Threatening Thresholds?

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

Florian K Diekert*

Work in progress... this version: May 9, 2014

Abstract

The potential occurrence of disastrous regime shifts is especially relevant for environmental goods and services. At the same time, the use of environmental goods and services is often prone to market failure due to their public nature and the ensuing free-rider incentives. This paper investigates how and when the threat of disastrous regime shifts works as a "commitment device" to enforce the socially optimal outcome in a non-cooperative setting. A general, but still analytically tractable, model of a dynamic game is developed. When the common-pool externality only pertains to the (endogenous) risk of passing the catastrophic threshold, the agents can coordinate on the first-best in a wide range of cases – if the location is known. If it is unknown, agents can coordinate on the first-best only if the socially optimal action is to use the environmental good at its current level, and if this status quo is sufficiently valuable. Any experimentation (both in the non-cooperative game and in social optimum) will be undertaken only in the first period, and the degree of experimentation is declining in the value of the initial set of safe consumption possibilities. (187 words)

Keywords: Dynamic Games; Thresholds and Natural Disasters; Renewable Resources; Climate Change.

JEL-Codes: C73, Q20, Q54

*Department of Economics and Centre for Ecological and Evolutionary Synthesis, University of Oslo, PO-Box 1095, 3017 Oslo, Norway. Email: f.k.diekert@ibv.uio.no. This research is funded by NorMER, a Nordic Centre of Excellence for Research on Marine Ecosystems and Resources under Climate Change. I would like to thank Geir Asheim, Scott Barrett, Matti Liski, Eric Nævdal, Volker Diekert and seminar participants at the Beijer institute, at Statistics Norway, at Oslo- and Helsinki university for constructive discussion and comments.

1 Introduction

Everyday-experience teaches that many things tolerate some stress, but if this is driven too far or for too long, they break. Not surprisingly, regime shifts, thresholds, or tipping points are a popular topic in the scientific literature. In this paper, I investigate whether and how the existence of a catastrophic threshold may be *beneficial* in the sense that it enables non-cooperative agents to coordinate their actions, thereby improving welfare over a situation where this tipping point is not present or ignored. To this end, I develop a generic model of using a productive asset that loses (some or all) its productivity upon crossing some threshold.

The model is applicable to many different settings, but I would like to think of it broadly in terms of ecosystem services. Prime example are the eutrophication of lakes, the bleaching of coral reefs, or saltwater intrusion (Scheffer et al., 2001; Tsur and Zemel, 1995). On a somewhat larger scale, the collapses of the North Atlantic cod stocks off the coast of Canada, or the capelin stock in the Barents Sea can – at least in part – be attributed to overfishing (Frank et al., 2005; Hjermann et al., 2004). The most important application may however be the climate system, where potential drivers of a disastrous regime shift could be a disintegration of the West-Antartic ice sheet, a shutdown of the thermohaline circulation, or a melting of Permafrost (Lenton et al., 2008).

In order to isolate the effect of a catastrophic threshold on the ability to cooperate in exploiting a renewable resource, I abstract – as a first step – from the dynamic common pool aspect of non-cooperative resource use. This allows me to obtain tractable analytic solutions and feedback Nash-equilibria for general utility functions and general continuous probability distributions. I will first show that in the benchmark case when the location of the catastrophic threshold is known, the first-best can be achieved for any number of players if they are sufficiently patient. I will then analyze the effect of uncertainty: If the location of the threshold is unknown, agents can coordinate on the first-best only if the socially optimal action is to use the environmental good at its current level, and if this status quo is sufficiently valuable. If preserving the status quo is not sufficiently valuable, the players 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, and this consumption level is, *ex post*, socially optimal. If the value of preserving the status quo is too low (or conversely, the probability of causing the disastrous regime shift is deemed too high) the only non-cooperative equilibrium is the immediate extirpation of the resource.

A particular feature of thresholds that the model captures is that learning is only affirmative. That is, given that I know that the current state is safe, and I expand the current state, I will only learn whether the future state is safe or not. I will not obtain any new information on whether I have come very close to the threshold. I show that this implies that any experimentation will be undertaken in the first period, and that the degree of experimentation is is declining in the value of the initial set of safe consumption possibilities (the more valuable the current state the more cautious I am).

The remainder of this section places the current contribution in the context of the relevant literature. In section 2, I explain the general modeling approach. The paper concludes with a discussion of the model and directions for further research (section 5).

Relation to the literature

The pioneering contribution analyzing the economics of regime shifts in an ecosystem/pollution context was by Cropper (1976). There are by now a good dozen papers on the optimal management of renewable resources under the threat of a irreversible regime shift. Recently, Polasky et al. (2011) have summarized and characterized the literature at hand of a simple fishery model. They contrast whether the regime shift implies a collapse of the resource or merely a reduction of its renewability, and whether probability of crossing the threshold is exogenous or depends on the state of the system (i.e. it is endogenous). They show that resource extraction should be more cautious when crossing the threshold implies a loss of renewability and the probability of crossing the threshold depends on the state of the system. In contrast, exploitation should be more aggressive when a regime shift implies a collapse of the resource and the probability of crossing the threshold cannot be influenced. There is no change in optimal extraction for the loss-of-renewability/exogenous-probability case and the results are ambiguous for the collapse/endogenous-probability case.

Until now the literature 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. They find that, depending on the productivity of the rangeland, under- or overexploitation might occur. Kossioris et al. (2008) focus on feedback equilibria and analyze, with help of numerical methods, non-cooperative pollution of a "shallow lake". 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. However, they highlight that for some combinations of parameter values, a regime shift can be avoided in a feedback Nash equilibrium (although the value of the game is still substantially lower than in the first-best). 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. Finally, Sakamoto (2014) has, by combining analytical and numerical methods, analyzed a non-cooperative game with an endogenous regime shift hazard. He shows that this setting may lead to more precautionary management, also in a strategic setting.

An important aspect is how the threat of the regime shift is modeled. Most of the previous studies translate the uncertainty about the location of the threshold in state space into uncertainty about the 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, as time leads to infinity, the event occurs with probability 1. In other words, even if we were to totally stop extracting/polluting, the disastrous regime shift could not be avoided. 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 players can avoid the event by staying at or below that level. I therefore follow the modeling approach of e.g. Tsur and Zemel (1994, 1995); Nævdal (2001); Lemoine and Traeger (2014) where – figuratively speaking – the threshold is interpreted as the edge of a cliff and the management question is whether, when, and where to stop walking. This modeling approach leads, under certain conditions, to a socially optimal "safe minimum standard of conservation" as pointed out by Mitra and Roy (2006). I show in this paper that also players in a non-cooperative game may, depending on initial conditions, coordinate on such "safe minimum standard". However there are also initial conditions for which the potential of a disastrous regime shift becomes a "self-fulfilling prophecy" as it leads to an immediate extirpation of the resource.

Taken together, the contributions of Fesselmeyer and Santugini (2013), Sakamoto (2014), and the current paper show that the effect of a regime shift pulls in the same direction in a non-cooperative setting as under optimal management. This notwithstanding, I point out that the strategic setting implies a more aggressive resource use for the set of initial values for which conservation cannot be coordinated upon. Moreover, I highlight how the effect of a threatening threshold depends on the initial safe value: The more the players 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.[1] This aspect has, to the best of my knowledge, not yet been appreciated.

In addition to this literature, the current paper is closely related to two articles that discuss the role of uncertainty about the threshold's location on whether the regime shift 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. He shows that an increase in uncertainty leads to increased consumption, but that increased ambiguity may have the opposite effect. While Aflaki's and Barrett's models are static, I develop a general dynamic game. Furthermore, I place only minimal requirements on the utility function (concavity and boundedness) and the probability distribution of the threshold (continuity). I show how a certain threshold will always allow coordination for sufficiently patient players and how even under uncertainty, coordination induced by the threat of a threshold may be feasible in some cases.

Analyzing how strategic interactions shape the exploitation pattern of a renewable resource 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 basically 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.

[1] In some sense, the antithesis of this effect is well summarized by Janis Joplin's "Freedom is another word for nothing left to lose"

2 Modeling approach

To repeat, I consider a situation where several agents share a productive asset. Each agent obtains utility from using it and as long as total use is at or below a threshold, the asset remains intact. Contrarily, if total use in any period exceeds the threshold, a disastrous regime shift will occur. The players are unable to make binding agreements. In order to expose the effect of a threatening regime shift on the cooperative and non-cooperative use of environmental goods and services as clearly as possible, I consider a model where the common pool externality applies only the risk of triggering the regime shift. While this is clearly a simplification, it allows for a tractable solution of a general dynamic model.

