Tipping points, delays, and the control of catastrophes

Matti Liski[*] François Salanié[†]

August 27, 2019

Abstract

A decision-maker chooses a flow variable (e.g., consumption) contributing to a stock (e.g., greenhouse gases) that may trigger a catastrophe at each new untried level. A key assumption is that once triggered the catastrophe itself occurs only after a stochastic delay. We study the optimality of non-monotonic policies, with rapid experiments followed by more cautious, wait-and-see phases. History, i.e., how the current state of the world was reached, becomes a critical determinant of future policies. Starting with a history of rapid past increase in the stock leads to cautiousness and to less experimentation in total.

JEL Classification: C61, D81, Q54.
Keywords: catastrophes, experimentation, delays

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

[†]Toulouse School of Economics, INRA, 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).

Catastrophes that are triggered by a stock variable exceeding a tipping point are relevant in many fields of economics and biology. If market size is unknown, demand may tip after excessive entry to the market (Rob, 1991).[1] The introduction of a new technology may promise efficiency improvements but can also entail additional risks that, if materializing, can result in a breakdown of the market (Bonatti and Hörner, 2016). Because of climate change, the Greenland and the Antarctic Ice Sheets may melt leading to a rise in the sea level if temperatures are to rise above a threshold that is unknown to us (Kriegler et al., 2009; Naevdal, 2006). Fisheries and ecosystems may collapse when their size are reduced below a threshold (Polasky et al., 2001). Rockström et al. (2009) document nine potential tipping points for our planet.

All such situations can be conceptualized through a state variable triggering tipping when it takes values above a threshold. Typically the precise value of this threshold is unknown, reflecting scientific uncertainty or stochastic shocks. Moreover, once the tipping point is reached the catastrophe may only happen in the future after a considerable delay. For example, it has been argued that habitat fragmentation leads to an extinction debt that may take decades or even centuries to materialize, with the gradual extinction of numerous species (Tilman, 1994). Hence inference about catastrophes is difficult before they actually happen – and it is too late for cautious management once a catastrophe has happened. This paper studies such situations in a very simple but general model that can be thought of as a stylized model of environmental tipping points such as those related to climate change or ecosystems collapse. In this setting, we deliver results that generalize beyond these particular applications.

At any point in time a catastrophe might well be under way. The question we ask is whether such delays make the planner more or less cautious. Should, for example, in climate change temperature be stabilized until we are pretty sure we are on the safe side of the distribution, or do delays encourage more intensive experimentation? To answer this question, we develop a model of experimentation where the planner controls a flow variable (e.g., consumption) that contributes to a stock (e.g., greenhouse gases). An experiment is an increase in the stock variable. In contrast with the standard tipping point or threshold models (discussed below), the experiment is informative only with a

[1]There is also intensive research on macroeconomic disasters corresponding a large downward jump in GDP growth with unknown size and recovery time (Barro, 2006; Gouriou, 2008). However, learning dynamics is not the main focus of this research but rather how a given process generating catastrophes impacts the macroeconomic aggregates such as asset prices.

delay: each increase in the stock may or may not trigger the event and this becomes known to the planner as time goes by.

Consider a planner who stabilizes the stock and thereby stops experimenting. The past experiments have payoff implications only in the future, implying that even after stabilization the planner is continuously learning. Over time, the planner becomes more optimistic since past experiments do not seem to have triggered a catastrophe so far. The implications for policies that follow from this feature are unique, allowing us to draw lessons for different types of planning situations for catastrophes. For example, the ultimate end of experiments (stock stabilization) depends on initial beliefs on if a catastrophe is under way. Intuitively, the beliefs determine the risk of following the future experimentation path. In our model, the initial beliefs explicitly depend on the timing of historical experiments, and therefore the entire solution becomes history-dependent. Using specific examples, we show that a planner inheriting a more favorable history (more optimistic beliefs) is more cautious in the long run and experiments less in total. Such a planner can have more to lose than a planner who is almost sure that tipping has been triggered already.

The optimal policies are quite naturally nonmonotonic in this setting. Because the damage in case a catastrophe happens at time t is a function of the stock at the very same time t, the delay introduces an additional incentive to reduce the stock, so as to reduce the damage in case tipping was triggered in the past. However, no news is good news, so that by decreasing the stock the planner does not experiment and, conditional on not observing the event, becomes more optimistic over time that past increases in the stock did not trigger tipping. After some time, the planner must be willing to increase the stock again as he becomes almost sure that the way back to the initial stock level is safe.

Delays in learning are not only a source of caution, quite the opposite result can follow. The fact that all impacts are delayed contributes positively to welfare because of discounting: any given experiment is costly only in the future, so that the planner enjoys a benefit from living on "borrowed time". Therefore, the effect of introducing delays is to make catastrophes look more benign, and this tends to favor experimentation.

These results are new in the literature. Let us first remind two approaches that have been extensively used in the past literature to model the arrival of catastrophes. In the first approach, one may conceptualize the system under pressure as a machine: the more pressure is put, the more likely is the breakdown. The probability of the catastrophe happening during the next (short) period only depends on the current state of the system, through an exogenously assumed hazard rate function. This Markovian assumption features in many important recent applied papers, for example in the quantitative assessments of the optimal climate-change policies in Cai et. al. (2013) and Lontzek et al. (2015); for more stylized developments and other applications, see Clarke and Reed (1996), Polasky et al. (2011), and, for example, Sakamoto (2014). In this approach, delays are not taken into account, and the past history does not matter.

In the second approach, one introduces an unknown tipping point, or threshold, that the state variable should not exceed, otherwise a catastrophe immediately occurs. This approach originates back to Tsur and Zemel (1995, 1996), and it has been recently used in a quantitative assessment of the optimal climate-change policies (Lemoine and Traeger, 2014). In the absence of delays, experimenting with the thresholds is extremely informative: the planner is 100% certain that the threshold has not been exceeded if no catastrophe has occurred so far, a feature that matches the facts in most learning environments quite badly. For example, Roe and Baker (2007) argue forcefully 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. By contrast, in such threshold models the past history impacts the present only through the maximum level of the stock reached in the past: beliefs are simply truncated from below to exclude the thresholds' values lying below this maximum level.

Both approaches have their virtues. But the choice between them should not be made casually, as rather different outcomes can follow in applied quantitative works. For example, in Lontzek et al. (2015) the risk of a tipping point has qualitatively almost orthogonal implications for the optimal climate policy than in Lemoine and Traeger (2015).[2] One also has to keep in mind that in both approaches the impact of past history is reduced to a minimum.

Our approach is also different from the workhorse models (bandit models) used to study experimentation in various economic settings. As in Poisson bandit models, from not observing the event the planner updates beliefs on the arrival rate of a catastrophe

[2]However, one must bear in mind that quantitative assessments can differ in multiple dimensions. See Gerlagh and Liski (2016) for a more detailed discussion of the assumptions made in these two studies. One important difference is that in Lemoine and Traeger there is a finite support for the tipping point whereas in Lontzek et al. the hazard rate remains bounded.

(Malueg and Tsutsui, 1997; see also Keller, Rady, Cripps, 2005 and Bonatti and Hörner, 2011). However, because of delays the current beliefs explicitly depend on the timing of past experiments, and not only on the maximum threshold reached in the past. Another difference with typical bandit models is that the stock variable also has direct payoff implications; after all, the post-event penalty in most tipping point applications depends on how much pressure was put on the system during the experimentation.[3] In a bandit formulation, choosing no action (e.g., R&D effort) implies no change in the beliefs that the planner is holding.[4] In our formulation, the past experiments have payoff implications only in the future, implying that even after stabilization the planner is continuously learning, which is the feature leading to non-monotonic policies.

1 Model

Time t is a continuous variable. The planning date is $t = 0$, but the full past history will be relevant and thus we let $t \in (-\infty, \infty)$. Let Q_t denote a stock variable. An action $q_t \in [\underline{q}, \overline{q}]$, with $\underline{q} < 0 < \overline{q}$, allows to control the stock according to the following law of motion:

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

The planner in charge of managing this dynamic system also faces the possibility of a catastrophe. 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:

$$T(S) \equiv \inf\{t : Q_t > S\}. \tag{1}$$

Note that $T(S)$ is infinite if the stock never exceeds S, and that $Q_{T(S)} = S$ otherwise. 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 $u(q, Q)$, and it is discounted at the rate $\delta > 0$. We thereby allow for the stock level to directly impact

[3]Our setting is therefore more complicated than the standard Poisson bandits setting, which is why we concentrate on the case of a single decision-maker in this paper.

[4]This property is necessary for the Gittin's index policy characterization in bandit problems; see Bergemann and Välimäki, 2009.

payoffs. 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 stock S, or on the maximum level tried in the past $\overline{Q}_\kappa$, would eliminate this possibility by assumption. To illustrate, one may imagine a skater on thin ice. Immediate 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.

Overall, given S, τ, and a path $(Q_t)_{t\in(-\infty,+\infty)}$, one can apply (1) 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 in this model. First, the threshold value S is uncertain to the planner, who entertains prior (at time $t = -\infty$) beliefs characterized by a cdf F on the interval $[\underline{S}, \overline{S}]$. Second, the delay τ is also uncertain: it is drawn in a Poisson distribution with parameter $\alpha > 0$, independently from S. Hence the planner is not only unsure of the value of the tipping point, but also about whether the threshold has already been passed; this last feature is new in the literature. It 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.

