Accepted Manuscript

Threatening Thresholds? The effect of disastrous regime shifts on the non-cooperative use of environmental goods and services

Florian K. Diekert

PII: S0047-2727(17)30013-0
DOI: doi:10.1016/j.jpubeco.2017.01.004
Reference: PUBEC 3745

To appear in: *Journal of Public Economics*

Received date: 6 April 2015
Revised date: 24 September 2016
Accepted date: 24 January 2017

Threatening Thresholds?

The effect of disastrous regime shifts on the non-cooperative use of environmental goods and services

Florian K Diekert*

January 25, 2017

Abstract

This paper presents a tractable dynamic game in which agents jointly use a resource. The resource replenishes fully but collapses irreversibly if the total use exceeds a threshold. The threshold is assumed to be constant, but its location may be unknown. Consequently, an experiment to increase the level of safe resource use will only reveal whether the threshold has been crossed or not. If the consequence of crossing the threshold is disastrous (i.e., independent of how far the threshold has been exceeded), it is individually and socially optimal to update beliefs about the threshold's location at most once. The threat of a disastrous regime thereby facilitates coordination on a "cautious equilibrium". If the initial safe level is sufficiently valuable, the equilibrium implies no experimentation and coincides with the first-best resource use. The less valuable the initial safe value, the more the agents will experiment. For sufficiently low initial values, immediate depletion of the resource is the only equilibrium. When the regime shift is not disastrous, but the damage depends on how far threshold has been exceeded, experimentation may be gradual.

Keywords: Dynamic Games; Thresholds and Natural Disasters; Learning.

JEL-Codes: C73, Q20, Q54

*Department of Economics, Heidelberg University, Bergheimer Straße 20, 69115 Heidelberg, Germany and CEES, Department of Biosciences, University of Oslo, PO-Box 1095 Blindern, 3017 Oslo, Norway. E-mail: mailto:diekert@eco.uni-heidelberg.dediekert@eco.uni-heidelberg.de.

1 Introduction

Many ecosystems are threatened by collapse if overused. Examples include the eutrophication of lakes due to agricultural runoff (Scheffer et al., 2001), sudden shifts in vegetation cover due to land-use changes (Anderies et al., 2002; Dekker et al., 2007), and the collapse of fish stocks, such as Canadian cod or capelin in the Barents Sea (Frank et al., 2005; Hjermann et al., 2004). In the climate system, drivers of a potential regime shift could be a disintegration of the West-Antartic ice sheet (Feldmann and Levermann, 2015), a shutdown of the thermohaline circulation (Nævdal and Oppenheimer, 2007), or a melting of Permafrost (Lenton et al., 2008).

The danger that a disastrous regime shift occurs once a threshold – or tipping point – is crossed, obviously imperils the sustainable provision of ecosystem services. However, the existence of a catastrophic threshold may also be *beneficial* in the sense that it enables non-cooperative agents to coordinate their actions (Barrett and Dannenberg, 2012). This aspect is important because most real-world problems are characterized by the presence of many interacting agents and the absence of central enforcement. Moreover, a key feature of tipping points is that their exact location is almost always unknown. This threshold uncertainty may induce a "safe minimum standard of conservation" (Mitra and Roy, 2006), but, depending on the trade-off between the cost of control and the gain from risk reduction, it may also lead to less precaution (Brozović and Schlenker, 2011).

In this paper, I develop a dynamic game in which agents jointly use a replenishing resource that loses (some or all) its productivity upon crossing some (potentially unknown) threshold. In order to isolate the effect of threshold uncertainty on the ability to cooperate, I abstract – as a first step – from the dynamic common pool aspect of non-cooperative resource use.

The model is presented in section 2. It is general and applicable to many different settings, but to fix ideas, consider the problem of saltwater intrusion in a freshwater reservoir: The reservoir is used by several agents. Its overall volume is approximately known, and the annual recharge (due to rainfall or snowmelt) is sufficient to fully replenish it. However, the agents fear that saltwater may intrude and irreversibly spoil the resource once the water table falls too low. Further, suppose the geology is so complex that it is not known how much water must be left in the reservoir to avoid intrusion. Saltwater intrusion has not occurred in the past, so that the current level of total use is known to be safe. Thus, the agents now face the trade-off whether to expand the current consumption of water, or not. If they decide to expand the current level of use, by how much should extraction increase, and in how many steps should the expansion occur? Moreover, could it be in one agent's own best interest to empty the remaining reservoir even when all others take just their share of the historical use?

In section 3.1, I expose the underlying strategic structure of the game by considering the case where the location of the threshold is known. I show that there is a Nash equilibrium where the resource is conserved indefinitely and a Nash equilibrium where the resource is depleted immediately. In terms of the above example, the former equilibrium will only exist if sharing the amount of water that leaves just enough in the reservoir to avoid intrusion is sufficiently valuable compared to the incentives to deviate and empty the reservoir.

When the location of the threshold is fixed but unknown, any increase in resource use will

– in the absence of passive learning – only reveal whether the updated state is safe or not. The agents will not obtain any new information on how much closer they have come to the threshold.[1] I call this type of learning "affirmative". When the consequence of crossing the threshold is disastrous in the sense that it does not matter by how far the threshold has been overstepped, then there is no point in splitting any given increase in resource use in several steps. Any experimentation is – if at all – undertaken in the first period. Moreover, the degree of experimentation is decreasing in the value of current use that is known to be safe.

This means that both in the sole-owner's solution (section 3.2) and in the non-cooperative game (section 3.3), the steady-state consumption level will depend on history: When the current level of resource use is sufficiently valuable, coordination on not expanding the set of safe consumption values is a Nash equilibrium. If it is socially optimal to use the water reservoir at its current level, this Nash equilibrium will in fact coincide with the first-best resource use. If preserving the status quo is not sufficiently valuable, agents may still refrain from depleting the resource, but they will increase their consumption by an inefficiently high amount. However, provided that the increase in consumption has not caused the disastrous regime shift, the players can coordinate on keeping to the updated level of consumption, which is, *ex post*, socially optimal.

The "once-and-for-all" dynamics of experimentation and resource use under "affirmative learning" are robust to several extensions that are explored in section 4. While the threat of the threshold may no longer induce coordination on the first-best when the externality relates to *both* the (endogenous) risk of passing the threshold and resource itself, the threshold may still encourage coordination on a time-profile of resource use that is, in expected terms, Pareto-superior compared to the Nash equilibrium without a threshold. As I show in section 4.4, repeated experimentation will take place *only* if the post-threshold value depends negatively on the pre-threshold degree of experimentation, and if this effect is sufficiently strong.

Section 5 concludes the paper and points to important future applications of the modeling framework. All proofs are collected in the Appendix.

Relation to the literature

This paper links to three strands of the literature. First, it contributes to the literature on the management of natural resources under regime-shift risk by explicitly analyzing learning about the location of a threshold in a tractable dynamic model. Second, the paper extends the literature on coordination in face of a catastrophic public bad, that has hitherto been analyzed in a static setting. Third, it relates to the broader literature by characterizing optimal experimentation in a set-up of "affirmative learning".

The pioneering contributions that analyze the economics of regime shifts in an environmental/resource context were Cropper (1976) and Kemp (1976). There are by now a good dozen papers on the optimal management of renewable resources under the threat of an irreversible regime shift (see Polasky et al., 2011, for a summary). Most previous studies translate the uncertainty about the location of the threshold in state space into uncertainty about the

[1] Empiricists will agree that there is no learning without experiencing.

occurrence of the event in time. This allows for a convenient hazard-rate formulation (where the hazard rate could be exogenous or endogenous), but it has the problematic feature that, eventually, the event occurs with probability 1. In other words, even if the agents were to totally stop extracting/polluting, the disastrous regime shift would be inevitable. Arguably, it is more realistic to model the regime shift in such a way that when it has not occurred up to some level, the agents can avoid the event by staying at or below that level (Tsur and Zemel, 1994; Nævdal, 2003; Lemoine and Traeger, 2014). To the best of my knowledge, this paper is the first to apply this modeling approach to a non-cooperative game.

In general, the literature in resource economics has been predominantly occupied with optimal management, leaving aside the central question of how agent's strategic considerations influence and are influenced by the potential to trigger a disastrous regime shift. Still, there are a few notable exceptions: Crépin and Lindahl (2009) analyze the classical "tragedy of the commons" in a grazing game with complex feedbacks, focussing on open-loop strategies. Ploeg and Zeeuw (2015b) compare the socially optimal carbon tax to the tax in the open-loop equilibrium under the threat of a productivity shock due to climate change. Reverting to numerical methods, Kossioris et al. (2008) analyze feedback equilibria in a "shallow lake" model. They show that, as in most differential games with renewable resources, the outcome of the feedback Nash equilibrium is in general worse than the open-loop equilibrium or the social optimum. In this paper, I am able to solve for the feedback equilibrium analytically by simplifying the dynamics of resource use.

Fesselmeyer and Santugini (2013) introduce an exogenous event risk into a non-cooperative renewable resource game à la Levhari and Mirman (1980). As in the optimal management problem with an exogenous probability of a regime shift, the impact of shifted resource dynamics is ambiguous: On the one hand, the threat of a less productive resource induces a conservation motive for all players, but on the other hand, it exacerbates the tragedy of the commons as the players do not take the risk externality into account. As risk is exogenous in Fesselmeyer and Santugini (2013), they can obtain analytical solutions in the Levhari-Mirman framework, but their model does not allow learning or adaptions to an evolving regime-shift risk. Sakamoto (2014) analyzes a non-cooperative game with an endogenous regime shift hazard by combining analytical and numerical methods. He shows that the regime-shift risk may lead to more precautionary management, also in a strategic setting. Miller and Nkuiya (2016) also combine analytical and numerical methods to investigate how an exogenous or endogenous regime shift affects coalition formation in the Levhari-Mirman model. They show that an endogenous hazard rate increases coalition sizes and it allows the players, in some cases, to achieve full cooperation. Using a different model setup that allows analytic solutions, this paper corroborates that the effect of a regime shift is qualitatively the same in a non-cooperative setting as under optimal management: for some combinations of parameters it induces more caution and for some combinations it induces less caution. Moreover, both the literature on optimal resource management under regime-shift risk and its non-cooperative counterpart have not explicitly addressed learning about the unknown location of the tipping point, which is the main focus of the present work.

There is a related literature on strategic experimentation in one-armed bandit problems (e.g.: Bolton and Harris, 1999; Keller et al., 2005; Bonatti and Hörner, 2015) that differs from the current paper in that there are no structural irreversibilities. Learning is then "informative" in the sense that agents obtain a random sample on which they base their inference about the state of the world and it pays to obtain repeated samples (but only finitely many in most cases) as this improves the estimate. The public nature of information introduces free-rider incentives in a strategic setting, so that learning is often sub-optimally slow. Here, experimentation will be overly aggressive in most cases.

The current paper is closely related to three articles that discuss the role of uncertainty about the threshold's location on whether a catastrophe can be avoided. Barrett (2013) shows that players in a linear-quadratic game are (in most cases) able to form self-enforcing agreements that avoid catastrophic climate change when the location of the threshold is known, but not when it is unknown. Similarly, Aflaki (2013) analyzes a model of a common-pool resource problem that is, in its essence, the same as the stage-game developed in section 3. Aflaki shows that an increase in uncertainty leads to increased consumption, but that increased ambiguity may have the opposite effect. Bochet et al. (2013) confirm the detrimental role of increased uncertainty in the stochastic variant of the Nash Demand Game: Even though "cautious" and "dangerous" equilibria co-exist (as they do in my model), they provide experimental evidence that participants in the lab are not able to coordinate on the Pareto-dominant cautious equilibrium.[2] However, the models in Aflaki (2013), Barrett (2013), and Bochet et al. (2013) are all static. Here, I show that the sharp distinction between known and unknown location of a threshold does not survive in a dynamic context. More uncertainty still leads to increased consumption, but this is now partly driven by the increased gain from experimentation.

As noted above, a key result of my model is that it is a Nash equilibrium to experiment once or never. Although I am unaware of an earlier comparable application to a strategic setting, results on optimal experimentation in the context of affirmative learning have appeared at various places before. For example, the classical book of Dubins and Savage (1965) analyzes circumstances under which it is optimal for gamblers to expose themselves to uncertainty in as few rounds as possible. Riley and Zeckhauser (1983) discuss price-negotiation strategies where the seller does not know the valuation of the buyer. They find that "[a] seller encountering risk-neutral buyers one at a time should, if commitments are feasible, quote a single take-it-or-leave-it price to each." Another well-known study is from Rob (1991), who analyzes optimal and competitive capacity expansion when market demand is unknown. Rob finds that learning will take place over several periods. In his model, experimenting too much (in the sense of installing more capital than is needed to satisfy the revealed demand) is very costly compared to experimenting too little several times (so that the true size of the market remains unknown). Consequently, learning takes place gradually. Under competition, learning

[2] Bochet et al. (2013, p.1) conclude that a "risk-taking society may emerge from the decentralized actions of risk-averse individuals". Unfortunately, it is not clear from the description in their manuscript whether the participants were able to communicate. The latter has shown to be a crucial factor for coordination in threshold public goods experiments (Tavoni et al., 2011; Barrett and Dannenberg, 2012). Hence, it may be that what they refer to as "societal risk taking" is simply the result of strategic uncertainty.

is even slower due to the private nature of search costs but the public nature of information.

In an application to environmental economics, Costello and Karp (2004) investigate optimal pollution quotas when abatement costs are unknown. In their model, the initial quota is binding with probability 1, but an increased quota may be slack (which is inefficient). While the information gain from a marginal increase in quota is small, there is no additional harm from experimenting too much. In line with the baseline model of the current paper, this feature leads to the conclusion that any experimentation takes place in the first period only. Similarly, Groeneveld et al. (2013) show that the upper bound of the belief about the threshold's location is updated only once in their model of a reversible flow-pollution threshold.

Lemoine and Traeger (2014) find that learning occurs over several periods. In section 4, I analyze two features that are present in their climate-change application and that may both induce repeated experimentation: First, as in Rob's model, the damage of the regime shift is larger the farther the threshold has been overstepped. Second, the dynamics of capital accumulation in Lemoine and Traeger (2014) effectively imply a constraint on the choice set. This leads mechanically to repeated experimentation.

When analyzing learning in a strategic setting, I point out that there are three different forces at work: First, the immediate gains from experimentation are certain and private while the cost of experimentation in terms of an increased regime-shift risk are borne by all. These two forces lead to more experimentation than socially optimal, but they are, to some extent, attenuated by the public nature of information: all agents gain from an expansion of the set of safe consumption values, provided the experiment has not triggered the regime shift. I provide sufficient conditions for when non-cooperative learning is more aggressive than socially optimal. Furthermore, I show that experimentation is decreasing in the value of the state that is known to be safe: The more the agents know that they can safely consume, the less will they be willing to risk triggering the regime shift by enlarging the set of consumption opportunities. This aspect has, to the best of my knowledge, not yet been appreciated.

Analyzing how strategic interactions shape renewable resource use under the threat of a disastrous regime shift is important beyond mere curiosity driven interest. It is probably fair to say that international relations are characterized by an absence of supranational enforcement mechanisms which would allow to make binding agreements. But also locally, within the jurisdiction of a given nation, control is seldom complete and the exploitation of many common pool resources is shaped by strategic considerations. Extending our knowledge on the effect of looming regime shifts by taking non-cooperative behavior into account is therefore a timely contribution to both the scientific literature and the current policy debate.

2 The model

This section presents the basic model setup (resource dynamics; agents, choices, and payoff; regime-shift risk) and discusses a number of tractability assumptions.

Resource dynamics

- Time is discrete and indexed by $t = 0, 1, 2, ...$
- Each period, agents can, in total, consume up to the available amount of the resource. There are two regimes: In the *productive regime*, the upper bound on the available resource is given by R, and in the *unproductive regime*, the upper bound is given by r (with $r < R$).
- The game starts in the productive regime and will stay in the productive regime as long as total consumption does not exceed a threshold T. The threshold T is the same in all periods, but it may be known or unknown.
- To highlight the effect of uncertainty about the threshold, I define the state variable s_t, denoting the upper bound of the "safe consumption possibility set" at time t. That is, total resource use up to s_t has not triggered a regime shift before, and it is hence known that it will not trigger a regime shift in the future (i.e. $\text{Prob}(T \leq s_t) = 0$).

