## Early Warning Signals<sup>\*</sup>

Florian Diekert<sup>1,2</sup>, Daniel Heyen<sup>3</sup>, and Frikk Nesje<sup>1,4</sup>

<sup>1</sup>Alfred Weber Institute for Economics, Heidelberg University
 <sup>2</sup>CEES, Department of Biosciences, University of Oslo
 <sup>3</sup>Center of Economic Research, ETH Zürich
 <sup>4</sup>Department of Economics, University of Oslo

January 30, 2020

#### Abstract

Many dynamic systems exhibit tipping points – they fundamentally change their character once a critical value, or threshold, is crossed. The location of such thresholds is typically unknown. In this project, we consider a resource economic model where an infinitely lived agent can learn about the location of threshold. The utility of the agent is increasing in consumption. The threshold is catastrophic, however, since the resource collapses if consumption exceeds the threshold. At each point in time, the agent chooses how much to increase consumption. This enables two types of learning. First, if the resource does not collapse, the agent learns that the threshold is located at higher consumption. This leads to successive experiments being increasingly cautious. Second, the possibly to obtain informative signals on the location of the threshold through consumption gives an additional incentive to experiment. Counterintuitively, the presence of such "early warning signals" can make experiments more risky.

JEL Classification: C61, D83, Q54

*Keywords*: Catastrophic regime shifts, early warning signals, tipping points, experimentation, learning.

## 1 Introduction

Many economic and ecological systems have complex dynamics that feature tipping points where the system suddenly shifts from a desirable regime to an undesirable regime. While the exact location of these tipping points is almost always unknown, it may be possible to receive "early warning signals" (or EWS, for short) that herald an impending

<sup>\*</sup>This research has been funded by the European Research Council Project NATCOOP (ERC StGr 678049). Correspondence: dheyen@ethz.ch.

collapse. A key question is whether EWS can be received, deciphered, and processed in time before it is too late (Biggs et al., 2009). In other words, will EWS necessarily lead to better management, or could they lull us in false security and lead to a slippery slope in risk taking? This paper is a first attempt to answer how the possibility of receiving an EWS affects the optimal management of dynamic systems.

The most important system that may exhibit tipping points is arguably the climate system (Lenton et al., 2008, 2019). Suggested tipping points include the collapse of the West Antarctic ice sheet (Feldmann and Levermann, 2015), a breakdown of the Mid Atlantic Overturning Circulation (Hawkins et al., 2011), or an altered carbon cycle due to boreal wildfires (Walker et al., 2019), with potential interactions leading to "domino effects" (Lemoine and Traeger, 2016; Rocha et al., 2018). Ecological systems that are likely to exhibit tipping points are tropical rain forests,<sup>1</sup> fisheries,<sup>2</sup> and coral reefs.<sup>3</sup> But also economic systems may exhibit tipping points. Both the two largest economic meltdowns of modern times, the great depression of the 1930s and the great recession of 2007-08 were, at least in part, caused by a bank-run (Gorton, 2010): Once a critical mass of actors has lost their faith in the banking system, there was no holding back.

The common element in all these examples is that the exact location of the tipping point is unknown, which makes managing these systems very hard. However, an important recent discovery has been the fact that these dynamic systems may show indications of the impending regime shift before it happens (Wiesenfeld, 1985; Kleinen et al., 2003). This phenomenon is known as "critical slowing down", which may serve as an EWS (Dakos et al., 2015). Critical slowing down describes the increase of variance and auto-correlation of the dynamic system close to a tipping point. Mathematically, the dominant eigenvalue of the Jacobian matrix tends to zero as a bifurcation point is approached (Held and Kleinen, 2004; Dakos et al., 2008).

For an illustrative example, picture a student in class that leans back and balances on a chair: There are two stable steady states in this system – sitting orderly at the desk, and laying on the floor. In between these two, there is also an unstable steady state in which the chair just balances. A small force perturbing an initial state at either of the stable points will have little effect and lead to a rapid movement back towards the respective initial state. In contrast, a small force perturbing the state close to the

<sup>&</sup>lt;sup>1</sup>For the Amazon, scientists estimate that an irreversible dieback caused by a breakdown of the self-sustaining hydrological system could occur when deforestation exceeds 20-25% (Lovejoy and Nobre, 2018).

 $<sup>^{2}</sup>$ The Canadian fishery for cod, once the most productive and valuable fisheries in the world, has collapsed in the 1990 has not recovered since (Frank et al., 2005).

<sup>&</sup>lt;sup>3</sup>Ocean acidification, pollution, and global warming massively harm coral reefs ecosystems that may face an irreversible collapse unless corrective action is undertaken (Hoegh-Guldberg et al., 2007).

tipping point will lead to a slow movement and may have a large effect (and hence the student's excitement in balancing at the tipping point).

Empirically, critical slowing down has been shown to anticipate abrupt changes in several systems, such as the past climate using paleontological time-series (Dakos et al., 2008; Lenton et al., 2008), wetland ecosystems using remote sensing data (Alibakhshi et al., 2017), and in experimentally controlled settings for yeast populations (Dai et al., 2012), daphnia populations (Drake and Griffen, 2010), and whole lake ecosystems (Carpenter et al., 2011).

The key application that we have in mind is the climate system. We take the perspective of a global government (henceforth: the agent) that seeks to maximize welfare by choosing the optimal amount of global warming. Tolerating more warming means increased consumption and hence higher short-run utility, but obviously increases the chance of crossing a tipping point, and may hence be disastrous for long-run utility. We isolate the issue of a potential tipping point and the effect of an EWS system by assuming that all non-catastrophic consequences of increasing global warming are subsumed in the current utility function. As a consequence, it is natural to assume that there is an interior value of global warming that maximizes welfare, even if there were no tipping point. We call this point L.

To focus attention on the inherent exploitation-exploration trade-off, we cast the problem of choosing the optimal amount of global warming as a location choice l on the domain [0, L] where a higher l yields higher utility, but also higher risk. That is, there is an unknown tipping point value  $T \in [0, L]$  such that utility is increasing in l as long as  $l \leq T$ , but once the chosen location exceeds the tipping point (l > T) utility is zero. In other words, were the location of T known, the optimal location that maximizes the discounted sum of utility would be to choose l = T indefinitely. However, because T is not known, the agent can learn with every choice of l' that exceeds the historically highest choice of l, whether the T is located in the interval between l and l'. If T is in (l, l'), the decision maker knows that choosing any l smaller or equal to l' is safe and will not cause the system to tip.

The existence of an EWS system now means that the agent not only learns whether or not she has crossed the tipping point, but – in case the tipping point has not yet been crossed – she also updates her belief about the value of T due to the potential reception of an EWS. In other words, there are two types of learning. One is by experience – whether or not the tipping point has been crossed. Figuratively speaking, this shifts the agent's belief upwards, but does not change its shape: Upon extending location from l



Figure 1: Illustration of the timing and updating of the agent's belief about the location of T. The top part of the figure shows the belief of the agent when choosing the location for time t+1. If the new location does not exceed the threshold, the agent gets the corresponding utility  $u(l_{t+1})$  and updates his belief. First, not crossing the threshold means that the original belief is truncated. For concreteness, we say this happens at time t + 0.5 (illustrated in the middle part of the figure). Second, the agent updates his belief differently, depending on whether she receives an early warning signal, or not. As the bottom part of the figure illustrates, not receiving an EWS implies a more optimistic belief about the location of the threshold near  $l_{t+1}$ .

to l' and not crossing the tipping point, the agent's prior belief about the location of the tipping point is now truncated to the domain (l', L) as the tipping point is not in (l, l'). Correspondingly, it is more likely that T is at any place on (l', L), but the relative likelihood between two location within (l', L) does not change (see the middle part of Figure 1 for an illustration).