Resource dynamics

- Time is discrete and indexed by $t = 0, 1, 2, ...$
- Each period, the players 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 $r \ll R$).
- 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 a state variable s_t, denoting the upper bound of the "safe consumption possibility set" at time t. That is, a total extraction up to s_t has not triggered a regime shift before and is hence known to not trigger a regime shift in the future (i.e. $\text{Prob}(T \leq s_t) = 0$).

Players, choices, and payoff

- Players derive utility from consuming the resource according to some general function $u(c_t^i)$. I assume that this function is continuous and bounded below by $u(0) = b$ and above by $u(R) = B$ with $u' > 0$ and $u'' \leq 0$.
- Let c_t^i be the consumption of player i at time t. For clarity, I split per-period consumption in the productive regime in two parts: $c_t^i = \alpha^i s_t + \delta_t^i$. This means:

1. The players consume a share α^i of s_t (the amount of the resource that can be used safely).
2. The players may choose to consume an amount δ_t^i more than $\alpha^i s_t$, effectively pushing the boundary of the production 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 extension of the production set by all other players. 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 players is to choose that sequence of extension decisions $\Delta^i = \delta_0^i, \delta_1^i, ...$ which, for given strategies of the other players Δ^{-i}, and for a given initial value s_0, maximizes the sum of expected per-period utilities, discounted by a common factor β with $\beta \in (0, 1)$.

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 $F(x) = \int_0^x f(\tau)\mathrm{d}\tau$.
- 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 extracting the entire amount of the resource is actually safe and the presence of a critical threshold is immaterial. When $R = A$ extracting the entire amount of the resource will trigger the regime shift for sure. Both R and A are known with certainty.
- Knowing that a given a given exploitation level s is save, the updated density of T on $[s, A]$ is given by $f_s(\delta) = \frac{f(\delta+s)}{1-F(s)}$. The cumulative probability of triggering the regime shift when, so to say, taking a step of distance δ from the safe value s is:

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

So that $F_s(\delta)$ is the discrete version of the hazard function.

- The key expression that I will use in the remainder of the paper is $L_s(\delta)$, which I call the 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. We have $L(x) = 1 - F(x)$ and

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

- The survival function $L_s(\delta)$ has the following properties: It is declining in δ as $\frac{\partial L_s(\delta)}{\partial \delta} = \frac{-f(s+\delta)}{1-F(s)} < 0$ and it is bounded below by the conditional probability that T is not in the interval $[s, R]$, that is $L_s(\delta) \in \left[\frac{1-F(R)}{1-F(s)}; 1\right]$. Whether the survival function is declining in s depends on f: $\frac{\partial L_s(\delta)}{\partial s} = \frac{-f(s+\delta)(1-F(s))+(1-F(s+\delta))f(s)}{[1-F(s)]^2} < 0 \Leftrightarrow \frac{1-F(s+\delta)}{1-F(s)} < \frac{f(s+\delta)}{f(s)}$. As the RHS of the last inequality is smaller one, we know at least that the chances of surviving a given step δ decline with s when the density of T is non-decreasing on $[0, R]$.

Preliminary clarifications and tractability assumptions

- It is well known that the static non-cooperative game of sharing a pie has infinitely many Nash-equilibria. Here, I focus on symmetric equilibria. Moreover, the game of sharing a pie requires a statement about the consequences when the sum of players' consumption plans exceed the total available resource. In this case, I assume that the resource is rationed so that each players gets an equal share. Finally, it is clear that in the dynamic game of sharing a pie developed above, the only best-reply to the other players' strategy to immediately extirpate the resource is to grab as much of the resource as possible as well. I will, in the following, not explicitly discuss this Nash equilibrium.

- I consider the disastrous regime shift to be irreversible.

- Without loss of generality, I set $r = 0$ and $u(0) = 0$. That is, the regime shift implies a complete collapse of the resource.

- This model abstracts from the dynamic common pool problem in the sense that the consumption decision of a player 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.

- The simplicity of the resource dynamics means that the idea that the system is the more likely to experience a disastrous regime shift the more of the stock has been exploited can simply be integrated into the cdf of the threshold location. The core model feature that a safe state s is safe (no matter how far or close to the threshold) implies also that the players cannot learn about the potential proximity of the threshold by observing the current state. That is, my model does not allow for “early warning signals”.

Interpretation

Essentially, the players face the problem of sharing a magic pie: If they do not eat too much today, they will tomorrow find the full pie replenished, to be shared again. Each player however faces the temptation to eat a little more than what was eaten yesterday since he or she will have that piece for himself, whereas the future consequences of his voracity will have to be shared by all.

A literal example for this model could be whale-watching: Imagine a group of whales coming to a certain spot to play and this sight is being enjoyed by tourists. The boat operators know that the current intensity of whale-watching is tolerated by the animals. Pondering whether to offer more tours, they have to weigh the additional income against the fear that there may be some intensity beyond which the whales get disturbed and avoid the area. A second example could be saltwater intrusion: players can take fresh water from a reservoir and the water will replenish. However, once the water pressure in the reservoir has dropped too low, saltwater will enter and the reservoir will be spoiled.

However, I also point to two more abstract interpretations: First, one could think of the long-term dynamics of a renewable resource such as a fishery. The upper bound on the consumption set R would then be the respective equilibrium harvest of the N-player exploitation problem. That is, in absence of the threat of a disastrous regime shift each player would obtain a harvest share $\alpha^i R$. Second, one could think of the climate system with an implicit emission-consumption link. R would then be level of emissions where the marginal benefit of emissions equals their marginal cost in absence of any considerations of a possible regime shift.

3 Social optimum and non-cooperative extraction

In this main section of the paper, I will as a preliminary step analyze the situation when the threshold is known. This scenario highlights the disciplinary effect of a threatening threshold. I will then turn to the more realistic case when the location of the threshold is unknown. I point out that coordination on staying on the safe side is more difficult but not impossible to achieve. A particular feature of the model is that any experimentation – if at all – is undertaken in the first period. I show that non-cooperative experimentation is larger than cooperative experimentation and, as it is intuitive, players experiment less the more they already know (both in the social optimum and Nash equilibrium). Comparative statics are analyzed in section 3.3 and an instructive specific example is provided in section 3.4.

3.1 Known threshold location

What is the first-best outcome in a situation when the threshold T is known? When T is very close to zero, so that a large part of the available resource R must be foregone to ensure its continued existence, it will be socially optimal to cross the threshold and extirpate the resource. When T is sufficiently close to R this will not be socially optimal. Rather, the first-best is to extract exactly that amount of the resource which does not cause the regime shift. How close is close will obviously depend on the discount factor β: the less the future is discounted, the more is one willing to sacrifice today's consumption for future consumption possibilities.

For a given safe value of total consumption s, the value of the game for player i is:

$$V(s) = \max_{\delta^i} \left\{ u(\alpha^i s + \delta^i) + I \cdot \beta V(s) \right\}, \text{ where } \begin{cases} I = 1 & \text{when} \quad s + \delta^i + \delta^{-i} \leq T \\ I = 0 & \text{when} \quad s + \delta^i + \delta^{-i} > T \end{cases} \tag{3}$$

Due to the simplicity of the model structure, it is clear that if staying below 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 $\frac{u(\alpha^i T)}{1-\beta}$. Conversely, the payoff from deviating and immediately extirpating the resource when all other players' policy is to stay at the threshold is given by $u\left(R - \frac{N-1}{N}T\right)$. Staying below the threshold can thus be sustained as a Nash equilibrium whenever $\frac{u(\alpha^i T)}{1-\beta} \geq u\left(R - \frac{N-1}{N}T\right)$. Denote by $\bar{\beta}$ the value of β for which this condition just so holds with equality (i.e. $\bar{\beta}$ is the lowest discount factor for which staying below the threshold can be sustained for a given N, T, and R). We have:

$$\bar{\beta} = 1 - \frac{u(\alpha^i T)}{u\left(R - \frac{N-1}{N}T\right)} \tag{4}$$

Although it will always be a Nash-equilibrium to extirpate the resource, there will always be a parameter combination where there is a second equilibrium which supports the first-best (Proposition 1). Given these conditions, the game exhibits the structure of a coordination game. Here, as in the static game from Barrett (2013, p.236), "[e]ssentially, nature herself enforces an agreement to avoid catastrophe."

