Journal of Public Economics 147 (2017) 30–49

Contents lists available at ScienceDirect

Journal of Public Economics

journal homepage: www.elsevier.com/locate/jpube

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

Florian K. Diekert

Department of Economics, Heidelberg University, Bergheimer Straße 20, Heidelberg 69115, Germany
CEES, Department of Biosciences, University of Oslo, PO-Box 1095 Blindern, Oslo 3017, Norway

ARTICLE INFO

Article history:
Received 6 April 2015
Received in revised form 24 September 2016
Accepted 24 January 2017
Available online 31 January 2017

JEL classification:
C73
Q20
Q54

Keywords:
Dynamic games
Thresholds and natural disasters
Learning

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.

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

E-mail address: florian.diekert@awi.uni-heidelberg.de.

http://dx.doi.org/10.1016/j.jpubeco.2017.01.004
0047-2727/© 2017 Published by Elsevier B.V.

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.

1.1. 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 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.

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