To write the planning problem at time 0, we define initial conditions as the chronicle of past stocks $(Q_t)_{t\le 0}$. Beliefs at time zero are obtained by conditioning the prior beliefs by the event "no catastrophe happened until time zero", or equivalently $\kappa = T(S) + \tau \ge 0$. Therefore, given $(Q_t)_{t\le 0}$, the planner's problem is as follows:

$$\max_{(q_t,Q_t)_{t\geqslant 0}} \mathbb{E}\left[\int_0^\kappa u(q_t, Q_t)\exp(-\delta t)dt + \exp(-\delta\kappa)V(Q_\kappa) \;\middle|\; \kappa \ge 0\right] \tag{2}$$

$$\dot{Q}_t = q_t \in [\underline{q}, \overline{q}], T(S) = \inf\{t : Q_t > S\}, \qquad (Q_t)_{t\le 0} \text{ given.} \tag{3}$$

We allow for jumps in q, as long as they do not threaten measurability. We say that a path (Q_t) is monotonic if Q_t is everywhere weakly decreasing, or weakly increasing, with time. We now turn to some specific assumptions, then offer a few examples, before developing a tractable approach to solving (2)-(3).

1.1 Assumptions and Notations

There are only three primitive elements in this model, that we study in turn: functions $u(q, Q)$ and $V(Q)$, and distribution F. To begin with, we assume:[5]

Assumption 1 *Function u is twice 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.

Note that whether an increase in Q is good news or bad news is left unrestricted, so as to fit various settings. Function ν encapsulates the trade-off between instantaneous gains from an increase in q, and effects due to the associated increase in Q in all subsequent periods. ν is decreasing for any positive δ if one assumes both $u_{QQ} \leq 0$ and $u_{qQ} \leq 0$, but the above formulation highlights which property is really needed. To avoid discussing multiplicities, we also assume that there exists a unique solution $Q^* \in [-\infty, +\infty]$ to the equation $\nu(Q) = 0$. Note that Q^* need not be finite.

As a benchmark, we now define the No-Catastrophe Problem (NCP), as the case in which catastrophes do not occur. In the NCP, the planner maximizes

$$\int_0^{+\infty} u(q_t, Q_t) \exp(-\delta t) dt \tag{4}$$

under $\dot{Q}_t = q_t$, Q_0 given. As time enters only through geometric discounting, the problem is said to be autonomous. Lemma A.1 in the Appendix studies this problem and shows it admits a monotonic solution. The monotonicity of ν implies that for this solution, the planner chooses to gradually increase (if $Q_0 < Q^*$) or reduce (if $Q_0 > Q^*$) the stock until it reaches the level Q^*.

We define catastrophes as costly, irreversible events. Irreversibility means that the continuation value V is fixed, and beyond the control of the social planner. The catastrophe is costly if V is less than the value of forever stabilizing the stock at a safe level.

[5]Subscripts denote partial derivatives.

We therefore define the damage function D as

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

and we assume:

Assumption 2 *The damage function $D(Q)$ is continuously differentiable, weakly positive, weakly increasing, and weakly convex.*

In particular, this assumption implies that V is bounded from above. Finally, let the c.d.f. F be continuously differentiable on its support $[\underline{S}, \overline{S}]$, 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.*

1.2 Examples

The setting covers multiple interpretations, in addition to the skater's adventures sketched above. It is formulated to fit the management of a stock of pollution, such as carbon dioxide. Or one may instead think to a fishery, with obvious changes in signs; q is then interpreted as the catch from the stock. As in many stylized pollution-control models, the stock may follow the law of motion

$$\dot{Q}_t = e_t - \gamma Q_t.$$

By the change of variable $q = e - \gamma Q$, we are back to our model. The utility $v(e, Q)$ now becomes $u(q, Q) = v(q + \gamma Q, Q)$, and Assumption 1 holds for example if v is concave, with a negative cross-derivative.

The model can also be applied to the management of debt. A borrower allocates his revenue w between consumption c (yielding a utility $U(c)$, assumed concave) and net savings s, with $w = c + s$. His debt Q commands interest payments rQ. The law of motion of debt is thus

$$\dot{Q}_t = rQ_t - s_t.$$

The change of variable $q = rQ - s$ yields $\dot{Q} = q$, as in our model. Moreover, one has $c = w + q - rQ$, so that utility $u(q, Q) \equiv U(w + q - rQ)$ is indeed a function of both q and Q. Then $\nu(Q) = (1 - r/\delta)U'(w - rQ)$, and thus Assumption 1 holds if $r > \delta$.

Finally, the lender triggers bankruptcy when debt exceeds some uncertain level S, and liquidation occurs after some stochastic delay. Quite naturally, the continuation payoff after liquidation is computed as a function $V(Q_\kappa)$ of the debt level at the liquidation date κ. Note that in the NCP, liquidation never occurs by assumption, but interest charges have to be paid. Because $r > \delta$, the borrower should reduce its debt as fast as he can. Indeed, ν is negative everywhere, and therefore the desired level of the debt is $Q^* = -\infty$. Whether this conclusion also holds when liquidation is allowed after an (uncertain) delay is studied in this paper.

Finally, we study in section 5 a linear model in which consumption q is always beneficial, and does not create any damages. Therefore, in the NCP the planner should consume as much as he can, and we have $Q^* = +\infty$. Introducing catastrophes leads to richer conclusions, as we shall see.

2 Restating the problem

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 now define the survival probability at time t, computed at the beginning of times using the prior beliefs F:

$$p_t \equiv Prob(\kappa \geq t).$$

One may distinguish two possibilities. Either S is above $\overline{Q}_t$, and in that case a catastrophe cannot happen before time t. Or S is below $\overline{Q}_t$, which means that a catastrophe was triggered at time $T(S) < t$, but did not happen before time t because the delay τ is above $t - T(S)$. Overall, we obtain

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

The second term 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:

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

The survival probability can also be expressed as follows. At time 0 we have

$$p_0 = 1 - F(\overline{Q}_0) + K,$$

where

$$K \equiv \int_{S \leq \overline{Q}_0} \exp(\alpha T(S)) dF(S)$$

is the probability that a catastrophe has been triggered before time 0, but did not happen yet. We complement these by the law of motion:

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

which indeed allows us to recover (5). Hence, in any case p_t is above $(1 - F(\overline{Q}_t))$, because past experiments may have triggered a catastrophe; and p_t is nonincreasing, without surprise since it 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 $(1 - F(\overline{Q}_t))$, and the legacy from the past π_t is zero everywhere. Otherwise, if delays exist ($\alpha < +\infty$), then either planning starts with no experimentation in the history: $\overline{Q}_0 \leq \underline{S}$, $\pi_0 = 0$, and $p_0 = 1$. Or the planner inherits some experiments from the past: $\overline{Q}_0 > \underline{S}$, and then one has $\pi_t > 0$ and $p_t > 1 - F(\overline{Q}_t)$ 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 does not alter the maximization problem. The expected payoff (2) can be written as follows:

$$\mathbb{E}\left[\int_{t \geq 0} 1_{\kappa \geq t} u(q_t, Q_t) \exp(-\delta t) dt + \exp(-\delta \kappa) V(Q_\kappa)\right]$$

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

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

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

$$\dot{Q}_t = q_t \in [\underline{q}, \overline{q}], \qquad Q_0 \text{ given} \tag{8}$$

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

$$\dot{p}_t = \alpha(1 - F(\overline{Q}_t) - p_t), \qquad p_0 \text{ given} \tag{10}$$

where $\overline{Q}_0 \geq Q_0$ and $p_0 \geq 1 - F(\overline{Q}_0)$. One can see that only three values are relevant at the planning date: Q_0, $\overline{Q}_0$, and p_0. These values are informative enough to replace the full description of the past trajectory $(Q_t)_{t \leq 0}$, thanks to the assumption of a Poisson process for delays.

Notice also that the problem is autonomous: time enters only through geometric discounting. However, there are three state variables here, so that solutions need not be monotonic.[6] 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. We will pay careful attention to the logic of such experimentation policies.

3 Watch your step: the no-delay case ($\alpha \to \infty$)

In the absence of delays, a catastrophe occurs as soon as it is triggered: learning is immediate, and the legacy from the past in the survival probability is zero. This case was first studied in Tsur and Zemel (1994, 1995, 1996) and recently applied by Lemoine and Traeger (2014), and our contribution here is mainly to provide simple and general proofs. Because learning is immediate, we have $p_t = 1 - F(\overline{Q}_t)$ at all dates. We also know that Q_0 is below $\overline{S}$, since otherwise a catastrophe would already have occurred. Now, at the planning date exactly two values matter: Q_0 and $\overline{Q}_0$. In spite of this multiplicity, we first show that one can focus on monotonic candidates.

Lemma 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.*

Monotonicity results from two effects. Conjecturing a non-monotonic path as depicted in Figure 1, we see that Q and the record $\overline{Q}$ 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.

[6] As observed in, e.g., Benhabib and Nishimura (1979).