The other type of learning comes from the reception of the EWS. This information changes the shape of the agent's belief. A model of EWS needs to describe both the probability of receiving a signal and the posterior belief that the reception of an EWS induces. We capture the former issue by a signal function, f, and assume that the signal is at least partly informative. In particular, we assume that the closer one is to the actual realization of the threshold T, the more likely it is to receive a signal. This means that the reception of an EWS makes the decision-maker more pessimistic about the location of the threshold. Conversely, the absence of an EWS makes the decision-maker more optimistic about the location of the threshold (see the bottom part of Figure 1 for an illustration). The solution to the planner's problem therefore depends on the initial belief as well as the history of location choices and received signals.

We find that the potential to obtain an EWS by increasing consumption induces successive learning in a model in which it is optimal to learn only once in the absence of an EWS (Diekert, 2017). It may therefore become optimal to experiment several times before eventually settling on a location  $l^*$  at which exploration stops.

On the one hand, EWS is welfare enhancing because successive experimentation may induce the decision-maker to stop earlier, or go further, than they would have done in absence of an EWS. On the other hand, EWS can actually imply a more risky optimal plan than without EWS even though the initial location choice is smaller than without the EWS (depending on signal realizations), see Figure 2.



Figure 2: Example of optimal plan, without and with early warning signals (EWS). The potential to obtain an EWS can induce successive learning and a more risky optimal plan (higher  $l^*$ ).

The intuition is as follows. At the margin, there is more information to be gained by choosing an additional small increases in l. Conditional on survival, the decision-maker may become more optimistic for given signal realizations. This implies stopping at a higher location  $l^*$ , and may deliver a slippery slope in risk-taking.

The contribution of our paper is to provide the first analysis of how a decision maker optimally reacts to the possibility of receiving an early warning signal. While there is an active literature on EWS in the natural sciences, there are – to the best of our knowledge – only three applications of the concept in economics.

First, Richter and Dakos (2015) investigate whether EWS occur in a simulations of a stylized socio-ecological system. They find that indeed upcoming regime shifts may be detected using socio-economic outcomes, such as individual profits. Second, Diks et al. (2019) analyze the observed time series of historical financial crises and find statistical evidence for critical slowing down before Black Monday 1987 but no EWS for more recent crises. Finally, Barrett and Dannenberg (2014) report on an incentivized experiment in which reduced uncertainty about the location of a tipping point made it easier for the participants to coordinate on the cooperative outcome. While these studies are important in showing that EWS may play an important role in economic applications, they do not consider active learning.

As active learning is the central aspect of our paper, our study is related to the large literature on experimentation/bandit models in economics (Bergemann and Välimäki, 2008). However, our model differs in two important ways from the workhorse models. First, the realization of the signal depends on the amount of experimentation. This may lead to a distribution of the signal arrival rate that is hump-shaped (Boyarchenko, 2018). Second, learning in our model combines active learning by experimentation with "affirmative learning" by experience (Diekert, 2017).

In addition to the literature on experimentation, our paper is of course closely related to the literatures on regime shift models. We bring those two strands of literature together in a resource economic model, but our focus is different than in (Liski and Salaniè, 2019), the paper closest to ours. While Liski and Salaniè study optimal pattern of experimentation and resource use when the threshold could be crossed but the catastrophic regime shift occurs later (relating to the concept "extinction debt"); we study a setting where the threshold has not yet been crossed and the decision-maker can learn about the threshold location.

Crépin and Nævdal (2019) organize the literature on regime shift models by distinguishing between papers that consider "time-distributed catastrophes" (TDC) (such as Cropper, 1976; Polasky et al., 2011; Cai and Lontzek, 2019) and papers that consider "state-space-distributed catastrophes" (SDC) (such as Tsur and Zemel, 1995; Nævdal, 2006; Lemoine and Traeger, 2014).

In the "TDC" class of models, the regime shift risk is captured by a hazard rate that depends on time, and the catastrophe consequently occurs with probability one as  $t \to \infty$ . In the "SDC" class of models, it is possible to secure the sustainability of the system by freezing the state variable at its current state. Without a change in the state variable, the probability of a regime shift is zero. In this latter class of models, the optimal response to regime shift risk is often non-monotonic. Sakamoto (2014) shows that this non-monotonicity is amplified in a non-cooperative setting. In simple terms, agents try to grab what they can before it is too late when catastrophe avoidance becomes unlikely, but cooperation and caution increases when the catastrophe may be avoided.

Crépin and Nævdal (2019) then propose a hybrid of both standard approaches, that they call "inertia risk". Similar to (Liski and Salaniè, 2019), this intermediate case of "inertia risk" may produce path dependency and stabilization targets at which the occurrence of a catastrophe is neither ruled out nor occurs with probability one. Different from (Liski and Salaniè, 2019), who model a lag between the triggering the catastrophe and feeling the consequence, Crépin and Nævdal (2019) model a lag between the increase in the hazard rate and the regime shift.

Whether the catastrophe occurrences is distributed along the dimension of time, or state-space, has dramatic consequences for the prospects of learning. In the former class of models, there is no learning and the long-run steady states do not depend on initial conditions. For state-space-distributed catastrophes, in contrast, learning is extreme. Once a given location in state space has been explored and no regime shift has occurred, it is known to be safe.

## 2 Model

This section presents the general model that we use to explore the effect of EWS on the optimal management of a dynamic resource system. As stated in the introduction, the key application that we have in mind is climate change. The agent (a hypothetical world government) controls the optimal amount of global warming. Up to a point L, the marginal benefits of more warming exceed the marginal cost of more warming. However, the agent is worried that the climate system exhibits a catastrophic tipping point Tbefore the amount of global warming L is reached.

The location of the threshold is determined and constant, but unknown. In other words, the agents holds a belief p about the location of the tipping point, or threshold, T

in [0, L]. Define  $p: [0, L] \to \mathbb{R}$  and  $\int p(T)dT = 1$ . At each step in time t = 0, 1, 2, ..., the agent can choose an amount of warming for the next period  $l_{t+1}$  (henceforth: a location  $l_{t+1}$ ). Hence, the action space is  $\mathcal{L} = [0, L]$ . Choosing  $l_{t+1}$  yields a reward of  $u(l_{t+1})$  in the next period, where u is increasing and concave. However, when  $l_{t+1}$  exceeds the threshold T, the climate system tips and utility in the next and all subsequent periods is normalized to zero.<sup>4</sup> The agent maximizes the discounted sum of period utilities by choosing a sequence of locations  $l_{t+1}$  starting from  $l_0 = 0$ . To keep track of whether the catastrophe has occurred yet, we introduce the variable  $C_t \in \{0, 1\}$ , where  $C_t = 0$  indicates that the threshold has not been crossed at time t.

Importantly, the agent may receive an early warning signal about the tipping point. The signal itself is not informative about the exact location of T, but the probability of receiving a signal is higher, the closer the chosen location is to the threshold realization. We denote the signal realization at time t by  $s_t \in \{0, 1\}$  with

$$\Pr(s_t = 1 | T, l_{t+1}, l_t) = f(T, l_{t+1}, l_t) = \begin{cases} K e^{-\lambda(T - l_{t+1})} & \text{if } l_{t+1} > l_t, \\ k e^{-\lambda(T - l_{t+1})} & \text{if } l_{t+1} \le l_t, \end{cases}$$
(1)

where  $\lambda \in [0, \infty)$  and  $0 \le k \le K \le 1$ . The parameter K in equation (1) is the sensitivity of the EWS. In other words, the probability to hear a signal when the threshold is at the chosen location  $l_{t+1}$  is K, meaning that K is the true positive rate (*i.e.* the sensitivity), while 1-K is the false negative rate or "miss rate". This implies that when the agent does not receive a signal, she knows with probability K that the threshold is not at the chosen location ("no news is good news").

Moreover, we distinguish the case where k = 0, in which no early warning signal can be received except new and unknown area of the state space is explored, and the case when k > 0, in which it is possible to receive an early warning signal even if no active experimentation is undertaken in the current period.

