

September 2023

“Catastrophes, delays, and learning”

Matti Liski and François Salanié

Catastrophes, delays, and learning^{*}

Matti Liski[†] François Salanié[‡]

September 15, 2023

Abstract

We propose a simple and general model of experimentation in which reaching untried levels of a stock variable may, after a stochastic delay, lead to a catastrophe. Hence, at any point in time a catastrophe might well be under way, due to past experiments. We show how to measure this legacy of the past from prior beliefs and the chronicle of stock levels. We characterize the optimal policy as a function of the legacy and show that it leads to a new protocol for planning that applies to a general class of problems, encompassing the study of pandemics or climate change. Several original policy predictions follow, e.g., experimentation can stop but resume later.

JEL Classification: C61, D81, Q54.

Keywords: catastrophes, experimentation, learning, delays

^{*}We have benefitted from feedbacks in seminars in Ascona, Copenhagen, Gothenburg, Helsinki, Madrid, Montpellier, Munich, Rome, Rennes, Tilburg, Toulouse, and Tel Aviv. We also thank Anne-Sophie Crépin, Jan Knoepfle, Topi Hokkanen, Rick van der Ploeg, Bernard Salanié, Arthur Snow, Nicolas Treich, and Amos Zemel for comments on this work.

[†]Aalto University, Department of Economics, and Helsinki GSE. e-mail: matti.liski@aalto.fi. Matti Liski acknowledges funding from the Finnish Cultural Foundation.

[‡]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). This research has also benefitted from the financial support of the research foundation TSE-Partnership.

How should society manage dynamic systems that may suddenly collapse? 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 opposite policy conclusions. In this paper, we argue that the key question is how these approaches deal with the possibility that a catastrophe may already be under way.

Consider 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 '70s, 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 accumulation of greenhouse gases in the atmosphere) and the effect (melting) (e.g., Fitzpatrick and Kelly, 2017). Similar thresholds and delays are not unheard of in other situations. Is a virus outbreak on its way to cause a breakdown of the health system? Will habitat fragmentation lead to a collapse of biodiversity, or is it already too late?

When facing such threats, one may take it as advisable to act on the assumption that the catastrophe is on its way to be appropriately prepared for its occurrence. On reflection, however, one may consider it equally advisable to assume the opposite to focus on actions that avoid triggering the catastrophe in the first place. Both premises produce valuable insights, as the literature has shown, but we are left with a logical dilemma: the assumptions are mutually exclusive and the choice between them dictates to a large degree the nature of policy recommendations. Our formal framework is designed to address this dilemma, allowing us to develop a new protocol for planning under the threat of a catastrophe.

We develop a general model of experimentation in which a planner manages 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 stock variable with multiple interpretations (e.g., temperature, finite resource, infected population). The stock *triggers* a catastrophe when it exceeds an unknown threshold. Once triggered, the catastrophe itself *occurs* only after a stochastic delay. The key assumption is that the planner does not know whether a catastrophe has been triggered or not: only the

occurrence of a catastrophe is observable. Reaching a previously untried level is thus an experiment whose results may be learned only later on.

The delay between the triggering of the event and its occurrence leads to an information structure in which the planner evaluates potential threats pending from the past. Formally, for any date we define the legacy of the past as the probability that past experiments, whether planned or simply inherited, have triggered the catastrophe. As time goes by without any catastrophe occurring, we are more confident that nothing will follow from the past experimentations and the legacy goes down — unless we keep on experimenting, thereby causing an increase in future values of the legacy. Likewise, when evaluating the present-day legacy, it matters *how and when* we experimented in the past. For instance, a rapid increase in greenhouse gases in the recent past creates a legacy higher than if the same increase took place in a distant past.

Two thought-experiments prove useful. First, if the planner could learn without any delay the results from experiments, there would be no legacy. In this situation, what would be the long-run level of the stock, say Q^E , at which the planner optimally stops experimenting? Second, conversely, if one knew for sure that the catastrophe is pending, what would be the stock level, say Q^D , that one aims at before the catastrophe occurs? It turns out that the ordering of the two stock levels partitions possible planning situations into two very different classes.

To fix ideas, consider the management of a pandemic for which the classical trade-off is between economic activity, typically associated to young people, and mortality (or morbidity) risk, typically borne by older people. In addition, there is the risk that too many cases might lead to a collapse of the health system. Hence, in our model the stock is the number of infected people which, by reaching an unknown threshold, may trigger a catastrophe. The planner thus manages simultaneously this catastrophe risk and the classical trade-off. Our protocol recommends in a first step to evaluate and rank the values of Q^E and Q^D .