Proposition 1. *When the location of the threshold is known with certainty, then there exists, for every combination of N, T, and R, a $\bar{\beta}$ such that the first-best can be sustained as a Nash-equilibrium when $\beta \geq \bar{\beta}$. The larger is N, or the closer T is to 0, the larger has to be β.*

Proof. For a given N, first note that $\frac{\mathrm{d}\bar{\beta}}{\mathrm{d}T} = -\left([u'\alpha^i \cdot u + u \cdot u' \frac{N-1}{N}]/[u^2]\right) < 0$ so that the players need to be the less patient the more valuable it is to stay below the threshold (i.e. as T grows). Second, note that for $T \to 0$, the last term of (4) approaches 0 so that the right-hand-side of (4) approaches 1. But since it approaches 1 from below, we can always find some value of β that could still sustain the first-best. Finally, note that for $\beta > \bar{\beta}$ the sharing scheme α will be indeterminate, but as $\beta \to \bar{\beta}$ only the symmetric sharing $\alpha^i = \frac{1}{N}$ remains. The reason is that the binding constraint for coordination is the player that obtains the lowest share and as players are symmetric sharing will be equitable when there is no surplus to distribute. This implies that as $N \to \infty$, $u(\frac{1}{N}T) \to 0$. Again, $\bar{\beta}$ approaches 1 from below, allowing to find some value of β that could still sustain the first-best. □

3.2 Unknown threshold location

I now turn to the case when the location of the threshold is unknown. In spite of the uncertainty about T, the players do know that resource consumption up to s does not trigger the disastrous regime shift. Hence, the threat of a disastrous regime shift can, in principle, be eliminated by coordinating on consuming exactly the amount s. However, there will be cases where s is so low that it is not even socially optimal to stay at s, but rather to expand the set of safe consumption possibilities by some amount δ. When deciding on the size of this step δ, the gain from expanding the set of safe consumption possibilities to $s' = s + \delta$ has to be weighted against the probability of causing the disastrous regime shift. The effect of strategic non-cooperative interaction will manifest itself in the classic way that each player realizes that the gain from increased consumption is private while the cost in terms of a higher regime-shift hazard are borne by all.

To introduce some notation, let s^* be a value of s at which it is not socially optimal to expand the set of safe extraction values: The threat of a disastrous regime shift lures too large. Similarly let s^{nc} be a value of s at which not expanding the set of safe extraction values can be sustained as a Nash-equilibrium. Let $\overline{s}^*$ and $\overline{s}^{nc}$ be respectively the lowest member of these sets of values. Clearly, we have $\overline{s}^* \leq \overline{s}^{nc}$. Conversely, let $\underline{s}$ be the highest of all those values of s for which either the current safe consumption level is so low, or the threat of a potential future regime shift weighs so little that it is rational to extract the entire resource (i.e. choose $\delta(s) = R - s$). Again, we will have $\underline{s}^* \leq \underline{s}^{nc}$, but note that $\underline{s}$ may in fact not exist when it is optimal to choose $\delta(s) < R - s$ for all s.

The outcome when the location of T is unknown is illustrated in Figure 1. The blue dashed line plots the socially optimal extension δ of the safe consumption set s (on the y-axis) as a function of the safe consumption set on the x-axis (where obviously $s \leq R$ and $\delta \in [0, R - s]$). For very low values of s (below $\underline{s}^*$) it is optimal to consume the entire resource, i.e. to choose $\delta(s) = R - s$. For high values of s (above $\overline{s}^*$) it is optimal to remain standing, i.e. to choose $\delta(s) = 0$. The red solid line plots the non-cooperative equilibrium, clearly showing how $\underline{s}^* \leq \underline{s}^{nc}$ and $\overline{s}^* \leq \overline{s}^{nc}$ (in some cases we may even have $\underline{s}^{nc} < \overline{s}^*$), but also showing that the first-best and the non-cooperative outcome coincide for very low and very high values of s_0.

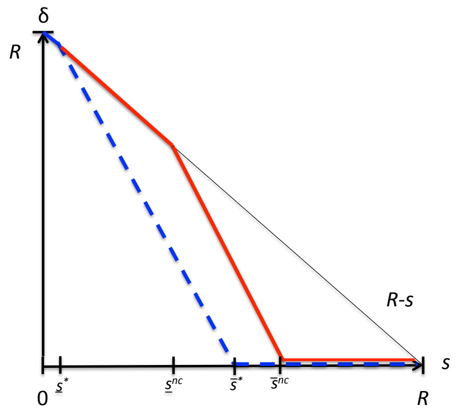

Figure 1: Illustration of policy function $\delta(s)$: blue dashed line is cooperative outcome, red line is non-cooperative outcome.

In the following, I will formally prove this dynamic consumption pattern, first for the first-best case in section 3.2.1 and then for the non-cooperative game in section 3.2.2. Finally, I will give a concrete example for a specific utility function and probability distribution in section 3.4. Before starting, let me – for simplicity and without loss of generality – equate the first-best with the consumption from a sole owner / social planner and postulate that the safe value of s is shared equitably among the symmetric players.

3.2.1 First-best outcome

Starting from a given s_0, the social planner has in principle two options: She can either stay at s_0 (choose $\delta = 0$), thereby ensuring the existence of the resource in the next period (as $L_s(0)$, the probability of not crossing the threshold, is 1). Alternatively, she can take a positive step into unknown territory (choose $\delta > 0$), potentially expanding the set of safe extraction possibilities to $s_1 = s_0 + \delta$, albeit at the risk of a resource collapse (as $L_s(\delta) < 1$ for $\delta > 0$). The social planner's "Bellmann equation" is thus:

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

The crux is, of course, that the value function $V(s)$ is *a priori* not known. However, we do know that once the planner has decided to not expand the set of safe extraction 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 ($s' = s$), so that the problem in the next period is identical to the problem in the current period. If moving in the next period would 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). Or seen from the other end, if one chooses $\delta = 0$ at some value of s, it must be optimal to choose $\delta = 0$ at the next period (nothing has changed except that time has passed), or it must not have been optimal to choose $\delta = 0$ in the first place.

Below, I will show that there indeed must be some value $\overline{s}^*$ so that for $s \geq \overline{s}^*$, it is optimal to choose $\delta = 0$. In this case, we know that $V(s) = \frac{u(s)}{1-\beta}$.

This leaves three possible paths when starting from values of s_0 that are below $\overline{s}^*$. The social planner could

a.) make one step and then stay,

b.) make several, but finitely many steps and then stay, c.) make infinitely many steps.

Suppose that a value at which it is socially 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 write down the function which is to be maximized as we know that the continuation value of staying at $s' = s + \delta$ is $\frac{u(s')}{1-\beta}$. Denote the social planner's valuation of taking exactly one step δ from the initial value s and then staying at s' forevermore by $\varphi(s)$ 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:

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

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. That is:

$$\delta^*(s) = \max \left\{ 0; \min \left\{ \arg\max \varphi(\delta; s); R - s \right\} \right\}. \tag{8}$$

With this explicit functional form in hand, I can show that it is better to traverse any given distance before remaining standing in one step rather than two steps. *A forteriori*, 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^*.

The intuition is simple: Given that it is optimal to eventually stop at some $s^* \geq \bar{s}$, the probability that the threshold is located on the interval $[s, s^*]$ is exogenous. Hence the probability of triggering the regime shift when going from s 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. In other words, given that one has to walk out into the dark, it is best to take a deep breath and get to it.

In short, the dynamics of the consumption pattern are stunted: For initial values of s below some threshold $\bar{s}^*$, it is optimal to make exactly one step and then stay at the updated value s' forever (provided T is not located between s and s', of course). For initial values of s above $\bar{s}^*$, it is optimal to never expand the set of safe consumption possibilities. In other words, any learning – if at all – is undertaken in the first period.

The first-best consumption is summarized by the following proposition:

Proposition 2. *The socially optimal total use of the resource is either* s_0 *for all* t *or* $s_0 + \delta^*(s_0)$ *for* $t = 0$ *and, if the resource has not collapsed,* s_1 *for all* $t \geq 1$.

Proof. First, I show that there exists some values s^* at which it is not optimal to further expand the set of safe consumption possibilities. In the second part, I show that for any $s \neq s^*$ it is better to make one step and then stay than to make two steps and then stay.