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 $\overline{Q}_{t_1}$. 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 1

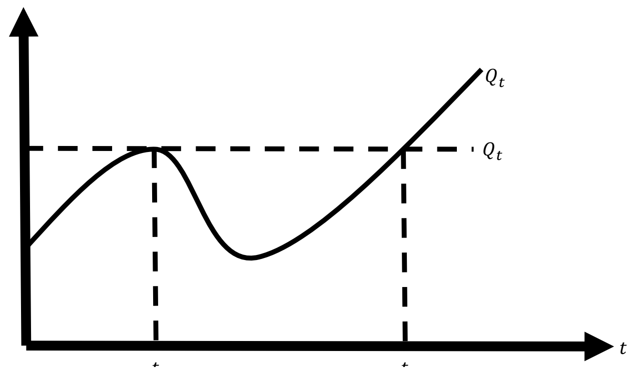

Monotonicity

Monotonicity thus results both from the fact that the problem is autonomous, and from the fact that catastrophes are costly. It ensures that the problem has a solution, and offers a simple manner to determine it. Indeed, either we impose that 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 we impose that the candidate is weakly increasing. In this latter case, if $Q_0 < \overline{Q}_0$, there is an initial phase without experiment, and then $\overline{Q}_t = Q_t$ 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 monotonic solution.[7] Overall, existence of a solution follows from the comparison of these two candidates.

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

We now turn to a characterization of solutions, as a function of the initial values Q_0 and $\overline{Q}_0$. One important reasoning is the following. When a path converges to the maximum value $\overline{Q}_\infty$, the option of experimenting more exists. It yields a gain measured by ν, 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 $Q_0^{**} \leq Q^*$ such that:[8]

$$\nu(Q_0^{**}) = \rho(Q_0^{**})D(Q_0^{**}). \tag{11}$$

For the sake of simplicity, assume that this value is unique in $[-\infty, +\infty]$.

Proposition 1 *In the absence of delay, for each value of* $(Q_0, \overline{Q}_0)$ *there exists a solution such that:*

(i) If $Q_0 \geq Q^*$*, then the planner never experiments, and the solution is the decreasing NCP path, converging to* Q^**.*

(ii) If $Q_0 < Q^*$*, then the solution is weakly increasing. Moreover:*

(ii.a) If $\overline{Q}_0 \geq Q^*$*, the solution is the NCP path, converging to* Q^**.*

(ii.b) If $Q_0^{**} < \overline{Q}_0 < Q^*$*, the solution is first increasing, then it is constant and equals* $\overline{Q}_0$*.*

(ii.c) If $\overline{Q}_0 \leq Q_0^{**}$*, the solution converges to* Q_0^{**}*.*

To convey some intuition, let us focus on the simpler case when the stock has been increasing in the past, so that $\overline{Q}_0 = Q_0$ 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^*. 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^{**}. 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^*, 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 $Q_0 > Q_0^{**}$ 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

[8] The case $Q_0^{**} = \infty$ leads to immediate adaptations to the statements that follow.

only when one begins with a value of the stock exactly equal to Q^*. 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^{**} and Q^*, 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 (1994, 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.[9] But the main conclusion is the same. To quote their 1996 paper, p. 1291:

> 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.

4 The case with delays

Let us now turn to the general problem with delays, as stated in (7)-(10). We begin by studying what can happen in the very long run. A first possibility is that the planner chooses a path that exceeds $\overline{S}$, thereby triggering a catastrophe with certainty. So let us call T a date at which $\overline{Q}_T \geq \overline{S}$, and let us study the planning problem after time T. Referring to (10), we immediately get that the survival probability is going to zero:

$$p_t = p_T \exp(-\alpha(t - T)). \tag{12}$$

Plugging this expression into (7), we have to maximize

$$\int_T^{+\infty} [u(q_t, Q_t) + \alpha V(Q_t)] \exp(-(\alpha + \delta)t) dt \tag{13}$$

under the constraint $\dot{Q}_t = q_t$. Notice the formal similarity with the NCP objective defined in (4). 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.1 to prove similar results. In particular, this problem is autonomous, and it admits a monotonic

[9] In particular, their proof of monotonicity relies on a complex assumption that we do not need here (1994, Assumptions U1-U2).

solution. Moreover, the behavior of the solution in the long-run is governed by the following function of Q, built from the objective function exactly as ν was built from $u\exp(-\delta t)$:

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

This function turns out to be equal to:

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

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 $\hat{Q}$ below Q^*. Once more, for simplicity we assume that this value is unique in $[-\infty, +\infty]$.

Lemma 2 *Suppose there exists T such that $Q_T \geq \overline{S}$. Then after T, there exists an optimal path $(Q_t)_{t \geq T}$ that is monotonic, and that converges to $\hat{Q}$.*

To understand the definition of $\hat{Q}$, recall that we are in the case when 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 that takes into account the stochastic delay before the catastrophe occurs.[10] A similar trade-off appears in Clarke and Reed (1996), Polasky et al. (2011), Cai et. al. (2013), Sakamoto (2014), or Lontzek et al. (2015.) In those works, the catastrophe happens at time t with a hazard rate $h(Q_t)$, where h is a given function, so that the survival probability writes:

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

Comparing with (12), 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: i) the delay before the catastrophe, which can be controlled by reducing the stock because h is assumed to be an increasing function of Q; this element is absent from our model, as by assumption the delay follows a process with a constant hazard rate α, independent from Q; ii) and the damage from the catastrophe, as in our model where $V(Q)$ is allowed to depend on Q. By assuming exogenous 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.

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

In particular, whether one would want to trigger a catastrophe with probability one is an intriguing question. We can show that this is not the case if $\hat{Q}$ is low enough, as follows:

Lemma 3 *Suppose $\hat{Q} \leq \overline{S}$, and $Q_0 \leq \overline{S}$. Then one can focus on paths that lie below $\overline{S}$: $Q_t \leq \overline{S}$, for all t.*

Indeed, assume $\overline{S}$ is reached at some finite time T for the first time. After T, because $\hat{Q} \leq \overline{S}$, Lemma 2 implies that the optimal path is weakly decreasing. This shows that the stock can never exceed $\overline{S}$, as claimed.

We can now give a useful definition. For a given path (Q), define $\overline{Q}_\infty$ as the supremum value for the stock. In general, this value will be finite, either because the stock never exceeds $\overline{S}$; or because it does exceed this level, but then converges to $\hat{Q}$. In fact, the only case where the supremum value can be infinite is when $\hat{Q}$ is itself infinite.

So far, we have studied the case when the supremum value is above $\overline{S}$. For the alternative case, a general result is as follows:

Lemma 4 *If $\overline{Q}_\infty < \overline{S}$, one can focus on paths that converge to $\min\{Q^*, \overline{Q}_\infty\}$.*

The idea here is very simple: since $\overline{Q}_t$ converges to a finite limit $\overline{Q}_\infty$, experimentation has to become negligible in the long-run. As time goes by, the legacy from the past goes to zero, and the objective becomes to maximize the integral of the instantaneous payoff u, under the constraint that the stock does not exceed $\overline{Q}_\infty$. Still, this last value is endogenous to the problem, and to determine it we have to dig deeper.

A simple argument applies here. Suppose that the path converges to $\overline{Q}_\infty$, and consider what happens in the very long-run. The planner could increase q by some small amount, during some small period of time. The gain in utility flow is then measured by ν. On the other hand, the associated experiment may trigger a catastrophe: the corresponding hazard rate is ρ, and compared to the constant path the damage is D. Notice additionally that this damage will only be incurred after a delay, and thus is discounted by $E\exp(-\delta\tau) = \frac{\alpha}{\alpha+\delta}$.

This intuitive reasoning motivates the following definition. Because ρ and D are weakly decreasing, there exists a value Q^{**}, with $Q_0^{**} \leq Q^{**} \leq Q^*$ such that:[11]

$$\nu(Q^{**}) = \frac{\alpha}{\alpha+\delta}\rho(Q^{**})D(Q^{**}). \tag{14}$$

[11]The case $Q^{**} = \infty$ leads to immediate adaptations to the statements that follow.

For the sake of simplicity, assume that this value is unique in $[-\infty, +\infty]$. The intuition given above indicates that when the stock is below Q^{**} one should experiment further, and indeed we have the following result:

Lemma 5 *One can focus on paths such that* $\overline{Q}_\infty \geq Q^{**}$.

Finally, our next result further restricts the set of paths to be considered.

Lemma 6 *One can focus on paths* (Q_t) *such that:*

(i) When the stock is strictly below $\hat{Q}$*, then it is weakly increasing.*

(ii) When the stock is strictly above Q^* *(but strictly below* $\overline{S}$*), then it is weakly decreasing.*

An intuition for the first result is as follows. Suppose the stock is below $\hat{Q}$. When a catastrophe was triggered in the past with probability one, Lemma 2 indicates that one should increase the stock. If more generally the probability of having triggered a catastrophe is less than one, then increasing the stock is even more valuable.

The proof of the second result is much more involved, but the intuition is similar. Suppose the stock is above Q^*. In the absence of catastrophe, the NCP solution indicates that one should adopt a decreasing path. If in addition a catastrophe may have been triggered in the past, then reducing the stock is even more valuable. On the other hand, when catastrophes are introduced one now has the possibility to experiment further, and this effect makes the proof more involved.[12]

Finally, an important consequence of this second result is that when the stock is initially lower than Q^*, it must remain below forever.

5 Applications

In this section, we provide a few applications of our general results.

[12]The difficulty comes from our definition of a catastrophe as a net loss compared to a constant path (see Assumption 2). This definition is simple and general, but it does not allow to rely on a catastrophe being sufficiently damaging. This is related to the Assumptions used in Tsur and Zemel (1994), see footnote 9.

