



FAERE

French Association
of Environmental and Resource Economists

Working papers

Catastrophes, delays, and learning

Matti Liski – François Salanié

WP 2020.20

Suggested citation:

M. Liski, F. Salanié (2020). Catastrophes, delays, and learning. *FAERE Working Paper*, 2020.20.

ISSN number: 2274-5556

www.faere.fr

Catastrophes, delays, and learning^{*}

Matti Liski[†] François Salanié[‡]

September 23, 2020

Abstract

How to plan for catastrophes that may be under way? In a simple but general model of experimentation, a decision-maker chooses a flow variable contributing to a stock that may trigger a catastrophe at each untried level. Once triggered, the catastrophe itself occurs only after a stochastic delay. Consequently, the rhythm of past experimentations determines the arrival of information. This has strong implications for policies in situations where the planner inherits a history of experiments, like climate change and pandemic crisis. The structure encompasses canonical approaches in the literature.

JEL Classification: C61, D81, Q54.

Keywords: catastrophes, experimentation, delays

^{*}We have benefitted from discussions in numerous seminars in Toulouse, Helsinki, Montpellier, Munich, Rennes, Copenhagen, Gothenburg, Ascona, Tel Aviv, and Madrid. In particular, we thank Anne-Sophie Crépin, Arthur Snow, and Nicolas Treich, for comments on this work.

[†]Aalto University, Department of Economics. e-mail: matti.liski@aalto.fi

[‡]Toulouse School of Economics, INRAE, Université Toulouse Capitole, Toulouse, France. e-mail: francois.salanie@inra.fr. François Salanié acknowledges funding from ANR under grant ANR-17-EURE-0010 (Investissements d’Avenir program).

How to plan for catastrophes? As economists, we are increasingly confronting this question. But when we study climate change, virus outbreaks turning to pandemics, or the collapse of fisheries and ecosystems, we encounter several approaches with different assumptions, sometimes yielding opposing conclusions. In particular, a key question is how these approaches deal with the possibility that a catastrophe may already be under way, as we now explain.

Consider for example the impact of climate change on the Greenland ice sheet. A catastrophic melting might well be under way, though no one knows exactly (e.g., Kriegler et al., 2009). We expect that *some* temperature increase will lead to a dramatic acceleration in melting, but this threshold is unknown, reflecting scientific uncertainty or stochastic shocks. Was this critical threshold exceeded already in the 70’s, or will it be reached in the near future? Evidently, we cannot tell the final effect of past actions because there is a considerable delay between the cause (the CO_2 concentration in the atmosphere) and the effect (melting) (Fitzpatrick and Kelly, 2017; van der Ploeg, 2018). Similar thresholds are not unheard of in other situations: we do not immediately observe if a virus outbreak is on its way to pandemic, or if habitat fragmentation will lead to a collapse of biodiversity.

In this paper, we develop a simple but general model of experimentation where the decision maker chooses both *how much to experiment* with an unknown threshold and *how to prepare* for the potential impacts from exceeding this threshold. The planner controls a state variable (“stock”) with multiple interpretations (i.e., temperature, infected population). The stock *triggers* a catastrophe when it exceeds an unknown threshold. Once triggered, the catastrophe itself *occurs* only after a stochastic delay. Reaching a previously untried level is thus an experiment whose results are only known later on. This distinction between triggering and occurrence allows for rich patterns of information arrival.

Going back to the Greenland ice sheet, consider low vs. high rhythms of experimentation in the past, leading to the same current CO_2 concentration or temperature level. Because of the delay between triggering and occurrence, a fast rhythm generates less information on the safety of the current standing than a slow rhythm. This legacy of the past is key to our analysis. Accordingly, the planner values the gain from further experimenting differently. If he is worried enough, he may also want to reduce the exposure to the catastrophe, in case it occurs. Our setting allows to do so, by allowing the

welfare losses borne when the catastrophe occurs to depend on the level of the stock: for example, a cooler climate may prevent a cascade of correlated events in case the ice cover is collapsing. This feature leads to different possibilities for the optimal path.

After a sustained rhythm of experimentation, the planner rationally fears that losses arrive any time soon. Interestingly, this can justify increasing or reducing the stock, depending on the properties of the welfare loss. If the loss is independent of the stock, then there is nothing the planner can do to mitigate the damage: he becomes fatalistic, which may mean experimenting even more if the instantaneous value of an increased consumption is high. In contrast, if the loss depends on the stock and is also high enough, reducing the stock is optimal because the priority is to control the damage. After a while, the fact that the catastrophe does not occur makes the planner more and more optimistic, until he chooses to increase the stock because damage control is not the priority anymore. Therefore, in such a case the optimal policy is non-monotonic. In a rapidly developing pandemic, the optimal policy might well be to impose first a strict lockdown, and then to accommodate monotonically increasing levels of infection up to the levels encountered in the past that triggered the lockdown in the first place: this is because in between we have learnt that the health system, or our institutions, were able to cope with these levels without collapsing.¹

After a low rhythm, the planner has more precise information, and is more certain that the current welfare is not threatened by a catastrophe. His behavior is then driven by a calculation of the pros and cons of experimenting further, i.e., a stronger economic activity vs. the risk of triggering a catastrophe. Hence, the risk of losing this secured welfare drives optimal actions, rather than damage control. History thus determines not only the information available but also the fundamental trade-off to be looked at in decision-making.

Two applications illustrate the results from the general model. First, our climate change illustration shows that the models in the literature come primarily from two opposite camps that emphasize distinct tradeoffs for determining the optimal policies. Our model unifies these approaches by showing how they follow as limiting cases of our

¹The idea of the hammer and the dance as characterization of the optimal policy during pandemic crises was entertained in an influential posting by Tomas Pueyo, March 19, 2020. <https://medium.com/@tomaspueyo/coronavirus-the-hammer-and-the-dance-be9337092b56>. On this topic, a recent work by Assenza et al. (2020) provides a complete literature review. Our work offers a rationale for the hammer-and-dance policies based on information and learning.

setting and how their comparison determines the optimal experimentation policy outside the limits. Second, the disease control and social distancing application contributes to the rapidly growing literature on virus outbreaks. We show how non-monotonic paths naturally emerge when the planner fears that a catastrophe may be under way.

Literature. Catastrophes, broadly interpreted, appear in a wide range of economic applications, including macroeconomics disasters (e.g., Barro, 2006; Gourio, 2008), technology breakdowns and demand tipping (e.g., Rob, 1991; Bonatti and Hörner, 2017), resource consumption (Kemp, 1976), nuclear accidents (Cropper, 1976), and pollution control (Clarke and Reed, 1994; Polasky et al., 2011; Sakamoto, 2014; van der Ploeg and de Zeeuw, 2017; Bretschger and Vinogradova, 2019; Cai and Lontzek, 2019).²

Our model does not cover all applications but it is general in the sense that it embeds two canonical approaches to modeling catastrophes in the literature. In the hazard rate approach, one may conceptualize the system as a machine that may break down under pressure. The probability of a catastrophe happening depends only on the current state of the system, typically through an exogenous hazard rate function. There is thus no memory of the past, and no learning over time. Moreover, the catastrophe has to happen sooner or later, just as any machine will ultimately break down. This assumption features in many important recent applied papers, for example in the quantitative assessments of the optimal climate-change policies (e.g., Besley and Dixit, 2019).

In the tipping point approach, the catastrophe occurs as soon as the stock exceeds a tipping point whose exact value is unknown. The formal approach appears in Kemp (1976), when studying the problem of “eating a cake of unknown size”; and, more closely in the catastrophe context, in Tsur and Zemel (1994) (see also Tsur and Zemel, 1995 and 1996.) It has also been used in a quantitative policy evaluation (Lemoine and Traeger, 2014). Learning occurs instantaneously over time: the planner is absolutely certain that the threshold has not been exceeded in the past if no catastrophe has occurred so far. Beliefs are thus revised through a simple truncation. This feature matches the facts in most learning environments quite badly. For example, Roe and Baker (2007) argue that the delays built into the feedback mechanisms governing climate change will prevent us from learning the true nature of the problem in the coming decades.³

²See Rheinberger and Treich (2017) for bibliometric analysis of the literature on catastrophes.

³Crépin and Nævdal (2019) extend the threshold approach as follows. The stock governs the rate of change of another state variable which makes the catastrophe to occur when it goes above an unknown tipping point. This introduces inertia to the path of this second state variable but learning is still

Our model embeds both approaches as special cases. When the delay in our model goes to zero, the optimal policy can be shown to converge to that in the tipping point approach; the latter was derived in Tsur and Zemel (1994), and we provide a more general study in Appendix B. On the other hand, if past actions are known to have triggered the event, the catastrophe has to occur, and the planner can only try to mitigate the damage from it. Hence, the policy in that case is conceptually equivalent to the one from the hazard rate approach. By introducing a delay, and by relaxing the assumption that the catastrophe was triggered in the past, we explore a more general case. In particular, in our model the decision maker remains uncertain if the current standing is safe, even if he stops experimenting. In fact, the exposure to this final uncertainty is a choice to be made.

Our approach is also different from the bandit models used to study experimentation in various economic settings. As in Poisson bandit settings, from not observing the event the planner updates beliefs on the arrival rate of a catastrophe (Malueg and Tsutsui, 1997; see also Keller, Rady and Cripps, 2005; and Bonatti and Hörner, 2011). In a sense, in our setting the decision maker runs an endogenous continuum of such bandits (thresholds tried), so that obtaining the information content of past actions requires aggregation over the experiments. The belief updating that follows from this aggregation is new to the experimentation literature.

A few recent papers on experiments are related to our work. Gerlagh and Liski (2017) consider an explicit climate-economy model with learning on potentially catastrophic damages to study if the path on carbon emissions is sensitive to the assumption that such damages are not currently observed but must be learned over time. The objective of that paper is to study the impact of speed of learning of given hidden state that determines if damages will ultimately arrive. However, it is not possible to make choices that impact the value of the hidden state, and therefore there is no experimentation with tipping in that paper.

Salmi, Laiho and Murto (2020) study “*Gradual learning from incremental actions*” which could also be the title of our paper. This model shares with our model the feature that the chronicle of past actions determines the speed of information arrival, and further actions have direct payoff consequences. In their model, the direct payoff consequence comes from capacity expansion and contemporaneous sales, running a risk of instantaneous.

overcapacity if the state turns out to be bad. In our model, the contemporaneous payoff relates to consumption utility. The main difference is that our contemporaneous choice of consumption utility is an experiment with the hidden state while in Salmi et al. the decision maker cannot influence the state.⁴

Guillouët and Martimort (2019) study the foundations of a precautionary principle in an environment where a catastrophe may happen after a delay when the stock exceeds an unknown tipping point, as in our paper. The focus is very different however, as they do not allow the planner to condition its policy on his beliefs. Hence, a time-consistency problem arises, and the best policy results from a Nash equilibrium between different selves. Our question is very different as we focus on the optimal policy only, so that beliefs inherited from the chronicle of past actions play a key role. We also have a payoff structure covering multiple planning situations, as illustrated by the applications.

1 The Model

1.1 The No-Catastrophe Problem (NCP)

We introduce the primitives by considering first the problem without catastrophes. The decision maker chooses a flow action $q_t \in [\underline{q}, \bar{q}]$, with $\underline{q} < 0 < \bar{q}$, to control a stock Q according to a simple law of motion:

$$\dot{Q}_t = q_t.$$

Hence q is a net flow, that can be positive or negative. The instantaneous payoff $u(q, Q)$ is allowed to depend both on the stock level Q and on the flow q , and it is discounted at the rate $\delta > 0$. We assume (subscripts denote partial derivatives):

Assumption 1 *Function u is twice continuously differentiable, bounded from above, and weakly concave in q . Moreover, the function*

$$\nu(Q) \equiv u_q(0, Q) + \frac{u_Q(0, Q)}{\delta}$$

is weakly decreasing with respect to Q .

⁴Models for reputation in dynamic games with incomplete information often have the feature that some agent can influence the hidden state such as the quality of a firm's product (Faingold and Sannikov, 2011, Board and ter Vehn, 2013, Bohren, 2019). These papers differ significantly in focus but, however, share the feature that past choices shape key trade-offs of current and future interactions in a capital-theoretic manner.

Function ν encapsulates the trade-off between instantaneous gains from an increase in q , and the long-run effects of the associated increase in Q . ν is decreasing for any positive δ if u_{QQ} and u_{qQ} are at most zero, but we use the above formulation to highlight the exact property that is really needed. To avoid multiplicities, we also assume that there exists a unique solution Q^N (where N stands for “No catastrophe”) to the equation $\nu(Q) = 0$, where by convention we let $Q^N = +\infty$ (resp. $-\infty$) if ν is everywhere positive (resp. negative.)

Note that the impact of Q on the payoff can be positive or negative, thereby allowing the model to fit various settings. For example, in Section 4 we shall study a model of climate change in which q stands for emissions associated to consumption (net of natural decay), while Q is the stock of greenhouse gases, that impacts the climate and therefore the production possibilities, and/or the representative agent’s payoff. In Section 5, our application to the control of a disease interprets q as a measure of the stringency of a social distancing policy, and Q as a measure of the number of infected people.

The optimal policy $(q_t, Q_t)_{t \geq 0}$ in the absence of catastrophes maximizes

$$\int_0^{+\infty} u(q_t, Q_t) \exp(-\delta t) dt \quad (1)$$

under $\dot{Q}_t = q_t$, Q_0 given. As time enters only through geometric discounting, the problem is autonomous. Lemma A.4 in the Appendix shows that solutions exist and are monotonic.⁵ Because ν is decreasing, the planner chooses to gradually increase (if $Q_0 < Q^N$) or reduce (if $Q_0 > Q^N$) the stock until it reaches the level Q^N . Q^N is thus to be interpreted as the long-term target, in the absence of catastrophes.

1.2 Introducing catastrophes and delays

We build on the primitives above to add the possibility of a catastrophe. The planning date is $t = 0$, but the full past history will be relevant and thus we let $t \in (-\infty, \infty)$. We say that a catastrophe is *triggered* when the stock Q exceeds a given threshold value S . Given a path $(Q_t)_{t \in (-\infty, +\infty)}$, the triggering time is a function of the threshold S :

$$T(S) \equiv \inf\{t : Q_t > S\}. \quad (2)$$

Note that $T(S)$ is infinite if the stock never exceeds S , and that $Q_{T(S)} = S$ otherwise.

⁵The Appendix actually studies a more general problem and presents results used in a later stage.

We also define the record stock at time t :

$$\overline{Q}_t \equiv \max_{t' \leq t} Q_{t'}$$

so that $T(S) < t$ if and only if $S < \overline{Q}_t$. The catastrophe itself *occurs* only after a delay $\tau \geq 0$, thus at date $\kappa = T(S) + \tau$. Before κ , the instantaneous utility is just $u(q, Q)$. At time κ , the catastrophe occurs, and the game stops, with the planner receiving a stopping payoff $V(Q_\kappa)$ that depends on the value of the stock at the catastrophe date. This means that the planner can mitigate the impact of a catastrophe by changing the level of the stock after the catastrophe was triggered, but before it occurs. Making the stopping payoff V instead depend on the threshold S , or on the maximum level tried in the past \overline{Q}_κ , would eliminate this possibility by assumption. To illustrate, one may imagine a skater on thin ice. Instantaneous utility flow increases with the distance from the shore, but the ice gets thinner and thinner. Once the first crack in the ice has appeared, the skater may turn back as long as the ice is still holding. When the ice finally breaks, the damage to the skater depends on the remaining distance to the shore, as assumed in the model.