(1) There exists some values s^* at which it is optimal to stay. At least one such a value exists because at $s = R$ there is no other choice but to remain standing. Furthermore, values of $s^* < R$ exist when the survival function $L_s(\delta)$ is sufficiently small close to R. The intuition is that the marginal benefit from staying is strictly increasing with s, but the marginal benefit from making a small step is then decreasing close to R: While the potential reward is small (since the value function is bounded above by $u(R)/(1-\beta)$), the likelihood of triggering the regime shift becomes high. When it is known that there is a catastrophic threshold on $[0, R]$, we have $L_s(\delta) \to 0$ as $\delta \to R - s$ and there will always be values of $s^* < R$. To show this more formally, compare the value from staying at some value of s close to R, that is $s = R - \varepsilon$, to the value of making a step towards R so that one stays at $R - \delta$ (with $\delta \in (0, \varepsilon]$). I claim that for some ε, we have:

$$\frac{u(R-\varepsilon)}{1-\beta} \geq u(R-\delta) + \beta L_s(\delta)\frac{u(R-\delta)}{1-\beta} \tag{9}$$

Clearly, $\lim_{\varepsilon \to 0}\left[\frac{u(R-\varepsilon)}{1-\beta}\right] = \frac{u(R)}{1-\beta}$ but since $\delta \in (0, \varepsilon]$ and $L_s(\delta) \to 0$ as $\delta \to R - s$, we have $\lim_{\varepsilon \to 0}\left[u(R-\delta) + \beta L_s(\delta)\frac{u(R-\delta)}{1-\beta}\right] = u(R) < \frac{u(R)}{1-\beta}$.

(2) For any $s \neq s^*$ it is better to make any one step and then stay than to make two steps and then stay. Take any s and find δ^* as defined in equation (8) above. I then show that at some $\tilde{s}$ below s (with $\tilde{s} = s - \delta_1$) the payoff from choosing $\delta_2 = \delta_1 + \delta^*$ (i.e. taking one step) exceeds the payoff from first taking one step from $\tilde{s}$ to s and then taking the second step δ^*:

$$\begin{aligned} &u(\tilde{s} + \delta_1 + \delta^*) + \beta L_{\tilde{s}}(\delta_1 + \delta^*)\frac{u(\tilde{s} + \delta_1 + \delta^*)}{1-\beta} \\ &\qquad \geq u(\tilde{s} + \delta_1) + \beta L_{\tilde{s}}(\delta_1)\Big(u(\tilde{s} + \delta_1 + \delta^*) + \beta L_{\tilde{s}+\delta_1}(\delta^*)\frac{u(\tilde{s} + \delta_1 + \delta^*)}{1-\beta}\Big) \end{aligned} \tag{10}$$

The important thing to note at this stage is that: $L_{\tilde{s}}(\delta_1)L_{\tilde{s}+\delta_1}(\delta^*) = \frac{L(\tilde{s}+\delta_1)}{L(\tilde{s})}\frac{L(\tilde{s}+\delta_1+\delta^*)}{L(\tilde{s}+\delta_1)} =$

$\frac{L(\tilde{s}+\delta_1+\delta^*)}{L(\tilde{s})} = L_{\tilde{s}}(\delta_1 + \delta^*)$. Hence, (10) can, upon inserting $\tilde{s} = s - \delta_1$, be written as:

$$u(s+\delta^*) + \beta \frac{L(s+\delta^*)}{L(s-\delta_1)} \frac{u(s+\delta^*)}{1-\beta} \geq u(s) + \beta \frac{L(s)}{L(s-\delta_1)} u(s+\delta^*) + \beta^2 \frac{L(s+\delta^*)}{L(s-\delta_1)} \frac{u(s+\delta^*)}{1-\beta}$$

$$\Leftrightarrow$$

$$u(s+\delta^*) - \beta \frac{L(s)}{L(s-\delta_1)} u(s+\delta^*) + \beta(1-\beta) \frac{L(s+\delta^*)}{L(s-\delta_1)} \frac{u(s+\delta^*)}{1-\beta} \geq u(s)$$

$$\Leftrightarrow$$

$$\left[1 + \beta \frac{L(s+\delta^*) - L(s)}{L(s-\delta_1)}\right] u(s+\delta^*) \geq u(s) \tag{10'}$$

Now by the definition of δ^* we know that:

$$u(s+\delta^*) + \beta L_s(\delta^*) \frac{u(s+\delta^*)}{1-\beta} > \frac{u(s)}{1-\beta}$$

$$\Leftrightarrow$$

$$(1-\beta)u(s+\delta^*) + \beta L_s(\delta^*) u(s+\delta^*) > u(s)$$

$$\Leftrightarrow$$

$$\left[1 - \beta + \beta \frac{L(s+\delta^*)}{L(s)}\right] u(s+\delta^*) > u(s)$$

$$\Leftrightarrow$$

$$\left[1 + \beta \frac{L(s+\delta^*) - L(s)}{L(s)}\right] u(s+\delta^*) > u(s) \tag{11}$$

Since $L'(s) < 0$, we know that $0 > \frac{L(s+\delta^*)-L(s)}{L(s-\delta_1)} > \frac{L(s+\delta^*)-L(s)}{L(s)}$. Therefore, combining (10') and (11) establishes the claim and completes the proof:

$$\left[1 + \beta \frac{L(s+\delta^*) - L(s)}{L(s-\delta_1)}\right] u(s+\delta^*) > \left[1 + \beta \frac{L(s+\delta^*) - L(s)}{L(s)}\right] u(s+\delta^*) > u(s)$$

□

3.2.2 Non-cooperative equilibrium

Similar to section 3.2.1, I define s^{nc} as a safe value of the consumption possibility set at which staying can be supported by a non-cooperative Nash equilibrium. Furthermore, I denote the aggregate step size that leads to a value s^{nc} from a given value $s \neq s^{nc}$ by $\delta^{nc}(s)$. Finally, $\delta_t^{i*}(s)$ is in general the individually optimal step size at some t for some given s.

For a given strategy of the other players Δ^{-i}, the "Bellman equation" of player i is given by:

$$V^i(s, \Delta^{-i}) = \max_{\delta^i \in [0, R-s]} \left\{ u(s+\delta^i) + \beta L_s(\delta^i + \delta^{-i}) V^i(s+\delta, \Delta^{-i}) \right\} \tag{12}$$

Also here, the crux is that V^i is *a priori* unknown. Similar to above we suspect that the equilibrium policy will be to extirpate, stay, or make one step and then stay. Therefore, I denote by ϕ the value for player i to take exactly one step of size δ^i and then remain standing when the other players' strategy $\Delta^{-i} = \{\delta^{-i}, 0, 0, 0, ...\}$ is to do the same:

$$\phi(\delta^i; \delta^{-i}, s) = u\left(\frac{s}{N} + \delta^i\right) + \beta L_s(\delta^i + \delta^{-i}) \frac{u\left(\frac{s+\delta^i+\delta^{-i}}{N}\right)}{1-\beta} \tag{13}$$

With these definitions in hand, I can prove the following Proposition that describes the feedback Nash-equilibrium of the game.

Proposition 3. *The set of Markov-strategies*

$$\delta^{i*}(\delta^{-i}, s) = \begin{cases} 0 & \text{if} \quad s \geq \overline{s}^{nc} & \text{(14a)} \\ g(\delta^{-i}, s) & \text{if} \quad s \in (\underline{s}^{nc}, \overline{s}^{nc}) & \text{(14b)} \\ R - s - \delta^{-i} & \text{if} \quad s \leq \underline{s}^{nc} & \text{(14c)} \end{cases}$$

constitutes a feedback Nash equilibrium. That is, for $s_0 \geq \overline{s}^{nc}$ coordination to stay at s_0 can be supported as a Nash equilibrium. For $s_0 < \overline{s}^{nc}$ taking one step and then staying at $s_1 = s_0 + \delta^{nc}$ can be supported as a Nash equilibrium.