5.1 A cake of unknown size, with delays

The question of how to eat a cake of unknown size was first proposed by Kemp (1976), and Kumar (2005) provides a fairly general study. The question may seem narrow, but in fact it encompasses questions such as the sustainability of an extraction effort under uncertainty, or the resilience of an ecosystem to the harvesting of a particular specie. In that model, the planner has to choose a consumption q at each date, until the cumulative consumption Q exceeds an uncertain threshold S, drawn in a distribution F. Then consumption is zero forever. We extend this framework by introducing a stochastic delay between the triggering date and the occurrence of the catastrophe, on one side; and by allowing for damages after the catastrophe that depend on the consumption stock Q at the date the catastrophe occurs.

So assume instantaneous utility $u(q)$ only depends on consumption q, where u is strictly increasing on $[\underline{q}, \overline{q}]$, and is normalized so that $u(0) = 0$. Then $\nu(Q) = u'(0)$ is strictly positive, and therefore $Q^* = +\infty$. This means that in the absence of catastrophe, the planner should forever extract as much as possible.

The tipping point is distributed on $[\underline{S}, \overline{S}]$, with $0 \leq \underline{S} < \overline{S}$, with a cdf F and an increasing hazard rate ρ. When a catastrophe occurs, the game stops, and the continuation value is assumed to be linear:

$$V(Q) = -(d_0 + d_1 Q) \qquad d_0 \geq 0, d_1 \geq 0.$$

This last assumption extends the literature, that has so far assumed $d_0 = d_1 = 0$. The values of these two parameters will turn out to be very important. For further reference, notice that the damage function $D(Q)$ is simply

$$D(Q) = d_0 + d_1 Q.$$

At this stage, one can apply Proposition 1 to solve the problem in the case without delays. Let Q_0^{**} be the value such that

$$u'(0) = \rho(Q)(d_0 + d_1 Q).$$

From the Proposition, we obtain that solutions are always nondecreasing. Only two cases (cases (ii).b and (ii).c) are possible. Either one inherits from the past a low level of experimentation ($\overline{Q}_0 < Q_0^{**}$), and then the optimal path is increasing and converges to Q_0^{**}. Or one has already experimented a lot ($\overline{Q}_0 \geq Q_0^{**}$), and then the optimal path consists in consuming zero forever. A higher damage function impacts this general picture only to the extent that it reduces the value of Q_0^{**}.

Let us now introduce delays that follow a Poisson distribution. Lemma 2 has shown that if the cumulative consumption Q exceeds $\overline{S}$ at some date, then it must converge toward the level $\hat{Q}$, such that $\nu(\hat{Q}) = \frac{\alpha}{\alpha+\delta} D'(\hat{Q})$. This definition suggests to distinguish two cases that are very different.

In **Case 1**, which is the case usually considered in the literature, the damage is a constant d_0, independent from cumulative extraction: so that $d_1 = 0$. Then it is easily shown that $\hat{Q}$ is infinite: once a catastrophe is triggered with certainty, the optimal management consists in extracting as much as possible before the catastrophe occurs. In such a case, from result (i) in Lemma 6 we obtain that the solution must be weakly increasing.

This result relies on the fact that when $d_1 = 0$, then the damage in case a catastrophe occurs does not depend on the present level of the stock. Hence, even though past experiments matter to present utility levels, they do not support a reduction of the present stock for precautionary reasons: the legacy of the past does not make the decision-maker more cautious.

To summarize: in Case 1, a monotonic solution exists (to be characterized soon!), and the introduction of delays makes the decision-maker less cautious (proof to be completed soon.)

In **Case 2**, the damage now depends on cumulative extraction: $d_0 = 0 < d_1$. Now assume

$$u'(0) < \frac{\alpha}{\alpha + \delta} d_1.$$

Then $\hat{Q} = -\infty$: once a catastrophe is triggered with certainty, the optimal management consists in reducing the stock as much as possible, thereby trying to mitigate damages before a catastrophe occurs. Now suppose that we start at time 0 after having experimented intensively in the recent past; so that the level of the stock is equal to the highest level on record ($Q_0 = \overline{Q}_0$), the level itself is quite high, 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 Case 2 non-monotonic policies are sometimes optimal. To show this formally, let us simplify this model by assuming that utility is linear:

$$u(q) = q \qquad q \in [\underline{q}, \overline{q}],$$

and that S is uniformly distributed on $[0, \overline{S}]$. An important threshold is Q^{**}, such that $\nu(Q) = \frac{\alpha}{\alpha+\delta}\rho(Q)D(Q)$. Here we obtain

$$Q^{**} = \frac{\overline{S}}{1 + \frac{\alpha}{\alpha+\delta}d_1}.$$

This threshold measures the incentives to experiment further, when there is no legacy from the past. The following claim proves that a non-monotonic path is optimal, under conditions that hold if marginal damages d_1 are high enough.

Lemma 7 *In Case 2, let* $\pi_0 \equiv 1 - \frac{1-F(\overline{Q}_0)}{p_0} < 1$ *be the legacy from the past at time zero. Suppose* $Q^{**} < Q_0 = \overline{Q}_0 < \overline{S}$, *and*

$$1 < \pi_0 \frac{\alpha}{\alpha + \delta} d_1. \tag{15}$$

Then 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* $t < t_1$;
- $q_t = \overline{q} > 0$ *for* $t_1 < t < t_2$;
- $q_t = 0$ *for* $t > t_2$.

Another interesting insight is that this phenomenon occurs if the legacy from the past is high enough, so that history matters. In particular, if one has experimented a lot in the recent past, then one should become more cautious in the near future. This is a form of hysteresis that did not appear in previous models.

6 Conclusion

Inferences about catastrophes are difficult before they actually happen. We have developed a novel approach for understanding the optimal experimentation with tipping points that come with delay and severity depending on past actions. History, i.e., how the current state of the world was reached, becomes a critical determinant of future policies.

References

[1] R.J. Barro (2006): "Rare Disasters and Asset Markets in the Twentieth Century", *The Quarterly Journal of Economics*, 121 (3): 823-866.

[2] A. Bonatti and J. Hörner (2011): "Collaborating", *American Economic Review*, 101(2): 632-63.

[3] A. Bonatti and J. Hörner (2016): "Learning to Disagree in a Game of Experimentation", working paper.

[4] Y. Cai, K.L. Judd, T.M. Lenton, T.S. Lontzek, D. Narita (2015b): "Environmental tipping points significantly affect the cost-benefit assessment of climate policies", *Proceedings of the National Academy of Sciences*, 112:4606-4611.

[5] H. Clarke and W. Reed (1994): "Consumption/pollution tradeoffs in an environment vulnerable to pollution-related catastrophic collapse", *Journal of Economic Dynamics & Control*, 18, pp. 991-1010.

[6] F. Gourio (2008): "Disasters and recoveries", *American Economic Review*, 98(2):68-73, May.

[7] M. Kemp (1976): "How to eat a cake of unknown size", in: M. Kemp (Ed.), Three Topics in the Theory of International Trade, North-Holland, Amsterdam, 1976.

[8] G. Keller, S. Rady, and M. Cripps (2005): "Strategic experimentation with exponential bandits", *Econometrica*, Vol. 73, No. 1, 39-68.

[9] R.C. Kumar (2005): "How to eat a cake with unknown size: A reconsideration", *Journal of Environmental Economics and Management* (50) 408-421.

[10] E. Kriegler, J.W. Hall, H. Held, R. Dawson, and H.J. Schellnhuber (2009): "Imprecise probability assessment of tipping points in the climate system", *Proceedings of the National Academy of Sciences*, Vol. 106: 5041-5046.

[11] A. Lange and N. Treich (2008): "Uncertainty, learning and ambiguity in economic models on climate policy: some classical results and new directions", *Climatic Change*, 89:7-21.

[12] D.M. Lemoine and C.P. Traeger (2014): “Watch Your Step: Optimal Policy in a Tipping Climate”, *American Economic Journal: Economic Policy* 6(1): forthcoming.

[13] D. Malueg and S. O. Tsutsui (1997): “Dynamic R&D Competition with Learning”, *The RAND Journal of Economics*, Vol. 28, No. 4, pp. 751-772.

[14] E. Naevdal (2006): “Dynamic optimization in the presence of threshold effects when the location of the threshold is uncertain, with an application to a possible disintegration of the Western Antarctic Ice Sheet”, *Journal of Economic Dynamics and Control*, 30, 1131-1158.

[15] S. Polasky, A. de Zeeuw, and F. Wagener (2011): “Optimal management with potential regime shifts”, *Journal of Environmental Economics and Management*, 62, pp. 229-240.

[16] R. Rob, “Learning and Capacity Expansion under Demand Uncertainty”, *The Review of Economic Studies*, Vol. 58, No. 4, June, pp. 655-675.

[17] J. Rockström, W. Steffen, K. Noone, A. Persson, F.S. Chapin III, E.F. Lambin, T.M. Lenton, M. Scheffer, C. Folke, H.J. Schellnhuber, B. Nykvist, C.A. de Wit, T. Hughes, S. van der Leeuw, H. Rodhe, S. Sörlin, P.K. Snyder, R. Costanza, U. Svedin, M. Falkenmark, L. Karlberg, R.W. Corell, V.J. Fabry, J. Hansen, B. Walker, D. Liverman, K. Richardson, P. Crutzen, and J.A. Foley (2009): ”A safe operating space for humanity”, *Nature*, 461: 472-475. doi:10.1038/461472a

[18] G.H. Roe and M.B. Baker (2007): “Why Is Climate Sensitivity So Unpredictable”, *Science*, Vol. 318, No.5850, October, pp. 629-632.