The parameter  $\lambda$  is the "EWS shape parameter" or "EWS distance specificity". It can indeed be shown that the higher  $\lambda$ , the higher the specificity of a certain hypothesis test (this specificity also depends on K though). The limit K=1 and  $\lambda \to \infty$  corresponds to an EWS system of maximal sensitivity (=1) and maximal specificity (=1). The relationship of EWS to hypothesis testing is further explained in Appendix A-2.

Before we can formulate the value function corresponding to this problem, we detail

<sup>&</sup>lt;sup>4</sup>Note that the assumption that choosing  $l_{t+1} = T$  does not cause the regime shift allows us to model the case where there is no tipping in climate system simply by having positive mass of the probability distribution of T at L.

the timing of events, the updating of beliefs, and the probability of receiving an EWS.

#### 2.1 Timing of events and Bayesian updating

At time t, the agent's state is fully described by  $(C_t, l_t, p_t)$ , where  $C_t = 0$  indicates that the threshold has not been crossed yet. If  $C_t = 1$  for any  $t = \theta$ , we have  $C_t = 1$  for all  $t > \theta$  (the regime shift is irreversible) and the agent can neither make any choices nor receive any utility. Thus, we only look the case  $C_t = 0$  in the following.

The timing is as follows. At the beginning of period t, the agent is at location  $l_t$  and receives utility  $u(l_t)$ . Moreover, the agent holds a belief  $p_t$  about the threshold position T. Clearly,  $p_t(T) = 0$  for all  $T < l_t$ . Next, the agent chooses a location  $l_{t+1}$ , with  $0 \le l_{t+1} \le L$ .

After choosing the location  $l_{t+1}$ , the agent gets essentially two pieces of information. For concreteness, say this happens at t + 0.5. At that time, she first learns whether she has crossed the threshold or not. The probability of crossing the threshold when moving the location to  $l_{t+1}$  is  $\int_{l_t}^{l_{t+1}} p_t(\tau) d\tau$  and we hence define the probability of not crossing the threshold:

$$P(C_{t+1} = 0; l_t, l_{t+1}, p_t) = 1 - \int_{l_t}^{l_{t+1}} p_t(\tau) \,\mathrm{d}\tau.$$
<sup>(2)</sup>

If the agent has crossed the threshold,  $C_{t+1} = 1$ , payoffs for the rest of the game are set to zero and no further action is possible. If she has not crossed the threshold,  $C_{t+1} = 0$ , and the agent updates her belief about threshold positions to:

$$\tilde{p}_{t+0.5}(T) = \begin{cases} 0 & \text{for } T < l_{t+1}, \\ \left( \int_{l_{t+1}}^{L} p_t(\tau) \, \mathrm{d}\tau \right)^{-1} p_t(T) & \text{for } T \ge l_{t+1}. \end{cases}$$
(3)

If the threshold has not been crossed, the second piece of information the decisionmaker gets at time t + 0.5 is to see whether she observes an early warning signal, s = 1, or not, s = 0. We denote the probability to receive an EWS, *conditional on the new location not crossing the threshold*, by q. We have:

$$q(s=1|C_{t+1}=0; l_t, l_{t+1}, \tilde{p}_{t+0.5}) = \int_{l_{t+1}}^{L} f(\tau, l_t, l_{t+1}) \tilde{p}_{t+0.5}(\tau) \,\mathrm{d}\tau, \tag{4}$$

where  $f(\tau, l_t, l_{t+1})$  is the early warning signal function, equation (1).

Depending on whether the agent receives an EWS (s=1) or nor (s=0), she updates

her belief  $\tilde{p}_{t+0.5}$  to  $p_{t+1}$ , where

$$p_{t+1}(T) = \begin{cases} \frac{f(T, l_t, l_{t+1})\tilde{p}_{t+0.5}(T)}{\int_{l_{t+1}}^{L} f(\tau, l_t, l_{t+1})\tilde{p}_{t+0.5}(\tau) \,\mathrm{d}\tau} & \text{if } s = 1, \\ \frac{(1 - f(T, l_t, l_{t+1}))\tilde{p}_{t+0.5}(T)}{\int_{l_{t+1}}^{L} (1 - f(\tau, l_t, l_{t+1}))\tilde{p}_{t+0.5}(\tau) \,\mathrm{d}\tau} & \text{if } s = 0. \end{cases}$$

$$(5)$$

Notice that, due to (3), this can also be written as

$$p_{t+1}(T) = \begin{cases} \frac{f(T-l_{t+1})p_t(T)}{\int_{l_{t+1}}^{L} f(\tau, l_t, l_{t+1})p_t(\tau) \, \mathrm{d}\tau} & \text{if } s = 1, \\ \frac{(1-f(T-l_{t+1}))p_t(T)}{\int_{l_{t+1}}^{L} (1-f(\tau, l_t, l_{t+1}))p_t(\tau) \, \mathrm{d}\tau} & \text{if } s = 0. \end{cases}$$
(6)

This completes one time step. The threshold has still not been crossed,  $C_{t+1} = 0$ , the agent is located at  $l_{t+1}$ , and she holds beliefs about the threshold position  $p_{t+1}(T)$ , where  $p_{t+1}(T) = 0$  for all  $T < l_{t+1}$ .

#### 2.2 Bellman equation and value function

At time t, the agent is in state  $C_t = 0$ , at location  $l_t$ , and has belief  $p_t(T)$ . We write the pre-event value function as  $V(l_t, p_t) \equiv V(C_t = 0, l_t, p_t)$  and normalize the constant post-event value  $V^C \equiv V(C_{t+1} = 1, l_t, p_t) = 0$ . Moreover, we express the posterior belief after choosing  $l_{t+1}$ , conditional on not crossing the threshold, by  $G(s = 1; l_{t+1}, p_t)$  if an EWS has been received, and by  $G(s = 0; l_{t+1}, p_t)$  if no EWS has been received.

Let  $\beta \in (0, 1)$  be the discount factor. The Bellman equation is then:

$$V(l_t, p_t) = \max_{l_{t+1} \in [0, R]} \left\{ u(l_t) + \beta P(l_t, l_{t+1}, p_t) \times \left[ q(l_t, l_{t+1}, p_t) V(l_{t+1}, G(s = 1; l_{t+1}, p_t)) + (1 - q(l_t, l_{t+1}, p_t)) V(l_{t+1}, G(s = 0; l_{t+1}, p_t)) \right] \right\}.$$
(7)

In Appendix A-1, we derive the corresponding first-order-condition and show that the continuation value of the Bellman equation depends only on  $p_t$  as the current (safe) location  $l_t$  is encoded in the belief. A closed-form solution of the problem, however, seems unattainable. Therefore, we use numerical methods to provide a first set of results.

## **3** Numerical solution

We solve the model numerically by backward recursion. For the implementation, it will be useful to normalize the lower value of the action space to 1 and to discretize the space. Hence,  $\mathcal{L} = \{1, 2, ..., L\}$ .

We approximate the original problem, value function (7) at t = 0, by considering a finite time horizon version of the problem. Our solution strategy is to solve for the original problem sequentially from the final time period problem. We denote the final time at which a decision is taken by  $t_{max}$  and the ultimate period in which no more decisions are taken by  $t_{\infty}$ .

In principle we need to consider all attainable paths of location choices and signal realizations (experienced histories from the initial point in time),  $\{l_t, s_t\}_{t=0}^{t_{\infty}}$ . To increase computational speed, we impose the restriction  $l_{t+1} \geq l_t$ , thereby ruling out location choices that clearly cannot be optimal.

It will be useful to distinguish between the problem of the agent before and after signals are realized. The way we account for this is by distinguishing between times t, t + 0.5 and t + 1. As discussed above, at the beginning of period t, utility from the location  $l_t$  is received, and then the location for next period  $l_{t+1}$  is chosen. Time t + 0.5refers to the time after the location choice is made but before the signals are realized. Time t + 1 refers to the segment of time after signals are realized (the beginning of the new period).