Our first theorem holds when $Q^E < Q^D$, a situation which follows when the planner puts a high weight on economic activity in comparison with the social costs of deaths. Then, all optimal policies allow infection levels to grow over time, as illustrated by path I in Figure 1. Moreover, a higher legacy (e.g., because there was a recent and fast increase in the number of cases before time t_0) leads to more experimentation and a higher total number of cases that the planner optimally tolerates: the idea is that since the occurrence

of the catastrophe is likely, it is better to reap the gains from economic activity while they still exist. Hence, a higher legacy of the past makes the planner *less* cautious. Our first theorem rationalizes such fatalism from a set of well-founded primitives.




Figure 1

The second theorem applies when $Q^E > Q^D$, a situation holding when the planner puts a high value to life vis-à-vis economic activity. A possible optimal policy is described by path *II* in Fig. 1. When facing the same legacy as the planner of the first theorem, it holds under intuitive conditions that the policy imposes a strict lockdown early to reduce the number of cases and thus the impact of the potential catastrophe on the health system. During the lockdown, the legacy of the past goes down as no new experiments take place, and therefore the planner becomes more and more optimistic over time. We prove that the first phase, in which the lockdown reduces the number of cases, is optimally followed by a second phase, in which the planner accommodates increasing levels of infections up to the level at which the lockdown started in the first place and possibly even above. Hence, this optimal policy is non-monotonic. It implies that, without changing the planner's preferences, a lockdown or increasing infections can be optimal depending on how the current level of infections was reached. Finally, and also in contrast to the first theorem, now a higher legacy of the past makes the planner more cautious in a sense that the lockdown becomes more stringent; however, the optimal asymptotic stock level does not depend on the legacy.

The disease control problem nicely illustrates the key stock-flow tradeoffs and contributes to the literature on virus outbreaks by adding a new learning-based rationale

for non-monotonic policies.¹ But these insights hold quite generally. Intuitively, when data indicate that the catastrophe is bound to happen and if, in addition, gains to mitigation are small, there is little reason to restrain actions that produce benefits prior to the occurrence. In the opposite case, gains to mitigation are high in the short run, but the concern regarding catastrophes pending from the past dwindles in the long run if no event occurs. This change in priority implies a non-monotonic trajectory for the stock.

In addition to pandemics, we illustrate the broad applicability of the results by “eating a cake of unknown size” (building on Kemp, 1976) which is a problem that captures the rudiments of many resource use settings and, finally, by a stylized climate-change application. Climate-change targets are commonly expressed as “budgets” for total CO₂ emissions; however, the budget deemed safe is highly uncertain (van der Ploeg, 2018; IPCC, 2021). We model this unknown budget as a threshold for cumulative emissions that should not be exceeded, lest a catastrophe be triggered. Several implications for the policies on CO₂ budgets follow. If the damage from a catastrophe cannot be much altered by reducing the present emissions, then the first theorem applies: the policies are monotonic and they imply a higher total budget when the legacy is higher (e.g., because the current emissions stock was reached quickly rather than slowly). Otherwise, the second theorem applies and then, for a legacy high enough, the policies aim at sharp emissions reductions early so that the budget is first not touched at all; only later can one move on to consume it. None of these implications follows from the existing literature.

Literature. Our model allows to tie together two canonical but distinct approaches to modeling catastrophes in the literature.² In the first approach, 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, while the catastrophe is bound to happen, action can be taken to delay its occurrence and severity. This assumption features in many recent applied papers (e.g., van der Ploeg and de Zeeuw, 2017), including quantitative assessments of the optimal climate-change policies (e.g., Besley and Dixit, 2019).

¹Assenza et al. (2020) provides a literature review on the so-called “hammer-and-dance” policies.

²Catastrophes, broadly interpreted, appear in a wide range of economic applications, including macroeconomic 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, de Zeeuw and Wagener, 2011; Sakamoto, 2014; van der Ploeg and de Zeeuw, 2017; Bretschger and Vinogradova, 2019; Cai and Lontzek, 2019). See Rheinberger and Treich (2017) for a bibliometric analysis of the literature on catastrophes.

In the second approach, the catastrophe occurs as soon as the critical variable exceeds a threshold whose exact value is unknown. The formal approach appears in Kemp (1976), who studied the problem of eating a cake of unknown size. In Rob (1991) the threshold is a kink in the demand curve. Tsur and Zemel (1994) focus on natural catastrophes (see also Tsur and Zemel, 1995 and 1996).³ In Chen (2020) firms face a common threshold that they search individually with a private cost from passing the threshold; in Diekert (2017) the cost of passing the threshold is common. The quantitative policy evaluation in Lemoine and Traeger (2014) uses the threshold approach. Learning occurs instantaneously in this literature: 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, after each step, through a simple truncation of the prior for the threshold. 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.⁴

Researchers in both camps end up working with a hazard rate for the event, one assumed exogenously and another derived from the threshold distribution. This choice may seem innocuous, but in fact its informational consequences could not be bigger: in one approach the catastrophe is pending for sure, while in the other one it is so far avoided with certainty. By introducing a delay, we explore a more general model where the planner remains uncertain if the current standing is safe, even if he stops experimenting. The approaches in the literature follow as special cases if the delay goes to zero or if past actions are known to have triggered the event. Neither of these canonical approaches is suitable for interpreting the information content of past experiments (planned or inherited) and thus they miss the mechanism that is key to our results.