[19] A. Seierstadt and K. Sydsaeter (1987): Optimal control theory with economic applications, New York, Elsevier Science Publishing.

[20] Y. Tsur and A. Zemel (1994): “Endangered species and natural resource exploitation: extinction vs. coexistence”, *Natural Resource Modeling*, Vol.8, No.4, Fall, pp.389-413.

[21] Y. Tsur and A. Zemel (1995): “Uncertainty and irreversibility in groundwater management”, *Journal of Environmental Economics and Management*, 29, 149-161.

[22] Y. Tsur and A. Zemel (1996): "Accounting for global warming risks: resource management under event uncertainty", *Journal of Economic Dynamics & Control*, 20, pp. 1289-1305.

A Proof Appendix

N.B.: to alleviate notations, we often write e for $\exp(-\delta t)$, and we omit arguments when there is no ambiguity.

We begin by studying the No-Catastrophe Problem (NCP), which consists in maximizing

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

under the constraints $\dot{Q}_t = q_t \in [\underline{q}, \overline{q}]$, Q_0 given. Let $W^*(Q_0)$ denote the value of this program.

Lemma A.1 *For every Q_0, there exists a monotonic solution to the NCP problem, that converges to Q^*. Moreover, for every path (Q, q), the function*

$$X^*(T) \equiv \int_0^T u(q_t, Q_t) \exp(-\delta t) dt + \exp(-\delta T) W^*(Q_T) \tag{A.1}$$

is weakly decreasing with T.

Proof of Lemma A.1: Existence follows from Theorem 15, p.237, in Seierstadt and Sydsaeter (1987). The NCP is an autonomous problem, and therefore at any time t the optimal decision for q can be specified so as to depend only on the value of the state variable Q_t. In turn, this implies that there exists a monotonic solution.[13] For this monotonic solution, because u is concave in q we have:

$$u(q_t, Q_t) \leq u(0, Q_t) + q_t u_q(0, Q_t) = u(0, Q_0) + \int_{Q_0}^{Q_t} u_Q(0, x) dx + q_t u_q(0, Q_t).$$

The second term can be integrated by parts:

$$\int_0^{+\infty} e \int_{Q_0}^{Q_t} u_Q(0, x) dx dt = \left[\frac{\exp(-\delta t)}{-\delta} \int_{Q_0}^{Q_t} u_Q(0, x) dx \right]_0^{+\infty} + \int_0^{+\infty} q_t \frac{u_Q(0, Q_t)}{\delta} edt.$$

[13] A general proof of this result is in Hartl (1987).

The bracketed term is zero. We have shown:

$$\int_0^{+\infty} u(q_t, Q_t)\exp(-\delta t)dt \leq \frac{u(0,Q_0)}{\delta} + \int_0^{+\infty} q_t \nu(Q_t)\exp(-\delta t)dt. \qquad \text{(A.2)}$$

Consider the case $Q_0 < Q^*$ (the case $Q_0 > Q^*$ can be handled in a symmetric way.) If the solution is weakly decreasing, then Q_t is everywhere belowQ^*, so that $\nu(Q_t) > 0$ from Assumption 1. Therefore, if in addition the solution is somewhere strictly decreasing, then from the above inequality the criterion on the left-hand side is strictly less than $u(0, Q_0)/\delta$, which is the feasible payoff associated to the constant path $(q = 0, Q = Q_0)$: we thus get a contradiction. This shows that the solution must be weakly increasing. Let $Q^+ \leq +\infty$ be its supremum.

If $Q^+ > Q^*$, then after some time T the stock is weakly increasing above Q^*, so that $q_t \geq 0$ and $\nu(Q_t) < 0$, and once more we get a contradiction thanks to (A.2).

If $Q^+ < Q^*$, or if Q^+ is finite and $Q^* = +\infty$, then for every positive ε there exists T such that $Q_t \geq Q^+ - \varepsilon$ for all $t \geq T$. After T, because ν is decreasing we have $0 < \nu(Q_t) \leq \nu(Q^+ - \varepsilon)$, which is finite. Moreover, we have $\int_{t \geq T} q_t dt = Q^+ - Q_T \leq \varepsilon$. This shows that the second term in the right-hand-side of (A.2) is proportional to ε, and finally that the right-hand-side of (A.2) can be made arbitrarily close to $u(0, Q^+)/\delta$. Therefore, for T high enough the best payoff the planner can hope for is also arbitrarily close to $u(0, Q^+)/\delta$.

But then one can add $\Delta q > 0$ to q during a short period of time Δt, with a change in payoff arbitrarily close to:

$$\Delta q \Delta t u_q(0, Q^+) + \Delta q \Delta t \frac{u_Q(0, Q^+)}{\delta} = \Delta q \Delta t \nu(Q^+).$$

If $Q^+ < Q^*$, this expression is strictly positive, thanks to Assumption 1, and we obtain a contradiction. Therefore the solution must converge to Q^*, as announced.

Finally, for a given path (q_t, Q_t), define $X^*(T)$ as in the Lemma. Then, for $T < T'$ we have:

$$X^*(T) - X^*(T') = \exp(-\delta T)W^*(Q_T) - \exp(-\delta T')W^*(Q_{T'}) - \int_T^{T'} u(q_t, Q_t)\exp(-\delta t)dt.$$

Let $T_0 = T' - T$, $q'_t = q_{t+T}$, $Q'_t = Q_{t+T}$. Then the above expression equals $\exp(-\delta T)$, times

$$W^*(Q'_0) - \left[\int_0^{T_0} u(q'_t, Q'_t)\exp(-\delta t)dt + \exp(-\delta T_0)W^*(Q'_{T_0})\right],$$

which is weakly positive by definition of W^*. This shows that $X^*(T)$ is weakly decreasing in T. ■

Proof of Lemma 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 ($\overline{Q}_0 = \overline{Q}_T$), 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 all similar cases, and in particular to those 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 $Q_T \leq Q^*$, the best thing to do is to make the last constraint bind everywhere,[14] 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 $Q_T > Q^*$. 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.1) for t above T. For $t < T$, because (Q_t) is weakly increasing we have $\overline{Q}_t = Q_t$. Therefore $p_t = 1 - F(Q_t)$, and the complete payoff from the candidate path is:

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

We can rewrite it as

$$(1 - F(Q_T)) X^*(T) + \int_0^T [u(q_t, Q_t)(F(Q_T) - F(Q_t)) + f(Q_t) q_t V(Q_t)] e dt \qquad \text{(A.3)}$$

where $X^*(T)$ is defined in Lemma A.1, and is weakly decreasing in T. The other terms are differentiable with respect to T, and their derivatives sum to:

$$q_t f(Q_t) \left[-X^*(T) + V(Q_T) \exp(-\delta T) + \int_0^T u e dt \right]$$
$$= q_t f(Q_T) \left[V(Q_T) - W^*(Q_T) \right] \exp(-\delta T)$$

[14] This is easily shown: this problem is autonomous, and therefore admits a monotonic solution; the solution cannot be somewhere strictly decreasing, by an argument similar to the one following the derivation of (A.2).

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 (A.3) is weakly profitable, at least until $Q_T = Q^*$ (and then the candidate becomes weakly increasing everywhere, as already observed), or until $T = 0$ (this happens when $Q_0 \geq Q^*$, and then the candidate is the NCP solution which is weakly decreasing overall). ■

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

$$W \equiv \int [pu - \dot{p}V]edt.$$

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

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

By concavity of u, this is less than

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

But the second term is negative, as $Q \geq Q^*$ 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 $p_0 u(0, Q_0)/\delta$.

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 $Q_0 < Q^*$, 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 $\overline{Q}_0$. In case (ii.a), $\overline{Q}_0$ is never reached, since the best thing to do is to adopt the NCP solution, as stated in the Proposition. Otherwise we have $\overline{Q}_0 < Q^*$, and therefore $\overline{Q}_0$ must be reached in finite time, say T. At T, one has to maximize

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

under the constraints $\dot{Q}_t = q_t \geq 0$, with an initial value $\overline{Q}_0$. As explained in the main text, a solution exists. The problem is autonomous, and we can proceed as in Lemma A.1 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^{**}. This is possible if $Q_0^{**} \geq \overline{Q}_0$ (case (ii.c)). Otherwise, the constraint $q \geq 0$ binds, as in case (ii.b). ■

Proof of Lemma 2: as explained in the text, the proof proceeds exactly as the proof of Lemma A.1. Indeed, problem (13) exhibits the same properties as the NCP: it is autonomous, and the objective $u + \alpha V$ is concave in q, and bounded from above thanks to Assumptions 1 and 2. ■

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

Lemma A.2 *For any T and $t \geq T$, one has:*

$$p_t = p_T \exp(-\alpha(t - T)) + \alpha \exp(-\alpha t) \int_T^t (1 - F(\overline{Q}_\tau)) \exp(\alpha\tau) d\tau. \tag{A.4}$$

In particular, when $\overline{Q}$ is a constant, we denote the survival probability by P, and one has:

$$P_t = 1 - F(\overline{Q}_T) + (p_T - (1 - F(\overline{Q}_T))) \exp(-\alpha(t - T)). \tag{A.5}$$

We now study the case when $\overline{Q}_\infty$ is reached in finite time. Suppose that Q reaches a maximum value $\overline{Q}_\infty$ at time $T < +\infty$. After time T, there is no new experimentation, and $\overline{Q}$ remains constant. We use (A.5) to compute the objective function:

$$pu - \dot{p}V = (1 - F(\overline{Q}_\infty))u + (p_T - 1 + F(\overline{Q}_\infty))(u + \alpha V) \exp(-\alpha(t - T)).$$

This motivates the study of the following problem: for $m \in [0,1]$ and Q_0 given, maximize

$$\int_0^{+\infty} [mu\exp(-\delta t) + (1-m)(u+\alpha V)\exp(-(\alpha+\delta)t)]\, dt$$

under the constraint $\dot{Q} = q$, Q_0 given. Let $W(Q_0, m)$ denote the value of this program; for $m = 0$ we are back to the program studied in Lemma 2.

