approach allows us to distinguish between the underlying motivations of resource constituents. Specifically, it allows us to separately identify the standard overconsumption problem from one created by individuals' strategic behavior to exhaust the resource.

Along these lines, the individual's first-stage optimal request level, conditional on being restricted to subset P_k, is the solution to the following problem:

$$\max_{r_j} E\big[\pi_j(r_j;R_{-j})\big] \quad \text{s.t. } r_j\in P_k \tag{4}$$

Let $\bar{r}_{j,k}$ denote the optimal request conditional on being restricted to subset P_k. A solution exists for all P_k that are not empty. For $k=1$, the solution is characterized by setting the marginal benefit of an additional unit equal to zero, $b'_j(\bar{r}_{j,k=1};R_{-j})=0$, if there exists an $r_j\in P_1$ that satisfies that condition, otherwise $\bar{r}_{j,k=1}=\alpha-R_{-j}$. For $k=2$, the solution is characterized by $b'_j(\bar{r}_{j,k=2};R_{-j})=\frac{D_j}{\beta-\alpha}$ if such an $\bar{r}_{j,k=2}$ exists.[1] For $k=3$, the solution is characterized by $b'_j(\bar{r}_{j,k=3};R_{-j})=0$.

The second stage of the optimization problem for each individual involves deciding which zone to choose. Denote $\overline{\pi}_{j,k}(R_{-j})$ as the maximum level of profit the individual can achieve conditional on the requests of others and being restricted to subset P_k; then, the individual's second-stage problem can be written as:

$$\max_{\gamma_{j,k}} \sum_k \gamma_{j,k}\overline{\pi}_{j,k}(R_{-j}) \tag{5}$$

where $\gamma_{j,k}\in\{0,1\}$ and $\sum_k\gamma_{j,k}=1$.[2] $\gamma_{j,k}$ represents the discrete choice made by each individual.

The solution to this problem is a set of $\gamma_{j,\tilde{k}}$ such that:

$$\gamma_{j,\tilde{k}}=\begin{cases} 1 & \text{if } \overline{\pi}_{j,\tilde{k}}(R_{-j})\ge\overline{\pi}_{j,k}(R_{-j})\ \forall k\ne\tilde{k} \\ 0 & \text{otherwise} \end{cases} \tag{6}$$

A NE exists if every player cannot improve their payoff unilaterally by changing r_j or $\gamma_{j,\tilde{k}}$. The model is parameterized for the experiment such that two such NE exist in each uncertainty (and certainty) treatment.

2.3. Social planner's decision problem

The social planner's problem is similar to the individual's problem in that she faces the same threshold uncertainty, represented by $\tilde{X}\sim U(\alpha,\beta)$. In calculating her solution, however, she accounts for the total damage incurred by the group, while the individual is only concerned with the damage she incurs. Thus, the social planner's problem is to choose a set of request levels that maximize the total expected payoff, which is defined as:

$$E\big[\pi(r_{1,2\ldots n})\big]=\begin{cases} \sum_{j=1}^n b_j(r_j) & \text{if } R\le\alpha \\ \sum_{j=1}^n b_j(r_j)-\frac{(\sum_{j=1}^n D_j)(R-\alpha)}{\beta-\alpha} & \text{if } \alpha<R<\beta \\ \sum_{j=1}^n b_j(r_j)-\sum_{j=1}^n D_j & \text{if } R\ge\beta \end{cases} \tag{7}$$

[1] $\frac{D_j}{\beta-\alpha}$ represents the expected marginal cost associated with increasing r_j between α and β.

[2] Note that $\sum_k\gamma_k$ can be greater than one if two $\overline{\pi}_k$ are equal. However, we ignore this special case given that it is not a possibility for the parameters/functions chosen for this experiment.

The social planner's continuous solution, conditional on discrete choices, is as follows: For $k=1$, the social planner's solution is characterized by setting the marginal benefit of an additional unit equal to zero for all users, $b_j'(\bar{r}^{sp}_{j,k=1})=0\ \forall j$, if there exists an $r_{1,\dots,n}\in P_1$ that satisfies that condition, otherwise $b_j'(\bar{r}^{sp}_{j,k=1})=b'_{-j}(\bar{r}^{sp}_{-j,k=1})\ \forall j$ and $R=\alpha$. For $k=2$, the solution is characterized by $b_j'(\bar{r}^{sp}_{j,k=2})=b_{-j}'(\bar{r}^{sp}_{-j,k=2})=\frac{\sum_{i=1}^{n}D_i}{\beta-\alpha}\ \forall j$. For $k=3$, the solution is characterized by $b_j'(\bar{r}^{sp}_{j,k=3})=0\ \forall j$. It is possible to show that the social planner's optimal total request level, R^{sp}, will never be larger than the sum of individual requests at either the *partially defect NE* or *fully defect NE*. While this result is generally true, the proceeding exposition focuses only on the cases relevant to our experiment.

The second-stage optimization problem and its solution are similar to above, such that, defining the social planner payoff for discrete choice, $\tilde{k}$ as $\overline{\pi}^{sp}_{\tilde{k}}$, the discrete choice social planner solution is characterized by:

$$\gamma^{sp}_{\tilde{k}}=\begin{cases}1 & \text{if } \overline{\pi}^{sp}_{\tilde{k}}\geq\overline{\pi}^{sp}_{k}\ \forall k\neq\tilde{k}\\ 0 & \text{otherwise}\end{cases}\tag{8}$$

2.4. Individual and social planner solutions under low and high uncertainty

Many possible cases exist in the general formulation of this problem, and it can be shown that the effect of uncertainty under the general formulation is ambiguous. We focus on specific cases that possess characteristics consistent with CPR settings. These cases have two NE and are defined by the following features: 1) there exists one symmetric NE in the uncertain zone (the *partially defect NE*) 2), there exists one symmetric NE in the destruction zone (the *fully defect NE*), 3) both NE have lower expected payoffs than the social planner's solution,[3] 4) the *fully defect NE* is socially inferior to the *partially defect NE*, and 5) conditional on others *fully defecting*, no individual who *fully defects* would unilaterally change to *partial defect*, and conditional on others *partially defecting*, no individual who *partially defects* would unilaterally change to *fully defect*. Within this subset of the general problem, we identify one case where the *SOR solution* is characterized by $\gamma^{sp}_1=1$ such that the continuous choice solution is characterized by $b_j'(\bar{r}^{sp}_{j,k=1})=b_{-j}'(\bar{r}^{sp}_{-j,k=1})\ \forall j$ and $R^{sp}=\alpha$ (the low uncertainty treatment), and one case where $\gamma^{sp}_2=1$ such that the optimal continuous-choice solution is characterized by $b_j'(\bar{r}^{sp}_{j,k=2})=b_j'(\bar{r}^{sp}_{-j,k=2})=\frac{\sum_{i=1}^{n}D_i}{\beta-\alpha}\ \forall j$ and $\alpha<R^{sp}\leq\beta$ (the high uncertainty treatment). These results imply that the social planner will not accept any probability of destroying the resource in the low uncertainty case, but will accept some risk in the high uncertainty case.