Introducing delays implies that any given welfare impact from triggering a catastrophe is delayed, which encourages to experiment more in total. This effect is somewhat trivial. But delays also give rise to the legacy which has an ambiguous impact on experiments. Under the first theorem, the legacy encourages to experiment more, and conversely under the second theorem. This opposition links closely the literature in a precise sense: the

³We discuss their contributions in detail in Section 7.

⁴Crépin and Nævdal (2019) extend the threshold approach. 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 in the path of this second state variable but learning is still instantaneous.

extreme informational assumptions of the literature define two stock-level targets whose comparison tells which one of the theorem applies. The theorems can be applied to a broad class of existing models, which we illustrate for a class of climate-change models.

Our approach is different from the bandit models used to study experimentation in various economic settings. As in Poisson bandit settings, the planner updates beliefs on the arrival rate of a catastrophe by not observing the event (as in Malueg and Tsutsui, 1997; see also Keller, Rady and Cripps, 2005; and Bonatti and Hörner, 2011). In a sense, our planner runs an endogenous continuum of such bandits (thresholds tried), and 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; even under a simplifying Poisson assumption for the distribution of the stochastic delay, this aggregation encapsulates not only the value of the highest stock on record but also the chronicle of past experiments.

A few recent papers on experiments are related to our work. Gerlagh and Liski (2018) consider an explicit climate-economy model with learning about potentially catastrophic damages. The objective of that paper is to study the impact of speed of learning on the optimal policy path when the legacy is strictly between zero and one (using the current terminology). In this sense, the paper is between the two canonical approaches to catastrophes in the literature. However, that model does not have a structure that connects the legacy to past experiments.

Laiho, Murto and Salmi (2023) shares with our paper the feature that the chronicle of past actions determines the speed of information arrival. In their model, stochastic flow gains are made possible by irreversible capacity expansions, but there is a risk of overcapacity if the profitability, given by an unknown state, turns out to be bad. In our model, the pay-off relevant stock level is reversible. Also, in our model, the chronicle of past actions is essential for beliefs; in their setting, the precise timing of past experiments does not matter.

Guillouët and Martimort (2020) 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 different however, as they do not allow the planner to condition the policy on his beliefs. Hence, a time-consistency problem arises, and they characterize policies resulting from a Nash equilibrium between different selves. Our question is very different as we focus on the optimal policy only,

and we develop a framework in which beliefs from a chronicle of past actions, potentially inherited, shape the policy. We also allow for non-monotonic policies and general payoff functions, and thereby cover a broad set of applications.

1 Model

A planner controls a dynamic system over time. Section 1.1 defines the planner's problem in the absence of catastrophes. Section 1.2 adds catastrophes and delays. Section 1.3 introduces uncertainty. With the model components at hand, Section 1.4 formulates the complete planning problem.

1.1 The Stock-Flow Problem (SFP)

Time t is a continuous variable in $(-\infty, +\infty)$. At each date $t \geq 0$, the planner chooses a flow action q_t to control a stock Q_t according to a simple law of motion:

$$\dot{Q}_t = q_t \in [\underline{q}, \bar{q}], \quad Q_0 \text{ given.} \quad (1)$$

We assume $\underline{q} < 0 < \bar{q}$, so that the stock may increase or decrease over time. The planner's objective function at date zero is the following discounted sum of payoffs:

$$\int_0^{+\infty} u(q_t, Q_t) \exp(-\delta t) dt. \quad (2)$$

We allow the instantaneous payoff $u(q, Q)$ to depend on both the stock level Q and the flow value q . Parameter $\delta > 0$ is the discount rate. The Stock-Flow Problem (SFP) is a classical calculus-of-variations exercise in which one maximizes (2) under the law of motion (1). To make this problem regular, 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*

$$\nu(Q) \equiv u_q(0, Q) + \frac{1}{\delta} u_Q(0, Q)$$

is weakly decreasing with respect to Q .

Function ν encapsulates the trade-off between instantaneous gains from an increase in the flow, and the long-run effects from the associated increase of the stock. This function

is decreasing for all values of the discount factor if both $u_{qQ}(0, Q)$ and $u_{QQ}(0, Q)$ are weakly negative, but the assumption highlights the exact property we will rely upon: increasing the stock is less valuable when the stock is higher.