Agents, choices, and payoff

- There are N identical agents. Each agent i derives utility from consuming the resource according to some general function $u(c_t^i)$, where c_t^i is the consumption of agent i at time t. I assume that u is continuous, increasing ($u' > 0$), concave ($u'' \leq 0$), and bounded below by $u(0) = b$.
- For clarity, I split the agent's per-period consumption in two parts: $c_t^i = \frac{s_t}{N} + \delta_t^i$. This means:
  1. The agents obtain an equitable share of the amount of the resource that can be used safely.
  2. The agents may choose to consume an additional amount δ_t^i, effectively pushing the boundary of the safe consumption possibility set at the risk of triggering the regime shift.
- In other words, δ_t^i is the effective choice variable with $\delta_t^i \in [0, R - s_t - \delta_t^{-i}]$, where δ_t^{-i} is the expansion of the safe consumption set by all other agents except i. I denote δ without superscript i as the total extension of the safe set, i.e. $\delta_t = \sum_{i=1}^N \delta_t^i$.
- The objective of the agents is to choose that sequence of state-dependent decisions $\Delta^i = \delta_0^i, \delta_1^i, ...$ which, for given strategies of the other agents Δ^{-i}, and for a given initial value s_0, maximizes the sum of expected per-period utilities, discounted by a common

factor $\beta \in (0,1)$. I concentrate on Markovian strategies because they are "the simplest form of behavior that is consistent with rationality" (Maskin and Tirole, 2001, p.193).

The probability of triggering the regime shift

- Let the probability density of T on $[0, A]$ be given by a continuous function f such that the cumulative probability of triggering the regime shift is *a priori* given by $F(x) = \int_0^x f(\tau)\mathrm{d}\tau$. $F(x)$ is the common prior of the agents, so that we are in a situation of risk (and not Knightian uncertainty).

- The variable A with $R \leq A \leq \infty$ denotes the upper bound of the support of T. When $R < A$, there is some probability $1 - F(R)$ that using the entire resource is safe and the presence of a critical threshold is immaterial. When $R = A$ using the entire resource will trigger the regime shift for sure. Both R and A are known with certainty.[3]

- Knowing that a given consumption level s is safe, the updated density of T on $[s, A]$ is given by $f_s(\delta) = \frac{f(s+\delta)}{1-F(s)}$ (see Figure 1). The cumulative probability of triggering the regime shift when, so to say, taking a step of distance δ from the safe value s is:

$$F_s(\delta) = \int_0^\delta f_s(\tau)\mathrm{d}\tau = \frac{1}{1-F(s)}\int_0^\delta f(s+\xi)\mathrm{d}\xi = \frac{F(s+\delta)-F(s)}{1-F(s)} \quad (1)$$

  So that $F_s(\delta)$ is the discretized version of the hazard rate. I assume that the hazard rate does not decrease in s.

- The (Bayesian) updating of beliefs is illustrated in Figure 1. Note that it is only revealed whether the state s is safe or not, but no new knowledge about the relative probability that the threshold is located at s_1 or s_2 (with $s_1, s_2 > s$) has been acquired.

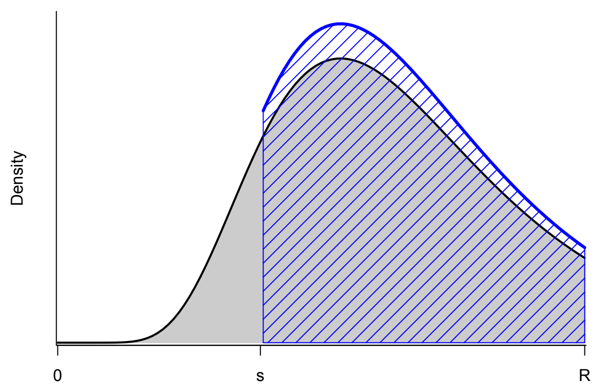

Figure 1: Updating of belief upon learning that $T > s$: Grey area is F, blue hatched area is F_s.

[3]The idea that a system is more likely to experience a disastrous regime shift the lower the amount of the resource that has been left untouched could simply be included in the belief $F(x)$. Additive disturbances, such as stochastic (white) noise, are independent of the current state and would not affect the calculations in a meaningful way. They could be absorbed in the discount factor.

- The key expression that I use in the remainder of the paper is $L_s(\delta)$, which I call the conditional survival function. It denotes the probability that the threshold is not crossed when taking a step δ, given that the event has not occurred up to s. Let $L(x) = 1-F(x)$:

$$L_s(\delta) = 1 - F_s(\delta) = \frac{1 - F(s) - (F(s+\delta) - F(s))}{1 - F(s)} = \frac{L(s+\delta)}{L(s)} \tag{2}$$

  The conditional survival function has the following properties:

  - It decreases with the step size δ: $\frac{\partial L_s(\delta)}{\partial \delta} = \frac{-f(s+\delta)}{1-F(s)} < 0$.
  - It decreases with s: $\frac{\partial L_s(\delta)}{\partial s} = \frac{-f(s+\delta)(1-F(s))+(1-F(s+\delta))f(s)}{[1-F(s)]^2} \leq 0 \Leftrightarrow \frac{f(s)}{1-F(s)} \leq \frac{f(s+\delta)}{1-F(s+\delta)}$ (as the hazard rate is non-decreasing).

Clarifications and tractability assumptions

- It is well known that the static non-cooperative game of sharing a given resource has infinitely many equilibria: Even when the agents are assumed to be symmetric, any given division of the total resource is an equilibrium. Moreover, the game requires a statement about the consequences when the sum of consumption plans exceeds the total available resource. Here, I assume that each agent gets an equal share. This assumption could be justified by relying on a cooperative bargaining solution such as Nash (1953) or as the outcome of a non-cooperative bargaining game where each agent is allowed to make a take-it-or-leave-it offer with equal probability (Harstad, 2012). The important assumption of symmetry is further discussed in section 5.

- The agent's prior $F(x)$ is fixed. The absence of any passive learning (an arrival of information simply due to the passage of time) is justified in a situation where all learning opportunities from other, similar resources have been exhausted. The only way to learn more about the location of the threshold in the specific resource at hand is to experiment with it.[4]

- The regime shift is irreversible. Moreover, I consider the regime shift to be disastrous, in the sense that crossing the thresholds breaks all links between the pre-event and the post-event regime. Because the post-event value function is then independent of the pre-event state, I set, for simplicity's sake, $r = 0$ and $b = 0$. In section 4.4, I discuss the case when the post-event value function depends on the pre-event state.

- The model abstracts from the dynamic common pool problem in the sense that the consumption decision of an agent today has no effect on the consumption possibilities tomorrow, *except* that a) the set of safe consumption possibilities may have been enlarged and b) the disastrous regime shift may have been triggered. This assumption is relaxed in section 4.2.

[4] An everyday example is blowing up a ballon: We all know that they will burst at some point, and we have blown up sufficiently many balloons, or seen our parents blow sufficiently many balloons to have a good idea which size is safe. But for a given balloon at hand, I do not know when it will burst.

3 Social optimum and non-cooperative equilibrium

In this main part of the paper, I will first expose the underlying strategic structure of the model by analyzing the situation when the threshold is known (section 3.1). In section 3.2, I describe the optimal course of action in absence of strategic interactions to highlight that any experimentation is – if at all – undertaken in the first period. Moreover, experimentation is decreasing with the value of the consumption level that is known to be safe. I then show that this feature of learning may allow for a cautious non-cooperative equilibrium: Either the resource is conserved with probability 1 or the agents experiment once (section 3.3). The degree of experimentation will be inefficiently large in most cases, but if the threshold has not been crossed, staying at the updated safe level is – *ex post* – socially optimal. In section 3.4, I analyze how optimal and non-cooperative resource use shifts with changes in the parameters. Finally, I provide an instructive example for which I derive closed-form solutions (section 3.5).

3.1 Known threshold location

When the threshold T is known, the first-best resource use, maximizing the sum of agent's utilities, is to equitably share just the amount of the resource that can be used safely if and only if $Nu(R/N) \leq N\frac{u(T/N)}{1-\beta}$.

Intuitively, when T is small, too much of the resource must be left untouched to ensure its future existence. As a consequence, it is socially optimal to cross the threshold and consume the entire resource immediately. When T is large, however, the per-period utility from staying at the threshold is sufficiently high so that the first-best is to indefinitely use exactly that amount of the resource which does not cause the regime shift. Whether a given T is large enough to induce conservation depends on the overall amount of the resource R and the discount factor β. The more of the resource must be left untouched, or the more the future is discounted, the less willing one is to sacrifice today's consumption of R to ensure continued consumption of T. Thus, I define the critical value T_c^* such that immediate depletion is first-best when $T < T_c^*$ and staying at T is first-best when $T > T_c^*$. That is, T_c^* is given by $u(R/N) - \frac{u(T/N)}{1-\beta} = 0$.

In the non-cooperative game with a known threshold, immediate depletion is always a Nash equilibrium. Clearly, an agent's best reply when the other agents cross the threshold is to demand the maximal amount of the resource as well. However, also here there will be a critical value T_c^{nc} so that staying at the threshold T is also Nash equilibrium when $T \geq T_c^{nc}$. In fact, as Proposition 1 states, there will always be a parameter combination so that the first-best of staying at T can be supported as a Nash equilibrium. Similarly, when $T < T_c^*$, the Nash-equilibrium of immediate depletion will again be socially optimal.

As the setup is stationary, it is clear that if staying at the threshold can be rationalized in any one period, it can be done so in every period. The payoff from avoiding the regime shift is $\frac{u(T/N)}{1-\beta}$. Conversely, the payoff from deviating and immediately depleting the resource when all other players intend to stay at the threshold is given by $u\big(R - \frac{N-1}{N}T\big)$. The lower T is, the lower the payoff from staying at the threshold, and the higher the payoff from deviating.

I can therefore define a function Ψ that captures agent i's incentive to grab the resource when all other agents stay at T:

$$\Psi(T, R, N, \beta) = u\left(R - \frac{N-1}{N}T\right) - \frac{u(T/N)}{1-\beta} \tag{3}$$

The function Ψ is positive at $T = 0$ and declines as T gets larger. Staying at the threshold can be sustained as a Nash equilibrium whenever $\Psi \leq 0$. The critical value T_c^{nc} is implicitly defined by $\Psi(T_c, R, N, \beta) = 0$. Note that $T_c^* < T_c^{nc}$ because $u\left(\frac{R}{N}\right) < u\left(R - \frac{N-1}{N}T\right)$ as $N > 1$ and $R > T$.

Proposition 1. *When the location of the threshold is known with certainty, then there exists, for every combination of β, N, and R, a value T_c^{nc} such that the first-best of staying at T can be sustained as a Nash equilibrium when $T \geq T_c^{nc}$, where T_c^{nc} is defined by $\Psi = 0$. The critical value T_c^{nc} is higher, the larger N or R are, or the smaller β is.*

Proof. The proof is placed in Appendix A.1 □

In other words, when T is known and $T \geq T_c^{nc}$, the game exhibits the structure of a coordination game with two Nash equilibria in symmetric pure strategies. Here, as in the static game from Barrett (2013, p.236), "[e]ssentially, nature herself enforces an agreement to avoid catastrophe." When staying at or below the threshold is not sufficiently valuable, immediate depletion is the only equilibrium.

Having exposed the underlying strategic structure of the game, I now turn to the situation when the location of the threshold is unknown: First, I disregard strategic interactions and study optimal experimentation of a single agent. Then, I analyze the non-cooperative game with unknown location of T.

3.2 Optimal experimentation when the location of T is unknown

Consider the problem of a single decision maker (a "sole-owner") with the following objective:

$$\max \sum_{t=0}^{\infty} \beta^t u(c_t) \quad \text{subject to: } R_{t+1} = \begin{cases} R_t & \text{if} \quad c_t \leq T \\ 0 & \text{if} \quad c_t > T \quad \text{or } R_t = 0 \end{cases} \quad ; \quad R_0 = R. \tag{4}$$

Starting from a historically given safe value s_t, and a belief about the location of the threshold, the sole-owner has in principle two options: She can either stay at s_t (choose $\delta = 0$), thereby ensuring the existence of the resource in the next period. Alternatively, she can take a positive step into unknown territory (choose $\delta > 0$), potentially expanding the set of safe consumption possibilities to $s_{t+1} = s_t + \delta$, albeit at the risk of a resource collapse. Recall that $L_s(\delta)$ is the probability of surviving (that is, not crossing the threshold when taking a step of size δ from the safe value s). We can thus write the sole-owner's Bellman equation as:

$$V(s) = \max_{\delta \in [0, R-s]} \left\{ u(s+\delta) + \beta L_s(\delta) V(s+\delta) \right\} \tag{5}$$

The crux is, of course, that the value function $V(s)$ is *a priori* not known. However, we do know that once the sole-owner has decided to not expand the set of safe consumption possibilities, it cannot be optimal to do so at a later period: If $\delta = 0$ is chosen in a given period, nothing is learned for the future ($s_{t+1} = s_t$), so that the problem in the next period is identical to the problem in the current period. If moving in the next period were to increase the payoff, it would increase the payoff even more when one would have made the move a period earlier (as the future is discounted).

To introduce some notation, let s^* be a member of the set of admissible consumption values $[0, R]$ at which it is not optimal to expand the set of safe consumption values (as the threat of a disastrous regime shift looms too large). Denote this set of values by S and let $\overline{s}^*$ be the smallest member of S. In Appendix A.2, I show that S must exist and that it is convex when the hazard rate is non-decreasing. Thus, for $s \geq \overline{s}^*$, it is optimal to choose $\delta = 0$. In this case, we know $V(s)$. It is given by $V(s) = \frac{u(s)}{1-\beta}$.

This leaves three possible paths when starting from values of s_0 that are below $\overline{s}^*$. The decision maker could: 1) make one step and then stay, 2) make several, but finitely many steps and then stay, and 3) make infinitely many steps. I now argue that 1) is optimal.

Suppose that a value at which it is optimal to remain standing is reached in finitely many steps. This implies that there must be a last step. For this last step, we can explicitly write down the objective function as we know that the value of staying at s^* forever is $\frac{u(s^*)}{1-\beta}$. Denote by $\varphi(\delta; s)$ the sole-owner's valuation of taking exactly one step of size δ from the initial value s to some value s^* and then staying at s^* forevermore, and denote by $\delta^*(s)$ the optimal choice of the last step. Formally:

$$\varphi(\delta; s) = u(s+\delta) + \beta L_s(\delta) \frac{u(s+\delta)}{1-\beta}. \tag{6}$$

This yields the following first-order-condition for an interior solution:

$$\varphi'(\delta; s) = u'(s+\delta) + \beta \left[L_s'(\delta) \frac{u(s+\delta)}{1-\beta} + L_s(\delta) \frac{u'(s+\delta)}{1-\beta} \right] = 0. \tag{7}$$

With these explicit functional forms in hand, I can show that it is better to traverse any given distance before remaining standing in one step rather than two steps (see Appendix A.2). *A fortiori*, this holds for any finite sequence of steps. Also an infinite sequence of steps cannot yield a higher payoff since the first step towards s^* will be arbitrarily close to s^* and concavity of the utility function ensures that there is no gain from never actually reaching s^*.

Let $g^*(s)$ be the interior solution to the first-order-condition (7). Note that we need not have an interior solution so that $\delta^*(s) = 0$ when $\varphi'(\delta; s) < 0$ for all δ and $\delta^*(s) = R - s$ when $\varphi'(\delta; s) > 0$ for all δ. The first corner solution arises when $s \geq \overline{s}^*$. Similarly, I define a critical value $\underline{s}^*$ so that the second corner solution arises when $s \leq \underline{s}^*$. (In most cases, this corner solution is not relevant.) That is, the optimal expansion of the set of safe consumption values is given by:

$$\delta^*(s) = \begin{cases} R - s & \text{if} \quad s \leq \underline{s}^* & \text{(8a)} \\ g^*(s) & \text{if} \quad s \in (\underline{s}^*, \overline{s}^*) & \text{(8b)} \\ 0 & \text{if} \quad s \geq \overline{s}^* & \text{(8c)} \end{cases}$$

The optimal consumption pattern is summarized by the following proposition:

Proposition 2. *There exists a set S so that for $s \in S$, it is optimal to choose $\delta^*(s) = 0$. That is, if $s_0 \in S$, the optimal use of the resource is s_0 for all t. If $s_0 \notin S$, it is optimal to experiment once at $t = 0$ and expand the set of safe values by $\delta^*(s_0)$. When this has not triggered the regime shift, it is optimal to stay at $s_1 = s_0 + \delta^*(s_0)$ for all $t \geq 1$.*

Proof. The proof is given in Appendix A.2. □

In other words, any experimentation – if at all – is undertaken in the first period. The intuition is the following: Given that it is optimal to eventually stop at some $s^* \in S$, the probability of triggering the regime shift when going from s_0 to s^* is the same whether the distance is traversed in one step or in many steps. Due to discounting, the earlier the optimal safe value s^* is reached, the better.[5]

Moreover, the degree of experimentation depends on history. When the second-order condition is fulfilled[6] it can be shown that the optimal step size $\delta^*(s)$ is declining in s (Proposition 3). The intuition for this effect is clear: The more valuable the current safe level of use, the less the sole-owner can gain from an increased use, but the more she can lose should the experiment trigger the regime shift. In other words, the more the decision maker knows, the less she wants to learn. In fact, this implies that the largest step is undertaken when $s = 0$, which is reminiscent Janis Joplin's dictum that "freedom is just another word for nothing left to lose".