These cases were chosen because they correspond to many real-world CPR settings and allow for meaningful analysis. For example, it is reasonable that the social planner would ensure resource survival, or that she is willing to accept some probability of destruction. However, if the socially optimal solution is to exhaust the resource ($\gamma^{sp}_3=1$), there would be no social dilemma. We also acknowledge that under the general set-up, it is possible that no NE exist, but these cases would not allow for meaningful analysis. The cases of interest also implicitly assume that individuals have the ability to cross the resource threshold (β is sufficiently small) and that resource survival creates more value than destruction—both assumptions are consistent with most CPR settings. Limiting our analysis to only these cases of interest, it is possible to investigate how uncertainty affects both the discrete and continuous choice, and how the NE compare to the *SOR solution*. First, we examine the effect of uncertainty on token requests.

Because this analysis is limited to cases where two NE exist, we cannot theoretically identify the effect of increasing uncertainty on individuals' discrete choices. The *partially defect NE* is superior to the *fully defect NE* in every treatment, but the nature of a Coordination Game does not allow us to determine which NE will be reached *ex ante*. Therefore, picking one discrete-choice NE over the other is a behavioral question and is tested using data from a laboratory experiment. However, this ambiguity does not exist for the continuous choice given a particular discrete choice. For example, conditional on ($\gamma_{j,2}=1$), individuals will request such that $b_j'(\bar{r}^{NE}_{j,k=2})=\frac{D_j}{\beta-\alpha}$. Because damage is constant, increasing $\beta-\alpha$ decreases the right-hand side of the equation. Given the concavity of $b_j(r_j)$, requests must increase to satisfy the equality. Thus, conditional on the discrete choice to *partially defect*, increasing uncertainty increases consumption. These theoretical results are empirically tested across the low and high uncertainty treatments.

2.5. Comparing the individual and social planner solutions

It is also possible to compare individual NE solutions to that of the social planner and to show that, for the cases of interest, individuals have an incentive to overconsume. To compare the social planner and individual decisions, note that if both solutions are in the uncertain zone ($\gamma_{j,2}=1,\gamma^{sp}_2=1$) an individual's request solution is characterized by $b_j'(\bar{r}^{NE}_{j,k=2})=\frac{D_j}{\beta-\alpha}$, whereas the social planner's solution is characterized by $b_j'(\bar{r}^{sp}_{j,k=2})=\frac{\sum_j D_j}{\beta-\alpha}$. This implies that, $b_j'(\bar{r}^{sp}_{j,k=2})>b_j'(\bar{r}^{NE}_{j,k=2})$. Given the concavity of the payoff function, it follows that $\bar{r}^{NE}_{j,k=2}>\bar{r}^{sp}_{j,k=2}$ and a social dilemma exists if individuals pick the

[3] Note that under the certainty treatment, the *SOR solution* is one of the possible Nash Equilibria.

partially defect NE. Moreover, this disparity in perceived marginal cost between the individual and social planner makes it possible that the social planner solution is to remain in the safe zone $\gamma_1^{sp} = 1$ when the NE discrete solution is $\gamma_{j,2} = 2$, which will be the case if $\overline{\pi}_{j,2}(R_{-j}^{sp}) \geq \overline{\pi}_{j,1}(R_{-j}^{sp})$ for any individual. These incentives create a Prisoner's Dilemma for the continuous choice and a Coordination Game when choosing the NE (the discrete choice).

2.6. Designing optimal tax and fine policies

Given the discrepancy between the individual NE and the *SOR solution*, policies may enhance the value derived from the resource. We consider two potential policies, including both a constant tax per unit requested as well as a fee paid by all if R exceeds the realized threshold, X. The optimal policy to address this CPR dilemma requires two things. First, it must induce individuals to pick the correct discrete choice, $\gamma_{j,k}$, and second, given that choice, it must also induce all individuals to request the same number of tokens as the social planner, $\overline{r}_{j,k}^{NE} = \overline{r}_{j,k}^{sp}$ $\forall j$. If the social planner's discrete solution is $\gamma_1^{sp} = 1$, then the solution is characterized by $R^{sp} = \alpha$ and $b'_j(\overline{r}_{j,k=1}) = b'_{-j}(\overline{r}_{-j,k=1})$ $\forall j$. To ensure a solution, a policy must dictate $\overline{\pi}_{j,1} > \overline{\pi}_{j,2}, \overline{\pi}_{j,3}$ for all players in addition to ensuring that $\overline{r}_{j,k=1}^{NE} = \overline{r}_{j,k=1}^{sp}$ $\forall j$. If instead, the social planner's solution is $\gamma_2^{sp} = 1$, then the solution is characterized by $\alpha < R^{sp} < \beta$ and policy must induce the solution, $b'_j(\overline{r}_j) = \frac{\sum_{i=1}^n D_j}{\beta-\alpha}$ and $\overline{\pi}_{j,2} > \overline{\pi}_{j,1}, \overline{\pi}_{j,3}$ for all players. We ignore the case where the social planner's solution is γ_3, since this solution would denote the absence of any social dilemma. Importantly, while a tax is charged per unit extracted, the fine is a lump-sum charge that only occurs if the threshold is exceeded.

The optimal tax in this setup depends on the specific equilibrium locations. In the simplest case, where the individual and the social planner problems both have solutions in the uncertain zone (the high uncertainty treatment), taxes are set such that $tax = [\frac{\sum_{j=1}^n D_{-j}}{\beta-\alpha}]$. When the social planner solution is below the uncertain zone (the low uncertainty treatment), the tax must be set differently since the social planner solution in this case is $b'_j(\overline{r}_{j,k=1}^{sp}) < \frac{\sum D_j}{\beta-\alpha}$. If the tax were set at the previous level, $[\frac{\sum_{j=1}^n D_{-j}}{\beta-\alpha}]$, individuals would be over taxed. In this case, it is optimal to set the tax such that $tax = b'_j(\overline{r}_{j,k=1}^{sp}) - \frac{D_j}{\beta-\alpha}$. In either case, an excise tax increases the marginal cost of extraction regardless of discrete choice, and it therefore effectively eliminates the *fully defect NE* and moves the *partially defect NE* to coincide with the *SOR solution* under both uncertainty treatments. The fine also shifts the *partially defect NE* to coincide with the *SOR solution*; however, it does not eliminate the *fully defect NE* since it only increases expected marginal value in the uncertain zone. Theoretically, the fine is equivalent to increasing the size of D_j to $\sum_{j=1}^n D_j$, and thus is set as $fine = \sum_{j=1}^n D_{-j}$. Because taxes increase marginal cost regardless of the discrete choice, they should lead to a reduction in total consumption. Fines, by comparison, only affect marginal cost in the uncertain zone, and may not induce efficient discrete choices, and thus, may not decrease consumption.