The solutions to the problem (1)-(2) are easily shown to converge monotonically toward the following long-run target:

Definition 1 Q^N (where N stands for “No catastrophe”) is the stock level at which $\nu(Q)$ is zero. We assume Q^N is unique. By convention, we set $Q^N = +\infty$ if ν is positive for all Q , and $Q^N = -\infty$ if ν is negative for all Q .

Indeed, because ν is decreasing from Assumption 1, if $Q < Q^N$ then the payoff ν from increasing the stock is positive, and conversely if $Q > Q^N$. Thus, Q^N is interpreted as the long-run target in the absence of catastrophes. We obtain:

Proposition 1 *There exists a solution to the Stock-Flow Problem (1)-(2). This solution is monotonic and converges to Q^N .*

1.2 Catastrophes and delays

Catastrophes are irreversible and costly events. An original feature of our model is that we distinguish when a catastrophe is triggered from when it actually occurs. We say that a catastrophe is *triggered* when the stock Q exceeds a threshold value S . Given a path $(Q_t)_{t \in (-\infty, +\infty)}$, the triggering time is a function of S :

$$T(S) \equiv \inf\{t : Q_t > S\}. \quad (3)$$

Note that $T(S)$ is infinite if the stock never exceeds S and that $Q_{T(S)} = S$ otherwise. We also define the highest stock on record at time t :

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

so that $T(S) < t$ if and only if $S < \bar{Q}_t$. The catastrophe itself *occurs* only after a delay $\tau \geq 0$, at date $\kappa = T(S) + \tau$. Note that, in contrast to the SFP, now the full past trajectory of the stock is relevant at time 0, as the catastrophe may have been triggered in the past without occurring yet.

Before the catastrophe occurs, the planner’s instantaneous utility is $u(q_t, Q_t)$ at each date t . At time κ , the catastrophe occurs, the game ends, and the planner receives a

continuation payoff $V(Q_\kappa)$ which depends on the value of the stock at the catastrophe date κ .⁵ Hence, 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 continuation payoff V to depend instead on the threshold S , or on the maximum level tried in the past \bar{Q}_κ , would eliminate this possibility by assumption.

We assume that the planner prefers stabilizing the stock forever over getting immediately the continuation payoff from the catastrophe, and that this preference is stronger when the stock is higher. Formally, define the damage $D(Q)$ as the difference between the discounted value of the no-catastrophe utility flow from stabilizing at Q and $V(Q)$:

$$D(Q) \equiv \frac{u(0, Q)}{\delta} - V(Q). \quad (4)$$

The following assumption is stated using this damage function D :⁶

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

Therefore, the planner can mitigate damages by reducing the stock value at the occurrence date. The last part of the assumption ensures that higher stock levels reduce the value of increasing the flow, taking into account the marginal damages associated to a catastrophe. This part of the assumption is weaker than the convexity of damages, but it is strong enough to imply that $\nu - kD'$ is decreasing for any $k \in [0, 1]$, and thus encompasses situations in which the catastrophe is discounted, or occurs with a probability below one.⁷

Overall, given S , τ , and a path $(Q_t)_{t \in (-\infty, +\infty)}$, one applies (3) to compute $T(S)$ and $\kappa = T(S) + \tau$, so that the planner's payoff from date $t = 0$ onward equals

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

⁵Applications to disease control and climate change will provide micro-foundations for V as the value function of a post-catastrophe problem.

⁶Equivalently, the assumption states that: V is twice differentiable; V is less than the value $\frac{u(0, Q)}{\delta}$ of stabilizing the stock forever, and this difference increases with the stock; and finally, $u_q(0, Q) + \dot{V}'(Q)$ decreases with the stock, a property which holds if $u_{qQ} \leq 0$ and V is concave.

⁷Indeed, for $k \in [0, 1]$ we have $\nu - kD' = (1 - k)\nu + k(\nu - D')$, and both terms are decreasing by Assumptions 1 and 2. The early literature on catastrophes often relies on much more complicated assumptions. For example, Tsur and Zemel (1994) relies on two assumptions U1 and U2, which involve the solution to a constrained dynamic program. We later provide a general proof of their results relying on our simpler assumptions.

1.3 Uncertainty

We now introduce uncertainty over both the threshold S and the delay τ . The planner has prior beliefs on S , characterized by a cumulative distribution function F on the interval $[\underline{S}, \bar{S}]$. We underline that these are beliefs held at the beginning of times, $t = -\infty$. We assume throughout that F is continuously differentiable on its support, with density f . We adopt a monotone hazard rate assumption, which makes the triggering of a catastrophe more likely conditional on reaching a higher stock level:

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

The delay τ is also unknown to the planner. We assume it follows an exponential distribution with parameter $\alpha > 0$, with the cumulative distribution function $1 - \exp(-\alpha\tau)$. Thus, τ and S are independent.