Catastrophes are costly, irreversible events. Irreversibility means that the continuation value V is fixed.⁶ The catastrophe is costly if V is less than the value of forever stabilizing the stock at a safe level. We therefore define the damage function D as

$$D(Q) \equiv \frac{u(0, Q)}{\delta} - V(Q).$$

Assumption 2 *The damage function $D(Q)$ is twice continuously differentiable, weakly positive, weakly increasing, and such that $\nu(Q) - D'(Q)$ is weakly decreasing with respect to Q .*

The last part of the assumption is a regularity condition that ensures that higher levels of the stocks reduce the value of increasing the flow, once the marginal damages associated to a catastrophe are taken into account. The assumption is strong enough to imply that $\nu - kD'$ is decreasing, for $k \in [0, 1]$, and thus accommodates easily the case when the catastrophe is discounted, or occurs with a probability below one.⁷

⁶Although we take this stopping payoff as a given and thus beyond the control of the social planner, the applications to climate change and disease control provide micro-foundations for V as the value function of the post-catastrophe problem.

⁷Indeed, $\nu - kD' = (1 - k)\nu + k(\nu - D')$, and both terms are decreasing by assumption. The early

Overall, given S , τ , and a path $(Q_t)_{t \in (-\infty, +\infty)}$, one can apply (2) to compute $T(S)$ and $\kappa = T(S) + \tau$, so that the planner's payoff from date $t = 0$ onward equals

$$\int_0^\kappa u(q_t, Q_t) \exp(-\delta t) dt + \exp(-\delta \kappa) V(Q_\kappa).$$

We now introduce uncertainty. First, the threshold value S is uncertain to the planner, who entertains prior beliefs characterized by a cumulative distribution function F on the interval $[\underline{S}, \bar{S}]$. At this point, it is important to realize that prior beliefs refer to beliefs held at the beginning of times, i.e. at $t = -\infty$, before any experimentation takes place. These prior beliefs are then revised continuously over time to take into account the fact that no catastrophe has occurred yet. Assume that F is continuously differentiable on its support, with density f . We adopt a monotone hazard rate assumption:

Assumption 3 *The hazard rate $\rho(Q) \equiv \frac{f}{1-F}(Q)$ is weakly increasing.*

Second, the delay τ is also uncertain, and drawn from a Poisson distribution with parameter $\alpha > 0$, independently from S . Hence the planner is not only unsure of the location of the tipping point, but also about whether the threshold has already been passed. This captures the idea that a catastrophe might well be under way, though no one knows. This ignorance precludes a precise adoption of preventive measures, possibly making the catastrophe even worse. In the biology literature, concepts like the extinction debt correspond to this idea (Tilman et al., 1994).

Beliefs at time zero obtain by conditioning the prior beliefs by the event “no catastrophe happened until time zero”, or equivalently $\kappa = T(S) + \tau \geq 0$. Such a conditioning is an original feature of our model: at any date, one has to take into account the various experimentations that took place in the past, and the possibility that they might have triggered a catastrophe that did not happen yet. Therefore, given $(Q_t)_{t \leq 0}$, the planner's problem is as follows:

$$\max_{(q_t, Q_t)_{t \geq 0}} \mathbb{E} \left[\int_0^\kappa u(q_t, Q_t) \exp(-\delta t) dt + \exp(-\delta \kappa) V(Q_\kappa) \mid \kappa \geq 0 \right] \quad (3)$$

$$\dot{Q}_t = q_t \in [\underline{q}, \bar{q}], T(S) = \inf\{t : Q_t > S\}, \kappa = T(S) + \tau. \quad (4)$$

literature on catastrophes often relies on much more complicated assumptions. For example, the seminal study in Tsur and Zemel (1994) relies on two assumptions U1 and U2 that involve the solution to a constrained dynamic program. Our contribution in Appendix B is to provide general proofs of their results relying on our simple assumptions, together with a solid connection to our main analysis.

While Q is continuous by construction, we allow for a countable number of jumps in q . We say that a solution $(Q_t)_{t \geq 0}$ is monotonic if Q_t is everywhere weakly decreasing, or everywhere weakly increasing, with respect to time. For any path $(Q_t)_{t \in (-\infty, +\infty)}$, we define $\bar{Q}_\infty \leq +\infty$ as the supremum value for the stock. We say that \bar{Q}_∞ is reached in finite time if there exists $T < +\infty$ such that $Q_T = \bar{Q}_\infty$, and otherwise we say that \bar{Q}_∞ is reached asymptotically; in that case, one has $Q_t \leq \bar{Q}_t < \bar{Q}_\infty$ for all t .

2 Beliefs

To study the above problem, we first examine how to deal with the planner's beliefs. Recall that it is the premise of planning that no catastrophe has happened at time $t = 0$. Given a path $(Q_t)_{t \in (-\infty, +\infty)}$, let us define the survival probability at time t , computed at the beginning of times using the prior beliefs F :

$$p_t \equiv \text{Prob}(\kappa \geq t).$$

One may distinguish two possibilities. Either S is above \bar{Q}_t , and in that case a catastrophe cannot happen before time t . Or S is below \bar{Q}_t : for each such S , survival means that though the catastrophe was triggered at time $T(S) < t$, it did not happen before time t because the delay τ is above $t - T(S)$. Because τ follows a Poisson process, the corresponding probability is $\exp(-\alpha(t - T(S)))$. Overall, we obtain

$$p_t = 1 - F(\bar{Q}_t) + \int_{S < \bar{Q}_t} \exp(-\alpha(t - T(S))) dF(S). \quad (5)$$

The second term thus measures the possibility that a catastrophe was triggered in the past, but did not happen yet. It is high if the record stock has been increasing quickly in the recent past, and low otherwise. The share π_t of this term in p_t can be interpreted as measuring the legacy from the past, compared to potential threats from the future:

$$1 - \frac{1 - F(\bar{Q}_t)}{p_t} \equiv \pi_t \in [0, F(\bar{Q}_t)]. \quad (6)$$

The survival probability can also be expressed as follows. At time 0, the initial value for the survival probability is a function of the past trajectory:

$$p_0 = 1 - F(\bar{Q}_0) + \int_{S < \bar{Q}_0} \exp(\alpha T(S)) dF(S).$$

We complement this initial value by the law of motion:

$$\dot{p}_t = \alpha[1 - F(\bar{Q}_t) - p_t]$$

which indeed gives (5). Hence, p_t is above $1 - F(\bar{Q}_t)$ and, without a surprise, it is nonincreasing as p_t is a survival probability. One important limit case is when delays are nil (α goes to infinity), so that catastrophes occur as soon as they are triggered. Then p_t is everywhere equal to $1 - F(\bar{Q}_t)$, and the legacy from the past π_t is zero everywhere. Otherwise, if delays exist ($\alpha < +\infty$), then either planning starts with no past experimentation: $\bar{Q}_0 \leq \underline{S}$, $\pi_0 = 0$, and $p_0 = 1$. Or the planner inherits some experiments from the past: $\bar{Q}_0 > \underline{S}$, and then one has $\pi_t > 0$ and $p_t > 1 - F(\bar{Q}_t)$ forever, as it is always possible that a catastrophe was triggered in the past but did not occur yet.

We now know the cumulative density function $(1 - p_t)$ of the catastrophe date κ . Conditioning on the event $\kappa \geq 0$ amounts to divide this cdf by p_0 , and therefore the expected payoff (3) is proportional to:

$$\begin{aligned} & \mathbb{E} \left[\int_{t \geq 0} 1_{\kappa \geq t} u(q_t, Q_t) \exp(-\delta t) dt + \exp(-\delta \kappa) V(Q_\kappa) \right] \\ &= \int_{t \geq 0} [\mathbb{E} 1_{\kappa \geq t}] u(q_t, Q_t) \exp(-\delta t) dt + \int_{\kappa \geq 0} \exp(-\delta \kappa) V(Q_\kappa) d(1 - p_\kappa). \end{aligned}$$

By relabelling κ into t in the second term, we end up with the following problem:

$$\max \int_0^\infty [p_t u(q_t, Q_t) - \dot{p}_t V(Q_t)] \exp(-\delta t) dt \quad (7)$$

$$\dot{Q}_t = q_t \in [\underline{q}, \bar{q}], \quad Q_0 \text{ given} \quad (8)$$

$$\bar{Q}_t = \max(\max_{0 \leq t' \leq t} Q_{t'}, \bar{Q}_0) \quad \bar{Q}_0 \text{ given} \quad (9)$$

$$\dot{p}_t = \alpha(1 - F(\bar{Q}_t) - p_t) \quad p_0 \text{ given} \quad (10)$$

where $\bar{Q}_0 \geq Q_0$ and $p_0 \geq 1 - F(\bar{Q}_0)$. Notice that the problem is autonomous: time enters only through geometric discounting. There are three state variables: the stock Q , the maximum stock recorded so far \bar{Q} , and the survival probability p . Their initial values (Q_0, \bar{Q}_0, p_0) provide a sufficient summary of the past trajectory $(Q)_{t < 0}$, thanks to the assumption that τ follows a Poisson process, although for later interpretations it should be borne in mind that p_0 depends on $(Q)_{t < 0}$. Note that p_0 may equivalently be replaced by the legacy of the past, $\pi_0 = 1 - \frac{1 - F(\bar{Q}_0)}{p_0}$, that measures the possibility that a catastrophe was triggered in the past.

Though the structure of the program is quite simple, the existence of solutions does not follow immediately from standard results. The difficulty lies with the maximum in (9), that creates non-convexities. We will be able to establish transparent conditions

under which the optimal paths must be monotonic, and the existence of solutions will then be easily proven.

However, due to the multiplicity of state variables not all relevant solutions are monotonic.⁸ This allows for scenarios where the planner may first experiment by increasing the stock variable, and then switch to reducing this variable so as to mitigate the loss in case a catastrophe occurs; after a while, the planner might become optimistic enough to experiment further, and so on. One interesting question is whether such policies might be optimal, which we will establish as well.

3 Searching for optimal policies

We begin by analyzing two extreme cases that provide important foundations.

3.1 Extremes: inherited catastrophe ($\pi_0 = 1$)

Let us begin with the case where a catastrophe was triggered with certainty in the past. Formally, this means that the stock has exceeded the maximum value for the tipping point before the planning date: $\overline{Q}_0 \geq \overline{S}$. Then the planner knows a catastrophe is going to occur, but he does not know when. Referring to (10), we immediately get that the survival probability is going to zero after time 0:

$$p_t = p_0 \exp(-\alpha t). \quad (11)$$

Plugging this expression into (7), the optimal policy maximizes

$$\int_0^{+\infty} [u(q_t, Q_t) + \alpha V(Q_t)] \exp(-(\alpha + \delta)t) dt \quad (12)$$

under the constraint $\dot{Q}_t = q_t$. The planner thus takes into account the trade-off between the payoff u before the occurrence of the catastrophe, and the continuation value V when the catastrophe occurs. As observed for example in Kemp (1976), the possibility of a catastrophe occurring modifies the apparent discount factor.

Notice the formal similarity with the NCP objective defined in (1). Here, the past triggering of a catastrophe modifies the payoffs and the discount factor, but we can use essentially the same proofs as in Lemma A.4 to prove similar results. In particular, this problem is autonomous, and it admits a monotonic solution. Moreover, the behavior of

⁸As observed in, e.g., Benhabib and Nishimura (1979).

the solution in the long-run is governed by the following function of Q , built from the objective function exactly as $\nu = u_q + \frac{1}{\delta}u_Q$ was built from $u \exp(-\delta t)$ in our NCP case:

$$u_q(0, Q) + \frac{u_Q(0, Q) + \alpha V'(Q)}{\alpha + \delta}.$$

This turns out to be equal to:

$$\nu(Q) - \frac{\alpha}{\alpha + \delta} D'(Q). \quad (13)$$

From Assumptions 1-2, this function is weakly decreasing in Q , and is below $\nu(Q)$. Therefore, it may reach zero only at a value $Q^D \leq Q^N$.

Definition 1 Q^D (where D stands for “Damages”) is the stock level at which (13) is zero, and for simplicity we assume it is unique. By convention, we set $Q^D = +\infty$ if (13) is positive for all Q , and similarly we set $Q^D = -\infty$ if (13) is negative for all Q .

To understand the definition of Q^D , recall that we are in the case where a catastrophe has been triggered, and will occur with certainty in the future. Then any increase in the stock yields an increase in future damages, discounted by a coefficient $\frac{\alpha}{\alpha + \delta}$ that takes into account the stochastic delay before the catastrophe occurs.⁹ Thus, Q^D measures the sensitivity of the expected damage to the stock level. If the damage D does not depend on Q , then Q^D will be high, and in fact equal to Q^N . In that case, the planner should behave in a fatalistic manner, since he cannot mitigate the damage. Conversely, if the damage is very dependent on the stock at the occurrence date, then the planner should actively mitigate the damage by targeting a low value Q^D for the stock. These intuitions are readily verified, as follows:

Lemma 1 Suppose $\bar{Q}_0 \geq \bar{S}$. Then there exists an optimal path $(Q_t)_{t \geq 0}$. Moreover, all such optimal paths are monotonic, and converge to Q^D .

This extreme case study shares some similarity with the hazard rate approach used in Clarke and Reed (1994), Polasky et al. (2011), Sakamoto (2014), van der Ploeg and de Zeeuw (2017), or Besley and Dixit (2019). In those works, the catastrophe happens at time t with hazard rate $h(Q_t)$, where h is a given function, so that the survival probability writes:

$$p_t = p_0 \exp\left(-\int_0^t h(Q_\tau) d\tau\right).$$

⁹Recall that the delay τ follows a Poisson process with parameter α , so that this coefficient is indeed $E \exp(-\delta \tau) = \frac{\alpha}{\alpha + \delta}$.

Comparing with (11), we see that these works can be interpreted as assuming that a catastrophe has been triggered in the past. They then focus on how to best manage two distinct elements. First, the delay before the catastrophe can be controlled by reducing the stock because h is assumed to be an increasing function of Q , which is an element absent from our model as, by assumption, the delay follows a process with a constant hazard rate α . Second, the damage from the catastrophe, as in our model $V(Q)$, is allowed to depend on Q . By assuming exogenous (random) delays our setting is thus less general but, on the other hand, it allows to deal with the question of whether to trigger a catastrophe in the first place.