3. Experimental setup

To test these theoretical predictions, an experiment was used in which policies and resource threshold uncertainty could be exogenously introduced and varied. The experiment was conducted at Colorado State University with 96 participants who received an average payment of $27.[4] It was conducted in the spring and fall semesters of 2015, using z-Tree (Fischbacher, 2007). The experiment consisted of six sessions—one session of 20 students, four with 16, and one with 12. Just over half of these participants (52%) were female, their average age was approximately 20 years, and the most common fields of study were "Other" (28%), Agricultural Sciences (24%), and Economics (13.5%). A full list of participant summary statistics is presented in Appendix A.

3.1. Experimental design and procedures

Prior to each session, participants were provided a set of instructions and asked to follow along as instructions were read aloud.[5] After instructions were completed, students were given a questionnaire to gauge and clarify their understanding of the experiment. At the beginning of the experiment, individuals were provided with a table showing their actual payoff per token requested as well as the damage amount they would incur if the sum of token requests surpassed the threshold. Half of the players in each group were assigned a low-productivity payoff schedule, in which the payoff per token was less than that of high-productivity participants (parameterization described in the following subsection). Individuals did not know others' true payoffs, only that they were heterogeneous and that individuals who receive higher payoffs per token also suffer larger damages when the threshold is exceeded. Heterogeneous payoff and damage functions were used to represent real-world settings in which resource users have differing levels of productivity (see Appendix C for a discussion of implications this experimental design feature may have for our results). Also consistent with many real-world settings,

[4] To avoid income effects, we randomly chose one round in each session to be payoff-relevant. A current debate about the appropriateness of "pay one" or "pay all" approaches is summarized in Charness et al. (2016).

[5] See Appendix B for a complete exposition of the instructions seen by experiment participants.

Table 1
Uncertainty range by treatment.

	α	β	Uncertainty range
Certainty	18	18	0
Low uncertainty	12	24	12
High uncertainty	8	28	20

Table 2
Payoff schedule and damage amount for each user.

Marginal payoff per token											
Tokens requested	0	1	2	3	4	5	6	7	8	9	10
Low productivity	.	\$11.00	\$7.50	\$4.50	\$2.80	\$1.75	\$0.90	\$0.10	−\$0.50	−\$1.20	−\$1.75
High productivity	.	\$13.00	\$9.50	\$6.50	\$4.80	\$3.75	\$2.90	\$2.10	\$1.40	\$0.80	\$0.25
Damage	Low productivity: \$16					High productivity: \$30					

we did not reveal the requests of other individuals in a group, and instead elected to only inform individuals if their group stayed below or exceeded the threshold.

In each session, individuals participated in 12 rounds made up of 8 static periods.[6] The 9 rounds of interest in this analysis differ by policy and uncertainty treatment, while the 8 periods within a round remained identical to one another except in group composition. We employed a "stranger" design (Andreoni, 1988), with groups of four (two low-productivity and two high-productivity) randomly changing their composition in each period.

At the start of each round, participants were informed about the level of uncertainty and the policy treatment for that round. Then, players independently requested specific numbers of tokens. After each player made a request, a payoff screen was displayed with key information from the period. This screen informed participants of how much they requested, if the group exceeded the threshold, the total money paid for potential damage and policy payments, and their net payoff for the period.

By introducing three uncertainty treatment levels and three policies, we can compare the instances of defecting under each permutation of policy and uncertainty level. The ordering of each policy and uncertainty treatment varied across sessions so that it would be possible to test for treatment ordering effects, which do not appear to be significant in our results.[7]

This experimental design allows us to systematically examine the relationship between threshold uncertainty and resource constituents' consumption and equilibrium choice. Note that the experiment was parameterized to have an easily identifiable *fully defect NE* token choice for each participant. The *fully defect NE* corresponds to the token amount at which marginal benefit equals zero on a participant's payoff table.

3.2. Model parameterization for the experiment

For the experiment, a group size of four is used to allow for symmetry and to create group dynamics that may not be present in a bilateral game. D_j and $b_j(r_j)$ are unchanged for each participant across each treatment to help with understanding. Thus, the only parameters that change across uncertainty treatments are those associated with the uncertainty range (α and β), although $E[\tilde{X}]$ remains constant throughout the experiment (18). α and β were chosen to reflect three uncertainty treatments: certainty, low uncertainty, and high uncertainty. The specific uncertainty parameterization of each treatment is presented in Table 1.

The payoff functions for the experiment participants were parameterized to represent high- and low-productivity resource extractors. Table 2 presents the payoffs associated with different levels of requests made by each user type. Notice that the high-productivity user has a higher payoff associated with each level of request and faces a larger damage penalty if the group exceeds the resource threshold, reflecting larger losses to high-productivity users from resource collapse.

The model was parameterized such that both the *partially* and *fully defect NE* exist in both uncertainty treatments. The individual *partially defect NE* is in the uncertainty zone under both low and high uncertainty, and the payoff from that equilibrium is always higher than that of the *fully defect NE*. Fig. 1, which will be explained in more detail in the next section, illustrates the expected payoff of high-productivity users, conditional on others' request rates under each uncertainty treatment, and the numeric solutions to every treatment is presented in Appendix C. Note that while the token request rates are continuous, the NE choice is discrete. Also note that the individually optimal request rate depends on the

[6] While students participated in 12 rounds, only 9 are used in this paper. The rounds with the Certainty-varying treatment were omitted from this analysis because they were used to answer a different question that will be analyzed in future work. The omitted rounds had different realizations of the threshold which were known to the participants and therefore cannot be compared to our uncertainty treatments. For example, in the context of this paper, there is not a meaningful way to compare a known threshold of 8 to the expected threshold of 18 under our uncertainty treatment.

[7] A schematic of the experiment, the order of treatments and a series of tests for ordering effects are presented in Appendix D.