We solve the original problem by backward recursion from time  $t_{\infty}$ . Write the terminal value as the utility of staying at  $l_{t_{\infty}}$  forever, which is the value of the discounted stream of  $u(l_{t_{\infty}})$ . Assuming logarithmic utility,

$$V_{t_{\infty}}(l_{t_{\infty}}) = \frac{\log l_{t_{\infty}}}{1 - \beta}.$$
(8)

The value when location choices needs to be made, that is at any time  $t \in \{1, 2, ..., t_{\text{max}}\}$ (for now ignoring the original problem t = 0), can be written

$$V_t(l_t, p_t) = \max_{l_{t+1} \in [l_t, L]} \Big\{ u(l_t) + \beta \Big[ P(l_t, l_{t+1}, p_t) \cdot V_{t+0.5}(l_{t+1}, \tilde{p}_{t+0.5}) \Big] \Big\}.$$

The value at time t is the utility of the location  $l_t$  as well as the discounted continuation. Conditional on survival, the discounted continuation is the time t + 0.5 value, the value of the problem before signals are realized.

Let us now express the value after location choices are made but before signals are

realized, that is any time  $t \in \{0.5, 1.5, \dots, t_{\max} - 0.5\}$ :

$$V_{t+0.5}(l_{t+1}, \tilde{p}_{t+0.5}(T)) = q(l_t, l_{t+1}, \tilde{p}_{t+0.5}) \cdot V_{t+1}(l_{t+1}, G(s=1; l_{t+1}, p_t))$$
(9)  
+  $\left(1 - q(l_t, l_{t+1}, \tilde{p}_{t+0.5})\right) \cdot V_{t+1}(l_{t+1}, G(s=0; l_{t+1}, p_t))$   
where:  $q(l_t, l_{t+1}, \tilde{p}_{t+0.5}) = \int_{l_{t+1}}^{L} f(T, l_{t+1}, l_t) \tilde{p}_{t+0.5}(\tau) \, \mathrm{d}\tau$ 

The value at time t+0.5 is the expected value at time t+1, weighted by the probabilities of signal realization and no signal realization.

Finally, the original problem is symmetric to the time t problem when t = 0. We thus repeat the same maximization process back to the original problem. We solve this problem using Matlab (the code is discussed in Appendix A-4).

## 4 Preliminary results

In this section, we present the numerical solutions to the problem described above. We will first present the optimal location choices given a specific EWS system, and then discuss the effect of the EWS system on the value function and the total risk taken by the agent.

#### 4.1 Optimal location choices

We plot the optimal control paths given the parameter values  $t_{\text{max}} = 5$ ,  $\beta = 0.95$ , and an unbiased uniform initial prior about the location of T on  $\mathcal{L} = \{1, 2, ..., 6\}$ . In doing so, we will distinguish between different EWS systems along two critical dimensions. For reference, we repeat the signal function here:

$$f(T, l_{t+1}, l_t) = \begin{cases} Ke^{-\lambda(T - l_{t+1})} & \text{if } l_{t+1} > l_t, \\ ke^{-\lambda(T - l_{t+1})} & \text{if } l_{t+1} \le l_t, \end{cases}$$

First, we set the parameter K in the signal function equal to 1, or to 0.5. Recall that the parameter K refers to the sensitivity of the signal. Setting K = 1 means that the "miss rate" is zero, that is, there are no false negatives. When the agent does not receive a signal in this case, she is certain that the next location is safe. This is of course an extreme assumption, which is why we also explore the case K = 0.5.

Second, we differentiate whether the agent can only learn when making an active

experiment (choosing  $l_{t+1} > l_t$ ), or when she can also receive a signal when not making an active experiment (choosing  $l_{t+1} \le l_t$ ). The former case of only learning under active experimentation is equivalent to setting the parameter k in the signal function equal to zero, and to explore the latter case we set k = K. In our presentation of the results, we will first start with the extreme case of no learning without active experimentation.

Every plot in Figures 3 to 6 shows the optimal location choice over time. In all figures, the three plots correspond to different trajectories of observing EWS or not: the leftmost plot shows the optimal location path when the agent never obtains an EWS (s=0 at all time steps t = 0.5, 1.5, 2.5, ...). Similarly, the rightmost plot shows the case in which the agent always obtains EWS (s=1 at all time steps t = 0.5, 1.5, 2.5, ...). The plot in the middle shows the case of an alternating sequence, where the agent first does not receive an EWS at t=0.5, then receives an EWS at t=1.5, does not receive an EWS at t=2.5, etc. In all plots we show the optimal trajectories under different values of the EWS shape parameter, ranging from  $\lambda=0$  (equivalent to the absence of an EWS system) to  $\lambda=\frac{3}{2}$  (where the probability to receive an EWS if the threshold is 5 steps away is already below 1% of the probability to receive an EWS if the threshold is at the current location).

To discuss the optimal paths in an intuitive way and in easy language, we will use words like "moving" or "walking" to describe situations in which the agent chooses  $l_{t+1} > l_t$  and "standing still" or "stopping" to describe situations in which the agent chooses  $l_{t+1} = l_t$ .

The first thing to note is that when the agent does not have an early warning signal system to inform the decision ( $\lambda=0$ , red curves), then the agent may move at t=0, but definitely not afterwards. This essentially replicates the finding in Diekert (2017). This is holds independent of whether k=0 (Figures 3 and 4) or k=K (Figures 5 and 6).

#### **Result 1** In absence of an EWS, any experimentation is undertaken in the first period.

Focussing on the case when the agent can only learn when making an active experiment (k=0; Figures 3 and 4), we first see that the agent never resumes walking once she has stopped. This is intuitive, as her belief cannot change without making a step into unknown territory, and if it was optimal to not walk further at some time t for a given belief, it must also be optimal to not walk at time t + 1 for the same belief.

Second, we see that the agent may take another step after not receiving an EWS; compare the move at time t=2 in plot (a) of Figure 3 with the corresponding move after receiving an EWS (plot (c) of Figure 3). This pattern reflects the fact that no news are good news. In fact, when the agent does not receive an EWS, she may take several



Figure 3: Optimal paths for model parameters K=1, k=0



Figure 4: Optimal paths for model parameters K=0.5, k=0

steps. Depending on the shape parameter  $\lambda$  and the signal sensitivity K, the initial step may be larger than subsequent steps.

It is however not the case that the agent always continues to move upon not receiving an EWS. As Figure 4 (a) illustrates, the agent stops after two steps, despite not receiving EWS. The reason is that the sensitivity of the signal is low, such that the agent is not sufficiently certain that the threshold is not located in the next step. The "miss rate" of the EWS is too high for K=0.5.

Irrespective of whether K=1 or K=0.5, we see that the agent stops whenever her belief gets more pessimistic due to the reception of an EWS. This can be clearly seen by comparing plot (b) of either Figures 3 and 4) with plot (c) of the same Figure. At time t=0.5, after having chosen  $l_1$ , the agent hears no news in plot (b) but hears a warning signal in plot (c). In plot (c) then, the optimal location choice at time 2 is then to choose  $l_2 = l_1$ , while in plot (b), the optimal choice is  $l_2 > l_1$  (except for  $\lambda=0$ , where the signal is not informative). At time t=1.5, however, the agent hears bad news also in plot (b) and does hence not move any further. In other words, once the agent receives an EWS she stops, to never walk again. The following statement summarizes our findings.

**Result 2** When EWS are only received after active experimentation (k = 0), the agent stops at the current (safe) location when she receives an EWS. When not receiving an EWS, the agent may continue to move if she is sufficiently optimistic, but once the agent stays, she stays forever.

We now discuss plots where learning is possible even when standing still (k > 0;Figure 5 and 6). Similar to the case when k = 0, it is always optimal to stay for the agent upon observing an EWS. However, it can happen that the agent moves again after staying. This happens when the good news of no EWS comes in, see for instance plot (b) in Figure 5 for  $\lambda \neq 0$ . Here, the agent receives an EWS at t=1.5 after moving to  $l_2$  and then chooses  $l_3 = l_2$ . At t=2.5, however, she does not receive an EWS, and as a consequence she chooses  $l_4 > l_3$ .