3.2 Extremes: no legacy of the past ($\pi_0 = 0$)

The opposite extreme case posits that the legacy of the past is negligible, either because experimentation stopped a long time ago, or because it did not begin at all yet. So assume $\underline{S} \leq Q_0 = \bar{Q}_0 < \bar{S}$, and $\pi_0 = 0$. Then beliefs on S are truncated at \bar{Q}_0 , and the interval $[\bar{Q}_0, \bar{S}]$ is *terra incognita*. In such a situation, one may stabilize the stock by playing $q = 0$ forever. One may also experiment a bit more before stabilizing the stock. To compare these policy options, one computes the instantaneous utility gain from experimenting, and subtract the expected damage of triggering a catastrophe, to obtain:

$$\nu(Q) - \frac{\alpha}{\alpha + \delta} \rho(Q) D(Q). \quad (14)$$

Noticeably, the second term involves the level of damage D , not its derivative as in the Poisson arrival rate model, and the hazard rate ρ measures the probability of triggering a catastrophe at this point. This expression is weakly decreasing in Q , under our assumptions.

Definition 2 Q^E (where E stands for “Experimentation”) is the stock level at which (14) is zero, and for simplicity we assume it is unique. By convention, we set $Q^E = \underline{S}$ if (14) is negative at \underline{S} , and similarly we set $Q^E = \bar{S}$ if (14) is positive at \bar{S} .

One consequence of this definition is that, if one prefers to stabilize the stock forever at some level Q , then it must be that $Q \geq Q^E$; otherwise, it would be strictly profitable to experiment a bit more. This effect is already present in the tipping point approach that we have mentioned in the Introduction (e.g., Lemoine and Traeger, 2014), and that

we study in Appendix B. We extend this approach by introducing a delay, so that the definition of Q^E includes a discounting term $\frac{\alpha}{\alpha+\delta}$.

3.3 Leaving extremes ($0 < \pi_0 < 1$)

We proceed next to a few general results on which our main analysis is based. Let us begin by asking whether one would want to trigger a catastrophe with probability one, so that the legacy of the past reaches its maximum ($\pi = 1$). We can show that this is not the case if Q^D is low enough, as follows:

Lemma 2 *Suppose $Q^D \leq \bar{S}$ and $Q_0 \leq \bar{S}$. Then every optimal path is such that $\bar{Q}_\infty \leq \bar{S}$.*

The idea here is that as soon as \bar{S} is reached, Lemma 1 implies that the path must converge to $Q^D \leq \bar{S}$ in a monotonic way, and thus this path cannot strictly exceed \bar{S} . Let us however underline that our focus in this study is on cases in which planning begins at low levels of the stock, a focus that preserves diverse possibilities.

To study this diversity of cases, we rely on the targets Q^E and Q^D . We know that both targets must lie below Q^N , but they cannot be ranked in general. In fact, we shall soon see that the comparison of these two values plays a key role for the characteristics of an optimal policy. Proofs rely on two methods. One method performs a simple calculus of variations: we replace part of a path by a constant value for the stock. This stabilization can be performed on a finite interval $[t_1, t_2]$, but one has to take care of the continuous pasting to the original path: here, this requires $Q_{t_1} = Q_{t_2}$, and also $\bar{Q}_{t_1} = \bar{Q}_{t_2}$, so that beliefs at time t_2 are not affected by the stabilization. These conditions disappear when one performs the stabilization on the whole future $[t_1, +\infty[$, and we use both cases in the proofs given in Appendix. The following result obtains using this technique:

Lemma 3 *Suppose (Q) is an optimal path. Then Q is weakly increasing when it is strictly below Q^D .*

Intuitively, when the legacy of the past is maximum ($\pi_0 = 1$), we know from Lemma 1 that the planner wants to increase the stock until it reaches Q^D , even though this means experimenting more. With a lower legacy from the past, the planner should be even more eager to increase the stock, because the associated increase in damages $D'(Q)$ is now weighted by a lower probability.

The second method is tailored to asymptotic studies. As time goes to infinity, the legacy of the past (i.e., the probability that a catastrophe was triggered in the past) goes to zero. This means that the asymptotic behavior of the stock is determined by the incentives to experiment further (the location of Q in comparison to Q^E), and not by the control of damages (whose importance is measured by Q^D). The following result relies on such an asymptotic reasoning:

Lemma 4 *Suppose an optimal path is such that $\bar{Q}_\infty < \min(Q^N, \bar{S})$. Then Q_t is weakly increasing, for t high enough. Moreover, Q_t converges to \bar{Q}_∞ , and $\bar{Q}_\infty \geq Q^E$.*

Indeed, stabilizing the stock strictly below Q^E would be bad policy, as it becomes strictly beneficial to experiment further, at least asymptotically. Interestingly, the Lemma also shows that the stock must ultimately be increasing with respect to time. Hence, the possibility of non-monotonic fluctuations in the stock seems to be linked to the interplay between experimentation on one hand, and the fear of the legacy of the past on the other hand. This interplay disappears asymptotically because the legacy of the past vanishes in the limit.

3.4 Monotonic paths

In this section and the following, we present the main results. Let us begin by a case where the initial level is low, and the various benchmarks are ranked as follows:

$$\bar{Q}_0 < Q^E < Q^D < \min(Q^N, \bar{S}).$$

This ranking of Q^E and Q^D is associated to two features. Firstly, the damage D associated to a catastrophe is high, or the hazard rate is high, so that experimenting is risky; hence Q^E is low. Secondly, the damage does not depend much on the stock level at the time of occurrence: $D'(Q)$ is low, so that Q^D is high. One may think for example to the sudden collapse of a productive ecosystem, that occurs after total catch has exceeded some threshold. The damage from the collapse is the existence value of the ecosystem, assumed high, but it does not depend on the aggregate catch. The following result characterizes the optimal policy:

Theorem 1 *Suppose $\bar{Q}_0 < Q^E < Q^D < \min(Q^N, \bar{S})$. Then there exists an optimal path. Moreover, all optimal paths are weakly increasing, and reach \bar{Q}_∞ at some time $T \leq +\infty$,*

with $Q_T \in [Q^E, Q^D)$ such that

$$\nu(Q_T) = \frac{\alpha}{\alpha + \delta} \left(\pi_T D'(Q_T) + (1 - \pi_T) \rho(Q_T) D(Q_T) \right). \quad (15)$$

Formula (15) allows for two possibilities: either experiments stop in infinite or finite time. In the first situation the planner stops experimenting asymptotically, $T = +\infty$, π_T goes to zero, and the formula implies that the stock converges to Q^E (see Definition 2). This is a case in which the legacy of the past is small, it remains small over time because the planner experiments very slowly, and thus the planner is quite certain on path that the followed policy secures the welfare against losses.

In the second situation, in contrast, T is finite and the planner stops experimenting while still worrying about the catastrophe that may be under way, $\pi_T > 0$. This second case highlights the role of the legacy of the past. If it is initially high, or if the planner decides to experiment intensively, then at any point in time the planner attributes a high probability to the fact that a catastrophe may be under way. Because by assumption the damage does not depend much on the stock level at the time of occurrence, otherwise we would not have $Q^E < Q^D$, the planner has incentives to reap as much instantaneous gains as possible before the catastrophe happens. This pushes the stock up in the direction of Q^D , as indicated by the first term on the right-hand side of the formula. For this to happen, it must be that π_T is high enough, and therefore Q_T is reached in finite time. After the stopping time T , π_t vanishes over time, and the planner becomes more and more optimistic; but because $Q_T > Q^E$ is already reached, there are no welfare gains to be made by reducing the stock, and stabilization continues to be optimal. Finally, and somewhat paradoxically, higher initial values of the legacy of the past encourages this fatalistic behavior and promotes even more experimentation, and the final value Q_T further increases. We offer next a simple example in which these insights are shown to hold quite generally.

Before proceeding to the example, we emphasize that in the situation with finite T experiments may stop under uncertainty which fits neither of the approaches in the literature: in the hazard rate approach, the planner knows that the catastrophe is pending; in the tipping point approach, the planner knows that the catastrophe is avoided by stopping. As already observed, the formula itself is a nice synthesis of the hazard rate approach (corresponding to the first term) and of the tipping point approach (associated to the second term in (15)). The formula includes in addition a delay that is responsible

for the factor $\frac{\alpha}{\alpha+\delta}$.

Consumption example: consider a consumer that consumes at each date a quantity $q_t \in [\underline{q}, \bar{q}]$. As in our model, a catastrophe is triggered when total cumulative consumption Q_t exceeds an unknown threshold S , and the catastrophe occurs after a stochastic delay. This simple example generalizes the analysis of how to eat a cake of unknown size (Kemp, 1976), but could as well be applied to the collapse of a productive ecosystem, or even to stochastic defaults in corporate finance when total borrowing exceeds some unknown tolerance level set by the lender. Our key assumptions here are, first, that instantaneous utility is a linear function of q only: $u(q, Q) = u_0 + u_1 q$, where $u_1 > 0$, and $u_0 > 0$. Second, when the catastrophe occurs the activity simply disappears, and the consumer ends up with a zero continuation payoff: $V(Q) = 0$.

When applied to this simple case, our definitions yield $\nu(Q) = u_1 > 0$, so that $Q^N = +\infty$: the consumer would consume over time without bounds in the absence of catastrophes. Then the damage associated to the occurrence of a catastrophe is simply $D(Q) = \frac{u_0}{\delta}$, the loss of safe utility flow. Because the damage does not depend on the level of the stock at the occurrence date, we immediately obtain $Q^D = +\infty$, meaning that the planner chooses to consume without limits over time if a catastrophe is known to have been triggered. Finally, the target Q^E is implicitly defined by the equality

$$u_1 = \frac{\alpha}{\alpha + \delta} \rho(Q^E) \frac{u_0}{\delta},$$

assuming that a solution $Q^E \in [\underline{S}, \bar{S}]$ exists. One easily checks that Assumptions 1-2-3 hold.

Assume now that at date zero, planning begins with $\underline{S} \leq Q_0 = \bar{Q}_0 < Q^E$, and some initial survival probability $p_0 \geq 1 - F(Q_0)$. Because $Q^D = +\infty$, Lemma 3 implies that optimal policies are weakly increasing. Consequently, in the problem in (7)-(10) simplifies because \bar{Q} equals Q everywhere, and constraint (9) has disappeared. Thanks to the assumption that utility and constraints are linear in q , the Pontryagin's principle becomes easy to follow: the optimal policy consist of setting q at its maximum level \bar{q} up to some time T , and then to stabilize the stock forever by setting $q = 0$. The associated payoff is thus a function of T only:

$$W(T) = (u_0 + u_1 \bar{q}) \int_0^T p_t \exp(-\delta t) dt + u_0 \int_T^{+\infty} p_t \exp(-\delta t) dt,$$

where the survival probability is computed by solving (10):

$$p_t = p_0 \exp(-\alpha t) + \alpha \exp(-\alpha t) \int_0^t (1 - F(Q_0 + \bar{q} \min(\tau, T))) \exp(\alpha \tau) d\tau.$$

One can check that the first-order condition $W'(T) = 0$ is equivalent to formula (15):

$$u_1 = \frac{\alpha}{\alpha + \delta} (1 - \pi_T) \frac{u_0}{\delta},$$

so that we get the value of π_T . But finding a maximizer of $W(T)$ raises an additional difficulty. Because $Q^D = +\infty$, the planner may find it profitable to trigger a catastrophe with certainty by choosing $T = +\infty$. Fortunately, we can prove the following results without having to fully characterize the maximum of W .

Proposition 1 *Assume $\underline{S} \leq Q_0 = \bar{Q}_0 < Q^E$ and $V = 0$ in the consumption example. Then, experiments stop at finite $T > 0$ if and only if*

$$u_0 \int_T^{+\infty} P_t \exp(-\delta t) dt \geq (u_0 + u_1 \bar{q}) \int_T^{+\infty} p_t \exp(-\delta t) dt,$$

where P_t on the left-hand side is survival probability p_t for which experimentation has stopped. Moreover, Q_T is strictly above Q^E , and T and Q_T are strictly increasing in the legacy π_0 .

Intuitively, the planner prefers stabilizing in finite time if safe utility u_0 is high enough, or if marginal utility u_1 is low enough. Low legacy, perhaps because the current stock standing was reached slowly, makes the planner cautious and experimentation stops early. We prove this result here in the text.

First, Lemma 4 implies that Q_T is at least equal to Q^E . Since we have assumed $Q_0 < Q^E$, this shows that T must be strictly positive. In other words, it is always profitable to increase the stock at least up to Q^E .

Second, by playing $\bar{q} > 0$ at each period one must reach Q^E in finite time. At that point, there must exist a strictly positive legacy of the past. As a consequence, an optimal path must increase the stock strictly above Q^E .

Third, an increase in p_0 has a very simple effect on survival probabilities:

$$\frac{\partial p_t}{\partial p_0} = \exp(-\alpha t).$$

This makes it easy to compute the derivative of $W(T)$ with respect to p_0 , and to show that this derivative is itself increasing with respect to T .¹⁰ Hence, by a supermodularity

¹⁰Indeed, we get $(\alpha + \delta) \frac{\partial W(T)}{\partial p_0} = u_0 + u_1 \bar{q} - u_1 \bar{q} \exp(-(\alpha + \delta)T)$, which is increasing wrt T .

argument we immediately obtain that the set of maximizers of W is increasing with p_0 in the strong set order. Now, recall the definition (6) of π : $\pi_0 = 1 - (1 - F(Q_0))/p_0$. We thus have shown that when the legacy of the past π_0 increases, for example because the rhythm of experimentation before reaching Q_0 at time zero is higher, then the planner's optimal response is to increase experimentation. Once more, this paradoxical effect occurs because the best way to adapt to a looming catastrophe is to behave fatalistically, by reaping as much consumption as possible before the date of occurrence.

Finally, one can check $W(T)$ is above $W(+\infty)$ if and only if the condition in the Proposition holds: the two policies are identical before time T .

3.5 Non-monotonic paths

In this Section, we turn to the case where the damage strongly depends on the stock level at the date of occurrence. Formally, this means that we reverse the key ranking of the target values, so that we assume $Q^D < Q^E$. In this case, damage control becomes an important priority when the legacy of the past is high. However, asymptotically this legacy vanishes, and the limit value of optimal policies is entirely determined by experimentation. A general result follows:

Theorem 2 *Suppose $Q^D < Q^E < \min(Q^N, \bar{S})$. If an optimal path is such that $\bar{Q}_\infty < \min(Q^N, \bar{S})$, then it converges to \bar{Q}_∞ , and $\bar{Q}_\infty = \max(\bar{Q}_0, Q^E)$.*

An important remark is that there is every reason to believe that optimal paths should remain below $\min(Q^N, \bar{S})$, because one does not want to experiment further above Q^E , and one would rather reduce the stock in case a catastrophe was triggered in the past, as Q^D is even lower. While we acknowledge that this intuition does not constitute a proof, we shall carry on however.

The Theorem leads to two possibilities, depending on the size of the legacy of the past at the planning date. If this legacy is small, then the main target is Q^E . Starting from $\bar{Q}_0 < Q^E$, the planner should aim at Q^E in a monotonic way. If on the other hand $\bar{Q}_0 > Q^E$, then the target is modified into $\max(\bar{Q}_0, Q^E)$, but monotonicity should remain unaffected.

By contrast, when the legacy of the past is high, the first priority is to control the damage, by reducing the level of the stock. This is only after some time that the planner becomes optimistic enough to switch to a new phase in which he will increase the stock

level. We are then back to the above phase, in which a monotonic path is likely to be used until the target $\max(\bar{Q}_0, Q^E)$ is reached.