Proof. First note 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. Again, there will be some s^{nc} at which staying is a Nash equilibrium, because at least at $s = R$, there is no other choice. But parallel to the argument in Proposition 2, there will also be some $s^{nc} < R$ when s close enough to R and $L_s(\delta)$ becomes sufficiently small. Also here, there will always be values of $s^{nc} < R$ when it is known that there is a catastrophic threshold on $[0, R]$. Suppose all other players stay at $s = R - \varepsilon$, then for ε small, the value from staying at $s = R - \varepsilon$ is at least as large as the value of making a step towards R so that the updated value is $R - \delta$ (with $\delta \in (0, \varepsilon]$):

$$\frac{u\left(\frac{R-\varepsilon}{N}\right)}{1-\beta} \geq u\left(\frac{R-\varepsilon}{N}+\delta\right)+\beta L_s(\delta)\frac{u\left(\frac{R-\delta}{N}\right)}{1-\beta} \tag{15}$$

Parallel to the social optimum we have $\lim_{\varepsilon\to 0}\left[\frac{u((\frac{R-\varepsilon}{N})}{1-\beta}\right] = \frac{u(R/N)}{1-\beta}$. Again, since $\delta \in (0,\varepsilon]$ and $L_s(\delta) \to 0$ as $\delta \to R-s$ we have $\lim_{\varepsilon\to 0}\left[u\left(\frac{R-\varepsilon}{N}+\delta\right)+\beta L_s(\delta)\frac{u\left(\frac{R-\delta}{N}\right)}{1-\beta}\right] = u(R/N) < \frac{u(R/N)}{1-\beta}$.

Now, as there is some s^{nc} at which staying is a Nash equilibrium, there will be a last step at which this value is reached. Take some value s at which staying is not a Nash equilibrium. Suppose the strategy of the opponents is to take some step $\delta_1^{-i} < \delta^{nc}(s)$ and then some step $\delta_2^{-i*}(s+\delta_1^{-i}+\delta_1^i)$. The following calculations show that the best-reply from player i is to take only one step δ_1^{i*}. Hence $\delta_2^{-i*}(s+\delta_1^{-i}+\delta_1^i) = 0$ and the equilibrium will be to reach a value at which staying is a Nash equilibrium in one step.

For player i the payoff from making one step $\delta_1^{i*} = s^{nc}-s_0-\delta_1^{-i}$ exceeds the payoff from making two steps $\delta_1^i < s^{nc}-s_0-\delta_1^{-i}$ and $\delta_2^{i*} = s^{nc}-s_1-\delta_2^{-i*}$ when:

$$\begin{aligned} &u\left(\frac{s_0}{N}+\delta_1^{i*}\right)+\frac{\beta}{1-\beta}L_{s_0}(\delta_1^{i*}+\delta_1^{-i})u\left(\frac{s^{nc}}{N}\right) \\ &\quad\geq u\left(\frac{s_0}{N}+\delta_1^{i}\right)+\beta L_{s_0}(\delta_1^{i}+\delta_1^{-i})\left[u\left(\frac{s_1}{N}+\delta_2^{i*}\right)+\frac{\beta}{1-\beta}L_{s_1}(\delta_2^{i*})u\left(\frac{s^{nc}}{N}\right)\right]\end{aligned} \tag{16}$$

As for the coordinated case, $L_{s_0}(s_1-s_0)L_{s_1}(s^{nc}-s_1) = L_{s_0}(s^{nc}-s_0)$ so that (16) implies:

$$u\left(\frac{s_0}{N}+\delta_1^{i*}\right)-u\left(\frac{s_0}{N}+\delta_1^{i}\right) \geq \beta\left[\frac{L(s_1)}{L(s_0)}u\left(\frac{s_1}{N}+\delta_2^{i*}\right)-\frac{L(s^{nc})}{L(s_0)}u\left(\frac{s^{nc}}{N}\right)\right] \tag{17}$$

For clarity, write this inequality as $A-a \geq B-b$. This inequality holds because both $A > B$ and $a < b$. To see that $A > B$ note that u is an increasing and concave function so that $u(\frac{s_0}{N}+\delta_1^{i*}) > u(\frac{s_1}{N}+\delta_2^{i*})$ when $\frac{s_0}{N}+\delta_1^{i*} > \frac{s_1}{N}+\delta_2^{i*}$. Inserting $\delta_1^{i*} = s^{nc}-s_0-\delta_1^{-i}$, $\delta_2^{i*} = s^{nc}-s_1-\delta_2^{-i*}$ and $s_1 = s_0+\delta_1^i+\delta_1^{-i}$ in this inequality simplifies to $(N-1)(\delta_1^i+\delta_1^{-i}) > 0$, which is true. By the same argument, $a < b$ when $\frac{s_0}{N}+\delta_1^i < \frac{s^{nc}}{N}$. Re-write this as $N\delta_1^i < s^{nc}-s_0$. This inequality holds because it is implied by the definition that $\delta_1^i < s^{nc}-s_0-\delta_1^{-i}$ and $\delta_1^{-i} < \delta^{nc}(s)$.

The best-reply function $g(\delta^{-i}, s)$ in equation (14b) is therefore defined by the interior solution to the first-order-condition of maximizing $\phi(\delta^i; \delta^{-i}, s)$:

$$\begin{aligned}\phi'(\delta^i; \delta^{-i}, s) &= u'\left(\frac{s}{N} + \delta^i + \delta^{-i}\right) \\ &+ \frac{\beta}{1-\beta}\left[L_s'(\delta^i + \delta^{-i})u\left(\frac{s}{N} + \delta^i + \delta^{-i}\right) + \frac{1}{N}L_s(\delta^i + \delta^{-i})u'\left(\frac{s}{N} + \delta^i + \delta^{-i}\right)\right]\end{aligned} \tag{18}$$

For a symmetric step size $\delta^{-i} = (N-1)\delta^i$, we have:

$$\begin{aligned}\phi'(\delta^{nc}, s) = u'\left(\frac{s}{N} + \delta^{nc}\right) & \\ + \frac{\beta}{1-\beta}&\left[L_s'(N\delta^{nc})u\left(\frac{s}{N} + \delta^{nc}\right) + \frac{1}{N}L_s(N\delta^{nc})u'\left(\frac{s}{N} + \delta^{nc}\right)\right]\end{aligned} \tag{19}$$

The value of $\underline{s}^{nc}$ is defined by $\delta^{nc} = \frac{R-s}{N}$, which is the largest value of s at which equation (19) does not yet have an interior solution but $\phi'(\delta, s) > 0$ for all $\delta \in [0, R-s)$. Similarly, the value of $\overline{s}^{nc}$ is defined by $\delta^{nc} = 0$, which is the smallest value of s at which equation (19) no longer has an interior solution but $\phi'(\delta, s) < 0$ for all $\delta \in (0, R-s]$ □

Aggregate extraction in equilibrium is therefore given by:

$$\begin{array}{rll} R-s & \text{if} & s \leq \underline{s}^{nc} \\ N\delta^{nc}(s) & \text{if} & s \in (\underline{s}^{nc}, \overline{s}^{nc}) \\ 0 & \text{if} & s \geq \overline{s}^{nc} \end{array}$$

3.3 Comparative statics

In order to analyze how the extraction pattern changes with changes in the parameters, I will first show that δ^{nc}, the equilibrium expansion of the set of safe values is monotonically decreasing in s (Lemma 1). That is, the aggregate extraction pattern as a function of the prior knowledge about the set of safe consumption possibilities indeed looks qualitatively like it is shown in Figure 1 on page 12. This implies that the effect of an increase in the fundamentals β, $L_s(\delta)$, and R can be analyzed by investigating changes to $\phi'(\delta^{nc}, s)$.

Lemma 1. *The equilibrium step size δ^{nc}, defined by the interior solution to* (19)*, is decreasing in s.*

Proof. The intuition is simple: The more I know, the more cautious I am as I have more to lose and less to gain. The formal proof is somewhat tedious and therefore placed in the Appendix (page 32). □

We then have the following comparative statics results:

- The boundaries $\underline{s}^{nc}$ and $\overline{s}^{nc}$, and aggregate extraction for $s \in [\underline{s}^{nc}, \overline{s}^{nc}]$, is larger the more impatient the players are.
- The higher the maximum potential reward R, the larger the range where there a Nash-equilibrium that does not imply immediate extirpation.
- The more unlikely the regime shift (in terms of a first-degree stochastic dominance), the smaller the region at which no further experimentation is a Nash equilibrium.
- The "tradedy of the commons" manifests itself in the fact that aggregate extraction is increasing in N as long as u is sufficiently concave or N is sufficiently large.