Proposition 3. *The optimal step size $\delta^*(s)$ is decreasing in s for $s \in (\underline{s}^*, \overline{s}^*)$.*

Proof. The proof is placed in Appendix A.3. □

With this characterization of the optimal experimentation in absence of strategic interactions in place, I turn to the non-cooperative game.

[5]The astute reader will wonder whether the adopted timing "action - consumption - reaction" is critical for the result of immediate experimentation. In Appendix A.2, I show that immediate experimentation is also optimal under the alternative timing assumption of "action - reaction - consumption" (i.e. when utility in the first period is only obtained when the regime shift has not occurred).

[6]The second-order condition is fulfilled when $\left(\frac{1-\beta}{\beta} + L_s(\delta^*)\right)u'' + 2L_s'(\delta^*)u' + L_s''(\delta^*)u < 0$. Note that while the first term is negative because $\beta \in (0,1)$, $L_s(\delta^*) \geq 0$, and $u'' \leq 0$, and the second term is also negative because $u' > 0$ and $L_s'(\delta^*) < 0$, the third term $L_s''(\delta^*)u$ may be positive.

3.3 Non-cooperative equilibrium when the location of T is unknown

For a given value of the total consumption that is known to be safe, and a given state-dependent strategy of the other players that extends, in sum, the set of consumption values by δ^{-i}, the Bellman equation of agent i is:

$$V^i(s,\delta^{-i}) = \max_{\delta^i \in [0,R-s-\delta^i]} \left\{ u(s/N+\delta^i) + \beta L_s(\delta^i+\delta^{-i}) V^i(s+\delta,\delta^{-i}) \right\} \tag{9}$$

Also here, the crux is that agent i's value function V^i is *a priori* unknown. However, as the analysis in the previous section has highlighted, we do know that s divides the state space into a safe region and an unsafe region. Moreover, due to the stationarity of the problem, we know that if the agents can coordinate to stay in the safe region once, they can do so forever. Below, I will show that there indeed exists a set S^{nc} where for any $s \in S^{nc}$ staying at s is an equilibrium. However, just as in the case when the threshold's location is known, immediate depletion is always also a Nash equilibrium. But different from the case when the threshold's location is known, immediate depletion need not be the best-reply when $s \notin S^{nc}$. Rather, the agents may coordinate on expanding the set of safe consumption values by some amount δ^{nc} and this experiment need not trigger the regime shift. Provided that the regime shift has not occurred, the set of safe consumption possibilities will be expanded up to a level where it is a Nash equilibrium to not expand it further. Parallel to the socially optimal experimentation pattern, it will be a Nash equilibrium to reach the set S^{nc} in one step. This "cautious" pattern of non-cooperative resource use is summarized by the following proposition.

Proposition 4. *There exists a set S^{nc} such that for $s_0 \in S^{nc}$, it is a symmetric Nash equilibrium to stay at s_0 and consume $\frac{s_0}{N}$ for all t. For $s_0 \notin S^{nc}$, it is a Nash equilibrium to take exactly one step and consume $\frac{s_0}{N} + \delta^{nc}(s_0)$ for $t = 0$ and – when this has not triggered the regime shift – to stay at $s_1 = s_0 + N\delta^{nc}(s_0)$, consuming $\frac{s_1}{N}$ for all $t \geq 1$.*

Proof. The proof is given in Appendix A.4 □

The key intuition for the existence of this "cautious equilibrium" is that 1) for high values of s, staying at s is individually rational when all other agents do so, too, and 2) that when $s \notin S^{nc}$, no agent has an incentive to deviate from a one-step experimentation that expands the set of safe consumption values into the region in which staying is optimal. Of course, there will always also exist an "aggressive equilibrium" in which the resource is depleted immediately, simply because the best-reply for player i when all other players plan to expand the consumption set by $\frac{R-s}{N}$ is to choose $\frac{R-s}{N}$ as well. Note that, for a given s, both the "cautious" and the "aggressive equilibrium" are unique.[7]

[7] Uniqueness of the latter type of equilibrium simply follows from the assumption that in case of incompatible demands, the resource is shared equally among the players. Uniqueness of the symmetric "cautious equilibrium" (should it entail $\delta^{nc}(s) < \frac{R-s}{N}$) can be established by contradiction. Suppose all other players $j \neq i$ choose to expand the consumption set to a level at which – should the threshold have not been crossed – no player would have an incentive to go further. Player i's best-reply cannot be to choose $\delta^i = 0$ in this situation as the gain from making a small positive step (which are private) exceed the (public) cost of advancing a little further. Hence, the only equilibrium at which the players expand the consumption set once is the symmetric one.

Let ϕ denote the payoff for agent i when she takes exactly one step of size δ^i and then remains standing and the strategy of the other agents, $\Delta^{-i} = \{\delta^{-i}, 0, 0, 0, ...\}$, is also to take only one step (of total size δ^{-i}):

$$\phi(\delta^i; \delta^{-i}, s) = u\Big(\frac{s}{N} + \delta^i\Big) + \beta L_s(\delta^i + \delta^{-i}) \frac{u\Big(\frac{s+\delta^i+\delta^{-i}}{N}\Big)}{1-\beta} \tag{10}$$

The corresponding first-order-condition for an interior maximum is:

$$\begin{aligned}\phi'(\delta^i; \delta^{-i}, s) = u'\Big(\frac{s}{N} + \delta^i\Big) \\ + \beta L'_s(\delta^i + \delta^{-i}) \frac{u\Big(\frac{s+\delta^i+\delta^{-i}}{N}\Big)}{1-\beta} \\ + \beta \frac{1}{N} L_s(\delta^i + \delta^{-i}) \frac{u'\Big(\frac{s+\delta^i+\delta^{-i}}{N}\Big)}{1-\beta} \quad = 0\end{aligned} \tag{11}$$

Denote the interior solution to the first-order-condition (if it exists) by $g(\delta^{-i}, s)$. Three forces determine g: The first term represents the gain from a marginal increase in current utility. For a given s, this term is larger the more agents there are (as $u'' \leq 0$). The second term represents the marginal decrease in the probability of surviving, which is evaluated at the updated safe consumption value. As agent i obtains only $\frac{1}{N}$th of the updated safe consumption value, these cost weigh less the more agents there are. Third, conditional on survival, there is the marginal utility gain from an expanded safe consumption set. As this benefits all agents equally, it is devalued by the factor $\frac{1}{N}$.

The first two terms capture the "tragedy of the commons" with respect to the regime shift risk in the sense that the current gains from an experiment are private but the cost in terms of increased risk are public and shared by all. Therefore, the first two terms push for a sub-optimally large expansion. However, the third term pulls in the opposite direction as the agents do not take the informational value that their experiment has for the other agents into account. Without further assumptions on functional forms, one cannot exclude the possibility that there may be cases where non-cooperation implies too little experimentation. A sufficient condition for when the first two terms outweigh the informational externality is $\frac{N}{N+1} \geq u'(\frac{R}{N})/u'(\frac{R}{N+1})$; see Proposition 5(b). Moreover, section 3.5 highlights how the non-cooperative expansion of the set of safe consumption possibilities is inefficiently large for the illustrative example. Nevertheless, experimentation is still "cautious" in the sense that it does not trigger the regime shift with probability 1.

Clearly, for a given s and δ^{-i} there need not be an interior solution. When the gain from expanding the set of safe consumption values is small, but the threat of triggering the regime shift is large, it may be individually rational to choose $\delta^i = 0$. Conversely, when the gain from expanding the set of safe consumption values is large and/or it is unlikely that there is a regime shift, it may be individually rational to choose $\delta^i = R - s - \delta^{-i}$.

For a symmetric step size $\delta^{-i} = (N-1)\delta^i$, we can write equation (11) as follows:

$$\phi'(\delta^{nc}; s) = u'\Big(\frac{s}{N} + \delta^{nc}\Big) + \beta\left[L_s'(N\delta^{nc})\frac{u\big(\frac{s}{N} + \delta^{nc}\big)}{1-\beta} + \frac{1}{N}L_s(N\delta^{nc})\frac{u'\big(\frac{s}{N} + \delta^{nc}\big)}{1-\beta}\right] = 0 \quad (12)$$

Let $g^{nc}(s)$ be the individual symmetric interior non-cooperative expansion. It is implicitly defined by $\phi'(\delta^{nc}; s) = 0$. Noting the similarity of (12) to (7) when replacing δ^* with $N\delta^{nc}$, it is possible to show that $g^{nc}(s)$ is decreasing in s. We can therefore define $\overline{s}^{nc}$, the smallest member of the set S^{nc}, by $g^{nc}(\overline{s}^{nc}) = 0$. In other words, for $s \geq \overline{s}^{nc}$, the threat of triggering a disastrous regime shift is sufficiently large so that the agents find it in their own best interest to stay at s when all other agents do so, too. Conversely, we can define the value $\underline{s}^{nc}$ by the other corner solution $g^{nc}(\underline{s}^{nc}) = \frac{R-s}{N}$. In other words, for $s \leq \underline{s}^{nc}$, the threat of triggering a regime shift is so small compared to the gains from increasing one's own consumption that it is individually rational to use the resource up to its maximal capacity R.

To sum up, in the non-cooperative game when the location of T is unknown, there is a "cautious equilibrium" that is described by the following set of Markov-strategies:

$$\delta^{nc}(s) = \begin{cases} \dfrac{R-s}{N} & \text{if} \quad s \leq \underline{s}^{nc} & (13a) \\ g^{nc}(s) & \text{if} \quad s \in (\underline{s}^{nc}, \overline{s}^{nc}) & (13b) \\ 0 & \text{if} \quad s \geq \overline{s}^{nc} & (13c) \end{cases}$$

Figure 2 illustrates the aggregate expansion of the set of safe consumption possibilities in the cautious equilibrium and contrasts it with the optimal expansion of a sole-owner.

In short, the game has the structure of a coordination problem. Clearly, the "cautious equilibrium" Pareto-dominates the "aggressive equilibrium".[8] Without strategic uncertainty, the cautious equilibrium would thus be the outcome of the game. But what happens when the agents are uncertain about the other agents' behavior? As the disastrous regime shift is irreversible, there is no room for dynamic processes that lead agents to select the Pareto-dominant equilibrium (Kim, 1996). Therefore, I turn to the static concept of risk-dominance (Harsanyi and Selten, 1988).

Since the game is symmetric, applying the criterion of risk-dominance has the following intuitive interpretation: The cautious equilibrium is selected if the expected payoff from playing cautiously exceeds the expected payoff from playing aggressively when agent i assigns probability p to the other agents playing aggressively. Whether the cautious or the aggressive equilibrium is risk-dominant depends both on this probability p as well as on the safe value s. We can, for a given safe value s, solve for the probability p^* at which agent i is just indifferent between playing cautiously or aggressively:

[8]This follows immediately from the fact that, by definition, $\delta^{nc}(s)$ is the interior solution to the symmetric maximization problem (9) (with $\delta^{-i} = (N-1)\delta^{nc}$) where the policy $\delta(s) = R - s$ was an admissible candidate.

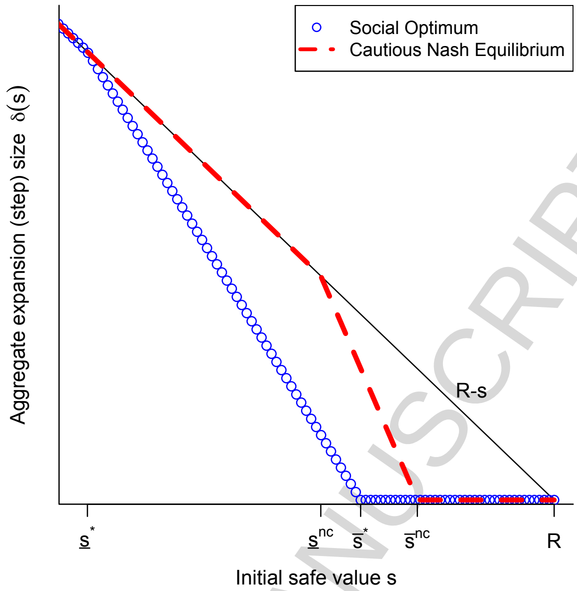

Figure 2: Illustration of policy function $\delta(s)$. The blue circles represent the optimal expansion δ of the safe consumption set s (on the y-axis) as a function of the safe consumption set (on the x-axis) when N=1 (where obviously $s \leq R$ and $\delta \in [0, R-s]$). For values of s below $\underline{s}^*$, it is optimal to consume the entire resource (choose $\delta(s) = R-s$). For values of s above $\overline{s}^*$, it is optimal to remain standing (choose $\delta(s) = 0$). The red dashed line plots the cautious non-cooperative equilibrium, showing how $\underline{s}^* \leq \underline{s}^{nc}$ and $\overline{s}^* \leq \overline{s}^{nc}$ (in some cases we may even have $\underline{s}^{nc} < \overline{s}^*$). It illustrates how even the "cautious" experimentation under non-cooperation implies excessive risk-taking. The figure also shows that the non-cooperative outcome may coincide with the sole-owner's choice for very low and high values of s.

$$p^* \cdot \pi_{[all\ aggressive]} + (1-p^*) \cdot \pi_{[only\ i\ aggressive]} = p^* \cdot \pi_{[only\ i\ cautious]} + (1-p^*) \cdot \pi_{[all\ cautious]}$$
$$\Leftrightarrow$$
$$p^* = \frac{\pi_{[all\ cautious]} - \pi_{[only\ i\ aggressive]}}{(\pi_{[all\ cautious]} - \pi_{[only\ i\ aggressive]}) - (\pi_{[only\ i\ cautious]} - \pi_{[all\ aggressive]})}$$

In the above calculation, $\pi_{[all\ aggressive]}$ refers to the payoff of playing aggressive when all other agents play aggressively, $\pi_{[only\ i\ aggressive]}$ refers to the payoff of playing aggressive when all other agents play cautiously, etc. In order to explicitly solve for the value of p^*, we need to put more structure on the problem. For the specific example developed in section 3.5 below, we can calculate and plot p^* as a function of s (see Figure 3). The grey area below the line drawn by p^* shows the set of values for which agent i prefers to play cautiously. Figure 3 illustrates how robust the cautious equilibrium is in this example: Even when the agents think that there is a 50% chance that all other agents play the aggressive strategy, it still pays to play cautiously for a wide range of initial values s. (Clearly, p^* is not defined for $s < \underline{s}^{nc}$ when the cautious and the aggressive equilibrium coincide.)

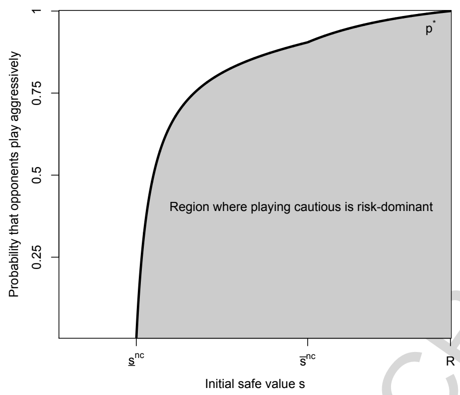

Figure 3: The black line plots p^* as a function of s for $u(c) = \sqrt{c}$, $f = \frac{1}{A}$ and $\beta = 0.8$, $A = R = 1$ and $N = 10$. It shows, for a given value of s the maximum value that agent i can assign to the probability that all other agents play aggressively and still prefer to play cautiously.

3.4 Comparative statics

In this section, I analyze how the consumption pattern in the cautious equilibrium shifts with changes in the parameters. Recall that g^{nc} is defined as the interior solution $\phi' = 0$ where ϕ' is given by (12). The effect of an increase in a parameter a in the interior range $s \in (\underline{s}^{nc}, \overline{s}^{nc})$ is given by $\frac{\mathrm{d}g^{nc}}{\mathrm{d}a} = -\frac{\partial\phi'/\partial a}{\partial\phi'/\partial g^{nc}}$. Further, recall that I assume that the second-order condition holds for $s \in (\underline{s}^{nc}, \overline{s}^{nc})$. Thus, to show that aggregate experimentation (the total expansion of the set of safe consumption values) is larger the higher the parameter a, it is sufficient to show that $\frac{\partial\phi'}{\partial a} > 0$ (since the second-order condition implies that $\frac{\partial\phi'}{\partial g^{nc}} < 0$). Because g^{nc} is monotonically decreasing in s, it is also sufficient to show that, for a given value of R, neither boundary $\underline{s}^{nc}$ or $\overline{s}^{nc}$ decreases and at least one boundary increases with a. The reason is that for a given value of R, an upward shift of $\underline{s}^{nc}$ or $\overline{s}^{nc}$ (and no downward shift of the respective other boundary) implies that all new values of g^{nc} must lie above the old values of g^{nc} (see Figure 2).