Therefore, non-monotonic paths may be optimal. But in any case, the asymptotic value of the stock remains the same, because it is set by the experimentation trade-off when the legacy of the past vanishes. We verify below these results in a variant of the consumption example.

A variant of the consumption example: we stick to the linear utility $u_0 + u_1 q$ with $u_0 = 0$ and $u_1 > 0$, but we now assume that the continuation value in case of a catastrophe is decreasing with Q with a sufficient slope, namely

$$V(Q) = -v_0 Q \quad v_0 > \frac{\alpha + \delta}{\alpha} u_1.$$

Thus, the function $\nu(Q) = u_1$ is unaffected, Q^N is still infinite, but the damage function becomes $D(Q) = v_0 Q$. Because of the condition on v_0 , Q^D is now $-\infty$: one should reduce the stock as much as possible when it is known that a catastrophe has been triggered for sure. Finally, Q^E is given implicitly by the equality

$$u_1 = \frac{\alpha}{\alpha + \delta} \rho(Q^E) v_0 Q^E,$$

assuming an interior solution in (\underline{S}, \bar{S}) . Now suppose that we start at time 0 after having experimented intensively in the recent past; the level of the stock is equal to the highest level on record ($Q_0 = \bar{Q}_0$), the level itself is quite high ($Q_0 > Q^E$), and so is the legacy from the past. Then one would like to reduce the stock so as to mitigate the damage in case a catastrophe was triggered in the past. After some time, the fact that a catastrophe did not occur makes the planner more optimistic, and he will switch to increasing the stock level again. We therefore expect that in this situation non-monotonic policies are optimal, and this is proven in the following proposition:

Proposition 2 *In the variant of the consumption example, suppose $Q^E < Q_0 = \bar{Q}_0 < \min(Q^N, \bar{S})$. If the legacy of the past is low enough:*

$$u_1 \geq \pi_0 \frac{\alpha}{\alpha + \delta} v_0, \tag{16}$$

then the optimal policy consists in stabilizing the stock forever: $q_t = 0$ for all t . Otherwise, there exists an optimal path, and it is such that, for some t_1, t_2 , with $0 < t_1 < t_2 < +\infty$:

- $q_t = \underline{q} < 0$ for $t < t_1$;
- $q_t = \bar{q} > 0$ for $t_1 < t < t_2$;
- $q_t = 0$ for $t > t_2$.

With low legacy, (16) holds and the planner could consider experimenting further but the inherited stock standing Q_0 is already too high for such a policy: from Theorem 2, we must have $\bar{Q}_\infty = \max(\bar{Q}_0, Q^E)$. Therefore, and because the same reasoning holds at any later date, it is optimal to freeze for good at the current stock standing.

However, with high legacy, condition (16) fails at the outset of planning and the planner turns to mitigation in the fear of damages arriving soon. But such fears must ultimately fade away, justifying the non-monotonic policy stated.

3.6 A simple comparison

One may summarize the above results by saying that a lot depends on the ranking of Q^E and Q^D . From their respective definitions, one can easily derive sufficient conditions to order these two targets. In particular, the function $Q \mapsto D(Q)(1 - F(Q))$ plays a key role.

Suppose first that this function is decreasing at $Q = Q^D$. This corresponds to a situation in which the damage does not depend much on the stock level at the date of occurrence, or the hazard rate is high. Then we immediately obtain $Q^E < Q^D$.¹¹ Consequently Theorem 1 applies: if one begins planning at a low initial level of the stock, then optimal paths are monotonic, and converge to a limit in $[Q^E, Q^D]$. Moreover (but this intuition was only verified in a simple example), the legacy of the past should intuitively determine how much to experiment, with low final stock level associated with low legacy and high final level with high legacy. The legacy thus determines if the planner adopts a cautious or fatalistic experimentation strategy.

Suppose now that this function is increasing at $Q = Q^D$. This corresponds to a situation in which the damage depends a lot on the stock level at the date of occurrence, or the hazard rate is low. Then we immediately obtain $Q^D < Q^E$. Consequently, Theorem 2 applies: if one begins planning at a low initial level of the stock, and the legacy of the past is small, then optimal paths are monotonic, and converge to a limit in Q^E . By contrast, if the legacy of the past is high, a first phase appears in which one may want

¹¹Indeed, by Definition 1 we have $\nu(Q^D) = \frac{\alpha}{\alpha+\delta}D'(Q^D) > \frac{\alpha}{\alpha+\delta}\rho(Q^D)D(Q^D)$ since $D(1 - F)$ is decreasing at Q^D . From Definition 2, Q^E must then be above Q^D .

to reduce the stock, until the legacy of the past is sufficiently reduced, before switching to the monotonic phase.

We now propose two applications in which this ranking can be studied more precisely.

4 Climate change

The two basic approaches to modeling climate catastrophes are the threshold approach (e.g., Lemoine and Traeger, 2014) and the hazard rate approach (e.g., Besley and Dixit, 2019). Lessons for policies depend crucially on the approach chosen, which we now demonstrate with a simple climate-change model inspired by Golosov et al. (2014). We let Q denote the CO₂ stock in the atmosphere, that follows a simple law of motion:

$$\dot{Q}_t = E_t - \gamma Q_t, \quad (17)$$

where E are total emissions, and $\gamma > 0$ is the constant decay rate.¹² In the absence of catastrophes, total output, denoted by Y_t , is

$$Y_t = \exp(-\theta Q) K^{1-\beta} E_t^\beta \quad (18)$$

where K stands for capital, which we shall set to 1 in this illustration, E_t is the fossil-fuel energy use, and $\beta \in (0, 1)$ is its factor share. The first term corresponds to the productive damages due to the accumulation of carbon in the atmosphere. Production is entirely consumed at each date, so that $C_t = Y_t$. Instantaneous utility of consumption is $U(C) = \ln C$.¹³

We are back to our model if we set $q = E - \gamma Q$. Then the instantaneous utility is

$$u(q, Q) = \beta \log(q + \gamma Q) - \theta Q,$$

and we obtain

$$\nu(Q) = \frac{\beta}{Q} \frac{\gamma + \delta}{\gamma \delta} - \frac{\theta}{\delta}.$$

Then the target in the absence of catastrophes is

$$Q^N = \frac{\beta}{\theta} \frac{\gamma + \delta}{\gamma}.$$

¹²The precise measure is the atmospheric temperature. We cut short the details of the emissions-temperature response of Golosov et al. (2014) and Gerlagh and Liski (2018). Some papers argue that focusing on cumulative emissions is enough for the essence of policy analysis. See Dietz and Venmans (2019).

¹³The analytical climate-economy papers often make this assumption; Traeger (2019) studies more general functional forms. We could also include a utility loss linear in Q , without affecting the results.

As with our general model, when Q exceeds the unknown threshold S , this triggers a catastrophe that occurs after a stochastic delay. In the climate change literature, there are numerous components of the Earth system that are susceptible to experiencing tipping events leading to irreversible processes (Lenton et al., 2008), with considerable variation in how long the catastrophes may be pending before they actually occur (van der Ploeg and de Zeeuw, 2017). Greenland ice-sheet is such a component for which the melting, after a critical temperature, is the irreversible process. As, for example, in Cai and Lontzek (2019), when occurring the catastrophe irreversibly changes the production possibility frontier. In our model, we model this impact by making θ increase by a factor k , and we assume that this shock is important enough:

$$k > \frac{\gamma + \delta}{\delta}.$$

After the catastrophe has occurred, planning goes on, and we can endogeneize the continuation value $V(Q)$ as follows:

$$V(Q_0) = \max \int_0^\infty \ln C_t \exp(-\delta t) dt \quad (19)$$

subject to $C_t = Y_t = \exp(-k\theta Q_t)(q_t + \gamma Q_t)^\beta$, and $\dot{Q} = q$, Q_0 given. This is a simple exercise in optimal control, whose solution leads to:

$$V(Q) = \frac{-k\theta}{\delta + \gamma} Q + \frac{\beta}{\delta} [\log(\gamma Q^N) - \log k - 1]. \quad (20)$$

Then the damage $D(Q) = \frac{u(0, Q)}{\delta} - V(Q)$ equals:

$$D(Q) = \theta Q \left(\frac{k}{\gamma + \delta} - \frac{1}{\delta} \right) + \frac{\beta}{\delta} \left(\log \frac{Q}{Q^N} + \log k + 1 \right).$$

We can now recall the approaches in the literature. Firstly, van der Ploeg and de Zeeuw (2017) provides a thoughtful analysis of how to prepare for catastrophes that are pending. The paper is explicit about the idea that the ultimate arrival of tipping is evident. In our setting, the assumption implies that the event was triggered in the past ($\bar{Q}_0 > \bar{S}, \pi_0 = 1$) but did not yet occur. From Definition 1, the optimal stock level converges monotonically to the long-run target Q^D :

$$Q^D = Q^N \frac{\gamma + \delta + \alpha}{\gamma + \delta + k\alpha},$$

which indeed is less than Q^N .

Secondly, the alternative approach assumes there is no delay between triggering and occurrence, so that there is no legacy from the past. Then experimentation determines the optimal policy, and the relevant target can be computed thanks to Definition 2, at least implicitly:

$$Q^N = Q^E \left[1 + \frac{\alpha}{\alpha + \delta} \rho(Q^E) \left(Q^E \left(\frac{\delta k}{\gamma + \delta} - 1 \right) + \frac{\beta}{\theta} \left(\log \frac{Q^E}{Q^N} + \log k + 1 \right) \right) \right]. \quad (21)$$

Increasing Q is optimal if $Q < Q^E$. Any such increase may trigger tipping and if this happens without a delay, we arrive at the threshold model with immediate learning: Expression (21) continues to hold but with $\frac{\alpha}{\alpha + \delta} = 1$. Then, this climate-economy model of ours has an information structure that is no different from that in Lemoine and Traeger (2014).

The illustration helps in interpreting the literature. When k is low and the prior ρ is high even for low Q , we get $Q^D > Q^E$. This means that the planner is prepared to increase emissions further if the “ice breaks”. In the literature, there are models where emissions increase after tipping occurrence (Bretschger and Vinogradova, 2019; van der Ploeg and de Zeeuw, 2017).¹⁴ In such models, past emissions do not have to be “undone” after tipping so one may say that priority is put to *ex ante* damage control. When k is large and the prior ρ remains low even for large Q , we get $Q^D < Q^E$. This seems to be the main case in the literature (Gerlagh and Liski, 2017; Lemoine and Traeger, 2014). One may say the planner puts priority to post-tip damage control in such models.

Proposition 3 *It holds for optimal climate policies with $\alpha < \infty$, $\bar{Q}_0 = Q_0 < Q^E$ and $\pi_0 \in (0, 1)$ that*

- (i) *Priority to ex ante damage control (low k , ρ high): Emissions continually increase until stopping in $[Q^E, Q^D)$.*
- (ii) *Priority to post-tip damage control (high k , ρ low): Emissions may first decline but ultimately converge to Q^E .*

The first part of the result is an application of Theorem 1. The overall emissions increase may depend on the legacy of the past. The second part follows from Theorem 2 where the total increase in \bar{Q} does not depend on the legacy. The planner may be initially

¹⁴In the latter model, the outcome arises in the case where the damage from the event is a one-time time drop of output.

almost certain that the calamity is pending if π_0 is close to unity. Then, in contrast with the benchmarks from the literature, where the planner returns to low emissions *after the tipping occurrence*, our planner may reduce emissions *before the event occurrence* if $Q^D < Q_0 < Q^E$ and finally chooses come back with Q and reach Q^E conditional on not observing the event.

5 Disease control and social distancing

We consider a population of agents whose mass is normalized to 1, facing a pandemic at its early stages. The proportion of infected at time t is I_t , with I_0 given, and it follows a simple law of motion:

$$\dot{I}_t = (R_t - (r + d))I_t.$$

The recovery rate r and the death rate d are given. We normalize the time unit so that parameters r and d add up to 1. Hence, from now on d is the proportion of deaths among those agents that exit the set of infected agents at every period. Thus,

$$r = 1 - d, \quad 0 < d < 1.$$

$R_t \in [0, R_0]$ is the random matching rate, and thanks to the above normalization it is also the reproduction rate, that measures how many agents are infected by a single agent. Its maximum value $R_0 > 1$ obtains when people behave as in the absence of the pandemic. By mandating social distancing, the social planner can reduce the value of R at each period. This comes at an economic cost, because the social monetary value of production at time t is an increasing and concave function $Y(R)$ of R . On the other hand, by reducing the growth of infections one can eventually reduce the number of deaths, each evaluated using a value of statistical life $w > 0$.

We can now define

$$Q = \log I \quad q = R - 1,$$

and we are back to our model, with

$$u(q, Q) = Y(q + 1) - wd \exp(Q), \quad \dot{Q} = q \in [-1, R_0 - 1],$$

and an initial value $I_0 > 0$. One important assumption is that the stock of infected is never zero, for example because some infections occur due to foreign travels, or the existence of a natural reservoir. Hence, there is no absorbing state here. As we have

seen, in the absence of a catastrophe the control of the disease is characterized by the function

$$\nu(Q) = u_q(0, Q) + \frac{u_Q(0, Q)}{\delta} = Y'(1) - \frac{wd}{\delta} \exp(Q).$$

In particular, we get that the optimal number of final infections is

$$I^N = \exp(Q^N) = \frac{\delta Y'(1)}{wd}.$$

For example, if $I_0 < I^N$, then the optimal path has $R_t > 1$, and asymptotically R_t goes to 1. The stock of infected grows monotonically until it reaches I^N . The production is below the level $Y(R_0)$ that it would reach in the absence of the disease, and decreases until it reaches $Y(1)$. The situation we model here is thus one in which the society accommodates a permanent level of infections, together with some social distancing. This corresponds to the idea of flattening the curve, hoping that in the long run some (unmodeled) discovery of a vaccine will allow a return to a “normal” situation.