Lemma A.3 *There exists a solution to the program defining $W(Q_0, m)$. Moreover, W is weakly increasing in m. Finally, for any path (q, Q), and any $(p_0 > 0, \overline{Q}_0)$, define p_t by (10). Then, if $\overline{Q}_t = Q_t$ on an interval $[t_0, t_1]$, the function*

$$X(T) \equiv \int_{t_0}^{T} [pu - \dot{p}V]edt + p_T \exp(-\delta T) W(Q_T, \frac{1-F(Q_T)}{p_T}) \tag{A.6}$$

is weakly decreasing with respect to T on $[t_0, t_1]$.

Proof of Lemma A.3: Once more, existence of a solution follows from Theorem 15, p. 237, in Seierstadt and Sydsaeter (1987). W is convex in m, as a maximum of linear functions. Moreover, by the envelope theorem its derivative equals

$$\int u\exp(-\delta t)dt - \int (u+\alpha V)\exp(-(\alpha+\delta)t)dt.$$

Thanks to the convexity of W in m, to show that this derivative is positive, it is sufficient to show that it is positive when computed using the solution to $W(Q_0, m=0)$. Along this path, we have

$$W(Q_0, 0) = \int_0^T (u+\alpha V)\exp(-(\alpha+\delta)t)dt + \exp(-(\alpha+\delta)T)W(Q_T, 0),$$

and by differentiating the right-hand-side with respect to T, we get:

$$u(q_t, Q_t) = (\alpha+\delta)W(Q_t, 0) - \alpha V(Q_t) - q_t \frac{\partial W}{\partial Q_0}(Q_t, 0). \tag{A.7}$$

By integrating over the whole path, we get

$$\int u\exp(-\delta t)dt = \int [(\alpha+\delta)W(Q_t,0) - \alpha V(Q_t)]\exp(-\delta t)dt - \int q_t \frac{\partial W}{\partial Q_0}(Q_t, 0)\exp(-\delta t)dt$$

$$= \alpha \int [W(Q_t, 0) - V(Q_t)]\exp(-\delta t)dt - \int \frac{d}{dt}[W(Q_t, 0)\exp(-\delta t)]dt$$

$$= \alpha \int [W(Q_t, 0) - V(Q_t)]\exp(-\delta t)dt + W(Q_0, 0).$$

Finally, because in the problem $W(Q,0)$ one can choose a constant path we have

$$W(Q,0) \geq \frac{u(0,Q)}{\delta} = V(Q) + D(Q) \geq V(Q)$$

and therefore we have shown that $\int u \exp(-\delta t) dt \geq W(Q_0, 0)$, as needed. This shows that $W(Q_0, m)$ is increasing in m at the right of zero, and therefore that $W(Q_0, m)$ is increasing in m on $[0, 1]$, by convexity.

Again proceeding as in Lemma A.1, for any $m \in [0,1]$ and T, define:

$$m'_T = \frac{m}{m + (1-m)\exp(-\alpha T)}.$$

Then for any path (q, Q), we can write:

$$W(Q_0, m) \geq \int_0^T [mu \exp(-\delta t) + (1-m)(u + \alpha V)] \exp(-(\alpha+\delta)t)] d\tau$$

$$+ \int_T^{+\infty} [mu \exp(-\delta t) + (1-m)(u + \alpha V)] \exp(-(\alpha+\delta)t)] d\tau.$$

where the second integral is computed using the solution to $W(Q_T, m'_T)$. Now, this second integral is easily seen to be equal to

$$(m + (1-m)\exp(-\alpha T)) \exp(-\delta T) W(Q_T, m'_T).$$

Therefore:

$$W(Q_0, m) \geq \int_0^T [mu \exp(-\delta t) + (1-m)(u + \alpha V) \exp(-(\alpha+\delta)t)] d\tau$$

$$+(m + (1-m)\exp(-\alpha T)) \exp(-\delta T) W(Q_T, m'_T).$$

Consequently, the right-derivative of the right-hand-side at $T = 0$ must be weakly negative, at any q_0, Q_0, m:

$$u + \alpha(1-m)V - (\delta + \alpha(1-m))W(Q_0, m) + q\frac{\partial W}{\partial Q_0}(Q_0, m) + \alpha m(1-m)\frac{\partial W}{\partial m}(Q_0, m) \leq 0. \tag{A.8}$$

Keeping this in mind, the derivative of $X(T)$ with respect to T is, after simplification by $\exp(-\delta T)$, and taking into account the equality $Q_T = \overline{Q}_T$:

$$pu - \dot{p}V + (\dot{p} - \delta p)W(Q_T, m) + pq\frac{\partial W}{\partial Q_0}(Q_T, m) - \frac{1}{p}(fqp + (1-F)\dot{p})\frac{\partial W}{\partial m}(Q_T, m)$$

where $m \equiv \frac{1-F(\overline{Q}_T)}{p_T}$. Accordingly, since $\dot{p} = \alpha(1 - F(Q) - p)$ we get $\dot{p}/p = -\alpha(1 - m)$. Simplifying by p, the above expression becomes

$$u + \alpha(1 - m)V - (\delta + \alpha(1 - m))W + q\frac{\partial W}{\partial Q_0} + (\alpha m(1 - m) - \frac{fq}{p})\frac{\partial W}{\partial m}.$$

Finally, from (A.8) at $q_0 = q_T$, $Q_0 = Q_T$, $m = \frac{1-F(\overline{Q}_T)}{p_T}$, we get that this derivative is less than $-\frac{fq}{p}\frac{\partial W}{\partial m}$, which is below zero as $q \geq 0$, and W has just been shown to be increasing in m. ■

Proof of Lemma 4: in the long-run, we know that $\overline{Q}_t$ converges to $\overline{Q}_\infty < \overline{S}$. Because F is continuous, this implies that $F(\overline{Q}_t)$ converges to $F(\overline{Q}_\infty)$. From (10), p_t converges to a strictly positive value $1 - F(\overline{Q}_\infty)$, while $\dot{p}_t$ converges to zero. From (7), the objective function is

$$pu - \dot{p}V = u[1 - F(\overline{Q}_\infty)] + u[p - (1 - F(\overline{Q}_\infty))] - \dot{p}V.$$

Maximizing the integral of the first term is the NCP problem. Because u and V are bounded from above, the last two bracketed terms are arbitrarily small for t high enough. Hence, these two last terms are transitory, and in the long-run only the first term matters. This shows that the planner should try to make the path as close as possible to Q^* in the long run. Taking into account the constraint $Q_t \leq \overline{Q}_\infty$, we get that the trajectory must converge to $\min\{Q^*, \overline{Q}_\infty\}$. ■

Proof of Lemma 5: Let us proceed by contradiction. Suppose $\overline{Q}_\infty < Q^{**}$. Then $\overline{Q}_\infty < Q^*$, and from Lemma 4 the path must converge to $\overline{Q}_\infty$, necessarily from below. By concavity of u in q, u converges to $u(0, \overline{Q}_\infty)$. p converges to $(1 - F(\overline{Q}_\infty))$, and $\dot{p}$ converges to zero. Total payoff is

$$W \equiv \int [pu - \dot{p}V]edt = p_0\frac{u(0, Q_0)}{\delta} + \int p[u - u(0, Q) - qu_q(0, Q)]edt + \int pq\nu edt + \int \dot{p}Dedt.$$

Suppose that for t high enough we modify q during a short period of time Δt by adding a small quantity $\Delta q > 0$. In the above expression, the first term is unaffected, and the second term can be neglected because q is close to zero and u is concave. For the third term, because q is close to zero only the initial variation of q matters, so that the variation of the third term is $\Delta q \Delta t$, times

$$(1 - F(\overline{Q}_\infty))\nu(\overline{Q}_\infty).$$

For the fourth and last term, from (A.4) we have:

$$p_t = p_0 \exp(-\alpha t) + \alpha \exp(-\alpha t) \int_0^t (1 - F(\overline{Q}_\tau)) \exp(\alpha \tau) d\tau.$$

Therefore, the variation of p is:

$$\Delta p = \alpha \exp(-\alpha t) \int_0^t (-f(\overline{Q}_\tau)\Delta q \Delta t) \exp(\alpha \tau) d\tau = -f(\overline{Q}_\infty)\Delta q \Delta t(1 - \exp(-\alpha t)).$$

Using (10), we have:

$$\Delta \dot{p} = -\alpha f(\overline{Q}_\infty)\Delta q \Delta t - \alpha \Delta p = -\alpha f(\overline{Q}_\infty)\Delta q \Delta t \exp(-\alpha t).$$

Because $\dot{p}$ converges to zero, the variation of the fourth term is:

$$-\alpha f(\overline{Q}_\infty)\Delta q \Delta t D(\overline{Q}_\infty)\frac{1}{\alpha + \delta}.$$

Finally, the total variation is $\Delta q \Delta t$, times

$$(1 - F(\overline{Q}_\infty))\nu(\overline{Q}_\infty) - \frac{\alpha}{\alpha + \delta} f(\overline{Q}_\infty) D(\overline{Q}_\infty).$$

This is positive if $\overline{Q}_\infty < Q^{**}$, and thus we get a contradiction. ■

The following Lemma is used repeatedly in the proofs below.