Proposition 5 summarizes the comparative statics results with respect to β, N, R and the agent's prior belief about the location of the threshold.

Proposition 5.

(a) *The boundaries $\underline{s}^{nc}$ and $\overline{s}^{nc}$, and aggregate experimentation in the cautious equilibrium, Ng^{nc}, decrease with β.*

(b) *A sufficient condition for aggregate experimentation to increase with N is that $\frac{N}{N+1} \geq u'(\frac{R}{N})/u'(\frac{R}{N+1})$.*

(c) *The more likely the regime shift (in terms of a first-order stochastic dominance), the larger the range where a separate cautious Nash-equilibrium exists and the lower aggregate experimentation.*

(d) *An increase of R to $\tilde{R}$ for an unchanged risk of the regime shift (i.e. $R < \tilde{R} \leq A$) decreases $\underline{s}^{nc}$ and leads to a larger range where a separate cautious equilibrium exists.*

Proof. The proofs are given in Appendix A.5. □

The first comparative static result conforms with basic intuition: The more patient the agents are, the more they value the annuity of staying at s, and the more cautious they are.

The second result provides a sufficient condition for when an increase in the number of agents exacerbates the "tragedy of the commons" in terms of aggregate experimentation. As discussed in relation to equation (11) above, there are three effects that an increase in N has on a given agent's incentives to expand the set of safe consumption values: First, a larger N implies that the marginal utility from a larger δ^i today increases. Second, a larger N means that the cost of an experiment in terms of an increased regime shift risk are diluted. Third, also the gain in marginal utility from an experiment that did not trigger the regime shift is shared among more agents. While the first two effects push towards a larger expansion, the last effect pulls in the other direction. When $\frac{N}{N+1} \geq u'(\frac{R}{N})/u'(\frac{R}{N+1})$, it is guaranteed that the first two effects dominate. Technically, this is shown by arguing that the range where a separate cautious equilibrium exists must shrink when $\frac{N}{N+1} \geq u'(\frac{R}{N})/u'(\frac{R}{N+1})$.

The third comparative static result also conforms with basic intuition: The more dangerous any step is, the more cautiously the agents experiment.

The last comparative statics result highlights the difference to the situation when the location of the threshold is known with certainty. In that situation, an increase in R leads to an increase in T_c, which shrinks the range in which the socially optimal outcome is a Nash equilibrium (Proposition 1). Here, immediate depletion is not necessarily the dominant strategy. An increase in R essentially means that the scope for an interior solution is widened so that the range for which immediate depletion is the only Nash equilibrium shrinks.

3.5 Specific example

For a given utility function and a given probability distribution of the threshold's location it is possible to solve for $\delta^*(s)$, $\delta^{nc}(s)$ and calculate the value function $V(s)$. To obtain closed form solutions, I assume that $u(c) = \sqrt{c}$ and that the agents believe that every value in $[0, A]$ is equally likely to be the threshold, i.e. $f = \frac{1}{A}$, and accordingly $L_s(\delta) = \frac{A-s-\delta}{A-s}$.

I first define the first-best. The problem of maximizing the sum of agent's utilities is:

$$\max_{\delta^i} \quad \sum_{i=1}^{N} \left\{ \sqrt{\frac{s}{N} + \delta^i} \; + \beta \frac{A - s - \sum \delta^i}{A - s} \cdot \frac{\sqrt{\frac{s}{N} + \delta^i}}{1 - \beta} \right\}$$

Because the agents are assumed to be identical, we can write the optimal total expansion of the set of safe consumption possibilities as:

$$\sum \delta^i = N\delta^* = \frac{A - (1 + 2\beta)s}{3\beta}$$

Note that in this specific example, the socially optimal experimentation is invariant to N, i.e.

it is the same as the optimal experimentation of a sole-owner. Clearly, δ^* is decreasing in β and s. There will only be an interior solution to (7) when $s \in [\underline{s}^*, \overline{s}^*]$. We have:[9]

$$\underline{s}^* = \max\left\{0, \frac{A - 3\beta R}{(1-\beta)}\right\}$$

$$\overline{s}^* = \min\left\{\frac{A}{1+2\beta}, R\right\}$$

Let us now consider the cautious equilibrium of the non-cooperative game. Solving (12) for the interior equilibrium g^{nc}, we have that total non-cooperative expansion is given by:

$$N\delta^{nc}(s) = \begin{cases} R - s & \text{if} \quad s \leq \underline{s}^{nc} \\ Ng^{nc} = \dfrac{\big((1-\beta)N + \beta\big)A - \big((1-\beta)N + 3\beta\big)s}{3\beta} & \text{if} \quad s \in (\underline{s}^{nc}, \overline{s}^{nc}) \\ 0 & \text{if} \quad s \geq \overline{s}^{nc} \end{cases}$$

where

$$\underline{s}^{nc} = \max\left\{0, \frac{\big((1-\beta)N + \beta\big)A - 3\beta R}{(1-\beta)N}\right\}$$

$$\overline{s}^{nc} = \min\left\{\frac{\big((1-\beta)N + \beta\big)A}{(1-\beta)N + 3\beta}, R\right\}$$

The closed form solutions make it easy to confirm the comparative statics results.

First, an increase in the discount factor implies that each agent is more patient and values the preservation of the resource for future consumption more. Thus, the boundaries $\underline{s}^{nc}$ and $\overline{s}^{nc}$ and aggregate experimentation decrease with β, which can be readily confirmed by noting that the denominator of $\frac{\partial g^{nc}}{\partial \beta}$ is $3N(s - A)$, which is negative because $s \leq A$.

The second comparative static result provides a sufficient condition for when an increase in N leads to more experimentation. This condition is not fulfilled for this specific example as $\frac{N}{N+1} < u'(\frac{R}{N})/u'(\frac{R}{N+1}) = \sqrt{\frac{N}{N+1}}$. However, $\frac{N}{N+1} < u'(\frac{R}{N})/u'(\frac{R}{N+1})$ is not a necessary condition. In fact, it is straightforward to check for this specific example that $\frac{\partial N g^{nc}}{\partial N} = g^{nc} + N\frac{1-\beta}{3\beta}(A - s) > 0$ (as $A \geq s$ by definition). In other words, any non-cooperative experimentation will be inefficiently large, and more so the larger N.

Third, a monotonic increase in the risk of a regime shift is here equivalent to a decrease in A (say, from A to $\hat{A}$). Clearly, when R is unchanged (so that $R < A$ and $R \leq \hat{A}$ when $\hat{A} < A$), a lower A means a lower $\underline{s}^{nc}$, a lower $\overline{s}^{nc}$, and a decreased expansion of the set of safe consumption values when $s \in (\underline{s}^{nc}, \overline{s}^{nc})$ because $((1-\beta)N + \beta) > 0$ for $N \geq 1$. The same holds when $A = R$ and $\hat{R} = \hat{A}$ with $\hat{A} < A$, because R only appears in the condition for $\underline{s}^{nc}$

[9] At $\underline{s}^*$ it is optimal to consume the entire resource, so that $\underline{s}^*$ is found by solving $R - \underline{s}^* = \frac{A-(1+2\beta)\underline{s}^*}{3\beta}$. At $\overline{s}^*$ it is optimal to remain standing, so that $\overline{s}^*$ is found by solving $0 = \frac{A-(1+2\beta)\overline{s}^*}{3\beta}$.

and $\underline{s}^{nc} > 0 \Leftrightarrow (1-\beta)N - 2\beta > 0$.

Finally, when R increases to $\tilde{R}$ but A remains unchanged (so that $R < A$ and $\tilde{R} \leq A$), it only has an effect on $\underline{s}^{nc}$ (provided that $\bar{s}^{nc} < R$). Provided that $\underline{s}^{nc} > 0$, it is plain to see that $\frac{\partial \underline{s}^{nc}}{\partial R} = -\frac{3\beta}{(1-\beta)N} < 0$. The range where a separate cautious equilibrium exists is larger.

Figure 4 plots the value function of a given agent for a uniform prior (with $A = R = 1$) and a discount factor of $\beta = 0.8$, illustrating how it changes as the number of agents increases. The more agents there are, the greater the distance of the non-cooperative value function (plotted by the red dashed line) to the socially optimal value function (plotted by the blue open circles). In particular when $N = 10$, one sees the region (roughly from $s = 0$ to $s = 0.2$) where there is no separate cautious equilibrium. Furthermore, the value when it becomes individually rational to remain standing, $\bar{s}^{nc}$, is relatively large (roughly 0.62). All in all however, this example shows that the threat of a irreversible regime shift is very effective when the externality applies only to the risk of crossing the threshold.[10] In particular, for $s \geq \bar{s}^{nc}$, the cautious equilibrium coincides with the social optimum.

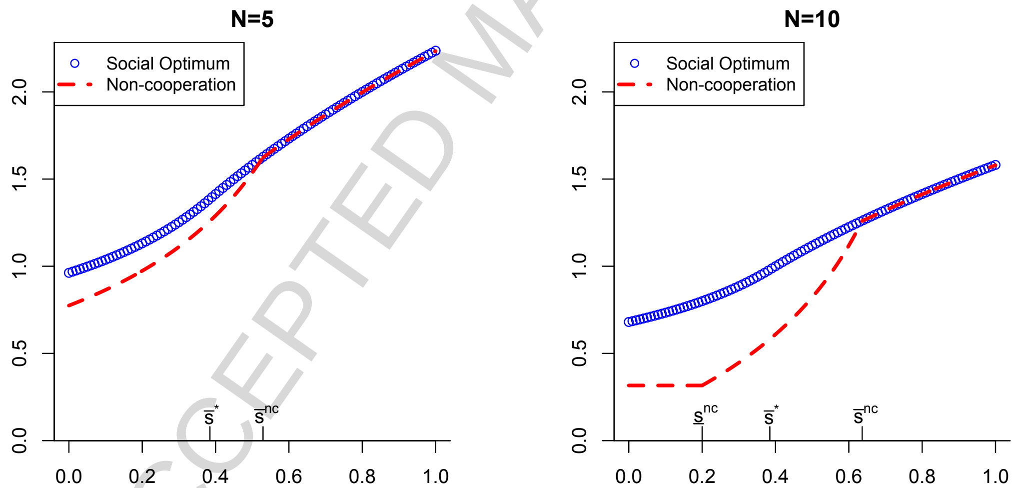

Figure 4: Illustration of the value a given agent derives from the socially optimal and non-cooperative use of the resource with $u = \sqrt{\frac{s}{N} + \delta^i}$, $\beta = 0.8$, and A=R=1 for N=5 and N=10. Note that the individual value, also in the social optimum, is lower when N=10 than N=5, simply because the resource is shared among more agents.

4 Extensions

The paper's main results do not rely on specific functional forms for utility or the probability distribution of the threshold's location. Tractability is achieved by considering extremely simple resource dynamics, namely the resource remains intact and replenishes fully in the next period as long as resource use in the current period has not exceeded T. In other words, there is no common-pool externality relating to the resource dynamics itself. In this part of the

[10]At least for this specific utility function and these parameter values. Note that $\beta = 0.8$ implies a unreasonably high discount rate, but it was chosen to magnify the effect of non-cooperation for a small number of agents.

paper, I explore to what extent the main results are robust to more general resource dynamics (section 4.1 and 4.2). Moreover, I show that the result of once-and-for-all experimentation does not rely on the assumption that the regime shift occurs immediately if the threshold is crossed (section 4.3). However, once-and-for-all experimentation may no longer be optimal when the regime shift is not disastrous in the sense that the post-event value function depends on the extent to which the threshold has been overstepped (section 4.4).

4.1 Growing R and constraints on the choice set

In this subsection, I shall relax the assumption that R is constant. Instead, I consider the case that R grows according to some function $G(R)$. However, I maintain the assumption that f and A are fixed, so that R_t converges to some $R_\infty \leq A$ with time. This situation may mechanically lead to repeated experimentation as long as the upper bound of the feasible choice set in the respective period is binding. As a consequence, the overall consumption plan will be more cautious.

Formally, the resource dynamics can be expressed as:

$$R_{t+1} = \begin{cases} R_t + G(R_t) & \text{if} \quad \sum_i c_t^i \leq T \\ 0 & \text{if} \quad \sum_i c_t^i > T \quad \text{or } R_t = 0 \end{cases} \tag{15}$$

where $G(R) > 0$ for $R \in (0, R_\infty)$, $G(R_\infty) = 0$, and $R_0 < R_\infty \leq A$.

Let us first consider the social optimum. As the set of safe values at which no more experimentation is optimal, S, depends only on the belief about the location of the threshold F and not on R, it will be optimal to expand the set of safe consumption values as long as $R_t \notin S$. Specifically, in this initial phase, it will be optimal to choose $\delta_t^* = R_t - s_{t-1}$. Once $R_t \in S$ for some t (say at $t = \tau$), it will be optimal to choose one last step $\delta^*(s_\tau)$ (which may be of size zero) and to remain at $s_{\tau+1} = s_\tau + \delta^*(s_\tau)$ for all remaining time.

Note that constraints on the choice set (such that $\delta \in [0, \delta_{max}]$ where $\delta_{max} < R - s_t$ for some period $t = 0, ..., \tau$) will lead to repeated experimentation for the same mechanistic reason: When the first-best unconstrained expansion is $\delta^*(s_0)$, but δ_{max} is such that it requires several steps to traverse this distance, then the safe value s will be updated sequentially (conditional on not causing the regime shift, of course). The optimal plan prescribes choosing δ_{max} for some period $t = 0, ..., \tau$ and then choose $\delta^*(s_\tau)$. Because $\frac{\partial \delta^*(s)}{\partial s} < 0$ (Proposition 3), this implies an overall more cautious plan (that is: $\sum_{t=0}^{\tau} \delta_{max} + \delta^*(s_\tau) < \delta^*(s_0)$).

The exact same reasoning applies in the non-cooperative game. Given that the agents coordinate on the cautious equilibrium, the set S^{nc} does not depend on R. Consequently, the equilibrium path prescribes choosing $\delta_t^{nc} = \frac{R_t - s_{t-1}}{N}$ for some period $t = 0, ..., \tau$ and then staying at $s_{\tau+1}$ for the remaining time (even though R_t may continue to grow). Note that this rests on the assumption that all agents rationally anticipate the evolution of R_t. Analyzing the effect of uncertainty about $G(R)$ could be very interesting, but is left for future work. That said, even with perfect knowledge about $G(R)$, strategic uncertainty will matter a lot in the real world. Given that a real-world agent knows that the incentive to grab is increasing through time and he or she is uncertain whether the other agents will actually stick to the cooperative choice, he or she will have strong incentives to pre-empt the other agents.

The discussion in this subsection highlights how the result of once-and-for-all experimentation is linked to the assumption of an unconstrained choice set $\delta \in [0, R-s]$. The discussion further sheds light on the difference of this model to e.g. the climate change application of Lemoine and Traeger (2014): One reason for the gradual approach in their model is that their assumed capital dynamics implicitly translate into a constrained choice set (as it is very reasonable in their setting, capital cannot be adjusted instantly and costlessly).

4.2 Non-renewable resource dynamics

So far, I have assumed that the resource replenishes fully every period unless the threshold has been crossed. In other words, the externality was only related to the risk of crossing the threshold. Here, I study the opposite case of a non-renewable resource to analyze the effect of a disastrous regime shift when the externality relates both to the risk of crossing the threshold and to the resource itself. Specifically, I consider the following model of extraction from a known stock of a non-renewable resource:

$$\max_{c_t^i} \sum_{t=0}^{\infty} \beta^t u(c_t^i) \quad \text{subject to: } R_{t+1} = \begin{cases} R_t - \sum_i c_t^i & \text{if } \sum_i c_t^i \leq T \\ 0 & \text{if } \sum_i c_t^i > T \end{cases} ; \quad R_0 \text{ given} \tag{16}$$

A simple interpretation of this model could be a mine from which several agents extract a valuable resource. If aggregate extraction is too high in a given period, the structure of the shafts may collapse, making the remainder of the resource inaccessible.[11]

I assume that the utility function is of such a form, that in a world without the threshold, there is a non-cooperative equilibrium in which positive extraction occurs in several periods. Due to discounting, it is clear that the extraction level will decline as time passes, both in the social optimum and in the non-cooperative equilibrium. Due to the stock externality, it is clear that the extraction rate in the non-cooperative equilibrium is inefficiently large (see e.g. Harstad and Liski, 2013). To introduce some notation, let $\tilde{c}^{nc}(R_t)$ be the total non-cooperative extraction level (as a function of the resource stock R_t) in absence of the regime shift risk.