Nevertheless, the scenario with the above smooth trade-offs assumes that society is able to withstand a permanent pressure on the health system, on the institutions in charge of imposing social distancing, or simply on the population’s morale. One might also want to take into account a possibility that these measures break down, or even that the infectious agent mutates into something more dangerous. We assume here that these events are triggered by a high enough number of infected J , and that this number is uncertain. We thus let F_J denote the cdf of the distribution of J , on the support $[0, \bar{J}]$. Equivalently, there is an uncertain threshold for $Q = \log I$, with a distribution F for the threshold $S = \log J$. If the hazard rate for J is ρ_J , the corresponding functions for S are:

$$F(S) = F_J(\exp(S)), \quad \rho(S) = \frac{f(S)}{1 - F(S)} = \exp(S) \rho_J(\exp(S)).$$

Hence, ρ is increasing wrt S when $J \rho_J(J)$ is increasing wrt J . When the catastrophe eventually occurs, the death rate increases from d to d' , while the recovery rate becomes r' . Society loses control, and the matching rate is set at some given value R' . Overall, we assume the following property on the rate of increase γ' of the number of infected:

$$0 < \gamma' \equiv R' - (r' + d') < \delta.$$

The second inequality allows to compute a meaningful welfare after the catastrophe starts

with an initial stock of infected $I = \exp(Q)$. We have¹⁵

$$V(Q) = \int_0^{+\infty} [Y(R') - wd'I_t] \exp(-\delta t) dt$$

where $I_t = I_0 \exp(\gamma't)$, and therefore,

$$V(Q) = \frac{Y(R')}{\delta} - \frac{wd'}{\delta - \gamma'} \exp(Q).$$

Consequently, the damage function is

$$D(Q) = \frac{u(0, Q)}{\delta} - V(Q) = \frac{Y(1) - Y(R')}{\delta} + w\mu \frac{d}{\delta} \exp(Q).$$

The first term is the actualized production loss, and the second term is associated to the increased mortality, with a dimensionless, multiplicative factor

$$\mu \equiv \frac{\frac{d'}{\delta - \gamma'} - \frac{d}{\delta}}{\frac{d}{\delta}} > 0.$$

Production loss $Y(1) - Y(R') > 0$ may be interpreted as the cost of a lockdown or other tough measures: the society must push R' down below one when the epidemic transmutes. This concludes the exposition of the model. We can now turn to determining the optimal management of this infectious disease.

Our method requires to compute two benchmark values. Firstly, Q^D is the social target when one knows that a catastrophe has been triggered in the past; for example, the disease has been uncontrolled and the tipping of healthcare is a question of time. We have

$$\nu(Q^D) = \frac{\alpha}{\alpha + \delta} D'(Q^D)$$

and we get, after some calculus:

$$I^D = \exp(Q^D) = I^N \frac{1}{1 + \frac{\alpha}{\delta + \alpha} \mu} < I^N. \quad (22)$$

Secondly, Q^E is the value of Q that balances the cost and benefits of further experimentation, in the absence of legacy from the past. We have

$$\nu(Q^E) = \frac{\alpha}{\alpha + \delta} \rho(Q^E) D(Q^E)$$

¹⁵Alternatively, one could assume that a vaccine is discovered after some (possibly stochastic) date T , independently from all other choices. Implicitly, the discount rate captures these ending considerations of the calamity.

and we now get an implicit expression:

$$I^N = \frac{\alpha}{\alpha + \delta} \rho(Q^E) \frac{Y(1) - Y(R')}{wd} + \exp(Q^E) \left[1 + \frac{\alpha}{\alpha + \delta} \mu \rho(Q^E) \right].$$

This defines a unique value $I^E < I^N$, such that:

$$I^N = I^E \left[1 + \frac{\alpha}{\alpha + \delta} \rho_J(I^E) \left(\frac{Y(1) - Y(R')}{wd} + \mu I^E \right) \right].$$

Comparing to (22), we immediately get that I^D is above I^E if and only if

$$\mu wd(1 - I^E \rho_J(I^E)) < \rho_J(I^E)(Y(1) - Y(R')). \quad (23)$$

Recall that μwd measures the increase in mortality after the catastrophe, while the term on the right-hand side measures the loss in production. As in the climate change application, we have now a simple measure of the loss, μwd , which together with the planner's beliefs as captured by ρ_J determine the ranking of I^D and I^E .


Proposition 4 *It holds for optimal disease control policies with current infections record at $I_0 < I^E$ and some legacy from the past $\pi_0 \in (0, 1)$ that*

- (i) *For $I^E < I^D$ (μwd low, ρ_J high), the planner chooses $R > 1$ for all I until reaching a value in $[I^E, I^D]$ and stops infections growth by choosing $R = 1$ thereafter;*
- (ii) *For $I^E > I^D$ (μwd high, ρ_J low), the planner may choose early containment $R < 1$ but finally chooses $R > 1$ so that infections converge to $I^E > I^D$.*

As in the climate change application, results (i) and (ii) are immediate consequences of Theorems 1 and 2, respectively. We illustrate next a policy path with a period of containment followed by relaxed distancing measures, consistent with the second result of the Theorem. Assume linear a relationship between R and output: $Y(1+q) = Y_0(1+q)$ with $Y_0 > 0$. Moreover, for the sake of illustration, assume that planning starts so late that $I_0 > I_N$. In the Appendix we compute the payoff for any feasible path $q_t, t \in [0, \infty)$, under $I_t \leq I_0$. This payoff turns out be linear in our distancing measure q :

$$\int_0^\infty q_t [G_t + H_t] dt + \text{constant}$$

Figure 1



The virus outbreak with linear production $Y(q) = Y_0(1 + q)$. The figure depicts the optimal time path of log of infections, starting with $\log I_0 = 10$. The parameters are: $\delta = .1, \gamma = .02, w = 1, d = .1, d' = .15, \alpha = .2, Y_0 = 1000, m_0 = .5$

where

$$G_t = m_0 \left(Y_0 - \frac{wd \exp(Q_t)}{\delta} \right) \exp(-\delta t),$$


$$H_t = (1 - m_0) \left(Y_0 - \left(\frac{\alpha w d'}{\delta - \gamma} + wd \right) \frac{\exp(Q_t)}{\alpha + \delta} \right) \exp(-(\delta + \alpha)t),$$

and $m_0 = \frac{1 - F(Q_0)}{p_0}$. Importantly, $G = 0$ gives the optimal level of infections in absence of catastrophes, $I^N = \exp(Q^N)$. Because $I_0 > I^N$, both terms G_0 and H_0 are strictly negative when the planning starts.¹⁶ It follows that complete lockdown $q = -1, R = 0$, is optimal at $t = 0$ and in fact for all $t \geq 0$ with $G_t + H_t < 0$. But the lockdown must end: it can be readily verified that $G_t + H_t$ turns positive at some finite $t' > 0$ when policy $q = -1$ is followed for all $t < t'$. The optimal policy after t' is to relax social distancing so that $I = \exp(Q)$ grows back to I^N . In the end, the planner knows that path is safe, as if tipping was absent.

In Figs. 1-2, we show the optimal path for the parameters reported. The first figure shows the path of infections and the second depicts the implemented R_t . When infections grow we must have $[G_t + H_t]q_t \geq 0$ which, in this illustration, holds as equality. In the

¹⁶For $H_0 < 0$, we may write: $\left(Y_0 - \left(\frac{\alpha w d'}{\delta - \gamma} + wd \right) \frac{\exp(Q_0)}{\alpha + \delta} \right) < \left(\frac{wd I_N}{\delta} - \left(\frac{\alpha w d'}{\delta - \gamma} + wd \right) \frac{\exp(Q_0)}{\alpha + \delta} \right) < \left(\frac{wd I_N}{\delta} - \left(\frac{\alpha w d'}{\delta - \gamma} + wd \right) \frac{I_N}{\alpha + \delta} \right) < \left(\frac{1}{\delta} - \frac{1}{(\delta - \gamma)} \right) \frac{\alpha}{\alpha + \delta} w d' I_N < 0$.

Figure 2



The virus outbreak with linear production $Y(q) = Y_0(1 + q)$. The figure depicts the optimal time path of reproduction rate R_t , associated with figure 1. The parameters are: $\delta = .1, \gamma = .02, w = 1, d = .1, d' = .15, \alpha = .2, Y_0 = 1000, m_0 = .5$

long-run, infections approach I^N and $R = 1$. The illustration confirms the optimality of early containment followed by increasing infections, as in hammer-and-dance policies for Covid-19 (e.g., Assenza et al., 2020). We believe that ours is the first work to rationalize the hammer-and-dance approach to disease control by learning.

6 Conclusion

Inferences about catastrophes are difficult before they actually happen. This paper developed a novel approach for understanding optimal experimentation with a tipping point that comes with delay and severity depending on past actions. The model interprets historical data for beliefs on the gains and losses of further experiments: it highlights the importance of timing of past actions. Slow histories generate more information than fast histories, so the same current stock standing can come with different information contents and different immediate actions forward. In Covid-19 crises, late planning starting after an explosion of infections can justify the optimality of a lockdown, but the same infection level can justify further steps forward if the current level was approached slowly.

The lessons carry over to climate change: the “lockdown of emissions” may be optimal until unknowns can be ruled out.

Our results suggest an agenda for the applied quantitative research that evaluates the optimal climate policies quantitatively with detailed climate-economy descriptions: The models should quantify the information content of past (unplanned) experiments to give a structural interpretation to beliefs. Our model and applications illustrate the idea but remain stylized. Cutting-edge quantitative approaches, including Cai and Lontzek (2019) and Traeger (2019), offer the frameworks for exploring the question.

References

- Assenza, T., F. Collard, M. Dupaigne, P. Fève, C. Hellwig, S. Kankanamge, and N. Werquin.** 2020. “The Hammer and the Dance: Equilibrium and Optimal Policy during a Pandemic Crisis.” TSE Working Paper, n. 20-1099, May 2020.
- Barro, Robert J.** 2006. “Rare Disasters and Asset Markets in the Twentieth Century*.” *The Quarterly Journal of Economics*, 121(3): 823–866.
- Benhabib, Jess, and Kazuo Nishimura.** 1979. “The Hopf bifurcation and the existence and stability of closed orbits in multisector models of optimal economic growth.” *Journal of Economic Theory*, 21(3): 421 – 444.
- Besley, Timothy, and Avinash Dixit.** 2019. “Environmental catastrophes and mitigation policies in a multiregion world.” *Proceedings of the National Academy of Sciences*, 116(12): 5270–5276.
- Board, Simon, and Moritz Meyer ter Vehn.** 2013. “Reputation for quality.” *Econometrica*, 81(6): 2381–2462.
- Bohren, J. Aislinn.** 2019. “Persistence in a Dynamic Moral Hazard Game.” forthcoming Theoretical Economics.
- Bonatti, Alessandro, and Johannes Hörner.** 2011. “Collaborating.” *American Economic Review*, 101(2): 632–63.
- Bonatti, Alessandro, and Johannes Hörner.** 2017. “Learning to disagree in a game of experimentation.” *Journal of Economic Theory*, 169: 234 – 269.
- Bretschger, Lucas, and Alexandra Vinogradova.** 2019. “Best policy response to environmental shocks: Applying a stochastic framework.” *Journal of Environmental Economics and Management*, 97: 23 – 41. SURED 2016: Dynamic and strategic aspects of environmental economics.
- Cai, Yongyang, and Thomas S. Lontzek.** 2019. “The Social Cost of Carbon with Economic and Climate Risks.” *Journal of Political Economy*, 127(6): 2684–2734.

- Clarke, Harry R., and William J. Reed.** 1994. "Consumption/pollution tradeoffs in an environment vulnerable to pollution-related catastrophic collapse." *Journal of Economic Dynamics and Control*, 18(5): 991 – 1010.
- Cropper, M.L.** 1976. "Regulating activities with catastrophic environmental effects." *Journal of Environmental Economics and Management*, 3(1): 1 – 15.
- Crépin, Anne-Sophie, and Eric Nævdal.** 2019. "Inertia Risk: Improving Economic Models of Catastrophes." *The Scandinavian Journal of Economics*, n/a(n/a).
- Dietz, Simon, and Frank Venmans.** 2019. "Cumulative carbon emissions and economic policy: In search of general principles." *Journal of Environmental Economics and Management*, 96: 108 – 129.
- Faingold, Eduardo, and Yuliy Sannikov.** 2011. "Reputation in continuous-time games." *Econometrica*, 79(3): 773–876.
- Fitzpatrick, Luke G., and David L. Kelly.** 2017. "Probabilistic Stabilization Targets." *Journal of the Association of Environmental and Resource Economists*, 4(2): 611–657.
- Gerlagh, Reyer, and Matti Liski.** 2017. "Carbon Prices for The Next Hundred Years." *The Economic Journal*, 128(609): 728–757.
- Gerlagh, Reyer, and Matti Liski.** 2018. "Consistent climate policies." *Journal of the European Economic Association*, 16(1): 1–44.
- Golosov, Mikhail, John Hassler, Per Krusell, and Aleh Tsyvinski.** 2014. "Optimal Taxes on Fossil Fuel in General Equilibrium." *Econometrica*, 82(1): 41–88.
- Gourio, François.** 2008. "Disasters and Recoveries." *The American Economic Review*, 98(2): 68–73.
- Guillouët, L., and M. Martimort.** 2019. "Precaution, information and time-inconsistency: Do we need a precautionary principle?" Working paper Paris PSE.
- Keller, Godfrey, Sven Rady, and Martin Cripps.** 2005. "Strategic Experimentation with Exponential Bandits." *Econometrica*, 73(1): 39–68.

- Kemp, M.** 1976. *in: M. Kemp (Ed.), Three Topics in the Theory of International Trade.*
. Second ed., North-Holland.
- Kriegler, Elmar, Jim W. Hall, Hermann Held, Richard Dawson, and Hans Joachim Schellnhuber.** 2009. “Imprecise probability assessment of tipping points in the climate system.” *Proceedings of the National Academy of Sciences*, 106(13): 5041–5046.
- Lemoine, Derek, and Christian Traeger.** 2014. “Watch Your Step: Optimal Policy in a Tipping Climate.” *American Economic Journal: Economic Policy*, 6(1): 137–66.
- Lenton, Timothy M., Hermann Held, Elmar Kriegler, Jim W. Hall, Wolfgang Lucht, Stefan Rahmstorf, and Hans Joachim Schellnhuber.** 2008. “Tipping elements in the Earth’s climate system.” *Proceedings of the National Academy of Sciences*, 105(6): 1786–1793.
- Malueg, David A., and Shunichi O. Tsutsui.** 1997. “Dynamic R&D Competition with Learning.” *The RAND Journal of Economics*, 28(4): 751–772.
- Polasky, Stephen, Aart deZeeuw, Florian Wagener, Aart deZeeuw, and Aart deZeeuw.** 2011. “Optimal management with potential regime shifts.” *Journal of Environmental Economics and Management*, 62(2): 229 – 240.
- Rheinberger, Christoph M., and Nicolas Treich.** 2017. “Attitudes Toward Catastrophe.” *Environmental and Resource Economics*, 67(3): 609–636.
- Rob, Rafael.** 1991. “Learning and Capacity Expansion under Demand Uncertainty.” *The Review of Economic Studies*, 58(4): 655–675.
- Roe, Gerard H., and Marcia B. Baker.** 2007. “Why Is Climate Sensitivity So Unpredictable?” *Science*, 318(5850): 629–632.
- Sakamoto, Hiroaki.** 2014. “Dynamic resource management under the risk of regime shifts.” *Journal of Environmental Economics and Management*, 68(1): 1 – 19.
- Salmi, J., T. Laiho, and P. Murto.** 2020. “Gradual Learning from Incremental Actions.” Working paper.