A key assumption is that the planner does not observe whether a catastrophe was triggered: he only observes its occurrence. This allows us to capture the idea that a catastrophe might well be under way: the planner is unsure if the threshold was exceeded in the past; he is unsure about the delay between triggering and occurrence. These are uncertainties often invoked in biology, under the name of extinction debt (Tilman et al., 1994).

To illustrate, one may imagine a skater on thin ice. Instantaneous utility flow increases with the distance to the shore at a decreasing rate (Assumption 1), but the ice gets thinner and thinner (Assumption 3). The skater does not observe whether the first crack in the ice has appeared, but he may turn back as long as the ice is still holding. When the ice finally breaks, the journey finds an abrupt end, and the damage to the skater depends on the remaining distance to the shore (Assumption 2).

We may use the skater illustration to introduce two benchmarks important for later results. First, assume the skater knows that the threshold has been exceeded and the ice will ultimately break; only the timing of the event is uncertain. Where should the skater stop? Suppose the skater has stopped and considers increasing the distance from

the shore marginally (i.e., variable Q), giving the benefit measured by⁸

$$\nu(Q) - \frac{\alpha}{\alpha + \delta} D'(Q). \quad (5)$$

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

Definition 2 Q^D (where D stands for “Damages”) is the stock level at which (5) is zero, and for simplicity we assume it is unique. By convention, we set $Q^D = +\infty$ if (5) is positive for all Q , and $Q^D = -\infty$ if (5) is negative for all Q .

Second, assume that the skater is still at the shore, or is otherwise sure that no threshold has been exceeded when standing at Q_0 . In such a situation, one may safely stabilize the situation by playing $q = 0$ forever. One may also experiment a bit more before stabilizing. To compare these policy options, one computes the instantaneous utility gain from experimenting and subtracts the expected discounted damage of triggering a catastrophe to obtain the net gain from a marginal experiment:

$$\nu(Q) - \frac{\alpha}{\alpha + \delta} \rho(Q) D(Q). \quad (6)$$

Noticeably, the second term involves the level of damage D , not its derivative; the hazard rate ρ measures the probability of triggering a catastrophe at the current level Q . Under our assumptions, expression (6) is weakly decreasing in Q and lies below $\nu(Q)$. Therefore, it may reach zero only at a value $Q^E \leq Q^N$:

Definition 3 Q^E (where E stands for “Experimentation”) is the stock level at which (6) is zero, and for simplicity we assume it is unique. By convention, we set $Q^E = \underline{S}$ if (6) is negative at \underline{S} , and $Q^E = \bar{S}$ if (6) is positive at \bar{S} .

As we will see in the next section, the ranking of targets Q^D and Q^E is key to our main theorems. The symmetry in equations (5)-(6) and our monotonicity assumptions make it easy to find sufficient conditions. For example, we have:

⁸The marginal damage is discounted by a coefficient that takes into account the stochastic delay before occurrence. Because this delay τ is distributed exponentially with parameter α , this coefficient is

$$E \exp(-\delta\tau) = \alpha \int_{\tau \geq 0} \exp(-(\alpha + \delta)\tau) d\tau = \frac{\alpha}{\alpha + \delta}.$$

Lemma 1 *If the function $D(Q)(1 - F(Q))$ increases (resp. decreases) at $Q = Q^D$, then $Q^D < Q^E$ (resp. $>$).*

The condition is simple, and its implication is clear: when the damage is relatively sensitive to Q , the planner who knows that the event has been triggered mitigates the expected losses from occurrence by aiming at low value Q^D for the stock. By contrast, the condition $Q^D > Q^E$ holds when the damage is a constant loss.

1.4 The planner's problem

At the planning date $t = 0$, the planner inherits historical data which consist of the past trajectory of the stock $(Q_t)_{t \leq 0}$. The plan is a contingency plan for survival: it conditions on the event “no catastrophe occurred before time zero”, or equivalently $\kappa = T(S) + \tau \geq 0$. Beliefs at time zero are obtained by this conditioning, i.e., they take into account the past experiments and the possibility that they might have triggered a catastrophe that did not occur yet. Therefore, the planner's problem is as follows:

$$\max_{(q_t)_{t \geq 0}} \mathbb{E} \left[\int_0^{\kappa} u(q_t, Q_t) \exp(-\delta t) dt + \exp(-\delta \kappa) V(Q_\kappa) \mid \kappa \geq 0, (Q_t)_{t \leq 0} \right] \quad (7)$$

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

While Q_t is continuous by construction, we only require q_t to be piecewise-continuous. We say that a policy $(Q_t)_{t \geq 0}$ is monotonic if Q_t is everywhere weakly decreasing, or everywhere weakly increasing, with respect to time. Given a 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 $Q_T = \bar{Q}_\infty$. Otherwise, we say that \bar{Q}_∞ is reached asymptotically, and in this case one has $Q_t \leq \bar{Q}_t < \bar{Q}_\infty$ for all t .