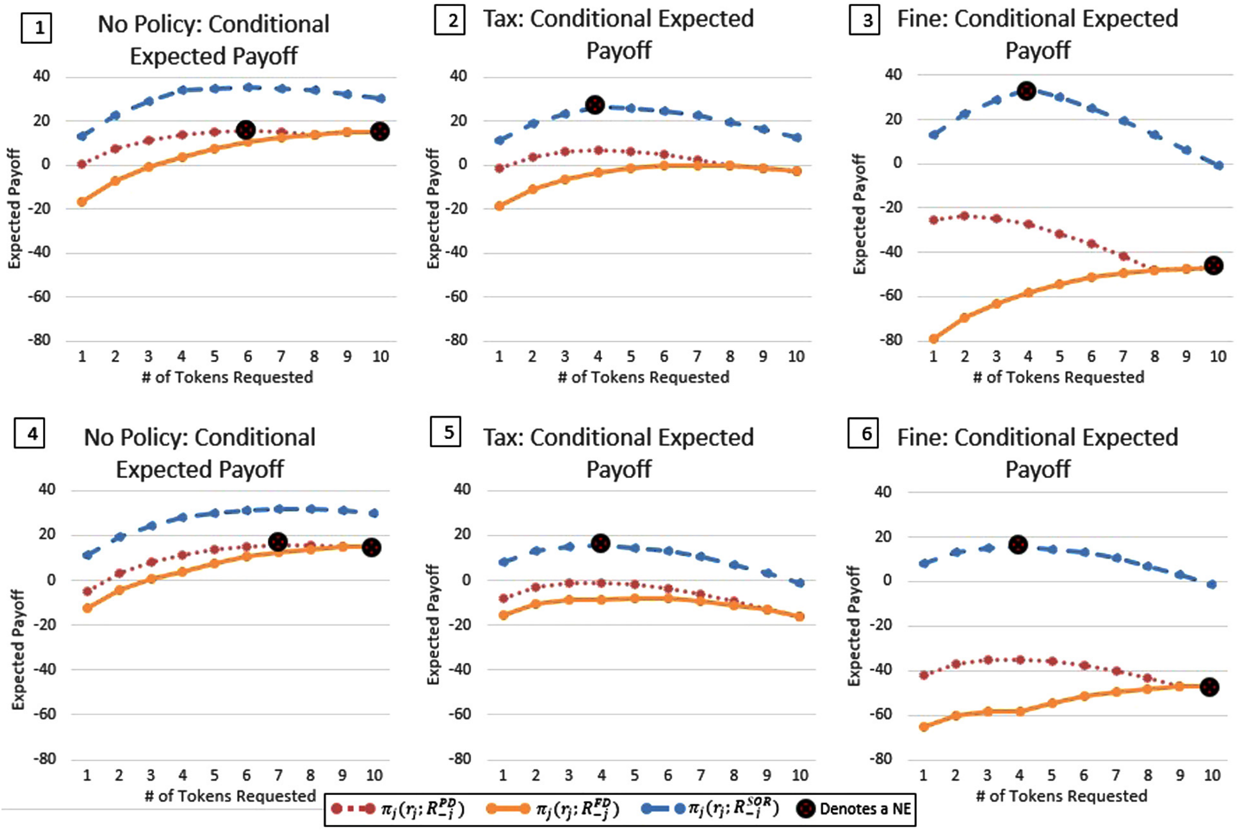

Fig. 1. Total expected payoff of high-productivity user under low and high uncertainty across policy treatments.

choices of others, and there may be a significant loss if an individual requests consistent with a *partially defect NE* while others in the group choose the *fully defect NE*.

3.3. Policy parameterization

Each of the three uncertainty treatments are investigated under two policies: an excise tax on extraction and a fine for exceeding the threshold. The optimal tax in the high uncertainty treatment is set such that $tax = [\frac{\sum_{j=1}^{n} D_{-j}}{\beta - \alpha}]$, or \$4.67 for low-productivity users and \$1.78 for high-productivity users. In the low uncertainty treatment, the tax is set to $tax = b'_j(\bar{r}^{sp}_{j,k=1}) - \frac{D_j}{\beta-\alpha}$, or \$3.80 for low-productivity users and \$3.10 for high-productivity users. The fine is set as the total damage others would incur if the threshold is exceeded. Numerically, this results in a \$76 fine for low-productivity users and a \$62 fine for high-productivity users. Under this parameterization, conditional on other group members *fully defecting* and pushing the group into the destruction zone, no single individual can unilaterally improve their expected payoffs by decreasing requests. In fact, if an individual chooses to *partially defect* when others in the group choose to *fully defect*, that individual would suffer substantial losses because she would incur both the damage and the fine without gaining the extra benefit from extracting $\bar{r}^{NE}_{j,k=3}$.

To visualize the decisions of resource constituents, Fig. 1 graphically illustrates the key features and solutions of the model. Each panel illustrates the expected payoff, $\pi_j(r_j; R_{-j})$, of a high-productivity user conditional on three different request combinations of others in the group: 1) others choose the requests consistent with the *partially defect NE* (labeled R^{PD}_{-j}), 2) others choose requests consistent with the *fully defect NE* (labeled R^{FD}_{-j}), or 3) others choose requests consistent with the *SOR solution* (labeled R^{SOR}_{-j}). Row one presents these payoffs for the low uncertainty treatment and row 2 presents them for high uncertainty. Requesting additional tokens along the same curve can be thought of as the continuous choice whereas moving from one curve to the other represents a shift in all other users' discrete strategies. For example, in Panel 1, the optimal choice, given that everyone else requests consistent with the *partially defect NE*, is to also *partially defect* and request 6 tokens (\$15.45). Accordingly, if everyone else *fully defects*, then expected payoff is highest when the individual also *fully defects* (\$15.00). If other users instead request tokens consistent with the *SOR solution*, the optimal solution for the individual is still to *partially defect*. This incentive reflects the Prisoner's Dilemma aspect of the CPR; because at the *SOR solution*, each user has an incentive to extract more, thus moving the group from $\pi_j(r_j; R^{SOR}_{-j})$ to $\pi_j(r_j; R^{PD}_{-j})$.

Table 3
Mean token requests by treatment and productivity type.

Productivity	No policy			Tax			Fine		
	Low	High	Difference (high-low)	Low	High	Difference (high-low)	Low	High	Difference (high-low)
Certainty	4.49	4.75	0.26**	3.44	4.32	0.88***	4.12	4.12	0.00
	(0.99)	(1.75)	(0.10)	(0.93)	(1.05)	(0.72)	(0.93)	(1.02)	(0.07)
Low uncertainty	4.64	5.51	0.87***	2.65	4.58	1.93***	3.58	3.71	0.13
	(1.57)	(2.67)	(0.16)	(1.17)	(1.65)	(0.10)	(1.18)	(1.24)	(0.09)
High uncertainty	4.78	5.42	0.64***	2.96	3.47	0.51***	3.51	3.37	−0.14
	(1.83)	(2.74)	(0.17)	(1.00)	(1.54)	(0.09)	(1.34)	(1.73)	(0.11)

Standard deviation is reported in parentheses.
$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$ denotes significance level of mean request difference across types as determined by a t-test.

In Column 1 of Fig. 1 (Panels 1 and 4), we see that without a policy, both NE exist and the *partially defect NE* is slightly superior to the *fully defect NE*. However, both NE are still inferior to the social planner's solution. In Column 2 (Panels 2 and 5), we observe that the introduction of a tax eliminates the *fully defect NE*, such that the only NE that exists is one consistent with the *SOR solution*. Column three (Panels 3 and 6) presents expected payoff functions when a fine is in place. There are a number of interesting features for this treatment. Note that the *fully defect NE* is not eliminated, even though it has substantially lower payoffs than the *partially defect NE*. While the fine does not eliminate the *fully defect NE*, it increases the expected damage such that the *partially defect* request rate aligns with the *SOR solution*. Thus, conditionally on *partially defecting*, the fine should induce optimal request rates among resource constituents.