Not receiving an EWS (s=0), however, is not sufficient for making a move: The confidence that the next step is safe can still be too low. This is illustrated in Figure 6 (b), which shows the optimal path for the same sequence of signal realizations as in Figure 5 (b), but for a lower sensitivity of the test.



Figure 5: Optimal paths for model parameters K=1, k=1

Figure 5 and 6 suggest that the first location choice (which only depends on prior and expectations, but not on learning) is fundamentally different from all other location choices (which can be conditioned on realizations of EWS) in terms of the dependence of the shape parameter  $\lambda$ . On the one hand, the initial move gets more cautious as the EWS structure gets more specific (i.e. with higher  $\lambda$ ). On the other hand, later location choices are more daring when  $\lambda$  is higher (cf. Figure 6 (a), where good news, i.e. the absence of an EWS, induces the agent only for a high value of  $\lambda$  to move). Note that this pattern is not monotonic: While the agent finds it optimal to take another step after



Figure 6: Optimal paths for model parameters K=0.5, k=0.5

receiving no news about the threshold at t=4.5 for  $\lambda=1$  (the agent chooses  $l_5>l_4$  in this case), she does not walk for  $\lambda=0.5$  or  $\lambda=1.5$  (the agent chooses  $l_5=l_4$  in these cases).

Note further how EWS sensitivity (i.e. parameter K) impacts optimal paths. What the comparison of Figure 5 with Figure 6 demonstrates is that lower EWS sensitivity (K=0.5) pushes the optimal paths towards the behavior that we know from the case of no EWS. Figure 6 (b) is a nice case in point: Under the alternating signal sequence (under which the decision-maker alternatingly moved and stayed in the maximal sensitive EWS structure K=1), the agent's optimal path coincides with the optimal path in the absence of an EWS structure. In fact, for k=0.5, neither the sequence of signal realizations, nor the parameter  $\lambda$  plays a role. In all cases, the agent finds it optimal to only experiment in the first period. The following statement summarizes our findings.

**Result 3** The more informative the EWS system, the more responsive is the agent to receiving or not receiving a signal.

Whether the agent can only learn when actively exploring new locations, or only receive EWS when standing still has an important effect on the optimal paths. When having to actively explore unknown locations to receive EWS, the agent becomes more cautious (makes a smaller initial step), which can be seen by comparing the corresponding plots of Figure 4 and 6. Here the agent always takes a step of size one under high  $\lambda$ and k=0, but a step of size 2 under k=0.5. The reason for this difference could be due to the interplay of time discounting and learning. Under k=0.5, the agent can learn also when standing still so that she takes the increased risk of moving to  $l_1=2$  immediately and, conditional on not crossing the threshold, she enjoys the payoff from this location for another period while waiting to observe another signal. Under k=0, this is not possible, so that the agent moves to location 2 in several steps. How exactly these aspects affect optimal paths in more general cases is an important area for future research.

#### 4.2 Total probability of crossing the threshold and the value of EWS

We calculate the risk of crossing the threshold as follows. We go through each potential threshold position T and calculate the probability that the agent crosses that threshold. The agent takes the optimal path through the tree that is jointly determined by deterministic location choices by the agent and stochastic realizations of the EWS system.

Focusing on a certain threshold T enables us to calculate the probability of a certain path: for each current location  $l_t$ , and the optimal next location  $l_{t+1}$ , we can calculate the probability of receiving an EWS as  $f(T, l_{t+1}, l_t)$ , cf. equation (1). Respectively, the probability of not receiving an EWS is given by  $(1 - f(T, l_t, l_{t+1}))$ . The probability of following a certain path is then the product of all those factors over all locations the agent has visited. For a given threshold T we know whether a path is safe, i.e. all  $l_t \leq T$ , or not. Accordingly we get, for a given threshold T, the probability of crossing the threshold. In a last step we calculate the total probability of crossing the threshold by weighing each threshold T by its probability  $p_0(T)$  and then summing over all possible threshold positions T.<sup>5</sup>



Figure 7: Contour plots showing the total probability of crossing T along the optimal path

Figure 7 shows the risk of crossing the threshold as a contour plot, where we vary the two relevant parameters of the EWS structure, namely the sensitivity K on the vertical axis and shape parameter  $\lambda$  on the horizontal axis. Cold and warm colors represent a small and high probability of crossing the threshold, respectively. We show results for three values of the discount factor,  $\beta=0.95$ ,  $\beta=0.8$  and  $\beta=0.5$ , and the numerical simulation is based on six locations and time steps (locations ranging from l=1 to L=6 and time ranging from 0 to  $t_{max}=5$ ) as in the plots of the optimal paths above. Here,

<sup>&</sup>lt;sup>5</sup>The underlying assumption here is that the agent's belief about the probability of different threshold positions is accurate in the sense that it matches the 'objective' probability. An analysis of the risk of crossing the threshold when the agent has a biased belief (either overly optimistic or overly pessimistic) is left for future research.

we focus on the case k=K and discuss the case k=0 in Appendix A-3.

We first note that if shape parameter or sensitivity is zero,  $\lambda = 0$  or K=0, then the probability of crossing the threshold does not depend on the respective other parameter. This is intuitive: A shape parameter of  $\lambda = 0$  implies that the EWS structure is not informative, it does not matter whether you always receive an EWS (K=1) or never (K=0) if the likelihood of receiving an EWS is not a function of the distance to the threshold. Similarly, an EWS structure with K=0 implies that the agent never gets an EWS, even if the distance resolution of the EWS structure, encoded in  $\lambda$ , is very high.

Comparing the plots from left to right illustrates the following result:

#### **Result 4** The risk of crossing the threshold increases as the discount factor $\beta$ decreases.

This result is expected: the more the agent discounts the future, the less they care about whether the threshold is crossed or not. Looking at each plot by itself, we then notice the following result:

# **Result 5** The probability of crossing the threshold is not a monotone function of the quality of the EWS structure, i.e. of sensitivity K and shape parameter $\lambda$ .

This finding is surprising and future research ought to verify whether these patterns hold more generally or whether they are an artifact of the numerical analysis. In general, however, this non-monotonicity is quite intuitive, it comes from the fact that the EWS introduces two aspects: On the one hand, the agent wants to advance to receive signals (this increases the risk the agents is exposed to), but on the other hand, the EWS also warns the agent, so that she is exposed to less risk. Depending on the parameter constellation, either effect may be stronger. For example, it is clear that for K=1 the former effect never increases the risk of crossing the threshold (as the miss rate is zero).

Consider the case of  $\beta=0.95$  first, cf. Figure 7 (a). Here, the basic pattern is that the higher  $\lambda$  and the higher K, the lower is the probability of crossing the threshold (but even here we already observe a non-monotonic behavior in the sensitivity parameter K).

The case  $\beta=0.8$  clearly shows non-monotonic behavior, cf. Figure 7 (b). While it is still true that very large values of K and  $\lambda$  are associated with a low probability of crossing the threshold (top-right corner), we see for intermediate values of sensitivity Kthat an increase in  $\lambda$  first increases the risk when  $\lambda$  is small, but later decreases the risk when  $\lambda$  is high. Similarly, for intermediate values of the shape parameter  $\lambda$ , an increase in K increases risk as long as K is below 0.2; for higher values of K, however, a further increase in K reduces the probability of crossing the threshold. Finally, consider the case of  $\beta=0.5$ , cf. Figure 7 (c). Not only do we observe the generally higher risk of crossing the threshold already mentioned above; we also see that we now see a monotonic behavior in the shape parameter, but in a strange and surprising way: Here, the higher  $\lambda$ , the *higher* the probability of crossing the threshold. This is the opposite pattern than the one found for a high value of beta, cf. Figure 7 (a). What persists is the non-monotonicity in sensitivity K: We find high risk for sensitivity parameter values around K=0.2 and K=0.6, and relatively low risk for extreme sensitivity values around K=0 and K=1.