2 Learning and the survival probability

The planner learns from past experiments by observing that a catastrophe did not yet occur: in this sense, no news is good news. Prior beliefs are thus revised over time by conditioning on survival. We now show how these beliefs can be summarized in a survival probability with simple dynamics. Given a path $(Q_t)_{t \in (-\infty, +\infty)}$, let us define the survival probability at time t as the decumulative density function of the catastrophe date κ , computed at the beginning of time using the prior beliefs F :

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

To characterize this probability, one may distinguish two possibilities for survival at time t . Either S is above \bar{Q}_t , so that no catastrophe could occur before time t , and survival is certain. Or S is below \bar{Q}_t , and in this case a catastrophe was triggered at time $T(S) < t$, but did not occur yet because the delay τ is above $t - T(S)$, an event that happens with probability $\exp[-\alpha(t - T(S))]$. Overall, we obtain

$$p_t = 1 - F(\bar{Q}_t) + \int_{S < \bar{Q}_t} \exp[-\alpha(t - T(S))] dF(S). \quad (9)$$

Hence, the survival probability at time t exceeds $1 - F(\bar{Q}_t)$, as a catastrophe may have been triggered in the past but did not occur yet. Define the *legacy of the past* π_t as the probability at time t that the event was triggered in the past, conditional on survival:

Definition 4 *For a given path, the legacy of the past at date t is*

$$\pi_t \equiv \frac{\int_{S < \bar{Q}_t} \exp[-\alpha(t - T(S))] dF(S)}{p_t} \in [0, F(\bar{Q}_t)].$$

Notice that π can also be computed directly from \bar{Q} and p , as follows:

$$\pi_t = 1 - \frac{1 - F(\bar{Q}_t)}{p_t}.$$

Let us underline that the legacy is a direct consequence of the delay between triggering and occurrence: in the limiting case without delay (α goes to infinity), p_t equals $1 - F(\bar{Q}_t)$, and then π_t is identically zero. When delays are introduced, as soon as some experimentation took place in the past, π_t is not zero anymore: it is a sum of terms which vanish over time, each term being associated to a possible value for the threshold $S < \bar{Q}_t$. Therefore, a past experiment contributes less to π_t if it took place a long time ago rather than just before t .

The dynamics of the survival probability can now be simplified, thanks to assuming an exponential law for the delay. Indeed, by applying (9) at $t = 0$, we get the information content of the data $(Q_t)_{t \leq 0}$ relevant for planning:

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

Moreover, by differentiating (9), we obtain a law of motion:

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

These two equalities in turn imply (9), at all future dates $t \geq 0$.

We are now in a position to rewrite the planner's problem defined in (7)-(8). Recall that the survival probability is the decumulative density function p_κ of the catastrophe date κ . It is the premise of planning that no catastrophe has happened at time $t = 0$, and therefore the planner should condition on the event $\kappa \geq 0$ by using the ratio p_κ/p_0 . His expected payoff can be written, leaving out the constant factor $1/p_0$,

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

By relabelling κ into t in the second integral, we obtain the following problem:

$$\max \int_0^\infty [p_t u(q_t, Q_t) - \dot{p}_t V(Q_t)] \exp(-\delta t) dt, \quad (10)$$

$$\dot{Q}_t = q_t \in [\underline{q}, \bar{q}], \quad Q_0 \text{ given}, \quad (11)$$

$$\bar{Q}_t = \max(\max_{0 \leq t' \leq t} Q_{t'}, \bar{Q}_0), \quad \bar{Q}_0 \geq Q_0 \text{ given}, \quad (12)$$

$$\dot{p}_t = \alpha(1 - F(\bar{Q}_t) - p_t), \quad p_0 > 0 \text{ given}, \quad p_0 \geq 1 - F(\bar{Q}_0). \quad (13)$$

As time enters only through exponential discounting, this problem is autonomous. The three state variables are the stock Q , the maximum stock on record \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 $(Q_t)_{t \leq 0}$, thanks to the assumption that the delay τ is exponential. In this triplet, one may equivalently replace p_0 by the legacy of the past, $\pi_0 = 1 - \frac{1 - F(\bar{Q}_0)}{p_0}$, which measures the probability that a catastrophe was triggered in the past, conditional on survival. Intuitively, \bar{Q}_0 measures total experimentation so far, while π_0 varies according to the past timing of these experiments; the data on the past enter through these variables.