First, as $\phi' = 0$ implicitly defines a monotonically decreasing function $\delta^{nc}(s)$ on $[\underline{s}^{nc}, \overline{s}^{nc}]$ (Lemma 1) and $\delta^{nc}(s)$ is bounded above by $R - s$ and below by 0, an increase in δ^{nc} will also lead to an increase in $\underline{s}^{nc}$ and $\overline{s}^{nc}$ respectively. Hence to prove the proposition's part with respect to β it is then sufficient to analyze $\frac{\mathrm{d}\phi'}{\mathrm{d}\beta}$. Note that the term in the squared brackets of (19) must be negative for $\phi' = 0$ to hold (since $u' > 0$). Therefore: $\frac{\mathrm{d}\phi'}{\mathrm{d}\beta} = \frac{2\beta[...]-\beta}{(1-\beta)^2} < 0$.

Second, to see the effect of an increase in R, note that this does not impact equation (19) directly, but it does have an effect on the first value $\underline{s}^{nc}$: As the diagonal line defining the upper bound of $\delta \in [0, R - s]$ shifts outwards, and $\delta^{nc}(s)$ is a downward sloping function steeper than $R - s$, the first value at which it is not optimal to extirpate the resource must be smaller.

Third, consider the condition for $\overline{s}^{nc}$, the smallest value at which it is a Nash equilibrium to remain standing and not experiment by enlarging the set of consumption possibilities (equation 15). Clearly, it cannot hold for a given $\overline{s}^{nc} = \hat{s}$ when the probability of a regime shift decreases (so that $\tilde{L}_s < L_s$) as this decreases the LHS but not the RHS of (15).

Finally, to show that aggregate extraction is non-decreasing with the number of players under certain conditions, I make the following argument: $\underline{s}^{nc}$, the largest value at which immediate extirpation is the only Nash equilibrium becomes larger when adding another player and $\frac{N}{N+1} > \frac{u'(\frac{R}{N})}{u'(\frac{R}{N+1})}$. Similarly, $\overline{s}^{nc}$ becomes smaller when subtracting a player and $\frac{N}{N-1} > \frac{u'(\frac{R}{N})}{u'(\frac{R}{N-1})}$ (where $\hat{s}$ is the first-value of no expansion in equilibrium for a given number of players N).

Consider the first case. For a given number of players N we have at a given $\underline{s}^{nc} = \hat{s}$ that $\phi'(\frac{R-s}{N}; \hat{s}) = 0$ and I show that $\phi'(\frac{R-s}{N+1}; \hat{s}) > 0$ when $\frac{N}{N+1} > \frac{u'(\frac{R}{N})}{u'(\frac{R}{N+1})}$:

$$
\begin{aligned}
&\phi'(\tfrac{R-s}{N+1}; \hat{s}) - \phi'(\tfrac{R-s}{N}; \hat{s}) > 0 \\
&\Leftrightarrow \\
&u'(\tfrac{R}{N+1}) - u'(\tfrac{R}{N}) + \frac{\beta}{1-\beta}\left[\left(u(\tfrac{R}{N+1}) - u(\tfrac{R}{N})\right) L'_{\hat{s}} + L_{\hat{s}}\left(\frac{1}{N+1}u'(\tfrac{R}{N+1}) - \frac{1}{N}u'(\tfrac{R}{N})\right)\right] > 0
\end{aligned}
$$

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

Now consider $\overline{s}^{nc}$, the first value of s at which the players can coordinate on remaining standing. For a given N we have that at a given $\overline{s}^{nc} = \hat{s}$, $\phi'_{[N]}(0; \hat{s}) = 0$ holds and I show that $\phi'_{[N-1]}(0; \hat{s}) < 0$ when $N-1$ (where the subscript denotes the number of players the condition is evaluated at):

$$
\begin{aligned}
&\phi'_{[N-1]}(0; \hat{s}) - \phi'_{[N]}(0; \hat{s}) > 0 \\
&\Leftrightarrow \\
&u'(\tfrac{\hat{s}}{N-1}) - u'(\tfrac{\hat{s}}{N}) + \frac{\beta}{1-\beta}\left[\left(u(\tfrac{\hat{s}}{N-1}) - u(\tfrac{\hat{s}}{N})\right) L'_{\hat{s}} + \frac{1}{N-1}u'(\tfrac{\hat{s}}{N-1}) - \frac{1}{N}u'(\tfrac{\hat{s}}{N})\right] > 0
\end{aligned}
$$

Again, the first part of the last line is positive due to concavity of u, the first term in the squared bracket is positive since $L'_s < 0$, but the last term in the squared bracket is positive only if $\frac{N}{N-1} > \frac{u'(\frac{R}{N})}{u'(\frac{R}{N-1})}$.

3.4 Specific example

For a given utility function and a given probability distribution of the threshold's location it is then possible to determine $\delta^*(s)$, $\delta^{nc}(s)$ and calculate the value function $V(s)$. Below, I do this for $u(c) = \sqrt{c}$ and a uniform probability distribution so that it is especially easy to explicitly solve (7) and (19). When the social planner thinks that every value in $[0, A]$ is equally likely to be the threshold, i.e. $f = \frac{1}{A}$, and accordingly $L_\delta(s) = \frac{A-s-\delta}{A-s}$, the optimal extension of the set of safe extraction possibilities is given by:

$$\delta^* = \frac{A - (1+2\beta)s}{3\beta} \tag{20}$$

There will only be an interior solution to (7) when $s \in [\underline{s}^*, \overline{s}^*]$. The values of $\underline{s}^*$ and $\overline{s}^*$ are given by:[2]

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

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

Exploiting the fact that due to symmetry we have $\delta^{-i} = (N-1)\delta^i$, I solve (19) for an interior equilibrium value δ^{nc}. Total non-cooperative expansion is then given by:

$$N\delta^{nc} = \frac{\big((1-\beta)N + \beta\big)A - \big((1-\beta)N + 3\beta\big)s}{3\beta} \tag{21}$$

Again, there will only be an interior equilibrium for $s \in [\underline{s}^{nc}, \overline{s}^{nc}]$ where these values are given by:

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

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

From inspection of (21) it becomes clear that the total non-cooperative consumption is increasing in N: $\frac{\partial[N\delta^{nc}]}{\partial N} = \frac{(1-\beta)(A-s)}{3\beta} > 0$. Furthermore, one can find the combination of parameters that would ensure a self-fulfilling prophecy of extirpation when staring from an initial value of $s = 0$ (no prior knowledge of a safe level of extraction). It is namely given by $N \geq \frac{\beta}{1-\beta}(3R - A)$, hence increasing in R and decreasing in β and A, as it is intuitive.

Figure 2 plots the value function for a uniform prior (with $A = R = 1$) and a discount factor of $\beta = 0.8$, illustrating how it changes as the number of player increases. The more players there are, the greater the distance of the non-cooperative value function (plotted by the blue solid diamonds) to the socially optimal value function (plotted by the black open circles). In particular when $N = 9$, one sees the region (roughly from $s = 0$ to $s = 0.12$) where it is individually rational to immediately extirpate the resource, and the large value of $\overline{s}^{nc}$ (roughly 0.62) when it first becomes individually rational to

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

remain standing. All in all however, this example shows that the threat of a irreversible regime shift is very effective when the common pool externality applies only to the risk of crossing the threshold. (At least for this specific utility function and these parameter values. Note that $\beta = 0.8$ in fact implies a unreasonably high discount rate, but it was chosen to magnify the effect of non-cooperation for a small number of players.)

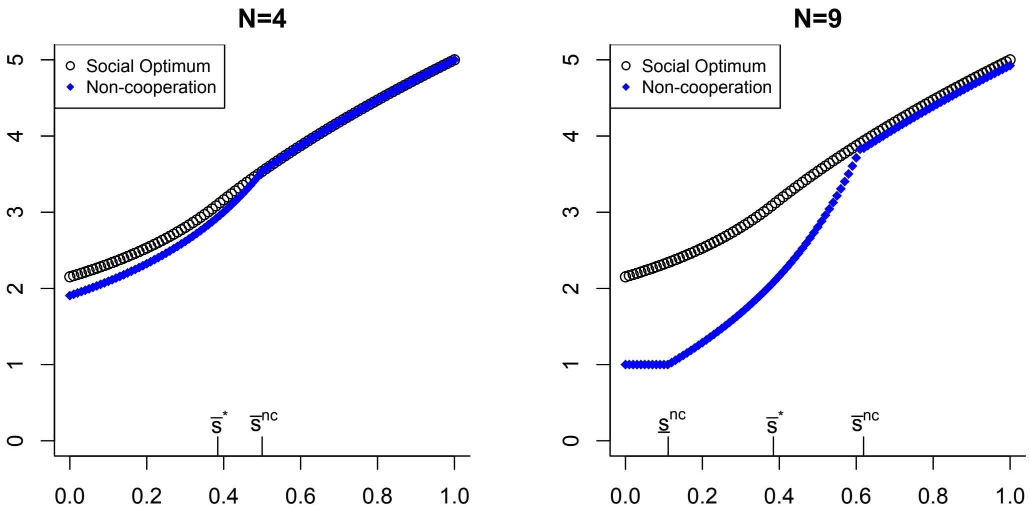

Figure 2: Illustration of cooperative and non-cooperative use of renewable resource: Value function $V(s)$ with $u(c) = \sqrt{c}$, $\beta = 0.8$, and $A = R = 1$ for $N = 4$ and $N = 9$.