Fig. 1 also makes clear that there is a cost of conditioning one's choice on the incorrect assumption about what others are choosing. For example, in Panel 4, if the individual believes other group members will *partially defect*, she would request 7 tokens to maximize expected payoff at \$15.55. However, if the group actually *fully defects* her expected payoff would drop to \$12.55; the individual would have been better off *fully defecting* (requesting 10) with an expected payoff of \$15.00—less than the *partially defect NE* payoff, but higher than individually picking that strategy when everyone else chooses to *fully defect*. This incentive reflects the Coordination Game imbedded in the model. In the extreme case, the presence of a fine under low uncertainty, an individual may choose the *SOR solution* (now also the *partially defect* NE) anticipating that others to do the same. If others choose this NE, then the expected payoff is over \$33.80. However, if other players choose the *fully defect NE*, the payoff for the player choosing the *partially defect NE* decreases by approximately \$90 to −\$58.20. Thus, an individual's choice will be influenced by her belief about the group's choice of NE, which may be affected by the level of threshold uncertainty.

4. Experimental results

Table 3 summarizes the total number of tokens requested across each of the uncertainty and policy treatments for low and high-productivity users. Consistent with theoretical predictions, and providing evidence that participants understood the experiment, low-productivity users requested fewer tokens than high-productivity users and total requests fell when the tax and fine policies were in place. The effect of the policies appears to be larger when the resource threshold is uncertain.

To summarize the prevalence of threshold exceedance in the experiment, Fig. 2 illustrates the proportion of periods in which the threshold was exceeded (Panel A) and the proportion of individual requests consistent with the *fully defect NE* (Panel B). For this analysis, the *fully defect NE* choice is defined with a binary variable that takes the value of 1 if a participant requests tokens such that $b'_j(r_j) = 0$, and 0 otherwise.[8]

Without a policy, the threshold was exceeded 37.5% of the time when the threshold was known with certainty, 58.9% in the low uncertainty treatment, and 52.1% in the high uncertainty treatment (Fig. 2, Panel A). Further, without policies individuals chose to *fully defect* 5.6% of the time under the certainty treatment, 16.0% of the time under the low uncertainty treatment, and 15.5% of the time under the high uncertainty treatment (Fig. 2, Panel B).

By comparison, when a tax or fine is introduced, the instances of exceeding the threshold are almost entirely eliminated in the certainty treatment, significantly reduced in the low uncertainty treatment, and brought in-line with the *SOR solution* in the high uncertainty treatment.[9] With policies, we also observe a substantial decrease in the instances of *fully defecting* in every uncertainty treatment (Fig. 2, Panel B).

Throughout the remainder of this section we turn to a conditional analysis to explore if uncertainty increases token requests, and if so, we disentangle the effect of uncertainty on marginal increases in requests from the effect of a change in the probability of choosing to *fully defect*. The dependent variable in each of the models estimated, y_{jtcs}, is either the number of tokens requested or a binary variable meant to represent the choice to *fully defect* by individual j in period t, round c, and session s of the experiment. We model the number of tokens or the discrete choice as a function of several

[8] We also defined *fully defect* as any request within one token of the actual *fully defect* request rate and ran the same analysis presented below with no qualitative change in results. Results from these models can be provided upon request.

[9] Note that there is a 0.25 probability of resource destruction at the *SOR solution* in the high uncertainty treatment.

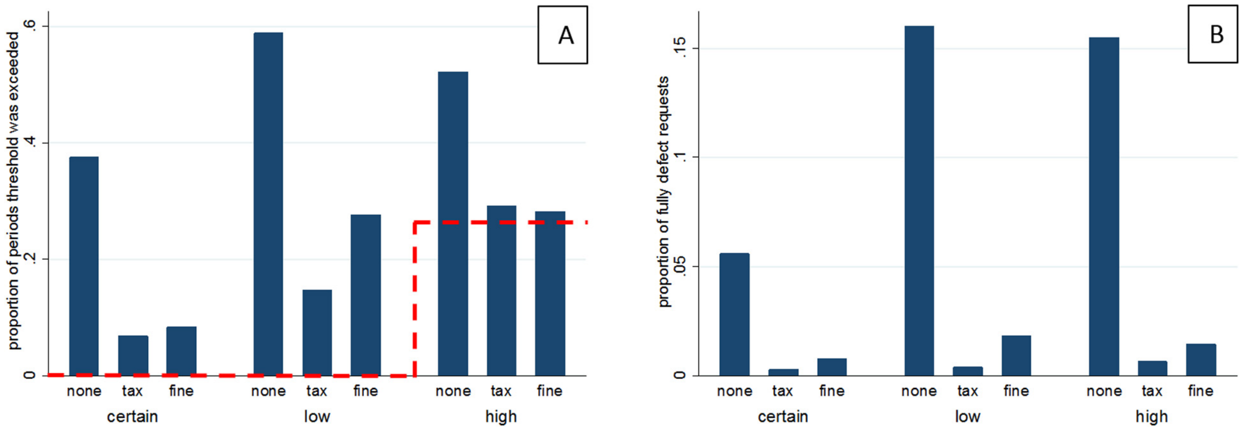

Fig. 2. Panel A illustrates the proportion of periods where the threshold was exceeded under each policy and uncertainty scenario. The dotted line in Panel A denotes the socially optimal expected number of threshold exceedances. Panel B illustrates the proportion of individual requests consistent with the *fully defect NE* under each policy and uncertainty scenario. A dotted line is not included in Panel B, because it is never socially optimal to fully defect.

independent variables that are specific to model periods but apply to all individuals within a period. First, x_{tcs} is a vector of dummy variables for each policy treatment observed in period t, round c, and session s of the experiment. Tax and fine dummy variables take a value of 1 if that policy is in place and 0 otherwise. w_{tcs} is a vector of dummy variables indicating the uncertainty treatment, ρ_{tcs} is a vector of policy-uncertainty interaction dummies and z_j is a dummy variable for risk aversion where $z_j = 1$ if an individual is risk averse and $z_j = 0$ otherwise.[10] If ϵ_{jtcs} is an idiosyncratic error term, our econometric models can be written as:

$$y_{jtcs} = f\left(\beta_{0j} + x'_{tcs}\beta_1 + w'_{tcs}\beta_2 + \rho'_{tcs}\beta_3 + z'_j\beta_4 + \epsilon_{jtcs}\right) \tag{9}$$

The coefficient vectors of interest include β_1, β_2, and β_3, which represent the effects of policy and uncertainty treatments on the dependent variable (token request of *fully defecting*). β_{0j} is a potentially individual-specific intercept. Equation (9) is used to estimate the effect of each treatment on both the discrete and continuous choices. A variety of alternative specifications are presented throughout the results section. In all cases, fixed and random effects are included at the subject level. Following Cherry et al. (2013a, 2013b), we present standard errors clustered at the subject level.[11]

4.1. The effect of uncertainty on token requests and the probability that participants choose the fully defect NE

Equation (9) is used to estimate the effect of introducing uncertainty on the number of tokens requested by each individual (OLS, both pooled and with individual fixed effects). Table 4 presents the results of Equation (9) for experimental periods when no policies are in place and y_{jtcs} is defined as the total number of tokens requested by individual j in period t, round c, and session s of the experiment. In Columns 2 and 4, only observations for subjects that did not choose the *fully defect NE* are included in the sample.