Lemma A.4 *Let $t_1 < t_2$, let (q, Q) be a feasible path, and let $N(Q) \equiv \int_{Q_{t_1}}^{Q} \nu(x)dx$. If (q, Q) is solution, then the following properties hold:*

(i) If $t_2 < +\infty$, $Q_{t_1} = Q_{t_2}$, $\overline{Q}_{t_1} = \overline{Q}_{t_2}$, then

$$\int_{t_1}^{t_2} \dot{p}(D - D(Q_{t_1}))edt \geq \int_{t_1}^{t_2} N\dot{\widehat{pe}}dt; \tag{A.9}$$

(ii) If $t_2 = +\infty$, then

$$\int_{t_1}^{+\infty} [\dot{p}(D - D(Q_{t_1}) - \frac{\alpha + \delta}{\alpha} N) + \delta(1 - F(\overline{Q}))N - \frac{\alpha\delta}{\alpha + \delta}(F(\overline{Q}) - F(\overline{Q}_{t_1}))D(Q_{t_1})]edt \geq 0. \tag{A.10}$$

Proof of Lemma A.4: Recall that by definition $V(Q) = u(0, Q)/\delta - D(Q)$. By integrating the planner's payoff by parts, we get:

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

The concavity of u in q implies:

$$W \leq -[p\frac{u(0,Q)}{\delta}e] + \int p[qu_q(0,Q) + q\frac{u_Q(0,Q)}{\delta}]edt + \int \dot{p}Dedt.$$

The first integral is exactly $\int pq\nu e$. Because $\dot{\widehat{N(Q)}} = q\nu(Q)$, we have:

$$\int pq\nu edt = [pNe] - \int N\dot{\widehat{pe}}dt.$$

Therefore:

$$W \leq [p(N - \frac{u(0,Q)}{\delta})e] + D(Q_{t_1})\int \dot{p}edt + \int \dot{p}(D - D(Q_{t_1}))edt - \int N\dot{\widehat{pe}}dt. \quad \text{(A.11)}$$

We now compute the payoff W_0 associated to a constant path. Hence, Q is constant, and q is zero. Therefore, repeating the above derivation now yields an equality; notice however 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.2.) We obtain:

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

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.11) is above W_0. To check this in case (i), notice that $\overline{Q}$ is the same constant for both paths, so that P and p are everywhere equal. Hence, the first bracketed term is the same in both expressions. The only terms that remain are those in (A.9).

In case (ii), the first bracketed terms are the same because they are both zero at infinity, and because they are equal at t_1. For the factors in $D(Q_{t_1})$, by applying (A.4) to both p and P, we get:

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

Then, we integrate by parts twice to get:

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

which is the last term in (A.10). The only remaining terms are the following ones:

$$\int \dot{p}(D - D(Q_{t_1}))edt - \int N\dot{\widehat{pe}}dt.$$

To simplify these terms, we notice that $\dot{\widehat{pe}} = (\dot{p} - \delta p)e$, and we can use (10) to remark that $p = -\dot{p}/\alpha + 1 - F(\overline{Q}_t)$; so that

$$\dot{\widehat{pe}} = [\dot{p}\frac{\alpha+\delta}{\alpha} - \delta(1 - F(\overline{Q}_t))]e.$$

This yields the first two terms in (A.10). ■

Proof of Lemma 6: We first show (i). Let us proceed by contradiction, and consider a path which is such that at some dates $t' < t''$ one has $Q_{t''} < Q_{t'} < \hat{Q}$. Because $\hat{Q} \leq Q^*$, from Lemma 4 Q must ultimately converge to a level above $Q_{t'}$. Therefore, there exists $t_1 < t_2$ such that $Q_t \leq Q_{t_1} = Q_{t_2} < \hat{Q}$ for all $t \in [t_1, t_2]$. Notice moreover that on this interval $\overline{Q}$ is a constant, and equals $\overline{Q}_{t_1}$. Therefore, the inequality (A.9) mut hold.

Now, because $Q_t < \hat{Q}$, we have $\nu(Q_t) \geq \frac{\alpha}{\alpha+\delta}D'(Q_t)$, and by integrating between Q_{t_1} and Q_t we get $N(Q_t) \leq \frac{\alpha}{\alpha+\delta}[D(Q_t) - D(Q_{t_1})]$.

Moreover, $\dot{\widehat{pe}} = (\dot{p} - \delta p)e \leq 0$, and we can use (10) to remark that $p = -\dot{p}/\alpha + 1 - F(\overline{Q}_{t_1})$; so that $\dot{\widehat{pe}} = (\dot{p}\frac{\alpha+\delta}{\alpha} - \delta(1 - F(\overline{Q}_{t_1})))e \leq \dot{p}\frac{\alpha+\delta}{\alpha}$. Overall, we get

$$N\dot{\widehat{pe}} \geq \frac{\alpha}{\alpha+\delta}[D(Q_t) - D(Q_{t_1})]\dot{\widehat{pe}} \geq [D(Q_t) - D(Q_{t_1})]\dot{p}e,$$

where the second inequality comes from the fact that $Q_t \leq Q_{t_1}$ and D is increasing. By integrating between t_1 and t_2, this contradicts the optimality condition (A.9) in Lemma A.4. So we can focus on paths which are weakly increasing with time when the stock is below $\hat{Q}$, as stated in (i) in the Lemma.

To show (ii), consider a path and a date at which Q is strictly above Q^*. Then $\overline{Q}_\infty > Q^*$, and from Lemma 4 Q must ultimately converge to Q^*. Therefore, $\overline{Q}_\infty$ must be reached in finite time, say T. We begin by showing that the path must be weakly decreasing for $t > T$.

After time T, the survival probability evolves as in (A.5), and the best the planner can hope is to get $p_T W(Q_T, m)$, where W is the function defined just before Lemma A.3, and $m = (1 - F(Q_T))/p_T > 0$. In fact, the planner gets exactly this payoff, because we now show that the solution to the problem $W(Q_T, m)$ must be weakly decreasing, as long as it is above Q^*.

Let us prove this result by contradiction. For $t > T$, suppose that the path (Q) solution to $W(Q_T, m)$ is somewhere strictly increasing while above Q^*. Because it ultimately converges to $\min(Q^*, \overline{Q}_\infty) = Q^*$, there must exist $t_1 < t_2$ such that

$$Q^* \leq Q_{t_1} = Q_{t_2} \leq Q_t \leq \overline{Q}_\infty$$

for all $t \in [t_1, t_2]$. Note that all paths verifying this property (and in particular the constant path $(q = 0, Q = Q_{t_1})$) share the same survival probabilities P on $[t_1, t_2]$, and the same values of Q and $\overline{Q}$ at $t = t_2$. Therefore, one can aim at maximizing

$$W = \int_{t_1}^{t_2} [Pu - \dot{P}V] e dt$$

over the set of such paths. We then check (A.9) in Lemma A.4. Because $Q_t > Q_{t_1} > Q^*$, we have $\nu(Q_t) \leq 0$ and $N(Q_t) \leq 0$. Because pe is decreasing with time, the right-hand-side in (A.9) is at least zero. Moreover, because D is increasing we have $D(Q_t) \geq D(Q_{t_1})$, and because p is decreasing we obtain that the left-hand-side is below zero. This contradicts (A.9), a contradiction. So we can focus on paths which are weakly decreasing with time after time T, as long as the stock is above Q^*.

Now, if $T = 0$, then the stock is weakly decreasing above Q^*, and this shows (ii). If $T > 0$, then there is an interval $[T', T]$ on which $q_t \geq 0$ and $Q_t = \overline{Q}_t > Q^*$. By assumption, the planner stops experimenting at $t = T$. Suppose he stops experimenting at $t_0 \in [T', T]$ instead, by choosing a path below Q_{t_0} after time t_0. From the above proof, the best path is weakly decreasing, as long as it remains above Q^*. Therefore, the payoff from stopping at t_0 is

$$\int_{T'}^{t_0} [pu - \dot{p}V] e dt + \exp(-\delta t_0) p_{t_0} W(Q_{t_0}, (1 - F(Q_{t_0}))/p_{t_0})$$

and the last result in Lemma A.3 states that this function is decreasing in t_0. This yields a contradiction, as the planner should stop experimenting earlier. This concludes the proof. ■

The following intermediate result is used below in the proof of Lemma 7:

Lemma A.5 *Suppose* $\max(Q^{**}, \overline{Q}_0) < \overline{Q}_\infty < \min(Q^*, \overline{S})$. *Then* $\overline{Q}_\infty$ *is reached in finite time, and* $\overline{Q}_\infty \leq \hat{Q}$.