Having discussed how the total probability of crossing the threshold depends on the EWS structure, we now show how the value function depends on the EWS (see Figure 8). Not surprisingly, the value increases with EWS quality, for all values of the discount factor. This finding, summarized in the final result, is essentially a reflection of the Bayesian principle that better information cannot harm.

**Result 6** The optimal value increases with the informativeness of the EWS, that is, the higher K and the higher  $\lambda$ .

Inspecting Figure 8, we see that the value function is more sensitive to a change in K than to a change in  $\lambda$ . Moreover, we note that the shape of the indifference curves is virtually independent of the discount factor  $\beta$  (while the attainable value of course decreases drastically with a lower  $\beta$ ).



Figure 8: Value function at  $l_0$  for different values of K and  $\lambda$ .

### 5 Discussion

In this paper, we study how the opportunity to receive an early warning about the location of a tipping point affects the optimal management of a resource that would collapse once the tipping point is crossed. Our model is abstract and generic, but the key application that we have in mind is the climate system.

Interpreting the location choice  $l_t$  as the target level of global warming of the agent, a hypothetical and benevolent world government, at time t, we find that the existence of an EWS would always improve welfare. However, we also find that an EWS may lead to a slippery slope in risk taking in the sense that the agent increases the level of global warming beyond the level that would be optimal in absence of an EWS. One reason is that the agent may allow more global warming in the belief that she would be warned when she would get close to the tipping point. Moreover, the EWS may itself induce an incentive to increase global warming in order to receive new information about the location of the tipping point. Both these reasons can then, for certain parameter combinations, lead to a situation where the EWS actually increases the total probability of experiencing the breakdown of the climate system.

Note that from a utilitarian perspective, this increase in risk taking is perfectly fine. There is no time-inconsistency or irrationality behind this result. It derives from backwards induction and maximizes discounted utility. From a conservationist's perspective, of course, the fact that EWS could lead to a higher probability of resource collapse is worrisome. Ruling out a trade-off between the utility of the discounted consumption stream and the value of preserving the resource, could therefore lead to calls for a "pre-cautionary principle" as a commitment device. Analyzing under which circumstances a precautionary principle would be selected by agents with different mixes of preferences in our model would be an interesting complement to the recent advances of Guillouet and Martimort (2019), but is clearly beyond the scope of the current paper.

Allowing for multiple agents with potentially non-aligned preferences is more generally an important avenue for future research. Here, we have assumed a single agent. How would results change if there is a sequence of agents, such as (overlapping) generations? Which contracts would avoid situations in which early generations either take excessive risks or explore too little from the perspective of later generations? Similarly, how would results change if the climate system is shared by several non-cooperative agents that live at the same time? Under which structures of informational spillovers would it be possible to agree on the optimal path, and under which structures would we be in a similar situation as Cassandra, where no one would believe the warning signal disclosed by another agent?

In this paper, we have taken the quality of the EWS system as given. That is, for a given parameter combination of K and  $\lambda$ , we have calculated the optimal path (contingent on signal realizations) and the resulting value and risk of crossing the threshold. However, the existence and the informativeness of generic early warning signals is not undisputed in the natural science literature. Boettiger and Hastings (2012, 2013), for example, point out that the detection of generic EWS is likely subject to the prosecutor's fallacy (false positives), because only systems with known regime shifts are being analyzed. Such a concern could be included in our model by allowing agents to hold a probabilistic belief about K. Future work could then also study how agents that have overly optimistic beliefs about K differ in their experimentation choices from agents that have overly pessimistic beliefs about K. In a setting where crossing the tipping point does not lead to an irreversible collapse, this may lead agents to conduct "metaexperiments", not to learn more about the potential location of T, but to learn about the quality of the EWS system.

Finally, an important distinction that we have made in our model and analysis is whether the agent could only learn when choosing a location in the unexplored state space, or whether she could receive an EWS also when not actively experimenting. Either case could be more realistic, depending on the specific situation that one would want to model, and maybe also on the general epistemological stance that the researcher takes. An idealistic researcher in the German tradition may favor an approach where learning is possible also without action, whereas a realist in the Anglo-Saxon tradition may share the view of Aristotle that there is no knowledge outside of experience.

In conclusion, early warning signals are an important and exciting topic for research, especially in light of the growing concerns about tipping points in the climate system (Lenton et al., 2019) Here we have presented a first model that allows us to explore the effects of EWS on optimal dynamic management. While we could only present a numerical solution in the present paper, we are optimistic to be able to prove the key results analytically. Stay tuned.

## References

- Alibakhshi, S., Groen, T., Rautiainen, M., and Naimi, B. (2017). Remotely-sensed early warning signals of a critical transition in a wetland ecosystem. *Remote sensing*, 9(4):352.
- Barrett, S. and Dannenberg, A. (2014). Sensitivity of collective action to uncertainty about climate tipping points. *Nature Clim. Change*, 4(1):36–39.
- Bergemann, D. and Välimäki, J. (2008). Bandit problems. In Durlauf, S. and Blume, L., editors, The New Palgrave Dictionary of Economics, pages 336–340. Macmillan Press, 2 edition.
- Biggs, R., Carpenter, S. R., and Brock, W. A. (2009). Turning back from the brink: Detecting an impending regime shift in time to avert it. Proceedings of the National Academy of Sciences, 106(3):826–831.
- Boettiger, C. and Hastings, A. (2012). Early warning signals and the prosecutor's fallacy. Proceedings of the Royal Society B: Biological Sciences, 279(1748):4734–4739.
- Boettiger, C. and Hastings, A. (2013). Tipping points: From patterns to predictions. *Nature*, 493(7431):157–158.
- Boyarchenko, S. (2018). Strategic experimentation with humped bandits. Available at SSRN: https://ssrn.com/abstract=3174107.
- Cai, Y. and Lontzek, T. S. (2019). The social cost of carbon with economic and climate risks. Journal of Political Economy, 127(6):2684–2734.
- Carpenter, S. R., Cole, J. J., Pace, M. L., Batt, R., Brock, W. A., Cline, T., Coloso, J., Hodgson, J. R., Kitchell, J. F., Seekell, D. A., Smith, L., and Weidel, B. (2011). Early warnings of regime shifts: A whole-ecosystem experiment. *Science*, 332(6033):1079–1082.
- Cropper, M. (1976). Regulating activities with catastrophic environmental effects. Journal of Environmental Economics and Management, 3(1):1–15.
- Crépin, A.-S. and Nævdal, E. (2019). Inertia risk: Improving economic models of catastrophes. The Scandinavian Journal of Economics, n/a(n/a):1–27.
- Dai, L., Vorselen, D., Korolev, K. S., and Gore, J. (2012). Generic indicators for loss of resilience before a tipping point leading to population collapse. *Science*, 336(6085):1175–1177.
- Dakos, V., Carpenter, S. R., van Nes, E. H., and Scheffer, M. (2015). Resilience indicators: prospects and limitations for early warnings of regime shifts. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 370(1659).
- Dakos, V., Scheffer, M., van Nes, E. H., Brovkin, V., Petoukhov, V., and Held, H. (2008). Slowing down as an early warning signal for abrupt climate change. *Proceedings of the National Academy of Sciences*, 105(38):14308–14312.
- Diekert, F. K. (2017). Threatening thresholds? the effect of disastrous regime shifts on the noncooperative use of environmental goods and services. *Journal of Public Economics*, 147:30 – 49.
- Diks, C., Hommes, C., and Wang, J. (2019). Critical slowing down as an early warning signal for financial crises? *Empirical Economics*, 57(4):1201–1228.
- Drake, J. M. and Griffen, B. D. (2010). Early warning signals of extinction in deteriorating environments. *Nature*, 467(7314):456–459.
- Feldmann, J. and Levermann, A. (2015). Collapse of the west antarctic ice sheet after local destabilization of the amundsen basin. *Proceedings of the National Academy of Sciences*, 112(46):14191–14196.
- Frank, K. T., Petrie, B., Choi, J. S., and Leggett, W. C. (2005). Trophic cascades in a formerly coddominated ecosystem. *Science*, 308(5728):1621–1623.
- Gorton, G. B. (2010). Slapped by the invisible hand: The panic of 2007. Oxford University Press.
- Guillouet, L. and Martimort, D. (2019). Precaution, Information, and time-inconsistency: Do we need a Precautionary Principle? Unpublished working paper, available at https://www.economics.

utoronto.ca/index.php/index/research/downloadSeminarPaper/110815.