Consistent with past research (Budescu et al., 1995; Rapoport et al., 1992; Walker and Gardner, 1992), we find that the presence of uncertainty has a positive and statistically significant effect on the number of tokens requested. However, the effect of uncertainty on token requests is not statistically different from zero when focusing only on those participants who chose not to *fully defect*.

We now utilize Equation (9) to estimate the effect of uncertainty on the probability that individuals choose to *fully defect*. Table 5 presents coefficient estimates using only periods when no policies are in place and y_{jtcs} is defined as being equal to 1 if individual j in period t, round c, and session s of the experiment chose to *fully defect*. If the individual does not *fully defect*, $y_{jtcs} = 0$. Three alternative specifications were estimated to test for robustness: logit, logit random-effects (logit RE), and OLS with individual fixed effects (FE). The logit random-effects model (Column 2) is the preferred model for estimating the impacts of policy and uncertainty because it utilizes the panel structure of the data while constraining predicted probabilities to be between zero and one.

Positive and significant coefficients on both uncertainty treatments indicate that the probability of *fully defecting* is significantly increased by the presence of threshold uncertainty. However, there is no significant difference in effect across low

[10] Risk preferences were included as a dummy variable because the marginal effect of a Likert scale is often problematic (Jamieson, 2004). However, qualitative results remain very similar when risk preferences are included in the regressions as a continuous variable and can be provided upon request.

[11] The results are robust across a wide range of approaches including calculating the standard errors using White's estimator and when clustering at the session level. The later includes, for the Fixed Effects and Pooled models, the bootstrapping approach developed by Esarey and Menger (2015) for samples with a small number of clusters.

Table 4
Token requests without policies.

Variables	(1) All participants (Pooled)	(2) Partially defect only (Pooled)	(3) All participants (FE)	(4) Partially defect only (FE)
Low uncertainty	0.456***	−0.0173	0.456***	0.0560
	(0.166)	(0.111)	(0.166)	(0.116)
High uncertainty	0.480***	0.0975	0.480***	0.129
	(0.169)	(0.147)	(0.169)	(0.143)
Risk aversion	−0.298	−0.309		
	(0.327)	(0.250)		
Constant	4.833***	4.581***	4.618***	4.327***
	(0.247)	(0.199)	(0.101)	(0.0740)
Observations	2,304	2,019	2,304	2,019
# of participants	96	96	96	96
R-squared	0.016	0.010	0.021	0.002

Standard errors (in parentheses) are clustered at the subject level $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table 5
Decision to fully defect.

Variables	(1) Logit	(2) Logit RE	(3) FE
Low uncertainty	0.124***	0.118***	0.104***
	(0.0356)	(0.0330)	(0.0278)
High uncertainty	0.120***	0.113***	0.0990***
	(0.0321)	(0.0296)	(0.0238)
Risk aversion	0.0180	0.00627	
	(0.0387)	(0.0361)	
Constant			0.0560***
			(0.0152)
Observations	2,304	2,304	2,304
# of participants	96	96	96
R-squared			.02

Standard errors (in parentheses) are clustered at the subject level.
Logit models are reported as average marginal effects.
$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

uncertainty and high uncertainty.[12] This suggests that moving from a world of certainty to one in which uncertainty exists around the resource threshold increases the chances of *fully defecting*, but the probability of *fully defecting* does not increase further with greater uncertainty.

Interestingly, risk aversion appears to have little effect on the number of tokens requested (Table 4, Columns 1 and 2) or the decision to *fully defect* (Table 5, Columns 1 and 2), even though theory suggests that extreme risk aversion should induce individuals to *fully defect* (see Appendix E for analysis of the impact of risk aversion). In general, extremely risk averse participants may choose to request fewer tokens to decrease the probability of incurring a loss, but they may also *fully defect* in order to prevent the severe personal loss that would occur if they *partially defect* but the resource is destroyed. When the *partially defect NE* creates an outcome in which the resource has some probability of destruction, resource constituents with strong risk aversion may intentionally destroy the resource to eliminate threshold uncertainty as well as the uncertainty around others' decisions to *fully defect*.

Taken together, the results presented in Tables 4 and 5 suggest that the increase in requests associated with the introduction of an uncertain threshold is largely driven by an increase in the number of participants who decide to *fully defect*. This result can be seen through statistically significant impacts of uncertainty on the probability of *fully defecting* but no impact of uncertainty on the number of tokens requested, conditional on not *fully defecting*. Making this distinction is important because it highlights a mechanism through which uncertainty increases the use of a CPR. When examining only the average effect of uncertainty, this mechanism remains hidden. The increase in average request under uncertainty obscures the underlying mechanism behind the CPR dilemma seen here: individuals are more likely to *fully defect* and choose assured resource destruction over a risky, but preferred NE. Viewing the social dilemma as a discrete shift rather than a simple rent dissipation problem has significant implications for the causes of and solutions to CPR problems.

[12] All coefficients comparisons presented herein were conducted using a Wald test. The finding that two coefficients were not significantly different refers to the 0.10 level. Unless noted otherwise, the reported result is consistent across each of the specifications presented.

Table 6
Token requests including policies.

Variables	(1) All participants (FE)	(2) Partially defect only (FE)
Low uncertainty	0.456***	0.0284
	(0.166)	(0.114)
High uncertainty	0.480***	0.117
	(0.169)	(0.145)
Tax	−0.737***	−0.525***
	(0.113)	(0.0935)
Fine	−0.497***	−0.298***
	(0.0970)	(0.0676)
Tax* low uncertainty	−0.724***	−0.305**
	(0.175)	(0.146)
Tax* high uncertainty	−1.148***	−0.805***
	(0.164)	(0.146)
Fine* low uncertainty	−0.931***	−0.536***
	(0.173)	(0.113)
Fine* high uncertainty	−1.165***	−0.845***
	(0.182)	(0.146)
Constant	4.618***	4.377***
	(0.0791)	(0.0636)
Observations	6,912	6,586
# of participants	96	96
R-squared	0.193	0.143

Standard errors (in parentheses) are clustered at the subject level.
$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

4.2. *The effect of policies on token requests and the probability that participants choose the fully defect NE*

Our results strongly suggest that uncertainty leads to overconsumption relative to the social optimum not just because individuals do not account for the damage additional token requests impose on others, but also because it increases the likelihood that they choose to *fully defect*. The theoretical model also suggests that policies such as a tax per token requested or a fine for threshold exceedance may improve the efficiency of outcomes. Empirical evidence of their effectiveness is presented in Tables 6 and 7.

We begin by estimating Equation (9), including observations from both policy and no-policy periods, as well as controls for each policy conditional on the level of uncertainty. Table 6 presents the results of this analysis, where the dependent variable, y_{jtcs}, is defined as the total number of tokens requested by each individual in a given period.