4 Extensions

The above section shows how non-cooperative extraction is more aggressive than the first-best for most initial values, but that the threat of a threshold does indeed allow for coordination on the first-best (no expansion of the set of safe consumption values) when the initial s is high enough. For s below a critical value, there is a Nash equilibrium where players make exactly one step. This step is sub-optimally large, but ex-post (conditional on survival) staying at the updated value is socially optimal. In other words, the dynamics of the game are stunted in the sense that any learning – if at all – is undertaken in the first period. The is because a particular feature of thresholds is that learning is only affirmative: In my model (and maybe also in reality), there is no learning without experiencing. After having made a step the players only realize whether they have crossed the threshold or not. As also the relative gain from expanding the set of safe consumption possibilities does not change, there is no point in not doing all experimentation at once.

In this section, I will therefore explore to what extent this result is robust to various extensions of the model's structure.

4.1 Delay in the occurrence of the regime shift

Consider a situation where the players, in a given period, observe only with some probability whether they have crossed the threshold. In fact, it is not unreasonable to model the true process of the resource as hidden and that it will manifest itself only after some delay (see Gerlagh and Liski (2014) for a recent paper that focusses on this effect in the context of optimal climate policies). Hence, as time passes the players will update their beliefs whether the threshold as been located on the interval $[s_t, s_t + \delta_t]$. How does this learning impact the optimal and non-cooperative strategies? This becomes an extremely difficult question as – due to the delay – the problem is no longer Markovian. However, at least when the probability of observing a catastrophe that has been caused in the past is not increasing through time, it is possible to show that a delay has no impact on the learning dynamics:

First, note that a constant exogenous hazard of a regime shift (denoted by h) has no impact on the learning behavior. Such a hazard would change the formulation of equation (12) to:

$$V^i(s, \Delta^{-i}) = \max_{\delta^i \in [0, R-s]} \left\{ u(s + \delta^i) + (1-h)\beta L_s(\delta^i + \delta^{-i}) V^i(s + \delta, \Delta^{-i}) \right\} \qquad (22)$$

That is, the structure remains unchanged only, the effective discount factor decreases. A declining exogenous hazard would lead to declining discount rates and could lead to time-inconsistency (players may regret having been so aggressive in the first period). This would however only strengthen the “one-step-then-stay” feature of learning.

Second, note that today’s decisions will be exogenous tomorrow (unless time-travel should be invented). Hence, they may lead to the players regretting what they have done in the initial period, but it will not induce them to expand the set of consumption possibilities any further. In other words, the fact that the learning dynamics are stunted is robust to a delaying 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. They almost surely will differ, as delaying the consequences of crossing the threshold should increase

It is quite intuitive that the general strategy of making one step and then staying is still optimal: As the players only learn that they have crossed the threshold when the disastrous regime shift occurs, they cannot capitalize on the information that they gain from the updating of beliefs.

4.2 Non-constant or uncertain R

Previously, the upper bound of the resource, R has been treated as known and constant. In this subsection, I shall depart from this assumption (but note that T remains constant).

First, I consider the case when R is uncertain, but its expected value is the same (and by construction, the lower bound of its support must be above s). This uncertainty will induce the players to be more cautious when expanding the set of safe consumption possibilities (as the strict concavity of u implies risk aversion) but will have no further implications. The reason is that for the players, staying at the safe value s or $s + \delta$ (if expansion has been positive and not caused the regime shift), the true value of R will be irrelevant; what matters is the the utility from sharing s or $s + \delta$.

Two cases will have to be distinguished in the discussion of the situations when R increases or decreases, respectively. Does the growth or decline of R level out or does it continue without bounds (when growing) or approach zero (when declining)? When R declines, but the decline stops at some value above the initial value s, the payoff from staying at s will remain unchanged. This is not the case when R declines below the level of s at some point in time. Similarly, the temptation to extirpate the resource will increase as R increases. When the growth of R converges to some R^U, there could be a possibility that the players can coordinate on staying at some safe value. In contrast, when R grows without bounds, also the temptation to turn the resource into a mine will grow without bounds. Consequently, there will some point at which the player can now longer resist this temptation and extirpation will be guaranteed. I discuss the implications of these scenarios in turn.

Declining R, converging to $R^L \geq s$ A declining R that converges to some $R^L \geq s$ will have no impact on the value of staying at a given s. But it will potentially diminish the gain from expanding beyond s, namely if the socially optimal / non-cooperative expansion in the original problem would have exceeded R^L. In this case, there are two forces at play: On the one hand, the gain from expansion is smaller, which induces caution. On the other hand, the set of values at which extirpation is the only Nash equilibrium, or even socially optimal, will increase, too.

Declining R, approaching 0 Compared to the situation above, the value of staying at s will be reduced in any case. Importantly, as R_t approaches zero, there will be some point at which it is socially optimal to no longer stay at s (or go down with R_t when $R_t < s$) but to extirpate the resource. This point will obviously earlier the lower is β. This – socially optimal – termination of the consumption stream has dramatic consequences when the resource is use non-cooperatively. The fact that there will be a finite period at which it will be collectively optimal to, so to say, turn the renewable resource into a mine, leads – by backward induction – to an unraveling of the game so that the resource is extirpated in the first period. It will be individually optimal to pre-empt the collective extirpation.

In contrast to the scenarios of a declining R, the scenarios of a growing R could lead to repeated experimentation, namely when the (individually or socially) optimal expansion of the set of consumption possibilities was stopped by the initial upper bound of the resource R_0. I will therefore focus on this case in my discussion below.

Increasing R, converging to R^U Consider first the social optimum. As expansion was stopped at the initial R_0, and R_t grows with the passing of time t, there is scope for continued experimentation. Up to some point, it will be optimal to continue expanding the set of consumption possibilities (so to say to move with the upper bound of this set). At some point τ however, this will become too risky and the optimal action is to stop at some interior $s_\tau \in [R_{\tau-1}, R_\tau)$. As R_t continues to grow but converges to some R^U, two situation can arise: Either it is socially optimal to stay at s_τ even in light of the temptation of turning the renewable resource R^U into a mine, or not. Then it will be optimal to stay at s_τ for some interval before R_t has grown so large that it pays off to consume it immediately.

Whether the players can coordinate on staying at some $s_\tau \leq R^U$ or not has again dramatic consequences for the game. In the first case, expansion will occur at the same pace as R_t grows until it stops and stays at s_τ.[3] In the second case, backward induction will unravel the game and the resource will be extirpated in the very first period. Consider some R_t at which the temptation of unilaterally extirpating the resource is still so small that the players could coordinate on sharing R_t in perpetuity, but at R_{t+1} this would no longer be the case. Even though the players would collectively better off letting the resource grow to R_{t+1} and then turn it into a mine, it will be in each player's individual best interest to pre-empt this and unilaterally extirpate the resource in the present period.

Increasing R, growing without upper bound The situation now is the same as in the scenario where R increased to a point where the temptation of extirpation could no longer be resisted. Backward induction will unravel the game. For the socially optimal extraction pattern, requirements will have to be placed on the growth function of R and the discount factor β for the problem to be well-defined. If β is too close to 1 or R grows too fast, the value function would have no maximum as it would always be better to wait yet another period before turning the renewable resource into a mine.

[3]This expansion will in fact be identical to the socially optimal expansion in the beginning, but it will in general stop at a later value of the state space.