- Hawkins, E., Smith, R. S., Allison, L. C., Gregory, J. M., Woollings, T. J., Pohlmann, H., and de Cuevas, B. (2011). Bistability of the atlantic overturning circulation in a global climate model and links to ocean freshwater transport. *Geophysical Research Letters*, 38(10).
- Held, H. and Kleinen, T. (2004). Detection of climate system bifurcations by degenerate fingerprinting. Geophysical Research Letters, 31(23).
- Hoegh-Guldberg, O., Mumby, P. J., Hooten, A. J., Steneck, R. S., Greenfield, P., Gomez, E., Harvell, C. D., Sale, P. F., Edwards, A. J., Caldeira, K., Knowlton, N., Eakin, C. M., Iglesias-Prieto, R., Muthiga, N., Bradbury, R. H., Dubi, A., and Hatziolos, M. E. (2007). Coral reefs under rapid climate change and ocean acidification. *Science*, 318(5857):1737–1742.
- Kleinen, T., Held, H., and Petschel-Held, G. (2003). The potential role of spectral properties in detecting thresholds in the earth system: application to the thermohaline circulation. *Ocean Dynamics*, 53(2):53–63.
- Lemoine, D. and Traeger, C. (2014). Watch your step: Optimal policy in a tipping climate. American Economic Journal: Economic Policy,, 6(1):137–166.
- Lemoine, D. and Traeger, C. P. (2016). Economics of tipping the climate dominoes. Nature Climate Change, 6(5):514–519.
- 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.
- Lenton, T. M., Rockström, J., Gaffney, O., Rahmstorf, S., Richardson, K., Steffen, W., and Schellnhuber, H. J. (2019). Climate tipping points – too risky to bet against. *Nature*, 575:592–595.
- Liski, M. and Salaniè (2019). Tipping points, delays, and the control of catastrophestipping points, delays, and the control of catastrophes. unpublished manuscript.
- Lovejoy, T. E. and Nobre, C. (2018). Amazon tipping point. Science Advances, 4(2).
- Nævdal, E. (2006). Dynamic optimisation in the presence of threshold effects when the location of the threshold is uncertain with an application to a possible disintegration of the western antarctic ice sheet. *Journal of Economic Dynamics and Control*, 30(7):1131–1158.
- Polasky, S., Zeeuw, A. d., and Wagener, F. (2011). Optimal management with potential regime shifts. Journal of Environmental Economics and Management, 62(2):229 – 240.
- Richter, A. and Dakos, V. (2015). Profit fluctuations signal eroding resilience of natural resources. *Ecological Economics*, 117:12–21.
- Rocha, J. C., Peterson, G., Bodin, O., and Levin, S. (2018). Cascading regime shifts within and across scales. *Science*, 362(6421):1379–1383.
- Sakamoto, H. (2014). Dynamic resource management under the risk of regime shifts. Journal of Environmental Economics and Management, 68(1):1–19.
- Tsur, Y. and Zemel, A. (1995). Uncertainty and irreversibility in groundwater resource management. Journal of Environmental Economics and Management, 29(2):149 – 161.
- Walker, X. J., Baltzer, J. L., Cumming, S. G., Day, N. J., Ebert, C., Goetz, S., Johnstone, J. F., Potter, S., Rogers, B. M., Schuur, E. A. G., Turetsky, M. R., and Mack, M. C. (2019). Increasing wildfires threaten historic carbon sink of boreal forest soils. *Nature*, 572(7770):520–523.
- Wiesenfeld, K. (1985). Virtual hopf phenomenon: A new precursor of period-doubling bifurcations. *Phys. Rev. A*, 32:1744–1751.

# Appendix

#### A-1 First-order condition

We show that the continuation value in (7) is independent of  $l_t$  (as this information is encoded in  $p_t(T)$ ). Assume an interior solution, giving the the first-order condition:

$$\begin{aligned} &(\frac{\partial}{\partial l}V(l_{t+1}, G(s=1;...) + \left(\frac{\partial}{\partial p}V(l_{t+1}, G(s=1;...)\frac{\partial}{\partial l}G(s=1;...)\right)\hat{q}(s=1, l_{t+1}, p_t(T)) \\ &+ \left(\frac{\partial}{\partial l}V(l_{t+1}, G(s=0;...) + \left(\frac{\partial}{\partial p}V(l_{t+1}, G(s=0;...)\frac{\partial}{\partial l}G(s=0;...)\right)\hat{q}(s=0, l_{t+1}, p_t(T)) \right. \\ &+ V(l_{t+1}, G(s=1;...)\frac{\partial}{\partial l}\hat{q}(s=1,...) + V(l_{t+1}, G(s=0;...)\frac{\partial}{\partial l}\hat{q}(s=0,...) = 0 \end{aligned}$$
(10)

Write  $\pi(l_t, p_t(T)) = l_{t+1}^*$ . Insert for  $\pi$  and differentiate with respect to  $l_t$ . It follows that

$$\frac{\partial}{\partial l_t} V(l_t, p_t(T)) = \pi'[...], \tag{11}$$

where the terms in the bracket equals 0 because  $l_{t+1}$  fulfills the first-order condition.

We can therefore write

$$V(p_t(T)) = \max_{l_{t+1}} \left\{ u(l_t) + \beta \times \left[ \hat{q}(s=1, l_{t+1}, p_t(T)) V(G(s=1; l_{t+1}, p_t(T)) + \hat{q}(s=0, l_{t+1}, p_t(T)) V(G(s=0; l_{t+1}, p_t(T))) \right] \right\},$$
(12)

with first-order condition

$$V'(G(s = 1; l_{t+1}, p_t(T))) \frac{\partial G(s = 1; l_{t+1}, p_t(T))}{\partial l_{t+1}} \hat{q}(s = 1, l_{t+1}, p_t(T))$$

$$+ V'(G(s = 0; l_{t+1}, p_t(T))) \frac{\partial G(s = 0; l_{t+1}, p_t(T))}{\partial l_{t+1}} \hat{q}(s = 0, l_{t+1}, p_t(T))$$

$$+ V(G(s = 1; l_{t+1}, p_t(T))) \frac{\partial \hat{q}(s = 1, l_{t+1}, p_t(T))}{\partial l_{t+1}}$$

$$+ V(G(s = 0; l_{t+1}, p_t(T))) \frac{\partial \hat{q}(s = 0, l_{t+1}, p_t(T))}{\partial l_{t+1}} = 0$$
(13)

#### A-2 Early warning signals and hypothesis testing

In this section we first outline the relationship between the agent's decision-problem of choosing a location and the general statistical concept of hypothesis testing. In a second step, we relate the parameters of the EWS structure to the well-known concepts of *sensitivity* and *specificity* of a test.

The link between the decision problem and hypothesis testing. For simplicity assume that the set of possible threshold locations is the unit interval [0, 1] and that the agent is located at l = 0. The agent contemplates whether to stay or move; and if to move, how far. The agent, therefore, is interested in hypotheses of the form "The interval  $\mathcal{I}(\varepsilon) = (0, \varepsilon)$  in front of me contains the threshold". If the latter hypothesis is true, then it is not safe to move to  $l + \varepsilon$ . The agent knows that the EWS structure provides information regarding the hypothesis.<sup>6</sup>. This can be regarded as a "test" about the hypothesis where an EWS, s = 1, corresponds to a "positive" test result and the absence of an EWS, s = 0, corresponds to a "negative" test result.