Proof of Lemma A.5: the assumptions in the Lemma imply that the path converges to $\overline{Q}_\infty$, necessarily from below. Let us proceed by contradiction, and suppose that $\overline{Q}_\infty$ is reached only asymptotically, so that for all t one has $Q_t \leq \overline{Q}_t < \overline{Q}_\infty$. We know that Q_t converges to $\overline{Q}_\infty$, and that $\dot{p}_t$ converges to zero. Therefore, for every $\varepsilon > 0$ there exists T such that

$$\forall\ t \geq T, \qquad |\overline{Q}_\infty - Q_t| < \varepsilon \qquad \text{and} \qquad |\dot{p}| < \varepsilon.$$

Moreover, there must exist[15] $t_1 \geq T$ such that $Q_1 = \overline{Q}_1$ (notice that to alleviate notations we write 1 instead of t_1.) From now on we focus on the time interval $[t_1, +\infty[$, with the aim of checking (A.10) in Lemma A.4. Referring to this result, we define two functions: first,

$$N(Q) = \int_{Q_1}^{Q} \nu(Q')dQ',$$

and second, function $A(Q, \overline{Q}, \dot{p})$:

$$\dot{p}\left[D(Q) - D(Q_1) - \frac{\alpha+\delta}{\alpha}N(Q)\right] + \delta\left[(1 - F(\overline{Q}))N(Q) - \frac{\alpha}{\alpha+\delta}(F(\overline{Q}) - F(\overline{Q}_1))D(Q_1)\right],$$

and we want to show that for $t \geq t_1$ we have $A(Q_t, \overline{Q}_t, \dot{p}_t) < 0$. Result (ii) in Lemma A.4 would then imply that the planner would be better off by adopting a constant path after time t_1, in contradiction with our assumptions.

Let us begin by the first bracketed term in the definition of A. Let

$$Z(Q) \equiv D(Q) - D(Q_1) - \frac{\alpha+\delta}{\alpha}N(Q).$$

Z is a convex function, and $Z(Q_1) = 0$. Therefore:

$$Z(Q) \geq (Q - Q_1)Z'(Q_1). \tag{A.12}$$

Moreover, because Z' is monotonic, and $\max(Q^{**}, \overline{Q}_0) < Q_1 < \overline{Q}_\infty$, we have

$$|Z'(Q_1)| \leq k_0 \equiv \max(|Z'(\max(Q^{**}, \overline{Q}_0))|, |Z'(\overline{Q}_\infty|).$$

Finally, we know that $|\dot{p}|$ is less than ε. This shows that the absolute value of the first term in A is below

$$k_0\varepsilon|Q - Q_1|. \tag{A.13}$$

We now turn to the second bracketed term in A. It can be written as

$$(1 - F(\overline{Q}))(N(Q) - N(\overline{Q})) + Y(\overline{Q}),$$

[15] Suppose otherwise. Then for every $t \geq T$ one has $Q_t < \overline{Q}_t$. But then $\overline{Q}_t$ must be a constant, in contradiction with the fact that $\overline{Q}_t$ converges to $\overline{Q}_\infty$, and reaches this value only asymptotically.

where

$$Y(\overline{Q}) = (1 - F(\overline{Q}))N(\overline{Q}) - \frac{\alpha}{\alpha+\delta}(F(\overline{Q}) - F(\overline{Q}_1))D(Q_1).$$

Notice first that

$$N(Q) - N(\overline{Q}) = \int_Q^{\overline{Q}} (-\nu(Q'))dQ' \leq -(\overline{Q} - Q)\nu(\overline{Q}_\infty) \tag{A.14}$$

because ν is decreasing, and $Q \leq Q' \leq \overline{Q} < \overline{Q}_\infty$. Turning to Y, for the same reason we have

$$N(\overline{Q}) \leq (\overline{Q} - Q_1)\nu(Q_1),$$

and because $1 - F$ is decreasing, and ρ is increasing, and $\overline{Q}_1 \leq \overline{Q}$,

$$F(\overline{Q}) - F(\overline{Q}_1) = \int_{\overline{Q}_1}^{\overline{Q}} f(Q')dQ' = \int_{\overline{Q}_1}^{\overline{Q}} \rho(Q')(1 - F(Q'))dQ'$$

$$\geq (1 - F(\overline{Q})) \int_{\overline{Q}_1}^{\overline{Q}} \rho(Q')dQ' \geq (1 - F(\overline{Q}))\rho(\overline{Q}_1)(\overline{Q} - \overline{Q}_1).$$

Because $\overline{Q}_1 = Q_1$, we obtain:

$$Y(\overline{Q}) \leq (1 - F(\overline{Q}))(\overline{Q} - Q_1)[\nu(Q_1) - \frac{\alpha}{\alpha+\delta}\rho(Q_1)D(Q_1)].$$

Let $M(Q_1)$ denote the last bracketed term. Because $Q_1 > Q^{**}$, it is strictly negative. Therefore:

$$Y(\overline{Q}) \leq (1 - F(\overline{Q}_\infty))(\overline{Q} - Q_1)M(Q_1).$$

Moreover, M is continuous. Hence, by choosing ε small enough $M(Q_1)$ can be made close to $M(\overline{Q}_\infty)$, so that

$$M(Q_1) \leq \frac{1}{2}M(\overline{Q}_\infty) < 0.$$

So we define

$$k_1 \equiv -\frac{1}{2}(1 - F(\overline{Q}_\infty))M(\overline{Q}_\infty) > 0$$

and overall, using also (A.13) and (A.14), we have shown:

$$A(Q, \overline{Q}, \dot{p}) \leq k_0\varepsilon|Q - Q_1| - \delta(1 - F(\overline{Q}))(\overline{Q} - Q)\nu(\overline{Q}_\infty) - \delta k_1(\overline{Q} - Q_1).$$

Finally, because

$$|Q - Q_1| \leq |Q - \overline{Q}| + |\overline{Q} - Q_1| = \overline{Q} - Q + \overline{Q} - Q_1,$$

we obtain

$$A(Q, \overline{Q}, \dot{p}) \leq -(\overline{Q} - Q)[\delta(1 - F(\overline{Q}))\nu(\overline{Q}_\infty) - k_0\varepsilon] - (\overline{Q} - Q_1)[\delta k_1 - k_0\varepsilon].$$

For ε small enough, the first bracketed term is strictly positive because $\overline{Q}_t < \overline{S}$ and $\overline{Q}_\infty < Q^*$, and therefore $(1 - F(\overline{Q}))\nu(\overline{Q}_\infty) > 0$; and so is the second term, because $k_1 > 0$. Comparing to (A.10) provides the contradiction we need.

This shows that $\overline{Q}_\infty$ is reached in finite time, say at time T. Moreover, $T > 0$, because $\overline{Q}_\infty > \overline{Q}_0$ by assumption. Therefore, there exists $T' < T$ such that Q is increasing on $[T', T]$, and therefore $Q = \overline{Q}$ on this interval. There remains to show the inequality $\overline{Q}_\infty \leq \hat{Q}$. To do so, we proceed by contradiction, and we assume $\overline{Q}_\infty > \hat{Q}$. This implies that we can choose $t_1 \in [T', T]$ so that $\hat{Q} < Q_1 < \overline{Q}_\infty$. For such a t_1, we can define N and A as above. Referring to the lines above, $Z(Q)$ is at least zero: indeed, $Q \geq Q_1$, $Z(Q_1) = 0$, and Z is increasing above Q_1 because $Q_1 > \hat{Q}$. Because $\dot{p} \leq 0$, the first term in A is therefore below zero. Finally, one shows that the second term in A is also below zero exactly as in the lines above. Overall, this shows once more that $A(Q_t, \overline{Q}_t, \dot{p}_t) \leq 0$ for all $t \geq t_1$, and once again we apply result (ii) in Lemma A.4 to get a contradiction. This shows $\overline{Q}_\infty \leq \hat{Q}$, as announced. ■

Proof of Lemma 7: from Lemma 3, we can focus on paths that are below $\overline{S}$. From Lemma 4, we can focus on paths that converge to their highest value $\overline{Q}_\infty$. Suppose $\overline{Q}_\infty > Q_0$. Then we can apply Lemma A.5, and we get a contradiction, as $\overline{Q}_\infty$ is supposed to be below $\hat{Q} = -\infty$. Therefore, we can focus on paths that do not exceed Q_0, and converge to Q_0. Therefore, the planner never experiments after time 0. Then p can be explicitly computed, and we end up maximizing

$$\int [m_0 u \exp(-\delta t) + (1 - m_0)(u + \alpha V) \exp(-(\alpha + \delta)t)]dt$$

where $m_0 \equiv (1 - F(Q_0))/p_0 > 0$, under the constraints $\dot{Q} = q$ and $Q \leq Q_0$. The objective function is

$$m_0 q \exp(-\delta t) + (1 - m_0)(q - \alpha(d_0 + d_1 Q)) \exp(-(\alpha + \delta)t).$$

Now we have, for any parameter $k > 0$:

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

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

$$\int q \left[m_0 + (1 - m_0) \exp(-\alpha t)(1 - \frac{\alpha}{\alpha + \delta} d_1) \right] \exp(-\delta t) dt.$$

Because $1 < \frac{\alpha}{\alpha+\delta} d_1$, the bracketed coefficient is increasing in t, and positive at $t = +\infty$. Because $1 - m_0 = \pi_0$, the same coefficient is negative at $t = 0$. The planner should thus behave as indicated in the claim. ■