- Seierstad, Atle, and Knut Sydsaeter.** 1987. *Optimal Control Theory with Economic Applications*. North Holland.
- Tilman, D., R.M. May, C.L. Lehman, and M.A. Nowak.** 1994. “Habitat destruction and the extinction debt.” *Nature*, 371: 65–66.
- Traeger, Christian.** 2019. “ACE - Analytic Climate Economy (with Temperature and Uncertainty) (Jan 9, 2019).” Available at SSRN: <https://ssrn.com/abstract=3307622> or <http://dx.doi.org/10.2139/ssrn.3307622>.
- Tsur, Yacov, and Amos Zemel.** 1994. “Endangered species and natural resource exploitation: extinction vs. Coexistence.” *Natural Resource Modeling*, 8(4): 389–413.
- Tsur, Yacov, and Amos Zemel.** 1995. “Uncertainty and Irreversibility in Groundwater Resource Management.” *Journal of Environmental Economics and Management*, 29(2): 149 – 161.
- Tsur, Yacov, and Amos Zemel.** 1996. “Accounting for global warming risks: Resource management under event uncertainty.” *Journal of Economic Dynamics and Control*, 20(6): 1289 –1305.
- van der Ploeg, F.** 2018. “The safe carbon budget.” *Climatic Change*, 147: 47–59.
- van der Ploeg, Frederick, and Aart de Zeeuw.** 2017. “Climate Tipping and Economic Growth: Precautionary Capital and the Price of Carbon.” *Journal of the European Economic Association*, 16(5): 1577–1617.

A Proof Appendix

N.B.: to alleviate notations, we often omit arguments when there is no ambiguity, and we write e for $\exp(-\delta t)$. We also make use of the function $N(Q, Q') = \int_Q^{Q'} \nu(x)dx$. $N(Q, Q')$ is concave in Q' under Assumption 1, and $N(Q, Q') - D(Q')$ is concave in Q' under Assumption 2.

For further reference, we state the following result, that follows from computing the unique solution to the differential equation (10):

Lemma A.1 For any T and $t \geq T$, one has:

$$p_t = p_T \exp(-\alpha(t - T)) + \alpha \exp(-\alpha t) \int_T^t (1 - F(\bar{Q}_\tau)) \exp(\alpha \tau) d\tau. \quad (\text{A.1})$$

In particular, when \bar{Q} is a constant on $[T, t]$, we denote the survival probability by P , and one has:

$$P_t = 1 - F(\bar{Q}_T) + (p_T - (1 - F(\bar{Q}_T))) \exp(-\alpha(t - T)). \quad (\text{A.2})$$

The following Lemma is also used repeatedly in the proofs below. The result stems from comparing the candidate path to a constant path, on a bounded or unbounded interval $[t_1, t_2]$.

Lemma A.2 If $(Q)_{t \geq 0}$ is an optimal path, then

$$\begin{aligned} & \int_{t_1}^{t_2} \left[\dot{p}_t (Q_t - Q_{t_1}) \left(D'(Q_{t_1}) - \frac{\alpha + \delta}{\alpha} \nu(Q_{t_1}) \right) \right. \\ & \left. + \delta(1 - F(\bar{Q}_t)) \left((Q_t - Q_{t_1}) \nu(Q_{t_1}) - \frac{\alpha}{\alpha + \delta} \rho(Q_{t_1}) D(Q_{t_1}) (\bar{Q}_t - \bar{Q}_{t_1}) \right) \right] \exp(-\delta t) dt \geq 0 \end{aligned} \quad (\text{A.3})$$

for all (t_1, t_2) such that one of the following two cases holds:

- Case (i): $t_1 < t_2 = +\infty$.
- Case (ii): $t_1 < t_2 < +\infty$, $Q_{t_1} = Q_{t_2}$, $\bar{Q}_{t_1} = \bar{Q}_{t_2}$.

Proof of Lemma A.2: Recall that by definition $V(Q) = u(0, Q)/\delta - D(Q)$, so that

$$W \equiv \int [pu - \dot{p}V] edt = \int [pu - \dot{p} \frac{u(0, Q)}{\delta} + \dot{p}D] edt.$$

By integrating the planner's payoff W by parts between t_1 and t_2 , we get:

$$W = -[p \frac{u(0, Q)}{\delta} e]_{t=t_1}^{t=t_2} + \int [p(u - u(0, Q) + q \frac{u_Q(0, Q)}{\delta}) + \dot{p}D] edt.$$

The concavity of u in q implies:

$$W \leq -[p \frac{u(0, Q)}{\delta} e]_{t=t_1}^{t=t_2} + \int \left[p(q u_q(0, Q) + q \frac{u_Q(0, Q)}{\delta}) + \dot{p}D \right] edt.$$

In the integral we recognize ν , and this expression can be rewritten as

$$W \leq -[p \frac{u(0, Q)}{\delta} e]_{t=t_1}^{t=t_2} + \int [pq\nu + \dot{p}(D - D(Q_{t_1}))] edt + D(Q_{t_1}) \int \dot{p} edt. \quad (\text{A.4})$$

We now compute the payoff W_0 associated to a constant $Q = Q_{t_1}$ between t_1 and t_2 . Because q is zero, the above computations now yield an equality. Notice also that, though the survival probability P associated to this constant path is initially the same ($P_{t_1} = p_{t_1}$), beyond t_1 it may differ from p (see Lemma A.1.) We obtain:

$$W_0 = -[P \frac{u(0, Q)}{\delta} e] + D(Q_{t_1}) \int \dot{P} e dt.$$

If (q, Q) is a solution, then it must be that the inequality $W \geq W_0$ holds, and therefore that the right-hand side in (A.4) is above W_0 . We now study this inequality.

The first bracketed terms are equal at $t = t_1$, and also at $t = t_2$ both in case (i) (because $t_2 = +\infty$ implies that the exponential is zero) and in case (ii) (because \bar{Q} is the same constant for both paths, so that P and p are everywhere equal.)

For the factor of $D(Q_{t_1})$, we know that $p = P$ in case (i), and we have in case (ii):

$$\int_{t_1}^{+\infty} (\dot{P} - \dot{p}) e dt = \delta \int_{t_1}^{+\infty} (P - p) e dt.$$

By applying (A.1) at $T = t_1$ to both p and P , we compute:

$$P_t - p_t = \alpha \exp(-\alpha t) \int_{t_1}^t (F(\bar{Q}_\tau) - F(\bar{Q}_{t_1})) \exp(\alpha \tau) d\tau.$$

We get, by integrating by parts once more:

$$\int_{t_1}^{+\infty} (\dot{P} - \dot{p}) e dt = \frac{\alpha \delta}{\alpha + \delta} \int_{t_1}^{+\infty} (F(\bar{Q}_t) - F(\bar{Q}_{t_1})) \exp(-\delta t) dt.$$

Moreover, Assumption 3 implies $\rho(\bar{Q}) \geq \rho(\bar{Q}_1) \geq \rho(Q_1)$, and therefore:

$$\begin{aligned} F(\bar{Q}) - F(\bar{Q}_1) &= \int_{\bar{Q}_1}^{\bar{Q}} f(S) dS = \int_{\bar{Q}_1}^{\bar{Q}} \rho(S) (1 - F(S)) dS \\ &\geq (1 - F(\bar{Q})) \int_{\bar{Q}_1}^{\bar{Q}} \rho(S) dS \geq (1 - F(\bar{Q})) \rho(Q_1) (\bar{Q} - \bar{Q}_1). \end{aligned}$$

We thus have recovered the last term in (A.3). Finally, we study the remaining terms in (A.4):

$$K \equiv \int [pq\nu + \dot{p}(D - D(Q_{t_1}))] e dt.$$

Let $N(Q_{t_1}, Q) = \int_{Q_{t_1}}^Q \nu(x) dx$. We have:

$$\int pq\nu e dt = [pNe]_{t=t_1}^{t=t_2} - \int N(\dot{p} - \delta p) e dt.$$

Notice that the bracketed term is zero at t_1 (because $N = 0$) and at t_2 (because the exponential is zero in case (i), or because $N = 0$ in case (ii)). Therefore:

$$K = \int [\delta p N + \dot{p}(D - D(Q_{t_1}) - N)] e dt.$$

Assumption 1 implies that N is concave, so that

$$N \leq (Q - Q_1)\nu(Q_{t_1}).$$

Assumption 2 implies that $D - N$ is convex, so that

$$D - N \geq D(Q_{t_1}) + (Q - Q_1)(D'(Q_{t_1}) - \nu(Q_{t_1})).$$

Replacing these elements in K yields the remaining terms in (A.3). ■

Proof of Lemma 4: For $t \geq t_1$, define the function

$$B(t_1, t) = \dot{p}_t(D'(Q_{t_1}) - \frac{\alpha + \delta}{\alpha}\nu(Q_{t_1})) + \delta(1 - F(\bar{Q}_t))\nu(Q_{t_1}).$$

$1 - F$ and ν are decreasing, and take positive values because Q and \bar{Q} lie strictly below \bar{S} and Q^N . Therefore, the second term is at least

$$k_0 \equiv \delta(1 - F(\bar{Q}_\infty))\nu(\bar{Q}_\infty) > 0.$$

Moreover, \dot{p} converges to 0. Therefore, there exists T such that, for $t \geq t_1 \geq T$, $B(t_1, t) > k_0/2 > 0$.

Now, let us proceed by contradiction, and suppose that there exists t_1 and T_2 such that $T \leq t_1 < T_2$ and $Q_{t_1} > Q_{T_2}$. A first possibility is that Q_t is below Q_{t_1} for all $t > t_1$. Then \bar{Q} is a constant, and the function under the integral in (A.3) at $(t_1, t_2 = +\infty)$ equals $(Q_t - Q_{t_1})B(t_1, t)$, and is everywhere weakly negative, and sometimes strictly negative by definition of T_2 . But this contradicts the inequality in (A.3). Therefore, Q_t cannot be below Q_{t_1} for all $t > t_1$, and there must exist t_2 such that $T \leq t_1 < t_2$, and $Q_{t_1} = Q_{t_2} \geq Q_t$ for all $t \in [t_1, t_2]$, and sometimes the last inequality is strict. Then \bar{Q}_t is a constant on this interval, and once more we get a contradiction with (A.3) at (t_1, t_2) .

We now know that after time T , Q_t is weakly increasing and bounded by \bar{Q}_∞ . Therefore, Q_t converges to some level $Q^+ \leq \bar{Q}_\infty$. Moreover, p converges to $(1 - F(\bar{Q}_\infty))$, and \dot{p} goes to zero. As a consequence, when computed at time T the planner's payoff verifies:

$$\lim_{T \rightarrow +\infty} \int_{t \geq T} (pu - \dot{p}V) \exp(-\delta(t - T)) dt = (1 - F(\bar{Q}_\infty)) \frac{u(0, Q^+)}{\delta},$$

so that the planner gets an instantaneous payoff arbitrarily close to $u(0, Q^+)$, conditional on reaching a high enough date T . Suppose that the planner deviates from T onward by playing a small quantity $\Delta q > 0$ during a small interval of time Δt , and then zero afterwards. If $Q^+ < \bar{Q}_\infty$, a simple expansion of the planner's payoff shows that the payoff from the deviation is at least

$$\frac{u(0, Q_T)}{\delta} + \frac{1}{2}\nu(Q^+)\Delta q\Delta t.$$

For T high enough, this is strictly above $u(0, Q^+)/\delta$, a contradiction. So we must have $Q^+ = \bar{Q}_\infty$. Then, if $\bar{Q}_\infty < Q^E$, by following a similar reasoning we get that the same deviation yields an additional payoff at least equal to $\frac{1}{2}[\nu(\bar{Q}_\infty) - \frac{\alpha}{\alpha+\delta}\rho(\bar{Q}_\infty)D(\bar{Q}_\infty)]\Delta q\Delta t > 0$, a contradiction once more. This concludes the proof. \blacksquare

Lemma A.3 *Suppose an optimal path is such that $Q^N < \bar{Q}_\infty < \bar{S}$. Then one of the two following cases hold:*

- *For t high enough, Q_t is weakly decreasing if Q_t is above Q^N . Moreover, Q_t converges to Q^N .*
- *Or Q_t fluctuates an infinite number of time, and each fluctuation reaches a higher stock level than the preceding one by beating the previous record \bar{Q} , and each fluctuation reaches either Q^N or a lower stock level than the preceding one.*

Proof of Lemma A.3: For $t \geq t_1$, define the function

$$B(t_1, t) = \dot{p}_t(D'(Q_{t_1}) - \frac{\alpha + \delta}{\alpha}\nu(Q_{t_1})) + \delta(1 - F(\bar{Q}_t))\nu(Q_{t_1}).$$

Suppose $Q^N < Q_{t_1}$ for some t_1 . Then $\nu(Q_{t_1}) < 0$, and $D'(Q_{t_1}) - \frac{\alpha+\delta}{\alpha}\nu(Q_{t_1}) > 0$, and $\dot{p} \leq 0$. Therefore, $B(t_1, t) < 0$ for all $t \geq t_1$. Now, the expression under the integral in (A.3) is less than $(Q_t - Q_{t_1})B(t_1, t)$. A first possibility is that $Q_t \geq Q_{t_1}$ for all $t \geq t_1$, the inequality being sometimes strict. But this would contradict the inequality (A.3). Therefore, if the stock is strictly increasing at $t = t_1$, then there must exist $t_3 > t_1$ such that $Q_{t_3} < Q_{t_1}$. In particular, there must exist t_2 such that $t_1 < t_2 < t_3$, and $Q_t \geq Q_{t_1} = Q_{t_2}$ for all $t \in [t_1, t_2]$, with sometimes a strict inequality. If $\bar{Q}_{t_1} = \bar{Q}_{t_2}$, then we obtain a contradiction with (A.3). Therefore, it must be that $\bar{Q}_{t_1} < \bar{Q}_{t_2}$.

Overall, what we have shown is that if the stock is strictly above Q^N , and is strictly increasing at t_1 , then the stock must first go strictly above \bar{Q}_{t_1} , before going strictly below Q_{t_1} . Now, there are two possibilities.

Either there are arbitrarily large dates at which the stock is strictly above Q^N and is strictly increasing. But then it must be that the stock fluctuates as we have just explained, with fluctuations that are increasing in size.

Or there exists a date such that beyond this date, the stock is weakly decreasing when it is above Q^N . This implies that \bar{Q}_∞ is reached in finite time, so that afterwards \bar{Q} is a constant. Moreover, \dot{p} goes to zero. Overall, the function in the integral in (A.3) has the sign of $(Q - Q_{t_1})\nu(Q_{t_1})$, and the inequality in (A.3) implies that the stock must be increasing when it is below Q^N . Overall, the stock must thus converge, and it is easily shown that it must converge to Q^N . This concludes the proof. ■

We can now study the No-Catastrophe Problem (NCP), which consists in maximizing

$$\int_0^{+\infty} u(q_t, Q_t) \exp(-\delta t) dt$$

under the constraints $\dot{Q}_t = q_t \in [\underline{q}, \bar{q}]$, Q_0 given.

Lemma A.4 *There exists a solution to the NCP. Moreover, all solutions are monotonic, and converge to Q^N .*

Proof of Lemma A.4: Existence of a solution to the NCP follows from Theorem 15, p.237, in Seierstad and Sydsaeter (1987). Consider such a solution. To study it, we can make use of the above Lemmas, taking into account that by definition catastrophes cannot happen: hence, we set $p = 1$, $\dot{p} = 0$, and $F = f = \rho = 0$. In particular, we have $Q^E = Q^N$ (see Definition 2.)