The threat of a regime shift has the potential to limit non-cooperative extraction below $\tilde{c}^{nc}(R_t)$ and thereby improve welfare. Also in the case of a non-renewable resource, with dynamics given by (16), it is possible to show that experimentation continues to exhibit once-and-for-all dynamics.[12] Thus, for a given value s_0, agents will either experiment once to learn

[11] Granted, in spite of this natural interpretation, two things are peculiar about this model setup: First, any player can extract any amount up to R_t. (The option to introduce a capacity constraint on current extraction – though realistic – would come at the cost of significant clutter without yielding any apparent benefit.) Second, the assumption that R_0 is known and that T is constant means that this is *not* a problem of eating a cake of unknown size. This problem has since long been dealt with in the literature (see e.g. Kemp, 1976; Hoel, 1978) and is not considered here.

[12] The proof is omitted because it follows the same steps as the proof of Proposition 4. In particular, the argument that agent i's payoff is higher when expanding the set of safe values in one step rather than two

whether they can safely extract $s_1^{nc} = s_0 + N\delta^{nc}(s_0)$, or they will not experiment at all. As the extraction path will decline, s_1^{nc} will only be binding on the per-period extraction for an initial phase (say until $t = \tau$, where τ is implicitly defined by $\tilde{c}^{nc}(R_\tau) = s_1^{nc}$. After time period τ, the extraction path will follow the same non-cooperative path as in the absence of a regime shift.

Consequently, the threat of a regime shift cannot induce coordination on the intertemporal first-best.[13] Nevertheless, the threat of a regime shift may induce an equilibrium that is, in expected terms, Pareto-superior to a situation without the regime shift risk: While the agents would obviously be better off without the regime shift risk when the initial experiment triggers the collapse, the agents benefit from the initial constraint on the per-period extraction (up to $t = \tau$) when the regime shift is not triggered by the initial experiment.

4.3 Delay in the occurrence of the regime shift

In this subsection, I depart from the assumption that all agents observe immediately whether last period's expansion has triggered the regime shift or not, but consider a situation where the agents observe only with some probability whether they have crossed the threshold. In fact, it is not unreasonable to model that the true state will manifest itself only after some delay. For example, the process of saltwater intrusion, though irreversible once the water table has fallen under a critical level, may take time to unfold (for a recent paper that focusses on this effect in the context of optimal climate policies, see Gerlagh and Liski (2014)). Hence, as time passes the agents will update their beliefs about whether the threshold has been located on the interval $[s_t, s_t + \delta_t]$. How does this learning affect the optimal and non-cooperative strategies? This becomes an extremely difficult question as the problem is no longer Markovian.

Nevertheless, it is possible to show that also when crossing the threshold at time t triggers the regime shift at some (potentially uncertain) time $\tau > t$, it is still socially and individually rational to experiment – if at all – in the first period only. The key is to realize that yesterday's decisions are exogenous today. This means that threat of a regime shift can be modeled as an exogenous hazard rate: Let h_t be the probability that the regime shift, triggered by events earlier than and including time t, occurs at time t (conditional on not having occurred prior to t, of course). The agent's problem in this situation can be formulated as:

$$V^i(s, \Delta^{-i}) = \max_{\delta^i \in [0, R-s]} \left\{ u(s + \delta^i) + (1 - h_t)\beta L_s(\delta^i + \delta^{-i}) V^i(s + \delta, \Delta^{-i}) \right\} \tag{17}$$

The structure of (17) is identical to the one in equation (9), only the effective discount factor decreases by $(1 - h_t)$. Agents anticipate how the effective discount factor changes with time, but their belief about the location of the threshold, given that it has not been crossed by the

rests on a comparison of the risk encountered by the two strategies and that staying after the first expansion is sub-optimal for agent i. Neither of these two arguments use the specific form of the continuation value (it is irrelevant when comparing the regime-shift risk and held constant in the second argument).

[13] Except in the very special case that $s_1^{nc} = s_0 = c^*(R_\mathcal{T})$ where $c^*(R_\mathcal{T})$ is the socially optimal extraction at the finite exhaustion period $\mathcal{T}$

current step, is not affected. Thus, the learning dynamics are unchanged.

In other words, the once-and-for-all dynamics of experimentation are robust to a delay in the occurrence of the regime-shift. This does of course not imply that the optimal decision under the two different models will be the same. It will almost surely differ, as delaying the consequences of crossing the threshold decreases the costs of experimentation. Yet, as the agents only learn that they have crossed the threshold when the disastrous regime shift actually occurs, they cannot capitalize on this delay by trying to expand the set of safe consumption possibilities several times.

4.4 Non-disastrous regime shift

A central feature of the baseline model was that the regime shift is disastrous: Crossing the threshold breaks any links between the state before and after the regime shift. The pre-event choices did not matter for the post-event value. This structure allowed me to simply normalize the continuation value in case of a regime shift to zero. For some applications, this independence of the post-event value is a fitting description. However, when the system under consideration is large, and the threshold effect on the damage is not truly catastrophic, but just one of many parts in the equation, a model with independent post-event value is not adequate. In such a setting, one would need to take into account how the continuation value depends on how far the set of consumption values has been expanded before the regime shift.

Denote the function that captures how the post-event continuation value depends on the pre-event expansion by $W(s_{t+1})$ (where $s_{t+1} = s_t + \delta_t$). How W would depend on the pre-event values of the state s_t and the choice variable δ_t is not generally clear. For example, Ren and Polasky (2014) discuss under which conditions regime-shift risk implies more cautionary or more aggressive management of renewable resources. In particular, they highlight the role of an "investment effect" that induces incentives for more aggressive management: Harvesting less (investing in the renewable resource stock) pays off badly should the regime shift occur. Ren and Polasky go on to show how these incentives are balanced (and potentially overturned) by the "risk reduction effect" and a "consumption smoothing effect" (that leads to more precaution in their application). Similarly, the capital stock in a climate change application likely has an ambiguous effect (Ploeg and Zeeuw, 2015a). On the one hand, it buffers against the adverse effects of the regime shift and hence smoothes consumption over regimes. On the other hand, a higher capital stock implies more intense use of fossil fuels, which aggravates climate damages.

Regardless of whether $W'(s_{t+1}) > 0$ or $W'(s_{t+1}) < 0$, the fact that the pre-event choices matter for the post-event value means that it is no longer immaterial by how much one has stepped over the threshold. I argue that 1) even when the regime shift is not disastrous, there will still be a set S or S^{nc} at which it is socially or individually rational to not experiment further, and 2) I point out that a necessary condition for a gradual approach to S is that the post-event value declines sufficiently strongly in s_{t+1}. As the analysis of the different forces at play is the same for the non-cooperative game, I concentrate on the sole-owner case for the general discussion below. With help of the concrete example, I then explicitly compare optimal experimentation in the absence of strategic interaction to the "cautious" equilibrium of the non-cooperative game.

To put some structure to the argument, I write down the Bellman equation of the sole owner before the regime shift has occurred:

$$V(s) = \max_{\delta \in [0, R-s]} \big\{ u(s+\delta) + \beta \big[L_s(\delta) \cdot V(s+\delta) + (1 - L_s(\delta)) \cdot W(s+\delta) \big] \big\} \tag{18}$$

The sole-owner seeks to choose that expansion of the set of consumption values that maximizes her current utility plus the discounted continuation value. With probability $L_s(\delta)$, the step of size δ turns out to be safe and the continuation value is given by $V(s+\delta)$. With probability $(1 - L_s(\delta))$, the threshold is located on the interval between s and $s+\delta$ and the continuation value is given by $W(s+\delta)$.

It will still be the case that there is a non-empty set S for which the optimal choice is $\delta^* = 0$. The reason is that, as long as the regime shift is a negative event ($V(s) > W(s)$), the gains from further expansion of the set of safe consumption values are bounded above, while the risk of triggering the regime shift grows exceedingly large as $s \to R$.

To obtain more insights about how the optimal expansion choice is changed by the existence of an endogenous continuation value, consider the derivative of the RHS of (18) (denoting this function by φ' again should not cause confusion):

$$\varphi' = u' + \beta \big[L_s'(\delta)(V - W) + L_s(\delta) \cdot V' + (1 - L_s(\delta)) \cdot W' \big] = 0 \tag{19}$$

The size of δ^* that solves equation (19) is determined by three factors: First, there is the gain in marginal utility u'. Second, there is the term involving $L_s'(\delta)$ which captures the increased risk of the regime shift. This term is negative as before. Third, the optimal choice of δ is affected by the marginal continuation value. Previously, only the event of not crossing the threshold mattered here. Now, the event of crossing the threshold also has to be evaluated explicitly.

Analyzing how an endogenous post-event value affects δ^*, I first note that the negative term $L_s'(\delta)(V - W)$ decreases in absolute value. When the second-order condition for an interior solution is satisfied, φ' is a decreasing function in the neighborhood of δ^*. A decrease in absolute value of the term $L_s'(\delta)(V - W)$ shifts the function upwards. Intuitively, a non-zero continuation value in case of a regime shift pushes for a larger current consumption. However, when $W' < 0$, the term $(1 - L_s(\delta)) \cdot W'$ is negative, which, *ceteris paribus*, leads to a lower value of δ^*. If, and only if, the post-event value declines sufficiently strongly in s_{t+1} will the optimal δ^* be so small that for $s_t \notin S$, we have $s_t + \delta_t^* = s_{t+1} \notin S$. The approach to S will then be gradual, implying periods of repeated experimentation.

To illustrate more concretely how these effects play out, I assume specific functional forms. As in section 3.5, I set $u(c) = \sqrt{c}$ and assume a uniform distribution for the location of the threshold so that $L_s(\delta) = \frac{A-s-\delta}{A-s}$ with $A = R$. For the post-event continuation value, I assume that the resource loses all its productivity once the regime shift occurs. In other words, W is the highest value that can be obtained when spreading the consumption of the now non-renewable resource $r_t = R - s_t - \delta_t$ over the remaining time horizon. We have $W^*(s_t + \delta_t) = \sqrt{\frac{R-(s_t+\delta_t)}{1-\beta^2}}$. Clearly, $W'(s_t + \delta_t) < 0$.

For a square-root utility function and without exogenous constraints on extraction, the number of agents cannot be too large in order to have an interior equilibrium with positive extraction over the entire time path in the non-cooperative game. Here I choose N=2. We have $W^{nc}(s_t + \delta_t^i + \delta_t^{-i}) = \sqrt{\frac{R-(s_t+\delta_t^i+\delta_t^{-i})}{1-\beta^2+\sqrt{1-\beta^2}}}$. The analytic closed form solutions are not particularly instructive. Instead, I present the results graphically.

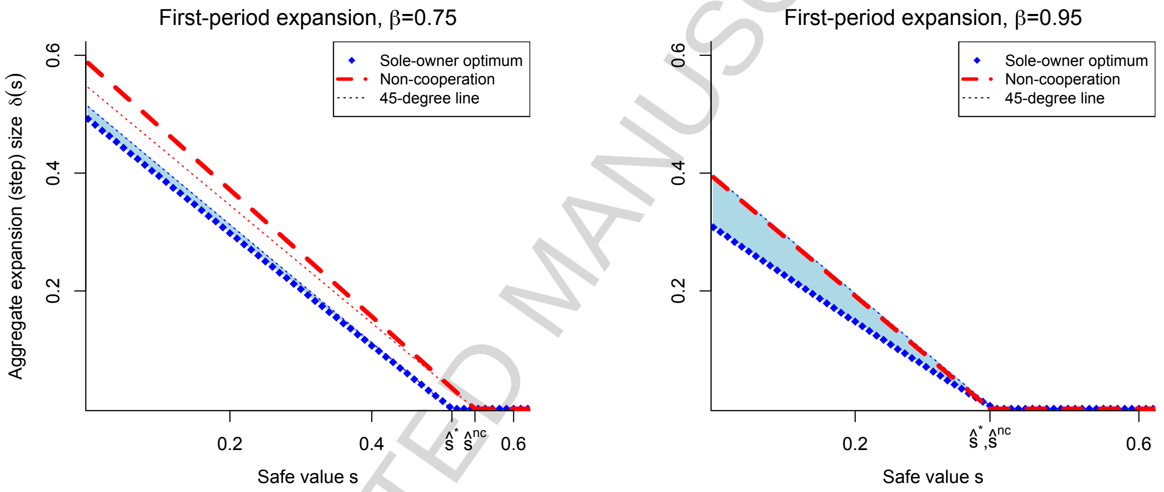

Figure 5: Illustration of optimal experimentation when pre-event choices matter for post-event value. Parameters and functional forms: $u(c) = \sqrt{c}$, N=2, A=R=1, $L_s(\delta) = \frac{1-s-\delta}{1-s}$; for $\beta = 0.75$ and β=0.95.

Figure 5 plots the first-period expansion δ as a function of s for a high and a low value of the discount factor. The blue dotted line shows the optimal expansion of a sole-owner and the light-blue area indicated, for a given s, by how much this first step falls short of the total step size (experimentation stops when $\hat{s}^*$ is reached). Note how the first experiment is a much larger fraction of the eventual area that is explored when $\beta = 0.75$ instead of $\beta = 0.95$.

The red dashed line shows the corresponding total expansion in the non-cooperative case. Note that here, no second step is taken – any experimentation is undertaken in the first period only. This is however not a general result: for very high values of β (0.998 and above; not shown here) the cautious non-cooperative equilibrium also implies repeated experimentation. Conversely, the lower is β, the steeper the first-period expansion as a function of s. Thus, non-cooperation has two effects: not only is the non-cooperative experimentation inefficiently large, the approach to the set at which it is optimal to cease experimenting is inefficiently fast. The latter aspect is caused by the fact that the sole-owner's marginal continuation value declines more steeply than its non-cooperative counterpart: The agents do not internalize the additional damage from the extent by which the threshold has been crossed.

5 Discussion and Conclusion

The threat of a disastrous regime shift can be beneficial because it allows coordination on a Pareto-dominant equilibrium, even when the location of the threshold is unknown. When the consequence of the regime shift is catastrophic, and learning is only affirmative, it is socially optimal and a Nash equilibrium to update beliefs only once (if at all). For a sufficiently valuable level of safe use, the threat of loosing the productive resource discourages any experimentation and effectively enforces the first-best consumption level. When the agents experiment, the expansion of the consumption set is in most cases inefficiently large. However, when the experiment has not triggered the regime shift, staying at the updated level is *ex post* socially optimal. In addition to this "cautious" Nash equilibrium there is always also an "aggressive" Nash equilibrium, in which the agents immediately deplete the resource. When the initial value of safe use is not valuable enough, immediate depletion will be the only equilibrium.

Empirically, we do not observe many resources where the safe level of use is updated once and the resource then either collapses or is used sustainably at the updated level. Whereas my model isolates the threat of a disastrous regime shift, many additional aspects that dilute the sharp once-and-for-all learning dynamics are likely to matter in the real world. In this paper, I explore the conditions under which the once-and-for-all learning dynamics and the existence of a "cautious" Nash equilibrium emerge. Importantly, I show that the threat of the disastrous regime shift loses importance when the externality applies both to the risk of triggering the regime shift and to the resource itself. Nevertheless, it can still act as a "commitment device" to at least dampen non-cooperative extraction.

These conclusions have been derived by developing a dynamic model that has placed only minimal requirements on the utility function and the probability distribution of the threshold. Nevertheless, there are a number of structural assumptions that warrant discussion.

First, a prominent aspect of this model is that the threshold itself is not stochastic. The central motivation is to isolate the effect of uncertainty about the threshold's location. This is arguably the core of the problem: We don't know which level of use triggers the regime shift. This assumption is consistent with Lemoine and Traeger (2014, p.28) who argue, "we would not actually expect tipping to be stochastic. Instead, any such stochasticity would serve to approximate a more complete model with uncertainty (and potentially learning) over the precise trigger mechanism underlying the tipping point." This being said, it would still be interesting to investigate how the choice between a hazard-rate formulation (as in Polasky et al., 2011 or Sakamoto, 2014) or a threshold formulation influences the outcome and policy conclusions in an otherwise identical model.

Second, to focus on the coordinating effect of the threat of the regime shift, I have assumed identical agents. One dimension along which players could differ is their valuation of the future. However, it is likely that a contract that gives a larger share of the gains to more impatient players could smooth out any such differences. One could also investigate the effect of heterogenous beliefs about the existence and location of the threshold. Agbo (2014) and Koulovatianos (2015) analyze this in the framework of Levhari and Mirman (1980). In the current set-up, such a heterogeneity could lead to interesting dynamics and possible multiple equilibria, where some players rationally do not want to learn about the probability distribution of T whereas other players do invest in experimentation. Another dimension along which players could differ is their size or the degree to which they depend on the resource. As larger players are likely to internalize a larger part of the externality than smaller players, different sets of equilibria may emerge. Especially in light of the discussions surrounding a possible climate treaty (Harstad, 2012; Nordhaus, 2015), it is topical to analyze situations where groups of players can form a coalition to ameliorate the negative effects of non-cooperation.