The optimal policy gives a path for the main variables that can be interrupted at any time by the occurrence of a catastrophe. An illustrative and simple situation is the one in which the legacy is one. Then, the planner knows that the catastrophe was triggered, but still does not know when the catastrophe will occur. From (13), the survival probability after time 0 is

$$p_t = p_0 \exp(-\alpha t). \quad (14)$$

Plugging this expression into (10), the optimal policy maximizes

$$\int_0^{+\infty} [u(q_t, Q_t) + \alpha V(Q_t)] \exp(-(\alpha + \delta)t) dt \quad (15)$$

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

Proposition 2 *Suppose $\bar{Q}_0 \geq \bar{S}$. Then there exists an optimal path, which solves problem (15). Moreover, this optimal path is monotonic, and converges to the value Q^D defined in Definition 2.*

Our main theorems focus on the remaining situations in which $\bar{Q}_0 < \bar{S}$.

3 Optimal policies

Characterizing the policy is not a simple task, as the problem involves three state variables, one being a record process, and additionally allows for non-parametric functions for both the payoffs and the belief distribution. For such a problem, methods from optimal control theory or calculus of variations methods that aim at deriving the policies from the first-order conditions do not readily apply. Instead, we derive the main qualitative properties that optimal paths must satisfy, and then the two main theorems.

The first qualitative property is the long-run behavior of an optimal path, as time goes to infinity. If the path has exceeded at some point the maximum possible threshold \bar{S} , thus triggering a catastrophe with certainty, then the legacy of the past is fixed at its maximum value 1, and convergence to Q^D follows, as shown in Proposition 2. Otherwise, the catastrophe might never be triggered. In such a case, one easily obtains from (13) that the survival probability converges to a finite positive value while the legacy of the past goes to zero. This means that the long-run value of experimenting further can be evaluated in a simple way. This allows us to determine the possible limits for the stock value, based on our definition of Q^E (see Lemma D.2).

The second class of properties relates to the monotonicity of an optimal path. Suppose that the stock is U-shaped, meaning that it takes the same value at two different dates while remaining below this value between these two dates. Then the maximum stock on record \bar{Q} is a constant on this interval, and the probability of survival evolves in a simple manner, following (13). One can then proceed to a simple experiment: replace on this interval the candidate path by a constant path for which the stock remains at its

initial level. Note that at the final date the two paths perfectly coincide with the same values for the three state variables, so that one can evaluate the difference in payoffs between these two paths by computing this difference only on the time interval under study. Optimality of the candidate path then implies that this difference is positive. We use this inequality repeatedly to derive monotonicity results (see Lemma B.1).

Finally, it is important to notice that standard continuity and convexity requirements for the constraints fail to hold for the general problem, because of the presence of a record process in (12). The consequence is that we are able to prove the existence of the optimum only after collecting enough qualitative properties for candidate paths. Under the assumptions of Theorem 1, the existence of the optimum can be verified from standard theorems, such as Theorem 15, p.237, in Seierstad and Sydsaeter (1987). For Theorem 2, we compute the optimum for specific examples.

All these arguments are derived step by step in the Appendix.

3.1 The first theorem: when $Q^E < Q^D$

The first theorem applies when

$$\bar{Q}_0 < Q^E < Q^D < \min(Q^N, \bar{S}). \quad (16)$$

In words, at the initial date, experimentation has barely begun, so the initial stock is low. A sequence of arguments shows how (16) leads to our first theorem.

First, it is a general property of optimal paths that they are monotonically increasing when they lie below Q^D . Intuitively, even in the worst case, in which the legacy is one, the policy would optimally increase the stock toward Q^D . Corollary C.1 in the Appendix shows is that this conclusion extends to lower levels of the legacy.

Second, given that $Q^E < Q^D$, it is not optimal to experiment further if one reaches Q^D . Intuitively, either the legacy is very small, and then one should not experiment any further if one is already above Q^E , or the legacy is very high, and then one should optimally come close to the long-run target Q^D (this is lemma F.1 in the Appendix).

We conclude that optimal paths must be increasing and bounded by Q^D , and therefore they must converge to some value $Q_T \leq Q^D$ at some date $T \leq +\infty$. Because the record-process disappears, existence of optimal paths is now easily proven, using standard results.

Finally, with the above preliminaries, one can proceed to a classical dynamic programming exercise: should the planner stop experimentation at date T , or a bit before T , or after T ? We outline next the key trade-off for the decision.⁹ If the planner stops and decides to stay at Q_T forever, the payoff is the continuation value conditional on survival at T , which we denote by z_T in:

$$z_T = \frac{u_T^0}{\delta} - \frac{\alpha}{\alpha + \delta} \pi_T D_T,$$

with $u_T^0 = u(0, Q_T)$ and $D_T = D(Q_T)$ for short. Intuitively, the legacy at T is crucial for the expected damage. Alternatively, the planner could continue experimenting for a short interval of time $[T, T + \Delta]$ with $q_T > 0$, and after this time stay at $Q_{T+\Delta}$ forever. Looking at the objective in (10), the flow gain from such a short experimentation period is $\frac{1}{p_T} [p_T u_T - \dot{p}_T V_T] \Delta$, in which we divide by p_T to condition on the fact that the planner has survived to T . Substituting for the definition of \dot{p} , and using $\pi = 1 - (1 - F)/p$ and $D = u^0/\delta - V$ in the flow gain, the welfare from the short experiment is

$$[u_T + \alpha \pi_T (\frac{u_T^0}{\delta} - D_T)] \Delta + \exp(-\delta \Delta) z_{T+\Delta}.$$