Then (A.3) becomes

$$\nu(Q_{t_1}) \int_{t_1}^{t_2} (Q_t - Q_{t_1}) \exp(-\delta t) dt \geq 0$$

for all (t_1, t_2) as in case (i) or case (ii) in Lemma A.2. Now, suppose there exists $T < T'$ such that $Q^N > Q_T > Q_{T'}$. A first possibility is that Q is weakly decreasing forever after T . Then we have both $\nu(Q_T) > 0$, and $Q_T \geq Q_t$ for all $t \geq T$, this inequality being strict for all $t \geq T'$. But this contradicts the above inequality at $(t_1 = T, t_2 = +\infty)$. Therefore, the stock must sometimes be increasing after time T , and this implies the existence of $t_1 < t_2$ such that $Q^N > Q_{t_1} = Q_{t_2} \geq Q_t$ for all $t \in (t_1, t_2)$, the last inequality being sometimes strict. But we obtain a similar contradiction at (t_1, t_2) , as $\nu(Q_{t_1}) > 0$ and $Q_t \leq Q_{t_1}$, the last inequality being sometimes strict.

Therefore, the stock Q is weakly increasing when it is strictly below Q^N . Symmetrically, Q is weakly decreasing when it is strictly above Q^N . This implies that Q never crosses Q^N , and that Q is monotonic, as announced.

Suppose $Q_0 < Q^N$. Then Lemma 4 implies that Q_t converges to a value at least equal to Q^E , which equals Q^N in the absence of catastrophe. Moreover, we just observed that Q_t cannot cross Q^N . Therefore, Q_t must converge to Q^N .

Suppose $Q_0 > Q^N$. Then we know that Q_t is weakly decreasing, so that we must be in the first case of Lemma A.3, and the stock thus converges to Q^N , as announced. ■

Proof of Lemma 1: the proof exactly follows the proof of Lemma A.4, as the two problems are formally identical, and u and $u + \alpha V$ share the same properties. ■

Proof of Lemma 2: Let us proceed by contradiction. If there exists t such that $Q_0 \leq \bar{S} < Q_t$, then \bar{S} is reached for the first time in finite time. At that time, Lemma 1 applies, and the solution must monotonically converge to $Q^D < \bar{S}$, and thus cannot exceed \bar{S} , in contradiction with our assumption. ■

Proof of Lemma 3: Let us proceed by contradiction. Consider a path such that $Q^D > Q_T > Q_{T'}$ at some dates $T < T'$. A first possibility is that $Q^D > Q_T \geq Q_t$ for all $t > T$, the last inequality sometimes being strict. Then we get a contradiction by applying (A.3) at $(t_1 = T, t_2 = +\infty)$ (case (i)), since the first product in (A.3) is negative as soon as $Q^D > Q_T > Q_t$, and the second product is weakly negative as \bar{Q} is a constant and $Q^N > Q_{T'}$.

Therefore, there exists $t_1 < t_2$ such that $Q^D \geq Q_{t_1} = Q_{t_2} \geq Q_t$ for all $t \in [t_1, t_2]$, the last inequality being sometimes strict. On this interval, one has $\bar{Q}_t = \bar{Q}_{t_1} = \bar{Q}_{t_2}$, so that case (ii) holds. But we get a similar contradiction as above by applying (A.3). ■

Lemma A.5 *Suppose $Q^E \leq Q^D$ and $\bar{Q}_0 \leq Q^D$. Then optimal paths cannot exceed Q^D .*

Proof of Lemma A.5: If $Q^D \geq \bar{S}$, Lemma 3 implies that if the stock reaches \bar{S} , it must converge to Q^D in a monotonic way, which shows the result.

Suppose now $Q^E \leq Q^D < \bar{S}$. Let us proceed by contradiction. Consider a path such that $\bar{Q}_0 \leq Q^D < Q_t < \bar{S}$ for some $t > 0$. Then there exists t_1 such that Q_t crosses Q^D

from below at t_1 . Moreover, from Lemma 3 we know that after t_1 the path must remain above Q^D . Therefore, we have $Q_t > Q_{t_1} = \bar{Q}_{t_1} = Q^D \geq Q^E$, and we apply (A.3) in case (i): the first term is zero, and the difference between the second term and the third term is strictly negative, as $1 - F(\bar{Q})$ is positive, $\bar{Q}_t - \bar{Q}_{t_1} \geq Q_t - Q_{t_1} > 0$, and $Q_{t_1} > Q^E$. ■

Lemma A.6 *Suppose an optimal path is such that $Q^E < \bar{Q}_\infty < \min(Q^N, \bar{S})$. Then \bar{Q}_∞ is reached in finite time.*

Proof of Lemma A.6: let us proceed by contradiction, and suppose that \bar{Q}_∞ is reached only asymptotically, necessarily from below. From Lemma 4, there exists T such that Q_t is weakly increasing for $t \geq T$, and converges to \bar{Q}_∞ . Moreover, because $\bar{Q}_\infty > Q^E$, we can choose T and $Q^0 > Q^E$ such that $Q_t > Q^0 > Q^E$ for $t \geq T$. Referring to Lemma A.2, we now define, for $t \geq t_1 \geq T$, the function

$$B(t_1, t) = \dot{p}(D'(Q_{t_1}) - \frac{\alpha + \delta}{\alpha} \nu(Q_{t_1})) + \delta(1 - F(\bar{Q}_t))(\nu(Q_{t_1}) - \frac{\alpha}{\alpha + \delta} \rho(Q_{t_1})D(Q_{t_1})).$$

Because $Q_{t_1} > Q^0 > Q^E$, the second term is strictly negative; in fact, it is strictly less than

$$k^- \equiv \delta(1 - F(\bar{Q}_\infty))(\nu(Q^0) - \frac{\alpha}{\alpha + \delta} \rho(Q^0)D(Q^0)) < 0.$$

Because \dot{p} goes to zero, the first term becomes negligible compared to the second one when t is high enough, say when $t \geq t_1 \geq T' \geq T$, so that $B(t_1, t) < k^-/2 < 0$. Finally, because the stock is weakly increasing, we have $Q_t = \bar{Q}_t$, and therefore the function in (A.3) equals $(Q_t - Q_{t_1})B(t_1, t)$. This function is everywhere weakly negative, and sometimes strictly negative since Q must grow up to \bar{Q}_∞ . So its integral in case (i) cannot be weakly positive, and we have a contradiction. This shows that \bar{Q}_∞ must be reached in finite time. ■

Lemma A.7 *Suppose \bar{Q}_∞ is reached in finite time, say T , and that Q_t remains constant after time T . Then the planner's payoff equals*

$$p_0 B(0) + \int_0^T p_t B(t) \exp(-\delta t) dt$$

where

$$B(t) \equiv u(q_t, Q_t) - u(0, Q_t) - q_t u_q(0, Q_t) + q_t C(t),$$

and

$$C(t) \equiv \nu(Q_t) - \frac{\alpha}{\alpha + \delta} [(1 - \pi_t)\rho(Q_t)D(Q_t) + \pi_t D'(Q_t)].$$

Proof of Lemma A.7: the planner's payoff is

$$W(T) \equiv \int_0^T [pu - \dot{p}V]edt + \int_T^{+\infty} [P_t u(0, Q_T) - \dot{P}_t V(Q_T)]edt$$

where the survival probabilities p and P are given in Lemma A.1. The function in the second integral can be computed as follows. First, we replace P_T by $1 - F(Q_T) - \dot{P}_T/\alpha$ from (10), and then we use (A.2) to compute \dot{P}/α . The second integral is thus

$$\int_T^{+\infty} [(1 - F(Q_T))u(0, Q_T) + (P_T - 1 + F(Q_T)) \exp(-\alpha(t - T))(u(0, Q_T) + \alpha V(Q_T))]edt$$

and thus equals $\exp(-\delta T)$, times

$$(1 - F(Q_T))\frac{u(0, Q_T)}{\delta} + (P_T - 1 + F(Q_T))\frac{u(0, Q_T) + \alpha V(Q_T)}{\alpha + \delta}.$$

Using the definition $V(Q) = u(0, Q)/\delta - D(Q)$, this can be simplified into

$$Z(T) \equiv P_T \frac{u(0, Q_T)}{\delta} - \frac{\alpha}{\alpha + \delta} (P_T - 1 + F(Q_T))D(Q_T).$$

The left-derivative of $W(T)$ is thus $\exp(-\delta T)$, times

$$p_T u(q_T, Q_T) - \dot{p}_T V(Q_T) + Z'(T) - \delta Z(T).$$

Clearly, we have $p_T = P_T$, and $\dot{p}_T = \alpha(1 - F(Q_T) - P_T)$. We obtain, with obvious notations:

$$\begin{aligned} & P_T [u(q_T, Q_T) - u(0, Q_T)] - \alpha(1 - F - P)(u/\delta - D) + \alpha(1 - F - P)u/\delta + Pqu_Q/\delta \\ & - \frac{\alpha}{\alpha + \delta} (\alpha(1 - F - P) + qf)D - \frac{\alpha}{\alpha + \delta} (P - 1 + F)qD' + \delta \frac{\alpha}{\alpha + \delta} (P - 1 + F)D. \end{aligned}$$

The u terms cancel each other, and so do almost all the D terms. We divide by $P_T = p_T > 0$ to get that the left-derivative of $W(T)$ is $p_t \exp(-\delta T)$, times

$$u(q_T, Q_T) - u(0, Q_T) + qu_Q/\delta - \frac{\alpha}{\alpha + \delta} qfD/P - \frac{\alpha}{\alpha + \delta} (1 - (1 - F)/P)qD'.$$

We finally use the definitions $\pi = 1 - (1 - F)/P$ and ν to get the result. ■

Proof of Theorem 1: from Lemma A.5, optimal paths cannot exceed Q^D ; and from Lemma 4, we obtain $\bar{Q}_\infty \in [Q^E, Q^D]$. From Lemma 3, optimal paths are monotonic, so

that $Q_t = \bar{Q}_t$ and $q_t \geq 0$. Existence of a solution then follows from Theorem 15, p.237, in Seierstad and Sydsaeter (1987).

If $\bar{Q}_\infty = Q^E$, then the Theorem holds, with $T = +\infty$ and $\pi_T = 0$.

Assume now $Q^E < \bar{Q}_\infty$. Because $\bar{Q}_\infty \leq Q^D < \min(Q^N, \bar{S})$, Lemma A.6 applies, and \bar{Q}_∞ is reached in finite time: there exists $T < +\infty$ such that $Q_T = \bar{Q}_\infty$. Finally, since $\bar{Q}_0 < Q^E$ it must be that $T > 0$.

We thus have to determine $T > 0$ to maximize the planner's payoff $W(T)$, as computed in Lemma A.7. Two optimality conditions must hold. Firstly, the first-order condition $W'(T) = 0$ yields $B(T) = 0$ at the optimal date T . Second, $W(T)$ is also maximized with respect to the path before T , and the value of q_T is free; therefore, $B(T)$ must be maximum with respect to q_T . These two conditions are:

$$u(q_T, Q_T) - u(0, Q_T) - q_T u_q(0, Q_T) + q_T C(T) = 0 \quad u_q(q_T, Q_T) - u_q(0, Q_T) + C(T) = 0.$$

They imply

$$u(q_T, Q_T) = u(0, Q_T) + q_T u_q(q_T, Q_T),$$

and because u is concave wrt q , the only possibility is that $u_q(q_T, Q_T) = u_q(0, Q_T)$ (so that smooth pasting occurs if u is strictly concave), and $C(T) = 0$. The last equality gives the formula in the Proposition. ■

Proof of Theorem 2: we know from Lemma 4 that every optimal path converges to $\bar{Q}_\infty \geq Q^E$. Moreover, we have $\bar{Q}_\infty \geq \bar{Q}_0$ by definition. Let us proceed by contradiction, and consider a path such that $\bar{Q}_\infty > \max(\bar{Q}_0, Q^E)$. Then the path converges to \bar{Q}_∞ , which is strictly above Q^E . From Lemma A.6, we get that \bar{Q}_∞ is reached in finite time, say T , and T is strictly positive because $\bar{Q}_\infty > \bar{Q}_0$. By the same reasoning as at the end of Theorem 1, we obtain (15). But this equality implies that Q_T lies between Q^E and Q^D , and this contradicts the inequalities $Q^D \leq Q^E < Q_T = \bar{Q}_\infty$. ■

Proof of Proposition 1: The proof in the text. ■

Proof of Proposition 2: From Theorem 2, we can focus on paths that converge to $\bar{Q}_\infty = \bar{Q}_0 = Q_0$. Therefore, the planner never experiments after time 0. Then p can be explicitly computed using (A.2):

$$P_t = 1 - F(Q_0) + (p_0 - 1 + F(Q_0)) \exp(-\alpha t).$$

Therefore, we have to maximize

$$W = \int_0^{+\infty} [P_t(u_1 q_t) + \dot{P} v_0 Q_t] \exp(-\delta t) dt$$

which is proportional to

$$\int_0^{\infty} [(1 - \pi_0)(u_1 q_t) \exp(-\delta t) + \pi_0(u_1 q_t - \alpha v_0 Q_t) \exp(-(\alpha + \delta)t)] dt$$

where $\pi_0 = 1 - (1 - F(Q_0))/p_0 \geq 0$ is the legacy of the past, under the constraints $\dot{Q} = q$ and $Q \leq Q_0$. Now we have, for any parameter $k > 0$:

$$\int_0^{\infty} Q_t \exp(-kt) dt = \frac{Q_0}{k} + \frac{1}{k} \int_0^{\infty} q_t \exp(-kt) dt.$$

We apply this formula to integrate the objective function. Ignoring constant terms, we obtain

$$\int_0^{\infty} q_t a_t \exp(-\delta t) dt,$$

where

$$a_t \equiv (1 - \pi_0)u_1 + \pi_0 \exp(-\alpha t)(u_1 - \frac{\alpha}{\alpha + \delta} v_0).$$

Now, if $u_1 \geq \pi_0 \frac{\alpha}{\alpha + \delta} v_0$, then a_t is positive for all t , and the planner would like to set q as high as possible, taking into account the constraint that the stock must converge to Q_0 . Hence, the solution indeed consists in stabilizing the stock from the start.