5 Discussion and Conclusion

The effect of potential regime shifts on the cooperative and non-cooperative use of environmental goods and services is polarizing. When the location of the threshold is known to be low, or if it is sufficiently likely that even a low level of consumption causes catastrophe, the game may exhibit prisoner-dilemma features: Although it would be optimal to sustain the resource at its current level of use, the only non-cooperative equilibrium will be the immediate extirpation of the resource.[4] In contrast, when the threshold is known to be high, or if it is sufficiently likely that the productive regime can be sustained even at a high level of consumption, the game changes into a coordination-problem: The threat of loosing the productive resource can effectively enforce the first-best consumption level. For intermediate values, the equilibrium will neither be extirpation nor status-quo consumption, but rather a one-time increase in consumption, expanding the set of safe consumption possibilities. This expansions will be inefficiently large, but if it has not caused the regime shift, the players will be able to coordinate on staying at the updated level. Staying at the updated level is *ex post* socially optimal when the externality applies only to the risk of a regime shift (i.e. any given level of safe consumption is efficiently shared among the agents).

These conclusions have been derived by using a general dynamic model that has placed only minimal requirements on the utility function (concavity and boundedness) and the probability distribution of the threshold (continuity). Nevertheless, there are a number of modeling assumptions that warrant discussion.

First, a prominent aspect of this model is that the threshold itself is not stochastic. The central motivation is that this allows concentrating on 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 modeling approach implies a clear demarcation between a safe region and a risky region of the state space. In particular, it implies that the edge of the cliff, figuratively speaking, is a safe – and in many cases optimal – place to be. The alternative approach, modeling the risk of a regime shift by a hazard rate acknowledges that, figuratively speaking, the edge of a cliff is often quite windy and not a particular safe place. This however implies, that the regime shift will occur with certainty as time goes to infinity, no matter how little of the resource is used; eventually there will be a gust of wind that is strong enough to blow us over the edge, regardless of where we stand. This is of course not very realistic either. But also on

[4]But note that extirpation is also the socially optimal course of action if the threshold is deemed to be very low and the status quo consumption would be of little value compared to the resource potential.

a deeper level one could argue that the non-stochasticity is in effect not a flaw but a feature: Let me cite Lemoine and Traeger (2014, p.28) who argue that "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, I have modeled the players to be identical. In the real world, players are rarely identical. One dimension along which players could differ could be their valuation of the future. However, *prima facie* it should not be difficult to show that any such differences could be smoothed out by a contract that gives a larger share of the gains from cooperation to more impatient players. Another dimension along which players could differ is their size or the degree to which they depend on the environmental goods or services in question. As larger players are likely to be able 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), it seems topical to analyze a situation where groups of players can form a coalition to ameliorate the negative effects of non-cooperation.

Third, one could investigate the effect of heterogenous beliefs about the existence and location of the threshold. Two recent papers have analyzed this (Agbo, 2014; Koulovatianos, 2010). 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 learning and experimentation.

Fourth, I have assumed the regime shift to be irreversible. This is obviously a considerable simplification. Groeneveld et al. (2013) have analyzed a problem of optimal learning about the location of a flow-pollution threshold in a setting where repeated crossings are allowed. In the current set-up it is however not clear whether the introduction of several thresholds would yield significant new insights in addition to making the analysis a lot more cumbersome. For example, if one presumes that crossing the threshold implies that one learns where it is,[5] the game turns into a repeated game. This may imply that cooperation is sustainable for sufficiently patient players (van Damme,

[5]Groeneveld et al. (2013) presume that crossing the threshold tells the social planner that the regime shift has occurred, but not where. This is a little bit like sitting in a car, fixing the course to a destination and then blindfolding oneself. Arguably conscious experimentation is more realistically described by saying that the course is set, but the eyes remain open.

1989), but 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.

Finally, the current model of uncertainty implies that it is, loosely speaking, pitch dark when the players take a step. It is only afterwards that they realize whether the disastrous regime shift has occurred or not. Exploring a threshold model where experimentation is not only affirmative but truly informative will be the task of future work. 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).

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*, forthcoming(0):1–15.

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

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

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.

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.

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

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.

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. (2010). A paradox of environmental awareness campaigns. MPRA Paper 27260, University Library of Munich, Germany.

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.

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. (2001). Optimal regulation of eutrophying lakes, fjords, and rivers in the presence of threshold effects. *American Journal of Agricultural Economics*, 83(4):972–984.

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

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.

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

Tsur, Y. and Zemel, A. (1995). Uncertainty and irreversibility in groundwater resource management. *Journal of Environmental Economics and Management*, 29(2):149 – 161.

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

Appendix

A.1 Proof that δ^{nc} is decreasing in s (Lemma 1).

Recall that the payoff from starting at s, taking a step δ^{nc}, and remaining standing thereafter when all other players follow the same strategy is given by:

$$\phi(\delta^{nc}, s) = u\left(\frac{s}{N} + \delta^{nc}\right) + \beta L_s(N\delta^{nc}) \frac{u\left(\frac{s}{N} + \delta^{nc}\right)}{1-\beta} \tag{13}$$

Accordingly the equilibrium choice of a positive expansion $\delta^{nc}(s)$ is the solution of $\phi'(\delta^{nc}, s) = 0$ where ϕ' is given by (19):

$$\phi'(\delta^{nc}, s) = u'\left(\frac{s}{N} + \delta^{nc}\right) + \frac{\beta}{1-\beta}\left[L'_s(N\delta^{nc})u\left(\frac{s}{N} + \delta^{nc}\right) + L_s(N\delta^{nc})u'\left(\frac{s}{N} + \delta^{nc}\right)\right] \tag{19}$$

with the second-order condition (where $\delta = N\delta^{nc}$ to simplify notation):

$$\phi''(\delta; s) = u'' + \frac{\beta}{1-\beta}\Big(NL''_s(\delta)u + (1+N)L'_s(\delta)u' + L_s(\delta)u''\Big) < 0 \tag{A-1}$$

To show that δ^{nc} is declining in s, I need to show that:

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

Since the denominator is negative when the second-order condition is satisfied, we have that $\frac{\mathrm{d}\delta^{nc}}{\mathrm{d}s} < 0$ when $\frac{\partial[\phi'(\delta^{nc};s)]}{\partial s} < 0$, so that the condition to check is (A-2):

$$\frac{\partial[\phi'(\delta^{nc}; s)]}{\partial s} = \frac{1}{N}u'' + \frac{\beta}{1-\beta}\left(\frac{\partial L'_s(\delta)}{\partial s}u + \frac{1}{N}L'_s(\delta)u' + \frac{\partial L_s(\delta)}{\partial s}u' + \frac{1}{N}L_s(\delta)u''\right) < 0 \tag{A-2}$$

Noting the similarity of (A-2) to the second-order condition (A-1), and realizing that (A-1) can be decomposed into a part A and a part B and that (A-2) can be decomposed into a part A and a part C, a sufficient condition for (A-2) to be satisfied is that $B > C$.

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

$$\underbrace{\frac{1}{N}\left[u'' + \frac{\beta}{1-\beta}\Big(L_s'(\delta)u' + L_s(\delta)u''\Big)\right]}_{A} + \underbrace{\frac{\beta}{1-\beta}\left(\frac{\partial L_s'(\delta)}{\partial s}u + \frac{1}{N}\frac{\partial L_s(\delta)}{\partial s}u'\right)}_{C} < 0 \tag{A-2'}$$

In order to show that $L_s''(\delta)u + \frac{1}{N}L_s'(\delta)u' > \frac{\partial L_s'(\delta)}{\partial s}u + \frac{1}{N}\frac{\partial L_s(\delta)}{\partial s}u'$, I use the first-order condition for an interior solution from (19) to write u' in terms of u:

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

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

$$L_s''(\delta) + L_s'(\delta)\left[\frac{-L_s'(\delta)}{\frac{1-\beta}{\beta} + L_s(\delta)}\right] > \frac{\partial L_s'(\delta)}{\partial s} + \frac{\partial L_s(\delta)}{\partial s}\left[\frac{-L_s'(\delta)}{\frac{1-\beta}{\beta} + L_s(\delta)}\right] \tag{A-3}$$

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-3) is indeed satisfied.

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

$\Leftrightarrow$

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

$\Leftrightarrow$

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

$\Leftrightarrow$

$$\left[N\frac{1-\beta}{\beta} + \frac{L(s+\delta)}{L(s)}\right]L(s) > L(s+\delta)$$

$\Leftrightarrow$

$$N\frac{1-\beta}{\beta}L(s) > 0$$