Third, I have assumed the regime shift to be irreversible. This is obviously a considerable simplification. Groeneveld et al. (2013) have analyzed the problem of how a sole-owner would learn about the location of a threshold in a setting where repeated crossings are allowed, but the exact location of the threshold remains unknown upon crossing it. If one presumes that crossing the threshold implies that one learns where it is, the game turns into a repeated game. This may imply that cooperation is sustainable for sufficiently patient agents (van Damme, 1989). However, there could also be cases where irreversibility emerges "endogenously" when it is possible – but not an equilibrium – to move out of a non-productive regime. The tractability of the current modeling approach may prove fruitful to further explore this issue.

A final, related, point is the fact that I have concentrated on Markovian strategies. When the agents are allowed to use history-dependent strategies, the threat of a threshold may allow them to coordinate on the social optimum in all phases of the game. They could simply agree on expanding the set of safe consumption possibilities by the socially optimal amount and threaten that if any agent steps too far, this triggers a reaction to deplete the resource in the next period. This obviously begs the question of renegotiation proofness, but it is plausible that a contract that is binding for two periods is already sufficient to achieve the first-best.

The threat of a disastrous regime shift is a very strong coordinating device. This is true irrespective of whether the threshold's location is known or unknown, because the agents learn only after the fact whether the disastrous regime shift has occurred or not. Would a catastrophic threshold lose its coordinating force when the agents can learn about its location without the risk of crossing it? Importantly, an extension of the model along these lines would speak to the recent debate on "early warning signals" (Scheffer et al., 2009; Boettiger and Hastings, 2013) and is the task of future work.

Acknowledgements

This research is funded by NorMER, a Nordic Centre of Excellence for Research on Marine Ecosystems and Resources under Climate Change. The constructive comments from two anonymous reviewers and the associate editor, Chris Knittel, have greatly improved the paper. Furthermore, I would like to thank Scott Barrett, Chris Costello, Bård Harstad, Svenn Jensen, Larry Karp, Matti Liski, Torben Mideksa, Eric Nævdal, and seminar participants at the Beijer Institute, Statistics Norway, in Oslo, Helsinki, Santa Barbara, and Berkeley for suggestions and discussion. Whitney Lyle provided valuable language assistance.

References

Aflaki, S. (2013). The effect of environmental uncertainty on the tragedy of the commons. *Games and Economic Behavior*, 82(0):240–253.

Agbo, M. (2014). Strategic exploitation with learning and heterogeneous beliefs. *Journal of Environmental Economics and Management*, 67(2):126–140.

Anderies, M. J., Janssen, A. M., and Walker, H. B. (2002). Grazing management, resilience, and the dynamics of a fire-driven rangeland system. *Ecosystems*, 5(1):23–44.

Barrett, S. (2013). Climate treaties and approaching catastrophes. *Journal of Environmental Economics and Management*, 66(2):235–250.

Barrett, S. and Dannenberg, A. (2012). Climate negotiations under scientific uncertainty. *Proceedings of the National Academy of Sciences*, 109(43):17372–17376.

Bochet, O., Laurent-Lucchetti, J., Leroux, J., and Sinclair-Desgagné, B. (2013). Collective Dangerous Behavior: Theory and Evidence on Risk-Taking. Research Papers by the Institute of Economics and Econometrics, Geneva School of Economics and Management, University of Geneva 13101, Institut d'Economie et Economtrie, Universit de Genve.

Boettiger, C. and Hastings, A. (2013). Tipping points: From patterns to predictions. *Nature*, 493(7431):157–158.

Bolton, P. and Harris, C. (1999). Strategic experimentation. *Econometrica*, 67(2):349–374.

Bonatti, A. and Hörner, J. (2015). Learning to disagree in a game of experimentation. Discussion paper, Cowles Foundation, Yale University.

Brozović, N. and Schlenker, W. (2011). Optimal management of an ecosystem with an unknown threshold. *Ecological Economics*, 70(4):627 – 640.

Costello, C. and Karp, L. (2004). Dynamic taxes and quotas with learning. *Journal of Economic Dynamics and Control*, 28(8):1661–1680.

Crépin, A.-S. and Lindahl, T. (2009). Grazing games: Sharing common property resources with complex dynamics. *Environmental and Resource Economics*, 44(1):29–46.

Cropper, M. (1976). Regulating activities with catastrophic environmental effects. *Journal of Environmental Economics and Management*, 3(1):1–15.

Dekker, S. C., Rietkerk, M., and Bierkens, M. F. P. (2007). Coupling microscale vegetationsoil water and macroscale vegetationprecipitation feedbacks in semiarid ecosystems. *Global Change Biology*, 13(3):671–678.

Dubins, L. E. and Savage, L. J. (1965). *How to gamble if you must : inequalities for stochastic processes*. McGraw-Hill, New York.

Feldmann, J. and Levermann, A. (2015). Collapse of the west antarctic ice sheet after local destabilization of the amundsen basin. *Proceedings of the National Academy of Sciences*, 112(46):14191–14196.

Fesselmeyer, E. and Santugini, M. (2013). Strategic exploitation of a common resource under environmental risk. *Journal of Economic Dynamics and Control*, 37(1):125–136.

Frank, K. T., Petrie, B., Choi, J. S., and Leggett, W. C. (2005). Trophic cascades in a formerly cod-dominated ecosystem. *Science*, 308(5728):1621–1623.

Gerlagh, R. and Liski, M. (2014). Carbon Prices for the Next Hundred Years. CESifo Working Paper Series 4671, CESifo Group Munich.

Groeneveld, R. A., Springborn, M., and Costello, C. (2013). Repeated experimentation to learn about a flow-pollutant threshold. *Environmental and Resource Economics*, forthcoming:1–21.

Harsanyi, J. C. and Selten, R. (1988). *A general theory of equilibrium selection in games*. MIT Press, Cambridge.

Harstad, B. (2012). Climate contracts: A game of emissions, investments, negotiations, and renegotiations. *The Review of Economic Studies*, 79(4):1527–1557.

Harstad, B. and Liski, M. (2013). Games and resources. In Shogren, J., editor, *Encyclopedia of Energy, Natural Resource and Environmental Economics*, volume 2, pages 299–308. Elsevier, Amsterdam.

Hjermann, D. Ø., Stenseth, N. C., and Ottersen, G. (2004). Competition among fishermen and fish causes the collapse of Barents Sea capelin. *Proceedings of the National Academy of Sciences*, 101:11679–11684.

Hoel, M. (1978). Resource extraction, uncertainty, and learning. *The Bell Journal of Economics*, 9(2):642–645.

Keller, G., Rady, S., and Cripps, M. (2005). Strategic experimentation with exponential bandits. *Econometrica*, 73(1):39–68.

Kemp, M. (1976). How to eat a cake of unknown size. In Kemp, M., editor, *Three Topics in the Theory of International Trade*, pages 297–308. North-Holland, Amsterdam.

Kim, Y. (1996). Equilibrium selection in n-person coordination games. *Games and Economic Behavior*, 15(2):203–227.

Kossioris, G., Plexousakis, M., Xepapadeas, A., de Zeeuw, A., and Mäler, K.-G. (2008). Feedback nash equilibria for non-linear differential games in pollution control. *Journal of Economic Dynamics and Control*, 32(4):1312–1331.

Koulovatianos, C. (2015). Strategic exploitation of a common-property resource under rational learning about its reproduction. *Dynamic Games and Applications*, 5(1):94–119.

Lemoine, D. and Traeger, C. (2014). Watch your step: Optimal policy in a tipping climate. *American Economic Journal: Economic Policy,*, forthcoming:1–47.

Lenton, T. M., Held, H., Kriegler, E., Hall, J. W., Lucht, W., Rahmstorf, S., and Schellnhuber, H. J. (2008). Tipping elements in the earth's climate system. *Proceedings of the National Academy of Sciences*, 105(6):1786–1793.

Levhari, D. and Mirman, L. J. (1980). The Great Fish War: An Example Using a Dynamic Cournot-Nash Solution. *The Bell Journal of Economics*, 11(1):322–334.

Maskin, E. and Tirole, J. (2001). Markov perfect equilibrium, i: Observable actions. *Journal of Economic Theory*, 100:191–219.

Miller, S. and Nkuiya, B. (2016). Coalition formation in fisheries with potential regime shift. *Journal of Environmental Economics and Management*, 79:189–207.

Mitra, T. and Roy, S. (2006). Optimal exploitation of renewable resources under uncertainty and the extinction of species. *Economic Theory*, 28(1):1–23.

Nævdal, E. (2003). Optimal regulation of natural resources in the presence of irreversible threshold effects. *Natural Resource Modeling*, 16(3):305–333.

Nævdal, E. and Oppenheimer, M. (2007). The economics of the thermohaline circulation – a problem with multiple thresholds of unknown locations. *Resource and Energy Economics*, 29(4):262–283.

Nash, J. (1953). Two-person cooperative games. *Econometrica*, 21(1):128–140.

Nordhaus, W. (2015). Climate clubs: Overcoming free-riding in international climate policy. *American Economic Review*, 105(4):1339–1370.

Ploeg, F. v. d. and Zeeuw, A. d. (2015a). Climate tipping and economic growth: Precautionary saving and the social cost of carbon. OxCarre Research Paper 118, OxCarre, Oxford, UK.

Ploeg, F. v. d. and Zeeuw, A. d. (2015b). Non-cooperative and cooperative responses to climate catastrophes in the global economy: A north-south perspective. OxCarre Research Paper 149, OxCarre, Oxford, UK.

Polasky, S., Zeeuw, A. d., and Wagener, F. (2011). Optimal management with potential regime shifts. *Journal of Environmental Economics and Management*, 62(2):229 – 240.

Ren, B. and Polasky, S. (2014). The optimal management of renewable resources under the risk of potential regime shift. *Journal of Economic Dynamics and Control*, 40(0):195–212.

Riley, J. and Zeckhauser, R. (1983). Optimal selling strategies: When to haggle, when to hold firm. *The Quarterly Journal of Economics*, 98(2):267–289.

Rob, R. (1991). Learning and capacity expansion under demand uncertainty. *The Review of Economic Studies*, 58(4):655–675.

Sakamoto, H. (2014). A dynamic common-property resource problem with potential regime shifts. Discussion Paper E-12-012, Graduate School of Economics, Kyoto University.

Scheffer, M., Bascompte, J., Brock, W. A., Brovkin, V., Carpenter, S. R., Dakos, V., Held, H., van Nes, E. H., Rietkerk, M., and Sugihara, G. (2009). Early-warning signals for critical transitions. *Nature*, 461(7260):53–59.

Scheffer, M., Carpenter, S., Foley, J. A., Folke, C., and Walker, B. (2001). Catastrophic shifts in ecosystems. *Nature*, 413(6856):591–596.

Tavoni, A., Dannenberg, A., Kallis, G., and Löschel, A. (2011). Inequality, communication, and the avoidance of disastrous climate change in a public goods game. *Proceedings of the National Academy of Sciences*, 108(29):11825–11829.

Tsur, Y. and Zemel, A. (1994). Endangered species and natural resource exploitation: Extinction vs. coexistence. *Natural Resource Modeling*, 8:389–413.

van Damme, E. (1989). Renegotiation-proof equilibria in repeated prisoners' dilemma. *Journal of Economic Theory*, 47(1):206–217.

Appendix

A.1 Proof of Proposition 1

Recall that Proposition 1 states that *when the location of the threshold is known with certainty, then there exists, for every combination of* β*,* N*, and* R*, a value* T_c^{nc} *such that the first-best of staying at* T *can be sustained as a Nash equilibrium when* $T \geq T_c^{nc}$*, where* T_c^{nc} *is defined by* $\Psi = 0$*. The critical value* T_c^{nc} *is higher, the larger* N *or* R *are, or the smaller* β *is.*

It is useful to replicate equation (3) that describes the gain from immediate depletion over staying at T when all other agents stay at T:

$$\Psi(T, R, N, \beta) = u\left(R - \frac{N-1}{N}T\right) - \frac{u(T/N)}{1-\beta} \tag{3}$$

To show that a value T_c, defined by $\Psi = 0$, always exists, I first note that Ψ declines monotonically in T: $\frac{\partial \Psi}{\partial T} = -\frac{N-1}{N}u'\big(R - \frac{N-1}{N}T\big) - \frac{1}{N}\frac{u'(T/N)}{1-\beta} < 0$ as $u' > 0, N \geq 1$ and $\beta \in (0,1)$. Then, I show that Ψ is larger than zero at $T = 0$: $\Psi(0, R, N, \beta) = u(R) > 0$. Finally, I show that Ψ is smaller than zero as $T \to R$: $\lim_{T \to R} \Psi = -\frac{\beta}{1-\beta}u(R/N) < 0$ as $\beta \in (0,1)$ and $u(R/N) > 0$. Thus, by the mean value theorem, for every combination of β, N, and R, there must be a value of T at which $\Psi = 0$.

Now, to show that staying at $T > T_c$ is indeed the socially optimal action (the first-best), I show that $\frac{\mathrm{d}T_c}{\mathrm{d}N} > 0$. This means that the critical value at which staying is a Nash equilibrium is higher the larger N is, which implies that T_c is smallest when $N = 1$ (the sole-owner case). As Ψ is monotonically declining in T, it will be socially optimal to stay at T for all values of T that are larger than the social-planner's T_c. T_c is implicitly defined by $\Psi = 0$ and $\frac{\mathrm{d}T_c}{\mathrm{d}N}$ is therefore given by $\frac{\mathrm{d}T_c}{\mathrm{d}N} = -\frac{\partial\Psi/\partial N}{\partial\Psi/\partial T}$.

We know that the denominator is negative, so that $\frac{\mathrm{d}T_c}{\mathrm{d}N} > 0$ when the numerator is positive. We have: $\frac{\partial \Psi}{\partial N} = \frac{T}{N^2}\big(u'\big(R - \frac{N-1}{N}T\big) - u'(T/N)\big) > 0$.

Finally, it remains to show that T_c is higher the larger R is and the smaller β is. Again, a sufficient condition for the former statement is $\frac{\partial \Psi}{\partial R} > 0$, which holds because $\frac{\partial \Psi}{\partial R} = u'\big(R - \frac{N-1}{N}T\big) > 0$. A sufficient condition for $\frac{\mathrm{d}T_c}{\mathrm{d}\beta} < 0$ is that $\frac{\partial \Psi}{\partial \beta} < 0$ which holds because $\frac{\partial \Psi}{\partial \beta} = -\frac{u(T/N)}{(1-\beta)^2} < 0$.

A.2 Proof of Proposition 2

Recall that Proposition 2 consists of two parts: First, it states that *there exists a set* S *so that for* $s \in S$*, it is optimal to choose* $\delta(s) = 0$*. That is, if* $s_0 \in S$*, the socially optimal use of the resource is* s_0 *for all* t*.* Second, the proposition states that *if* $s_0 \notin S$*, it is optimal to experiment once at* $t = 0$ *and expand the set of safe values by* $\delta^*(s_0)$*. When this has not triggered the regime shift, it is socially optimal to stay at* $s_1 = s_0 + \delta^*(s_0)$ *for all* $t \geq 1$.

Part (1) First, I show that there is a non-empty set $S \subset [0, R]$ at which it is optimal to stay. Assume for contradiction that for all s in $[0, R]$:

$$\max_{\delta\in(0,R-s]} \left\{u(s+\delta)+\beta L_s(\delta)V(s+\delta)\right\} > u(s)+\beta V(s) \tag{A-1}$$

Then there is a value of δ such that:

$$V(s) - \frac{L(s+\delta)}{L(s)}V(s+\delta) < \frac{u(s+\delta)-u(s)}{\beta} \tag{A-2}$$

Now as u is concave, positive, and bounded above by $u(R)$, we know that for an s sufficiently close to R, the numerator of the RHS of (A-2) is bounded above: $u(s+\delta)-u(s) < \beta K\delta$. Using this and multiplying both sides by $L(s)$ as well as dividing both sides by δ we have:

$$\frac{L(s)V(s) - L(s+\delta)V(s+\delta)}{\delta} < KL(s) \tag{A-3}$$

Now, take the limit as $s \to R$. Because $\delta \in (0, R-s]$, we have that $\delta \to 0$ when $s \to R$ so that the LHS of (A-3) is the negative of the derivative of $L(s)V(s)$: $\lim_{s\to R} -\frac{\partial[L(s)V(s)]}{\partial s} = f(R)V(R)$, which is positive while the RHS of (A-3) vanishes when $F(R) = 1$. Thus, we have a contradiction and there must be some s at which it is optimal to choose $\delta(s) = 0$. When there is a positive probability that there is no threshold on $[0, R]$ (that is, $F(R) < 1$), the RHS of (A-3) does not vanish. Nevertheless, there will always be value of s, namely $s = R$, at which it is optimal to stay – simply because there is no other choice.