Consistent with the analysis presented in Table 4, the results suggest that uncertainty leads to an increase in total requests when considering all rounds, but does not have a statistically significant effect when conditioning on the decision to not *fully defect*. The tax has a greater impact on requests than the fine when the threshold is certain (difference significant at .05). This result is consistent with theoretical predictions that the fine does not eliminate the *fully defect NE*. Both the tax and the fine are less effective at reducing request when we restrict the sample to only those individuals who do not *fully defect* (both with and without uncertainty). This is not surprising, given that this specification does not capture the effect of each policy on the decision to *fully defect*.

Interestingly, the difference between the effect of the tax and fine under uncertainty is not statistically different from zero in either the full sample or when restricting it to only those who do not *fully defect*. Using just the observations that do not *fully defect*, this result is expected because both policies are parameterized to incentivize the *SOR solution*. In the case of the full sample, this result is somewhat surprising because the fine does not eliminate the *fully defect NE*. This equivalency could be due to two counteracting effects. First, the certainty of the tax provides a strong incentive to reduce requests. On the other hand, conditional on exceeding the threshold, the magnitude of the fine is significantly larger than the total amount of taxes paid, assuming individuals choose the *SOR solution*. This effect, which depends on risk preferences, may result in a greater reduction in token requests in response to the fine. These two counteracting effects could explain the equivalency of the fine and tax policies in practice despite the theoretical difference between the policies' ability to eliminate the *fully defect NE*.

Finally, we examine if the policies reduce the probability of *fully defecting*. Table 7 presents coefficient estimates corresponding to the discrete choice for the full sample including controls for the various levels of uncertainty and policies. First note that the impact of uncertainty on the probability that individuals choose to *fully defect* is still positive and statistically significant. Consistent with theory, the tax has a greater negative impact (relative to the fine) on the probability that individuals choose to *fully defect* across all three models. This difference is statistically significant at the 10% level under the certain and low uncertainty treatments (for the Logit and Logit RE models), but not statistically different under the high uncertainty treatment. The tax is also more effective because it more effectively differentiates token requests across productivity types (Table 3).

Table 7
Decision to fully defect including policies.

Variables	(1) Logit	(2) Logit RE	(3) FE
Low uncertainty	0.124***	0.115***	0.104***
	(0.0356)	(0.0320)	(0.0278)
High uncertainty	0.120***	0.111***	0.0990***
	(0.0321)	(0.0286)	(0.0238)
Tax	−0.331***	−0.265***	−0.0534***
	(0.0871)	(0.0641)	(0.0157)
Fine	−0.214**	−0.175**	−0.0482***
	(0.106)	(0.0832)	(0.0178)
Tax* low uncertainty	−0.0808	−0.0820	−0.103***
	(0.0967)	(0.0757)	(0.0277)
Tax* high uncertainty	−0.0222	−0.0356	−0.0951***
	(0.0854)	(0.0662)	(0.0233)
Fine* low uncertainty	−0.0329	−0.0432	−0.0937***
	(0.119)	(0.0959)	(0.0315)
Fine* high uncertainty	−0.0548	−0.0599	−0.0924***
	(0.107)	(0.0848)	(0.0255)
Risk aversion	0.0223	0.0122	
	(0.0347)	(0.0333)	
Constant			0.0560***
			(0.0135)
Observations	6,912	6,912	6,912
# of participants	96	96	96
R-squared			0.090

Standard errors (in parentheses) are clustered at the subject level.
Logit models are reported as average marginal effects
$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

In summary, our experimental results suggest that uncertainty in a resource threshold incentivizes individuals to use a CPR more intensely. This effect is driven by uncertainty's impact on the probability that individuals choose the *fully defect NE*. Taxes on resource use as well as fines for resource destruction can be effective at decreasing the probability that an individual knowingly exceeds a resource threshold and the experimental evidence suggests that a tax is more efficient because it incentivizes reduced resource use while better differentiating the behavior of high and low-productivity resource constituents. These results suggest that, when considering effective CPR management, resource managers should construct policies that address both marginal overextraction and the discrete decision to cross resource thresholds because of the possibility that other users will do the same.

5. Conclusion and discussion

The theoretical and empirical results presented here suggest that the destruction of CPRs may result from a discrete choice to exhaust the resource as well as a failure to internalize the full costs of resource extraction. When uncertainty around a resource's ecological or physical tipping point exists, individuals are more likely to ignore potential resource degradation even when the corresponding consumption decisions lead to inferior outcomes and assured resource destruction. We find that while uncertainty causes individuals to more frequently choose the NE that results in certain resource destruction, it has little effect on those individuals who do not choose the *fully defect NE*. In addition, taxes and fines for resource extraction decrease the probability that individuals choose to exhaust the resource under the assumption that resource collapse is inevitable. Even though a fine may not eliminate the *fully defect NE*, its presence does significantly improve the chances that individuals coordinate to keep the resource healthy.

Previous studies have found that threshold uncertainty leads to increased consumption. Our results do not contradict these findings, but they suggest that the increase in consumption under uncertainty is partially driven by resource constituents' decisions to "rationally give up" and ignore impacts on resource stocks in anticipation that others will do the same. Our main contribution is to distinguish CPR over-use incentives in a context of threshold uncertainty—a distinction that is important from a behavioral and policy perspective. If little is known about a resource, users may take the mindset that it is likely to be destroyed regardless of their choices. In this context, the rational way to use the resource is to myopically maximize net value, ignoring any true impact on the resource. In this way, CPR dilemmas in the presence of tipping points create both a Prisoner's Dilemma and a coordination problem.

This research has several policy-relevant implications. First, regulation of CPRs with uncertain tipping points through taxes and fines has the potential to improve efficiency. This improvement comes through both marginal decreases in extraction as well as deterring individuals from choosing to deliberately cross a resource threshold. When strict tax and fine policies are politically infeasible, there may still be significant gains from interventions that eliminate the ability (or desire) of resource constituents to choose the *fully defect NE*.

In settings where outside policies are unlikely to be adopted and local institutions do not exist or lack the capacity to design and implement conservation policies, better information about the resource may provide an opportunity to improve the sustainability of that resource over time and increase the value obtained from scarce natural resources, even if the first-best solution cannot be reached. Reducing or re-characterizing uncertainty to affect public perception (and actions) has been explored in the field of climate science, where "identifying and reducing scientific uncertainty about climate change is a dominant theme of many scientific assessment[s]" (Shackley and Wynne, 1996). Climate scientists have chosen to largely focus public information campaigns on exact thresholds because such precise goals can encourage coordination.[13] As with greenhouse gas emissions, decreasing scientific uncertainty may change resource extraction behavior even without implementing additional policies. In this way, improving the physical understanding of vulnerable CPR systems can avoid resource exhaustion and enhance the livelihoods of people who depend on the resource.

While there are some shortcomings associated with condensing a time-dependent CPR scenario into a single-period experiment, our setup maintains much of the incentive structure present in CPR settings. Moreover, setting up this repeated static game allows us to easily differentiate between discrete choices in the presence of threshold uncertainty.