Sensitivity and specificity. In the following we restrict to the EWS structure

$$f(d) = K e^{-\lambda d},\tag{14}$$

where d = T - l is the distance between the actual threshold T and the location at which the EWS is observed / is not observed and  $\lambda \in [0, \infty]$  and  $K \in [0, 1]$  are parameters. It is those parameters that we want to link to properties of the hypothesis tests now.

For this section we make another simplifying assumption, namely that the decision-maker holds a uniform belief over all possible (remaining) threshold positions.

For a given  $\varepsilon > 0$ , the relevant hypothesis is that the threshold is less than  $\varepsilon$  away, i.e.  $T \in \mathcal{I}(\varepsilon)$ . Based on this hypothesis, we can characterize the 'test' EWS by

$\operatorname{TP}(\varepsilon) := \Pr\left(s = 1 \mid T \in \mathcal{I}(\varepsilon)\right)$	true positive
$\operatorname{FN}(\varepsilon) := \operatorname{Pr}\left(s = 0 \mid T \in \mathcal{I}(\varepsilon)\right)$	false negative
$\operatorname{FP}(\varepsilon) := \Pr\left(s = 1 \mid T \notin \mathcal{I}(\varepsilon)\right)$	false positive
$\operatorname{TN}(\varepsilon) := \Pr\left(s = 0 \mid T \notin \mathcal{I}(\varepsilon)\right)$	true negative

It suffices to calculate  $TP(\varepsilon)$  and  $TN(\varepsilon)$  as  $FN(\varepsilon) = 1 - TP(\varepsilon)$  and  $FP(\varepsilon) = 1 - TN(\varepsilon)$ . It is

$$TP(\varepsilon) = \Pr(s = 1 \mid T \in \mathcal{I}(\varepsilon))$$

$$= [\Pr(T \in \mathcal{I}(\varepsilon))]^{-1} \cdot \Pr(s = 1 \land T \in \mathcal{I}(\varepsilon))$$

$$= \frac{1}{\varepsilon} \int_{0}^{\varepsilon} K e^{-\lambda \tau} d\tau$$

$$= \frac{1}{\varepsilon} \frac{K}{\lambda} (1 - e^{-\lambda \varepsilon})$$
(15)

<sup>&</sup>lt;sup>6</sup>Actually, the EWS structure provides information simultaneously about all hypothesis indexed by  $\varepsilon \in (0, \infty)$ ).

Similarly, we get

$$TN(\varepsilon) = \Pr(s = 0 \mid T \notin \mathcal{I}(\varepsilon))$$

$$= [\Pr(T \notin \mathcal{I}(\varepsilon))]^{-1} \cdot \Pr(s = 0 \land T \notin \mathcal{I}(\varepsilon))$$

$$= \frac{1}{1 - \varepsilon} \int_{\varepsilon}^{1} (1 - Ke^{-\lambda\tau}) d\tau$$

$$= \frac{1 - \varepsilon}{1 - \varepsilon} - \frac{1}{1 - \varepsilon} \int_{\varepsilon}^{1} Ke^{-\lambda\tau} d\tau$$

$$= 1 - \frac{1}{1 - \varepsilon} \frac{K}{\lambda} \left( e^{-\lambda\varepsilon} - e^{-\lambda} \right)$$
(16)

The agent wonders whether to move at all and is accordingly interested in these tests for very small values of  $\varepsilon$ . We therefore analyse the limit of above expressions for  $\varepsilon \to 0$ . From (15) we get

$$TP(0) = \lim_{\varepsilon \to 0} TP(\varepsilon)$$
$$= \frac{K}{\lambda} \lim_{\varepsilon \to 0} \frac{\lambda e^{-\lambda}}{1}$$
$$= K$$
(17)

Similarly, we use (16) and get

$$TN(0) = \lim_{\varepsilon \to 0} TN(\varepsilon)$$
  
= 1 -  $\frac{K}{\lambda} (1 - e^{-\lambda})$  (18)

The true positive rate of a test is usually called *sensitivity* of the test. We therefore call K the **sensitivity of the EWS structure**, or short **EWS sensitivity**. The true negative rate of a test is called the *specificity* of the test. We here see that the specificity of the EWS structure depends on both K and  $\lambda$ . It is interesting to look at the extreme cases of  $\lambda$ ,

$$\lim_{\lambda \to 0} \operatorname{TN}(0) = 1 - K , \quad \lim_{\lambda \to \infty} \operatorname{TN}(0) = 1$$
(19)

We see that the EWS 'test' has maximal sensitivity and specificity when K = 1 and  $\lambda$  is large; this is essentially the case when, from hearing, the agent can infer that the threshold is close, and from not hearing they can infer that it is safe to walk. The EWS test has a high sensitivity and low specificity when K is near 1 and  $\lambda$  is near 0. This is the case when the agent nearly always receives an EWS, irrespective of whether the threshold is close or not, and which is hence of little use. On the other hand, an EWS structure with low K, essentially for all values of the shape parameter  $\lambda$ , displays a high specificity; if the threshold is not near the agent will typically not receive an EWS by mistake; the problem in this case is that even if the threshold is close they will often not receive an EWS. From (17) and (19) we see that it is impossible to have an EWS test with both low sensitivity and low specificity.

**Extensions.** What remains to be done is to generalize above analysis to arbitrary beliefs, i.e. establish the link between the EWS parameters K and  $\lambda$  and the properties of the hypothesis test for any belief the agent may hold. This is left for future research.

#### A-3 Total probability of crossing the threshold when k=0

Figure A-1 shows the contour plots that indicate the risk of crossing the threshold when the agent can only receive an EWS when actively experimenting (k=0). While the specific values differ from the situation when k = K shown in Figure 7 in the main text (as the optimal paths differ, notably they are simpler when k=0 because there is no stopping and walking again), we see the same non-monotonicity. In other words, the risk of crossing the threshold may be higher with an EWS than without.



Figure A-1: Contour plots showing the total probability of crossing T along the optimal path when  $k{=}0$ 

Similarly, the value function is qualitatively the same when k=0 or when k=K (compare Figure A-2 below to Figure 8 in the main text). That is, the value is monotonically increasing the more informative the EWS (higher K and higher  $\lambda$ ).



Figure A-2: Value function at  $l_0$  for different values of K and  $\lambda$  when k=0

#### A-4 Numerical implementation

We use Matlab to solve the model. We rely on @tree, which is a user generated Matlab class to represent tree data structures. The tree data structures enable a representation of the agent's original problem by backward recursion through paths of location choices and signal realizations that are clearly not suboptimal (see Section 3).

The structure of the code is perhaps clearest from considering the location tree, which stores the agent's current location through such paths. Suppose for ease of illustration that  $t_{\text{max}} = 2$ and L = 3. The root of the location tree is the initial location, that is  $l_0 = 1$ . The root of the tree has three children nodes. These correspond to the timing after a location choice is made  $(l_1 \in \{1, 2, 3\})$  but before signals are realized, and store the new location. We add two nodes to each of the children nodes of the root, corresponding to the timing after signals are realized. These nodes store the same location as their parent (the relevant children node of the root). At the relevant grandchildren node of the root, the agent makes a new location choice.



Figure A-3: Example of location tree, under assumptions  $t_{\text{max}} = 2$  and L = 3.

In addition to the location tree, we represent the following information in tree data structures:

- Signal tree: Stores whether the agent can observe and receive a signal at each node.
- Time tree: Stores the time at each node.
- Belief tree: Stores the belief at each node.
- Value tree: Stores the value at each node.
- Optloc tree: Stores the optimal location choice at each node.
- Optmove tree: Stores the optimal node to relocate to at each node.