Otherwise, if $u_1 < \pi_0 \frac{\alpha}{\alpha + \delta} v_0$, then a_t is initially negative, before becoming positive at some strictly positive time t_1 . The solution therefore consists in setting $q = \underline{q} < 0$ until t_1 , and then setting $q = \bar{q}$ until the stock is back to Q_0 , at time t_2 such that $\underline{q}t_1 + \bar{q}t_2 = 0$. The optimal policy is thus as stated in the claim. ■

Proof of Proposition 3: From Theorems 1 and Theorem 2. ■

Proof of Proposition 4: From Theorems 1 and Theorem 2. ■

Payoff expression in the social distancing illustration: Under constraint $I_t \leq I_0$, we can solve for p explicitly and, as in the proof of Proposition 2, write the general payoff as

$$\int_0^{\infty} m_0 u e^{-\delta t} + (1 - m_0)(u + \alpha V) e^{-(\delta + \alpha)t} dt$$

where $m_0 = \frac{1-F(Q_0)}{p_0}$. We note that

$$\begin{aligned} u &= Y_0(1+q) - wd \exp(Q) \\ u + \alpha V &= \left(Y_0(1+q) + \frac{\alpha}{\delta} Y_0 R' \right) - \left(\frac{\alpha w d'}{\delta - \gamma} + wd \right) \exp(Q) \end{aligned}$$

As in the proof of Proposition 2, we apply the following formula: for any parameter $k > 0$, we have

$$\int Q \exp(-kt) = \frac{Q_0}{k} + \int q \frac{1}{k} \exp(-kt) dt.$$

After some manipulations, we collect the terms that multiply q and drop the constants in the expression to obtain

$$\int_0^\infty q \left[m_0 \left(Y_0 - \frac{wd}{\delta} \exp(Q) \right) e^{-\delta t} + (1 - m_0) \left(Y_0 - \left(\frac{\alpha w d'}{\delta - \gamma} + wd \right) \frac{\exp(Q)}{\alpha + \delta} \right) e^{-(\delta + \alpha)t} \right] dt$$

which is the payoff as stated in the text.

B Watch your step: the no-delay case ($\alpha \rightarrow \infty$)

N.B.: The notice on notations from Appendix A applies.

This Appendix consolidates the results from the literature on the tipping point (or threshold) approach. In this approach, a catastrophe occurs as soon as it is triggered: learning is immediate after each untried increase in the stock. The situation was first studied in Tsur and Zemel (1994), Tsur and Zemel (1995), Tsur and Zemel (1996) and, for example, recently applied in Lemoine and Traeger (2014). Our contribution here is mainly to provide simple and general proofs.

Because learning is immediate, we have survival probability $p_t = 1 - F(\bar{Q}_t)$ at all dates. We also know that Q_0 is below S , since otherwise a catastrophe would already have occurred. Now, at the planning date exactly two values matter: Q_0 and \bar{Q}_0 . In spite of this multiplicity of stocks, we first show that one can focus on monotonic candidates.

Lemma B.1 *The value of the problem in the no-delay case is unaffected if one further imposes the constraint that the stock variable (Q_t) is monotonic.*

Proof of Lemma B.1: consider a candidate (q_t, Q_t) . Suppose there exist two arbitrary dates 0 and $T > 0$, such that for any $t \in [0, T]$, we have $Q_t \leq Q_0 = Q_T$. In such a case,

the record stock is a constant ($\bar{Q}_0 = \bar{Q}_T$), and therefore the problem at time zero and the problem at time T are identical. This proves that at time zero the planner could as well adopt the strategy he has planned to apply at time T .

This procedure can be applied to cases when Q is first decreasing, then increasing. Therefore, we can focus on paths that are first weakly increasing on some interval $[0, T]$, and then weakly decreasing on $[T, +\infty[$. If $T = 0$ or $T = +\infty$, we are done, so suppose $0 < T < +\infty$. Then Q_T is the maximum stock value. Therefore, after time T catastrophes cannot occur anymore, and one maximizes $\int_{t \geq T} u(q_t, Q_t) \exp(-\delta t) dt$ under the constraints $\dot{Q} = q$ and $Q_t \leq Q_T$. If $Q_T \leq Q^N$, the best thing to do is to make the last constraint bind everywhere,¹⁷ and therefore we are done, as the candidate path is weakly increasing on $[0, T]$ and constant over $[T, +\infty[$, and is thus monotonic.

The only remaining case is when $Q_T > Q^N$. Then the best thing to do after time T is to behave as in the NCP, and to adopt a path that is decreasing (see Lemma A.4) for t above T . For $t < T$, because (Q_t) is weakly increasing we have $\bar{Q}_t = Q_t$. Therefore $p_t = 1 - F(Q_t)$, and the complete payoff from the candidate path is:

$$\int_0^T [u(q_t, Q_t)(1 - F(Q_t)) + f(Q_t)q_t V(Q_t)] \exp(-\delta t) dt + \exp(-\delta T) W^*(Q_T)(1 - F(Q_T)),$$

where $W^*(Q)$ denotes the value of the NCP program at stock Q . We can rewrite the complete payoff as

$$(1 - F(Q_T))X^*(T) + \int_0^T [u(q_t, Q_t)(F(Q_T) - F(Q_t)) + f(Q_t)q_t V(Q_t)] \exp(-\delta t) dt \quad (\text{B.1})$$

where

$$X^*(T) \equiv \int_0^T u(q_t, Q_t) \exp(-\delta t) dt + \exp(-\delta T) W^*(Q_T). \quad (\text{B.2})$$

$X^*(T)$ is weakly decreasing in T . The other terms are differentiable with respect to T , and their derivatives sum to:

$$\begin{aligned} & q_t f(Q_t) \left[-X^*(T) + V(Q_T) \exp(-\delta T) + \int_0^T u \exp(-\delta t) dt \right] \\ &= q_t f(Q_T) [V(Q_T) - W^*(Q_T)] \exp(-\delta T) \end{aligned}$$

which is below zero because $q_t \geq 0$ for $t < T$, and $W^*(Q_T)$ is at least $\frac{u(0, Q_T)}{\delta}$ (because one can adopt a constant path), which itself is above $V(Q_T)$ from Assumption 2. So reducing T in the payoff (B.1) is weakly profitable, at least until $Q_T = Q^N$ (and then the


¹⁷This is easily shown: this problem is autonomous, and therefore admits a monotonic solution.

candidate becomes weakly increasing everywhere, as already observed), or until $T = 0$ (this happens when $Q_0 \geq Q^N$, and then the candidate is the NCP solution which is weakly decreasing overall). ■

Monotonicity results from two effects. Conjecturing a non-monotonic path as depicted in Figure 3, we see that Q and the record \bar{Q} are the same at t_1 and t_2 . For this reason, if it is optimal to increase Q at t_2 , the planner cannot gain anything from the detour to a lower level but rather should behave identically at t_1 . Note that this reasoning fails in the presence of delays: at t_1 the planner might want to reduce the stock in order to mitigate losses associated to a past triggering of a catastrophe. This effect is weaker at t_2 , since the planner has observed that no catastrophe has occurred so far, which makes less likely that a catastrophe was triggered in the past.

One further observes that at the left of t_1 the planner has chosen to increase the stock, even though such experimentation may trigger a costly catastrophe. However, this fear is irrelevant at the right of t_1 , as one cannot trigger a catastrophe by staying below the level \bar{Q}_{t_1} . Therefore the planner should not choose to reduce the stock, in contradiction to the Figure. And once more, this second reasoning also fails in the presence of delays, because it might be worthwhile to reduce the stock at the right of t_1 in case a catastrophe was triggered in the past.

Figure 3



Monotonicity

Monotonicity offers a simple manner to determine the optimal solution. Indeed, either

(i) the candidate is weakly decreasing: then catastrophes cannot occur, p_t is a constant $(1 - F(Q_0))$ forever, and we are back to the NCP case with the additional constraint $q_t \leq 0$, for which existence of a solution is easily proven. Or (ii) the candidate is weakly increasing. In this latter case, if $Q_0 < \bar{Q}_0$, there is an initial phase without experiment, and then $\bar{Q}_t = Q_t$ everywhere, and $p_t = 1 - F(Q_t)$. After this initial phase, the objective function is

$$\int_0^{+\infty} [u(q_t, Q_t)(1 - F(Q_t)) + f(Q_t)q_t V(Q_t)] \exp(-\delta t) dt$$

to be maximized under the constraint $\dot{Q}_t = q_t \geq 0$. This problem is autonomous, and once more our assumptions ensure the existence of a solution.¹⁸ Overall, a solution follows from the comparison of these two candidates.

We now turn to a characterization as a function of the initial values Q_0 and \bar{Q}_0 . One important reasoning is the following. When a path converges to the maximum value \bar{Q}_∞ , the option of experimenting more exists. It yields a gain measured by $\nu = u_q(0, Q) + \frac{u_Q(0, Q)}{\delta}$, but risks triggering a catastrophe with a hazard rate ρ , implying a damage D . This motivates the following definition. From our assumptions, functions ρ and D are weakly increasing, and therefore there exists a value $Q_0^E \leq Q^N$ such that:¹⁹

$$\nu(Q_0^E) = \rho(Q_0^E)D(Q_0^E). \quad (\text{B.3})$$

For the sake of simplicity, assume that this value is unique in $[-\infty, +\infty]$. Notice that due to the absence of delays, we have $Q_0^E \leq Q^E \leq Q^N$, with Q^E from Definition 2.

Proposition 5 *In the absence of delay, for each value of (Q_0, \bar{Q}_0) there exists a solution such that:*

- (i) *If $Q_0 \geq Q^N$, then the planner never experiments, and the solution is the decreasing NCP path, converging to Q^N .*
- (ii) *If $Q_0 < Q^N$, then the solution is weakly increasing. Moreover:*
 - (ii.a) *If $\bar{Q}_0 \geq Q^N$, the solution is the NCP path, converging to Q^N .*
 - (ii.b) *If $Q_0^E < \bar{Q}_0 < Q^N$, the solution is first increasing, then it is constant and equals \bar{Q}_0 .*
 - (ii.c) *If $\bar{Q}_0 \leq Q_0^E$, the solution converges to Q_0^E .*

¹⁸In particular, we need that f is bounded; see Theorem 15, p. 237, in Seierstadt and Sydsaeter, 1987.

¹⁹The case $Q_0^E = \infty$ leads to immediate adaptations to the statements that follow.

Proof of Proposition 5: the problem admits a monotonic solution for each (Q_0, \bar{Q}_0) . In case (i), suppose that the path Q_t is weakly increasing. Then, as long as $Q < \bar{Q}_0$ the flow objective function is $u(1 - F(\bar{Q}_0))$, and above \bar{Q}_0 the flow objective function is $u(1 - F(Q)) + f(Q)qV$. By setting $p = 1 - F(\max(Q, \bar{Q}_0))$, we are left with the maximization of

$$W \equiv \int [pu - pV]edt.$$

Now we use the inequality $V(Q) \leq u(0, Q)/\delta$ (from Assumption 2), and we integrate by parts to get

$$W \leq \int [pu - p\frac{u(0, Q)}{\delta}]edt = p_0\frac{u(0, Q_0)}{\delta} + \int p[u - u(0, Q) + q\frac{u_Q(0, Q)}{\delta}]edt.$$

By concavity of u , this is less than

$$p_0\frac{u(0, Q_0)}{\delta} + \int pqvedt.$$

But the second term is negative, as $Q \geq Q^N$ and thus $\nu \leq 0$, and as $q \geq 0$. Hence, the planner would be better off by choosing the constant path, since the latter yields the payoff $p_0u(0, Q_0)/\delta$.

We thus have reached a contradiction, and the planner should choose a weakly decreasing path. By construction, such a path involves no experiment. The best one is thus the NCP solution, as stated in the Proposition.

In case (ii), a weakly decreasing path would involve no experiment, and therefore would maximize $\int uedt$, with the additional constraint $q_t \leq 0$. But because $Q_0 < Q^N$, the solution to the NCP is weakly increasing, and therefore this additional constraint would be binding everywhere. Therefore, a weakly decreasing path would in fact be a constant path, so that we can focus on the case of a weakly increasing path.

Notice that such a path involves no experiment as long as it stays below \bar{Q}_0 . In case (ii.a), \bar{Q}_0 is never reached, since the best thing to do is to adopt the NCP solution, as stated in the Proposition. Otherwise we have $\bar{Q}_0 < Q^N$, and therefore \bar{Q}_0 must be reached in finite time, say T . At T , one has to maximize

$$\int_T^{+\infty} [u(q_t, Q_t)(1 - F(Q_t)) + f(Q_t)q_tV(Q_t)] \exp(-\delta t)dt$$

under the constraints $\dot{Q}_t = q_t \geq 0$, with an initial value \bar{Q}_0 . As explained in the main text, a solution exists. The problem is autonomous, and we can proceed as in Lemma

A.4 to show that it converges to a value Q such that $w_q(0, Q) + w_Q(0, Q)/\delta = 0$, where w is the function in the integral above. Here, this condition translates into

$$u_q(1 - F) + fV + \frac{u_Q(1 - F) - uf}{\delta} = 0$$

or equivalently $\nu(Q) = \rho(Q)D(Q)$, which is the definition of Q_0^E . This is possible if $Q_0^E \geq \bar{Q}_0$ (case (ii.c)). Otherwise, the constraint $q \geq 0$ binds, as in case (ii.b). ■

To convey some intuition, let us consider that the stock has been increasing in the past, so that $\bar{Q}_0 = Q_0$ at the planning date $t = 0$. In case (i), this initial value is high, so that it is best to reduce the stock, without experimenting. Then catastrophes do not play any role, and one simply has to follow the NCP solution, which is decreasing and converges to Q^N . In case (ii.c), this initial value is low. The best strategy is thus to increase the stock and to experiment, until one reaches Q_0^E . At that value, the gain ν from increasing the stock is exactly balanced by the loss ρD associated to a possible triggering of a catastrophe.

Case (ii.b) is remarkable. One begins at a value of the stock that is below Q^N , and thus one would like to increase the stock because $\nu(Q_0) > 0$. But this would imply experimenting, and experimenting is too costly because $Q_0 > Q_0^E$ and therefore $\nu(Q_0) < \rho(Q_0)D(Q_0)$. Then the solution is a constant path, equal to Q_0 . The apparition of an optimal constant path is an original feature. In the NCP, a constant path is optimal only when one begins with a value of the stock exactly equal to Q^N . In the problem with catastrophes and no delays, a constant path appears in case (ii.b), corresponding to a region with non-empty interior. Moreover, this region is relevant even if the stock was optimally managed in the past. In fact, if at the origin of time the initial stock is between Q_0^E and Q^N , then the optimal strategy is indeed to follow a constant path, equal to the initial stock Q_0 .

As already mentioned, similar results were already obtained in Tsur and Zemel, Proposition 2.1; 1995, Proposition 5.1; 1996, Proposition 4.1). Their models are slightly different, essentially because the damage function in case of a catastrophe takes a particular and different form in each of these papers. They also rely on more demanding assumptions.²⁰ But the main conclusion is the same.²¹

²⁰In particular, the proof of monotonicity relies on a complex assumption (1994, Assumption U2) that we do not need here.

²¹To quote their 1996 paper, p. 1291:

Finally, return to the consumption choice example where $u(q, Q) = u_0 + u_1 q$ and $V(Q) = -v_0 Q$. Now, $Q^N = +\infty$ as $\nu = u > 0$, and Q_0^E comes from

$$u = \rho(Q_0^E) \left(\frac{\bar{u}}{\delta} + v_0 Q_0^E \right). \quad (\text{B.4})$$

Then, by Proposition 5, $Q^N = +\infty$ implies that one should never reduce the stock. If $Q_0 \leq \bar{Q}_0 \leq Q_0^E$, one should increase Q until Q_0^E , and then stop experimentation; if $Q_0^E < \bar{Q}_0$, the stock increases until it becomes a constant at \bar{Q}_0 .

"The steady states of the optimal emission process form an interval, the boundaries of which attract the pollution process from any initial level outside the interval."