Given the strong effect of threshold uncertainty on individual decisions to choose the *fully defect NE*, our results suggest the need to consider this feature in future CPR modeling whenever such uncertainty exists. Future work in this area should evaluate the interaction of social and threshold uncertainty within this framework, the introduction of other types of policies such as communication, and the effect of changes in relative superiority on NE choice. Additionally, future work should include a rigorous investigation of individuals' beliefs about others and a thorough exploration into the behavioral and psychological components that cause individuals to "rationally give up" on conserving common pool resources.

Acknowledgments

This research was funded by the Rocky Mountain Research Station, U.S. Forest Service, funding was also provided by NSF Sustainability Research Network (SRN) Cooperative Agreement 1444758.

Appendix A. Supplementary material

Supplementary material related to this article can be found online at http://dx.doi.org/10.1016/j.geb.2017.06.009.

References

Adams, W.M., Brockington, D., Dyson, J., Vira, B., 2003. Managing tragedies: understanding conflict over common pool resources. Science 302 (5652), 1915–1916.

Aflaki, S., 2013. The effect of threshold uncertainty on the tragedy of the commons. Games Econ. Behav. 82, 240–253.

Anderies, J.M., Janssen, M.A., Lee, A., Wasserman, H., 2013. Environmental variability and collective action: experimental insights from an irrigation game. Ecolog. Econ. 93, 166–176.

Andreoni, J., 1988. Why free ride? Strategies and learning in public goods experiments. J. Public Econ. 37 (3), 291–304.

Apesteguia, J., Maier-Rigaud, F.P., 2006. The role of rivalry: public goods versus common-pool resources. J. Conflict Resolution 50 (5), 646–663.

Barrett, S., Dannenberg, A., 2012. Climate negotiations under scientific uncertainty. Proc. Natl. Acad. Sci. USA 109 (43), 17372–17376.

Biel, A., Gärling, T., 1995. The role of uncertainty in resource dilemmas. J. Environ. Psychol. 15 (3), 221–233.

Blanco, E., Lopez, M., Walker, J., 2015. The opportunity costs of conservation with deterministic and probabilistic degradation externalities. Environ. Resource Econ. 64 (2), 255–273.

Botelho, A., Dinar, A., Costa Pinto, L.M., Rapoport, A., 2014. Time and uncertainty in resource dilemmas: equilibrium solutions and experimental results. Exper. Econ. 17 (4), 649–672. http://dx.doi.org/10.1007/s10683-013-9388-2. Springer US.

Botsford, L.W., Castilla, J.C., Peterson, C.H., 1997. The management of fisheries and marine ecosystems. Science 277 (5325), 509–515.

Budescu, D.V., Rapoport, A., Suleiman, R., 1990. Resource dilemmas with threshold uncertainty and asymmetric players. Eur. J. Soc. Psychol. 20 (6), 475–487.

Budescu, D.V., Rapoport, A., Suleiman, R., 1995. Common pool resource dilemmas under uncertainty: qualitative tests of equilibrium solutions. Games Econ. Behav. 10 (1), 171–201.

Charness, G., Gneezy, U., Halladay, B., 2016. Experimental methods: pay one or pay all. J. Econ. Behav. Organ. 131, 141–150.

Cherry, T.L., Cotten, S.J., Kroll, S., 2013a. Heterogeneity, coordination and the provision of best-shot public goods. Exper. Econ. 16 (4), 497–510.

Cherry, T.L., Kallbekken, S., Kroll, S., McEvoy, D., 2013b. Cooperation in and out of markets: an experimental comparison of public good games and market with externalities. Econ. Letters 120, 93–96.

Dannenberg, A., Löschel, A., Paolacci, G., Reif, C., Tavoni, A., 2015. On the provision of public goods with probabilistic and ambiguous thresholds. Environ. Resource Econ. 61, 365–383.

Dawes, R.M., 1980. Social dilemmas. Annu. Rev. Psychol. 31 (1), 169–193.

Esarey, J., Menger, A., 2015. Practical and Effective Approaches to Dealing with Clustered Data. Unpublished Manuscript. Department of Political Science, Rice University.

Fischbacher, U., 2007. Z-Tree: Zurich toolbox for ready-made economic experiments. Exper. Econ. 10 (2), 171–178. Springer.

Gärling, T., Biel, A., Gustafsson, M., 1998. Different kinds and roles of threshold uncertainty. J. Environ. Psychol. 18 (1), 75–83.

Gordon, H.S., 1954. The economic theory of a common-property resource: the fishery. J. Polit. Economy 62 (2), 124–142.

Hanemann, W. Michael, 1984. Discrete/continuous models of consumer demand. Econometrica (Journal of the Econometric Society), 541–561.

Hardin, G., 1968. The tragedy of the commons. Science 162 (3859), 1243–1248.

Jamieson, S., 2004. Likert scales: how to (ab)use them. Med. Educ. 38 (12), 1217–1218.

[13] Two common examples are 350.org, which was named for its goal of reducing atmospheric carbon-dioxide to 350 parts per million and the Paris Climate Agreement's goal of maintaining global temperatures within 2 degrees of the pre-industrial average.

McBride, M., 2010. Threshold uncertainty in discrete public good games: an experimental study. Econ. Governance 11 (1), 77–99.
Messick, D.M., Allison, S.T., Samuelson, C.D., 1988. Framing and communication effects on group members' responses to environmental and social uncertainty. Appl. Behav. Econ. 2, 677–700.
Ostrom, E., 2006. The value-added of laboratory experiments for the study of institutions and common-pool resources. J. Econ. Behav. Organ. 61 (2), 149–163.
Poe, G.L., Schulze, W.D., Segerson, K., Suter, J.F., Vossler, C.A., 2004. Exploring the performance of ambient-based policy instruments when nonpoint source polluters can cooperate. Amer. J. Agr. Econ. 86 (5), 1203–1210.
Rapoport, A., Budescu, D.V., Suleiman, R., Weg, E., 1992. Social dilemmas with uniformly distributed resources.
Shackley, S., Wynne, B., 1996. Representing uncertainty in global climate change science and policy: boundary-ordering devices and authority. Sci. Technol. Human Values 21 (3), 275–302.
Suter, J.F., Duke, J.M., Messer, K.D., Michael, H.A., 2012. Behavior in a spatially explicit groundwater resource: evidence from the Lab. Amer. J. Agr. Econ. 94 (5), 1094–1112.
Van Dijk, E., Wilke, H., 2000. Decision-induced focusing in social dilemmas: give-some, keep-some, take-some, and leave-some dilemmas. J. Pers. Soc. Psychol. 78 (1), 92–104.
Walker, J.M., Gardner, R., 1992. Probabilistic destruction of common-pool resources: experimental evidence. Econ. J. 102 (414), 1149–1161.
Walker, J.M., Gardner, R., Ostrom, E., 1990. Rent dissipation in a limited-access common-pool resource: experimental evidence. J. Environ. Econ. Manage. 19 (3), 203–211.
Wit, A., van Dijk, E., Wilke, H., Groenenboom, A., 2004. The interplay between environmental and social uncertainty in social dilemmas. In: Contemporary Psychological Research on Social Dilemmas, pp. 376–398.