Thus, the set S is not empty. Moreover, when the hazard rate is not decreasing with s (that is when $\frac{\partial L_s(\delta)}{\partial s} = \frac{-f(s+\delta)(1-F(s))+(1-F(s+\delta))f(s)}{[1-F(s)]^2} < 0 \Leftrightarrow \frac{f(s)}{1-F(s)} < \frac{f(s+\delta)}{1-F(s+\delta)}$), it can be shown that the set S is convex, so that $S = [\overline{s}^*, R]$ where $\overline{s}^*$ is defined in the main text as the lowest value of s at which it is optimal to never experiment. First, note that convexity of S is trivial when $\overline{s}^* = R$. Consider then the case that $\overline{s}^* < R$. By definition, the first-order condition (equation (7) in the main text) must just hold with equality for $\overline{s}^*$:

$$\varphi'(\delta; s) = 0 \quad \Leftrightarrow \quad u'(\overline{s}^*) = \beta \frac{f(\overline{s}^*)}{1-F(\overline{s}^*)} \frac{u(\overline{s}^*)}{1-\beta}$$

S is convex if for any $s \in (\overline{s}^*, R]$ we have $\varphi' < 0$ (i.e. a boundary solution of $\delta^* = 0$). That is:

$$u'(\lambda\overline{s}^* + (1-\lambda)R) < \beta \frac{f(\lambda\overline{s}^* + (1-\lambda)R)}{1-F(\lambda\overline{s}^* + (1-\lambda)R)} \frac{u(\lambda\overline{s}^* + (1-\lambda)R)}{1-\beta} \quad \text{for all } \lambda \in (0,1] \tag{A-4}$$

Because $u' > 0$ and $u'' \leq 0$, the term on the LHS of (A-4) is smaller the larger λ is. Because $u' > 0$ the rightmost fraction of (A-4) is larger the larger λ is, and β is a positive constant. The term in the middle is the hazard rate, which is non-decreasing by assumption.

Part (2) When $s_0 \notin S$, it is not optimal to stay. Thus, it is optimal to expand the set of safe consumption values by choosing $\delta > 0$. Due to discounting, it cannot be optimal to approach S asymptotically but never actually reach it. Thus, there must be a last step from some $s_t \notin S$ to $s_{t+1} = s_t + \delta_t$ with $s_{t+1} \in S$. Below, I show that it is in fact optimal to take only one step. It then follows that when $s_0 \notin S$, it is optimal to choose $s_0 + \delta^*(s_0)$ for $t = 0$ and, if the resource has not collapsed, s_1 for all $t \geq 1$.

Denote $\delta^*(\tilde{s})$ the optimal last step when starting from some value $\tilde{s} \notin S$ and $s^* = \tilde{s} + \delta^*$ with $s^* \in S$. The following calculations show that going from some s to $\tilde{s}$ (by taking a step of size $\tilde{\delta}$) and then to s^* (by taking a step of size δ^*) yields a lower payoff than going from s to s^* directly (by taking a step of size $\hat{\delta} = \tilde{\delta} + \delta^*$; see the box below for a sketch of the involved step-sizes).

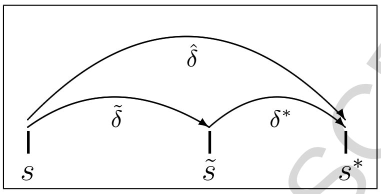

That is, I claim:

$$u(s+\tilde{\delta}) + \beta L_s(\tilde{\delta})\left(u(s+\tilde{\delta}+\delta^*) + \beta L_{s+\tilde{\delta}}(\delta^*)\frac{u(s+\tilde{\delta}+\delta^*)}{1-\beta}\right) \quad \leq \quad u(s+\hat{\delta}) + \beta L_s(\hat{\delta})\frac{u(s+\hat{\delta})}{1-\beta} \tag{A-5}$$

The important thing to note is that: $L_s(\tilde{\delta})L_{s+\tilde{\delta}}(\delta^*) = \frac{L(s+\tilde{\delta})}{L(s)}\frac{L(s+\tilde{\delta}+\delta^*)}{L(s+\tilde{\delta})} = \frac{L(s+\tilde{\delta}+\delta^*)}{L(s)} = L_s(\tilde{\delta}+\delta^*)$. Hence, (A-5) can, upon using $\hat{\delta} = \tilde{\delta} + \delta^*$ and splitting the RHS into three parts ($t = 0, t = 1, t \geq 2$), be written as:

$$u(s+\tilde{\delta}) + \beta L_s(\tilde{\delta})u(s+\hat{\delta}) + \beta^2 L_s(\hat{\delta})\frac{u(s+\hat{\delta})}{1-\beta} \quad \leq \quad u(s+\hat{\delta}) + \beta L_s(\hat{\delta})u(s+\hat{\delta}) + \beta^2 L_s(\hat{\delta})\frac{u(s+\hat{\delta})}{1-\beta}$$

which simplifies to:

$$u(s+\tilde{\delta}) \quad \leq \quad \left[1 + \beta\Big(L_s(\hat{\delta}) - L_s(\tilde{\delta})\Big)\right]u(s+\hat{\delta})$$

$$u(s+\tilde{\delta}) \quad \leq \quad \left[1 + \beta\frac{L(s+\hat{\delta}) - L(s+\tilde{\delta})}{L(s)}\right]u(s+\hat{\delta}) \tag{A-5'}$$

Because the term in the squared bracket is smaller than 1 (as $L(s+\hat{\delta}) < L(s+\tilde{\delta})$), it is not immediately obvious that the inequality in the last line holds. However, we can use the fact that because $\tilde{s} \notin S$, and because δ^* is defined as the optimal last step from $\tilde{s}$ into the set S, the following must hold:

$$\frac{u(\tilde{s})}{1-\beta} < u(\tilde{s}+\delta^*) + \beta L_{\tilde{s}}(\delta^*)\frac{u(\tilde{s}+\delta^*)}{1-\beta}.$$

Using the fact that $\tilde{s} = s + \tilde{\delta}$ and that $\tilde{s} + \delta^* = s + \hat{\delta}$, this can be re-arranged to give:

$$\frac{u(s+\tilde{\delta})}{1-\beta} < u(s+\hat{\delta}) + \beta\frac{L(s+\hat{\delta})}{L(s+\tilde{\delta})}\frac{u(s+\hat{\delta})}{1-\beta}$$

$$\Leftrightarrow$$

$$u(s+\tilde{\delta}) < \left[1 + \beta\frac{L(s+\hat{\delta}) - L(s+\tilde{\delta})}{L(s+\tilde{\delta})}\right]u(s+\hat{\delta}) \tag{A-6}$$

Since $L'(s) < 0$, we know that $\frac{L(s+\hat{\delta})-L(s+\tilde{\delta})}{L(s+\tilde{\delta})} < \frac{L(s+\hat{\delta})-L(s+\tilde{\delta})}{L(s)} < 0$. Therefore, combining (A-5') and (A-6) establishes the claim and completes the proof:

$$u(s+\tilde{\delta}) \quad < \quad \left[1 + \beta\frac{L(s+\hat{\delta}) - L(s+\tilde{\delta})}{L(s+\tilde{\delta})}\right]u(s+\hat{\delta}) \quad < \quad \left[1 + \beta\frac{L(s+\hat{\delta}) - L(s+\tilde{\delta})}{L(s)}\right]u(s+\hat{\delta}) \tag{A-7}$$

As pointed out in footnote 5, it may not immediately obvious that the assumption on the timing adopted in the model ("action - consumption - reaction") is innocuous. Below, I show that immediate experimentation is also optimal under the alternative timing assumption of "action - reaction - consumption".

Consider two plans, "A" and "B" (where plan A implies cautious experimentation and plan B immediate experimentation). Under the timing assumption of "action - reaction - consumption", I decide in plan A to expand the set of safe values by $\tilde{\delta}$, but before I obtain the utility from consuming $s+\tilde{\delta}$ I must first see whether the regime shift occurs or not (the latter event happens with probability $L_s(\tilde{\delta})$. The payoff from period 2 and the remaining periods follows the same logic. The expected payoff from plan A is therefore:

$$P^A = L_s(\tilde{\delta})u(s+\tilde{\delta}) + \beta L_s(\hat{\delta})u(s+\hat{\delta}) + \beta^2 L_s(\hat{\delta})\frac{u(s+\hat{\delta})}{1-\beta} \tag{A-8}$$

The payoff from "plan B" is almost identical, only that all uncertainty is revealed before any utility from consumption is obtained:

$$P^B = L_s(\hat{\delta})u(s+\hat{\delta}) + \beta L_s(\hat{\delta})u(s+\hat{\delta}) + \beta^2 L_s(\hat{\delta})\frac{u(s+\hat{\delta})}{1-\beta} \tag{A-9}$$

As can be clearly seen from (A-8) and (A-9) $P^A < P^B$ when $L_s(\tilde{\delta})u(s+\tilde{\delta}) < L_s(\hat{\delta})u(s+\hat{\delta})$. As $L_s(\tilde{\delta}) = \frac{L(s+\tilde{\delta})}{L(s)}$, and $s+\tilde{\delta} = \tilde{s}$, this is equivalent to:

$$u(\tilde{s}) < \frac{L(s^*)}{L(\tilde{s})}u(s^*) \tag{A-10}$$

Now, the same argument as above can be used. Because $\tilde{s} \notin S$, and because δ^* is defined as the optimal last step from $\tilde{s}$ into the set S, the following must hold:

$$\frac{u(\tilde{s})}{1-\beta} < L_{\tilde{s}}(\delta^*)u(s^*) + \beta L_{\tilde{s}}(\delta^*)\frac{u(s^*)}{1-\beta} \tag{A-11}$$

Simple reformulation then yields (A-10).

A.3 Proof of Proposition 3.

Proposition 3 states that *the optimal step size $\delta^*(s)$ is decreasing in s for $s \in (\underline{s}^*, \overline{s}^*)$*.

Recall that $\delta^*(s)$ is implicitly defined by the solution of $\varphi'(\delta^*; s) = 0$ for $s \in (\underline{s}^*, \overline{s}^*)$, where φ' is:

$$\varphi'(\delta^*; s) = u'(s+\delta^*) + \frac{\beta}{1-\beta}\left[L_s'(\delta^*)u(s+\delta^*) + L_s(\delta^*)u'(s+\delta^*)\right] \tag{7}$$

I assume that the second-order condition $\varphi'' < 0$ is satisfied:

$$\varphi''(\delta^*; s) = u'' + \frac{\beta}{1-\beta}\Big(L_s''(\delta^*)u + 2L_s'(\delta^*)u' + L_s(\delta^*)u''\Big) < 0 \tag{A-12}$$

To show that δ^* is declining in s, I can use the implicit function theorem and need to show that:

$$\frac{\mathrm{d}\delta^*}{\mathrm{d}s} = -\frac{\partial\big[\varphi'(\delta^*; s)\big]/\partial s}{\partial\big[\varphi'(\delta^*; s)\big]/\partial\delta^*} < 0$$

The denominator is negative when the second-order condition is satisfied. Therefore, a sufficient condition for $\frac{\mathrm{d}\delta^*}{\mathrm{d}s} < 0$ is that $\frac{\partial[\varphi'(\delta^*;s)]}{\partial s} < 0$ holds. In other words, we must have:

$$\frac{\partial[\varphi'(\delta^*; s)]}{\partial s} = u'' + \frac{\beta}{1-\beta}\left(\frac{\partial L_s'(\delta^*)}{\partial s}u + L_s'(\delta^*)u' + \frac{\partial L_s(\delta^*)}{\partial s}u' + L_s(\delta^*)u''\right) < 0 \tag{A-13}$$

Noting the similarity of (A-13) to the second-order condition (A-12), and realizing that (A-12) can be decomposed into a common part A and a part B, and that (A-13) can be decomposed into the common part A and a part C, a sufficient condition for (A-13) to be satisfied is that $B > C$.

$$\underbrace{\left[u'' + \frac{\beta}{1-\beta}\Big(L_s'L(\delta^*)u' + L_s(\delta^*)u''\Big)\right]}_{A} + \underbrace{\frac{\beta}{1-\beta}\Big(L_s''(\delta^*)u + L_s'(\delta^*)u'\Big)}_{B} < 0 \tag{A-12'}$$

$$\underbrace{\left[u'' + \frac{\beta}{1-\beta}\Big(L_s'(\delta^*)u' + L_s(\delta^*)u''\Big)\right]}_{A} + \underbrace{\frac{\beta}{1-\beta}\left(\frac{\partial L_s'(\delta^*)}{\partial s}u + \frac{\partial L_s(\delta^*)}{\partial s}u'\right)}_{C} < 0 \tag{A-13'}$$

In order to show that $L_s''(\delta)u + L_s'(\delta)u' > \frac{\partial L_s'(\delta)}{\partial s}u + \frac{\partial L_s(\delta)}{\partial s}u'$, I use the first-order condition for an interior solution from (7) to write u' in terms of u:

$$u' = \frac{-L_s'(\delta)}{\frac{1-\beta}{\beta} + L_s(\delta)}u$$

Upon inserting and canceling u, I need to show that:

$$L_s''(\delta) + L_s'(\delta)\left[\frac{-L_s'(\delta)}{\frac{1-\beta}{\beta} + L_s(\delta)}\right] > \frac{\partial L_s'(\delta)}{\partial s} + \frac{\partial L_s(\delta)}{\partial s}\left[\frac{-L_s'(\delta)}{\frac{1-\beta}{\beta} + L_s(\delta)}\right] \tag{A-14}$$

Recall that $L_s(\delta) = \frac{L(s+\delta)}{L(s)}$ and hence:

$$L_s'(\delta) = \frac{L'(s+\delta)}{L(s)} \qquad \frac{\partial L_s(\delta)}{\partial s} = \frac{L'(s+\delta)L(s) - L(s+\delta)L'(s)}{[L(s)]^2}$$

$$L_s''(\delta) = \frac{L''(s+\delta)}{L(s)} \qquad \frac{\partial L_s'(\delta)}{\partial s} = \frac{L''(s+\delta)L(s) - L'(s+\delta)L'(s)}{[L(s)]^2}$$

Tedious but straightforward calculations then show that (A-14) is indeed satisfied.

$$\underbrace{\left[\frac{1-\beta}{\beta}+\frac{L(s+\delta)}{L(s)}\right]}_{a}\frac{L''(s+\delta)}{L(s)}-\left[\frac{L'(s+\delta)}{L(s)}\right]^2>\underbrace{\left[\frac{1-\beta}{\beta}+\frac{L(s+\delta)}{L(s)}\right]}_{a}\frac{\partial L_s'(\delta)}{\partial s}-\frac{\partial L_s(\delta)}{\partial s}L_s'(\delta)$$

$$\Leftrightarrow$$

$$a\frac{L''(s+\delta)}{L(s)}-\left[\frac{L'(s+\delta)}{L(s)}\right]^2>a\frac{L''(s+\delta)L(s)-L'(s+\delta)L'(s)}{[L(s)]^2}-\frac{\partial L_s(\delta)}{\partial s}L_s'(\delta)$$

$$\Leftrightarrow$$

$$aL''(s+\delta)L(s)-[L'(s+\delta)]^2>a\big(L''(s+\delta)L(s)-L'(s+\delta)L'(s)\big)-\big(L'(s+\delta)L(s)-L(s+\delta)L'(s)\big)\frac{L'(s+\delta)}{L(s)}$$

$$\Leftrightarrow$$

$$-[L'(s+\delta)]^2L(s)>-aL'(s+\delta)L'(s)L(s)-L'(s+\delta)L(s)L'(s+\delta)+L(s+\delta)L'(s)L'(s+\delta)$$

$$\Leftrightarrow$$

$$aL'(s+\delta)L'(s)L(s)>L(s+\delta)L'(s)L'(s+\delta)$$

$$\Leftrightarrow$$

$$aL(s)>L(s+\delta)\quad\Leftrightarrow\quad\left[\frac{1-\beta}{\beta}+\frac{L(s+\delta)}{L(s)}\right]L(s)>L(s+\delta)\quad\Leftrightarrow\quad\frac{1-\beta}{\beta}L(s)>0.\quad\text{True because }\beta,L\in(0,1).$$

A.4 Proof of Proposition 4.

Recall that Proposition 4 states that *There exists a set S^{nc} such that for $s_0 \in S^{nc}$, it is a symmetric Nash equilibrium to stay at s_0 and consume $\frac{s_0}{N}$ for all t. For $s_0 \notin S^{nc}$, it is a Nash equilibrium to take exactly one step and consume $\frac{s_0}{N}+\delta^{nc}(s_0)$ for $t=0$ and – when this has not triggered the regime shift – to stay at $s_1 = s_0 + N\delta^{nc}(s_0)$, consuming $\frac{s_1}{N}$ for all $t \geq 1$.*

Preliminarily, note that the game's stationarity implies that if it is a Nash equilibrium to stay at some s in any one period, it will be a Nash equilibrium to stay at that s in all subsequent periods.

The first part of the proof, showing the existence of S^{nc}, is parallel to the first part of the proof of Proposition 2 and is not repeated here. It rests on the same argument, namely that there is some s at which the gains from increasing consumption are small compared to the expected loss, even when the short term gain does not have to be shared among all N agents.

To prove the second part of the proposition, I need to show that, for $s_0 \notin S^{nc}$, any agent i prefers to reach the set S^{nc} in one step rather than two when the strategy of all other agents is to first take one step of fixed size and then a second feedback step $\delta^{-i*}(s_1)$ that ensures reaching $s^{nc} \in S^{nc}$. Importantly, I use the symmetry of the agents, that is, the second step $\delta^{-i*}(s)$ is given by $\delta^{-i*}(s) = (N-1)\delta^{i*}(s)$ where $\delta^{i*}(s)$ is the state-dependent best reply defined by equation (11) in the main text. Moreover, I assume that all agents coordinate on staying at s^{nc} whenever it is reached.

The setup is the following: All agents stand at s_0 at the beginning of the first period. All agents except i take a step of fixed size and their combined expansion is given by $\tilde{\delta}^{-i}$. If agent i chooses the same symmetric fixed step, her expansion being denoted by $\tilde{\delta}^i$, the state $\tilde{s} \notin S^{nc}$ is reached (provided the regime shift has not occurred). That is, $s_0 < \tilde{s} < s^{nc}$ and $\tilde{s} - s_0 = N\tilde{\delta}^i$. Agent i can also choose her first step so that s^{nc} is already reached in the first period. Denote this step that expands the set of safe consumption values from s_0 to s^{nc} when the total expansion of all other agents taken together is $\tilde{\delta}^{-i}$ by $\hat{\delta}^{i*}(s_0)$. The size of the step $\hat{\delta}^{i*}(s_0)$ is thus $\hat{\delta}^{i*}(s_0) = s^{nc} - s_0 - \tilde{\delta}^{-i}$. Because all other agents remain at s^{nc} once it is reached, agent i's payoff is in this case:

$$\pi^{(1)} = u\Big(\frac{s_0}{N} + \hat{\delta}^{i*}(s_0)\Big) + \frac{\beta}{1-\beta} L_{s_0}(\hat{\delta}^{i*} + \tilde{\delta}^{-i}) u\Big(\frac{s^{nc}}{N}\Big) \tag{A-15}$$

When agent i takes two steps, first a step of size $\tilde{\delta}^i$ to the value $\tilde{s}$ (with $\tilde{s} \notin S^{nc}$) and then a second step $\delta^{i*}(\tilde{s})$ of size $\delta^{i*}(\tilde{s}) = s^{nc} - \tilde{s} - \delta^{-i*}(\tilde{s})$, her payoff is:

$$\pi^{(2)} = u\Big(\frac{s_0}{N} + \tilde{\delta}^i\Big) + \beta L_{s_0}(\tilde{\delta}^i + \tilde{\delta}^{-i}) \Big(u\Big(\frac{\tilde{s}}{N} + \delta^{i*}\Big) + \frac{\beta}{1-\beta} L_{\tilde{s}}(\delta^{i*} + \delta^{-i*}) u\Big(\frac{s^{nc}}{N}\Big)\Big) \tag{A-16}$$

I now show that $\pi^{(2)} < \pi^{(1)}$. For clarity, rewrite (A-15) and (A-16) by splitting it in three terms ($t = 0, t = 1$, and $t \geq 2$) and using the fact that $L_{s_0}(\hat{\delta}^{i*} + \tilde{\delta}^{-i}) = L_{s_0}(\tilde{\delta}^i + \tilde{\delta}^{-i}) L_{\tilde{s}}(\delta^{i*} + \delta^{-i*}) = \frac{L(s^{nc})}{L(s_0)}$:

$$\pi^{(1)} = u\Big(\frac{s_0}{N} + \hat{\delta}^{i*}(s_0)\Big) + \beta \frac{L(s^{nc})}{L(s_0)} u\Big(\frac{s^{nc}}{N}\Big) + \beta^2 \frac{L(s^{nc})}{L(s_0)} \frac{u(s^{nc}/N)}{1-\beta} \tag{A-15'}$$

$$\pi^{(2)} = u\Big(\frac{s_0}{N} + \tilde{\delta}^i\Big) + \beta \frac{L(\tilde{s})}{L(s_0)} u\Big(\frac{\tilde{s}}{N} + \delta^{i*}(\tilde{s})\Big) + \beta^2 \frac{L(s^{nc})}{L(s_0)} \frac{u(s^{nc}/N)}{1-\beta} \tag{A-16'}$$

Thus, $\pi^{(2)} < \pi^{(1)}$ if:

$$u\Big(\frac{s_0}{N} + \tilde{\delta}^i\Big) < u\Big(\frac{s_0}{N} + \hat{\delta}^{i*}(s_0)\Big) + \beta \Big[\frac{L(s^{nc})}{L(s_0)} u\Big(\frac{s^{nc}}{N}\Big) - \frac{L(\tilde{s})}{L(s_0)} u\Big(\frac{\tilde{s}}{N} + \delta^{i*}(\tilde{s})\Big)\Big] \tag{A-17}$$

First, by symmetry of the agents and the definition of $\delta^{i*}(\tilde{s})$, we have $\frac{\tilde{s}}{N} + \delta^{i*}(\tilde{s}) = \frac{\tilde{s} + N\delta^{i*}(\tilde{s})}{N} = \frac{s^{nc}}{N}$.

$$u\Big(\frac{s_0}{N} + \tilde{\delta}^i\Big) < u\Big(\frac{s_0}{N} + \hat{\delta}^{i*}(s_0)\Big) + \beta \frac{L(s^{nc}) - L(s_0)}{L(s_0)} u\Big(\frac{s^{nc}}{N}\Big) \tag{A-17'}$$

Similarly, we have $\frac{s_0}{N} + \tilde{\delta}^i = \frac{s_0 + N\tilde{\delta}^i}{N} = \frac{\tilde{s}}{N}$:

$$u\Big(\frac{\tilde{s}}{N}\Big) < u\Big(\frac{s_0}{N} + \hat{\delta}^{i*}(s_0)\Big) + \beta \frac{L(s^{nc}) - L(\tilde{s})}{L(s_0)} u\Big(\frac{s^{nc}}{N}\Big) \tag{A-17''}$$

Now, note that $\frac{s_0}{N} + \hat{\delta}^{i*}(s_0) = \frac{s_0 + N\hat{\delta}^{i*}(s_0)}{N}$ and that the step $\hat{\delta}^{i*}(s_0)$ is larger than the symmetric step that would be necessary to reach s^{nc} from s_0. Formally: $\hat{\delta}^{i*}(s_0) = s^{nc} - s_0 - \tilde{\delta}^{-i} > \frac{s^{nc} - s_0}{N} \Leftrightarrow Ns^{nc} - Ns_0 - N(N-1)\tilde{\delta}^i > s^{nc} - s_0 \Leftrightarrow s^{nc} - s_0 > N\tilde{\delta}^i = \tilde{s} - s_0$ which is true by construction. It follows that $\frac{s_0 + N\hat{\delta}^{i*}(s_0)}{N} > \frac{s^{nc}}{N}$ and we therefore have $u\Big(\frac{s^{nc}}{N}\Big) + \beta \frac{L(s^{nc}) - L(\tilde{s})}{L(s_0)} u\Big(\frac{s^{nc}}{N}\Big) < u\Big(\frac{s_0}{N} + \hat{\delta}^{i*}(s_0)\Big) + \beta \frac{L(s^{nc}) - L(\tilde{s})}{L(s_0)} u\Big(\frac{s^{nc}}{N}\Big)$ so that a sufficient condition for $\pi^{(2)} < \pi^{(1)}$ is:

$$u\left(\frac{\tilde{s}}{N}\right) < u\left(\frac{s^{nc}}{N}\right) + \beta\frac{L(s^{nc}) - L(\tilde{s})}{L(s_0)}u\left(\frac{s^{nc}}{N}\right) \tag{A-18}$$

Parallel to the argument in Proposition 2, we can use the fact that because $\tilde{s} \notin S^{nc}$ we must have

$$\frac{u(\tilde{s}/N)}{1-\beta} < u\left(\frac{\tilde{s}}{N} + \delta^{i*}(\tilde{s})\right) + \beta L_{\tilde{s}}(\delta^{i*}(\tilde{s}) + \delta^{-i*}(\tilde{s}))\frac{u(s^{nc}/N)}{1-\beta}$$

Using that agents are symmetric, we have $\frac{\tilde{s}}{N} + \delta^{i*}(\tilde{s}) = \frac{s^{nc}}{N}$ and re-arranging shows that (A-18) holds, so that $\pi^{(2)} < \pi^{(1)}$ as claimed.

A.5 Proof of Proposition 5.

Let me repeat the argument from the main text: The effect of an increase in a parameter a in the interior range $s \in (\underline{s}^{nc}, \overline{s}^{nc})$ is given by $\frac{\mathrm{d}g^{nc}}{\mathrm{d}a} = -\frac{\partial\phi'/\partial a}{\partial\phi'/\partial g^{nc}}$. Thus, to show that aggregate consumption is higher the higher the parameter a, it is sufficient to show that $\frac{\partial\phi'}{\partial a} > 0$ (because the second-order condition implies that $\frac{\partial\phi'}{\partial g^{nc}} < 0$). Because g^{nc} is monotonically decreasing in s, it is also sufficient to show that, for a given value of R, neither boundary $\underline{s}^{nc}$ or $\overline{s}^{nc}$ decreases and at least one boundary increases with a. The reason is that for a given value of R an upward shift of $\underline{s}^{nc}$ or $\overline{s}^{nc}$ (and no downward of the respective other boundary) necessarily implies that all new values of g^{nc} must lie above the old values of g^{nc}.

(a) *The boundaries $\underline{s}^{nc}$, $\overline{s}^{nc}$, and aggregate consumption in the cautious equilibrium, Ng^{nc}, decrease with β.* Here, it is simple to show that $\frac{\partial\phi'}{\partial\beta} < 0$. We have $\frac{\partial\phi'}{\partial\beta} = \frac{[...]}{(1-\beta)^2}$, where the term in the squared brackets [...] is the term in the squared brackets of equation (12). We know that this term must be negative for an interior solution because $u' > 0$.

(b) *An increase in N leads to higher resource use in the cautious equilibrium when $\frac{N}{N+1} \geq u'(\frac{R}{N})/u'(\frac{R}{N+1})$.* Here, I argue that both $\underline{s}^{nc}$ and $\overline{s}^{nc}$ increase when adding another player and $\frac{N}{N+1} \geq u'(\frac{R}{N})/u'(\frac{R}{N+1})$:

First, for a given number of players N we have at a given $\underline{s}^{nc} = \hat{s}$ that

$$\phi'(\frac{R-s}{N};\hat{s}) = u'\left(\frac{\hat{s}}{N} + \frac{R-\hat{s}}{N}\right) + \frac{\beta}{1-\beta}\left[L'_{\hat{s}}(N\delta^{nc})u\left(\frac{R}{N}\right) + \frac{1}{N}L_{\hat{s}}(N\delta^{nc})u'\left(\frac{R}{N}\right)\right] = 0$$

I now show that for $N+1$ we have $\phi'(\frac{R-s}{N+1};\hat{s}) > 0$ when $\frac{N}{N+1} \geq \frac{u'(\frac{R}{N})}{u'(\frac{R}{N+1})}$:

$$\phi'(\tfrac{R-s}{N+1};\hat{s}) - \phi'(\tfrac{R-s}{N};\hat{s}) > 0$$

$$\Leftrightarrow$$

$$u'(\tfrac{R}{N+1}) - u'(\tfrac{R}{N}) + \frac{\beta}{1-\beta}\left[\left(u(\tfrac{R}{N+1}) - u(\tfrac{R}{N})\right)L'_{\hat{s}} + L_{\hat{s}}\left(\frac{1}{N+1}u'(\tfrac{R}{N+1}) - \frac{1}{N}u'(\tfrac{R}{N})\right)\right] > 0$$

The first part of the last line is positive due to concavity of u, the first term in the squared bracket is positive since $L'_s < 0$ and $u(\frac{R}{N+1}) < u(\frac{R}{N})$, and the last term in the squared bracket is positive whenever $\frac{N}{N+1} \geq \frac{u'(\frac{R}{N})}{u'(\frac{R}{N+1})}$.

Thus, when $\frac{N}{N+1} \geq \frac{u'(\frac{R}{N})}{u'(\frac{R}{N+1})}$, it is guaranteed that $\phi'(\frac{R-s}{N+1};\hat{s}) > 0$, which implies that for $N+1$ the upper bound of the choice set is a binding constraint at $\hat{s}$ and that the corresponding smallest value of s at which the agents can coordinate on cautious experimentation is larger. Note that $\frac{N}{N+1} \geq \frac{u'(\frac{R}{N})}{u'(\frac{R}{N+1})}$ is not a necessary condition: Of course, we may have $\phi'(\frac{R-s}{N+1};\hat{s}) - \phi'(\frac{R-s}{N};\hat{s}) > 0$ also when $\frac{N}{N+1} < \frac{u'(\frac{R}{N})}{u'(\frac{R}{N+1})}$ as the specific example in section 3.5 shows.

Second, for a given number of players N we have at a given $\bar{s}^{nc} = \check{s}$ that

$$\phi'(0;\check{s},N) = u'\left(\frac{\check{s}}{N}\right) + \frac{\beta}{1-\beta}\left[L'_{\check{s}}(0)u\left(\frac{\check{s}}{N}\right) + \frac{1}{N}u'\left(\frac{\check{s}}{N}\right)\right] = 0$$

Clearly, we can make exactly the same argument as above to show that $\phi'(0;\check{s},N+1) > 0$ when $\frac{N}{N+1} \geq \frac{u'(\frac{R}{N})}{u'(\frac{R}{N+1})}$.

(c) *The more likely the regime shift (in terms of a first-order stochastic dominance), the larger the range where a separate cautious Nash-equilibrium exists and the lower aggregate consumption.*

Suppose that for some given value $\hat{s}$, equation (12) has an interior solution that defines g^{nc}:

$$\phi'(\delta^{nc};\hat{s},L) = u'\left(\frac{\hat{s}}{N}+\delta^{nc}\right) + \frac{\beta}{1-\beta}\left[L'_{\hat{s}}(N\delta^{nc})u\left(\frac{\hat{s}+N\delta^{nc}}{N}\right) + \frac{1}{N}L_{\hat{s}}(N\delta^{nc})u'\left(\frac{\hat{s}+N\delta^{nc}}{N}\right)\right] = 0 \qquad (12)$$

I now show that a more likely regime shift (in terms of a first-order stochastic dominance) means a change in L_s to $\tilde{L}_s$ in such a way that $\phi'(\delta^{nc};\hat{s},\tilde{L}) < 0$ so that for every $s \in (\underline{s}^{nc}, \bar{s}^{nc})$ we have that the resulting interior solution $\tilde{g}^{nc}$ is smaller than the orginal g^{nc}. As a consequence, the range where a separate cautious Nash-equilibrium exists will also be larger.

A first-order stochastic dominance means that $\tilde{F} \geq F$ (where the inequality is strict for at least one s). Because the hazard rate is non-declining, this means that $\frac{\tilde{F}(s+\delta)-\tilde{F}(s)}{1-\tilde{F}(s)} \geq \frac{F(s+\delta)-F(s)}{1-F(s)}$ and consequently $\tilde{L}_s \leq L_s$. This implies also that $\tilde{L}'_{\hat{s}} = \frac{-\tilde{f}(s+\delta)}{1-\tilde{F}(s)} < \frac{-f(s+\delta)}{1-F(s)} = L'_{\hat{s}} < 0$. Thus, both the negative first term and the positive second term in the squared bracket of (12) are smaller, which implies that $\phi'(\delta^{nc};\hat{s},\tilde{L}) < 0$.

(d) *An increase of R to $\tilde{R}$ for an unchanged risk of the regime shift (i.e. $R < \tilde{R} \leq A$) decreases $\underline{s}^{nc}$ and thus leads to a larger range where a separate cautious equilibrium exists.*

Note that R does not affect equation (12) when $R < \tilde{R} \leq A$, but it has an effect on the first value $\underline{s}^{nc}$: As the diagonal line defining the upper bound of δ shifts outwards, and $g^{nc}(s)$ is a downward sloping function steeper than $R-s$, the first value at which it is not optimal to deplete the resource, $\underline{s}^{nc}$, must be smaller when R increases to $\tilde{R}$.

Highlights

- Thresholds may benefit non-cooperative agents by facilitating coordination.
- Learning depends on history: higher initial safe level implies less experimentation.
- For high initial value, preserving the resource with certainty is a Nash equilibrium.