Now, if it is indeed optimal to stop at T , the planner must be indifferent between doing so and continuing as described. Manipulating this indifference leads to the condition presented in theorem 1:

Theorem 1 *Suppose (16) holds. Then there exists an optimal path. It is weakly increasing, and reaches its maximum value Q_T at some time $T \leq +\infty$, such that:*

$$\nu(Q_T) = \frac{\alpha}{\alpha + \delta} [(1 - \pi_T) \rho(Q_T) D(Q_T) + \pi_T D'(Q_T)]. \quad (17)$$

This condition nicely consolidates the definitions (5)-(6) of Q^D and Q^E , with weights given by the legacy at the time when the experimentation stops. The condition implies that a higher legacy π_T is associated to a higher long-run value of the stock. In this precise sense, higher legacies of the past promote more experimentation: immediate consumption becomes more of a priority when it is more likely that a catastrophe was triggered in the past because, by the assumption $Q^E < Q^D$, relatively little can be done to limit the damages from a potential catastrophe. This fatalism pushes the final value above

⁹The proof of theorem 1 is in Appendix F, and the dynamic programming interpretation in Appendix H.1.

Q^E , towards Q^D . This means that *a higher legacy of the past should make the planner experiment more, and thus be less cautious*. Proposition 3 in Section 4 formally proves this result in the case of a simple cake-eating problem.

Once Q_T is reached, as time goes by and no catastrophe occurs, the planner becomes more and more certain that no catastrophe was triggered at all. Then, the legacy of the past goes to zero. Now, since the stock is already above Q^E , there is no point in experimenting further; and since the stock is below Q^N , reducing the stock is also harmful. This is why the planner chooses to stabilize the stock forever after time T .

3.2 The second theorem: when $Q^E > Q^D$

We next reverse the key ranking of Q^E and Q^D , thus switching to a case when damages are sensitive to the stock level at the occurrence date of a catastrophe:

$$Q^D < Q^E < \min(Q^N, \bar{S}). \quad (18)$$

In this situation, a striking result is that the long-run target for the stock is easily computed. Indeed, if the stock remains below \bar{S} , then in the long-run the legacy of the past must go to zero. This implies that one should not stabilize below Q^E as further experimentation would be valuable. Conversely, further experimentation above Q^E is suboptimal both when the legacy is zero, by construction of Q^E , and when the legacy is high, because then one should aim at the lower value Q^D . We obtain:

Theorem 2 *Suppose (18) holds. If an optimal path is such that $\bar{Q}_\infty < \min(Q^N, \bar{S})$, then it converges to \bar{Q}_∞ , and $\bar{Q}_\infty = \max(\bar{Q}_0, Q^E)$.*

Once more, the interpretation is easier if one starts from a low level of the stock. Then every optimal path that remains below the long-run target for the stock level in the absence of catastrophe, and that does not trigger a catastrophe with certainty, must converge to the value of the stock that makes further experimentation valueless. This long-run target is in particular independent of the legacy at the initial date, contrary to what happened with the previous theorem.

On the other hand, in the short-run the path need not be monotonic. We are only able to provide a partial result: Lemma C.2 (in Appendix) shows that an optimal path may be decreasing at some date only if the legacy of the past is above a threshold at this date. The applications we now study will confirm that non-monotonic paths can indeed be optimal.

4 Eating a cake of unknown size

In a seminal paper, Kemp (1976) studies a cake-eating problem in which the size of the cake is initially unknown. Consider, for instance, the extraction of services from an ecosystem, such as a fishery, that may collapse due to overexploitation; or the viewpoint in climate change that one should not exceed a safe carbon budget whose value is uncertain (van der Ploeg, 2018), lest a catastrophe be triggered. We here extend this model by incorporating a delay between triggering and occurrence of a catastrophe. We additionally make strong assumptions on functional forms, so as to be able to perform some comparative statics with respect to the initial legacy of the past π_0 .

At each date t , a decision-maker with discount rate $\delta > 0$ chooses a net consumption $q_t \in [\underline{q}, \bar{q}]$ and receives an instantaneous payoff $u_0 + u_1 q_t$, where $u_0 \geq 0$ is the existence value of the cake, and $u_1 > 0$ is the value of a unit of the cake. In contrast to Kemp's model, we allow net consumption to be negative: this might be for example because the resource is at least partially renewable. The cumulative consumption is Q_t .

The catastrophe is triggered when Q_t exceeds the unknown threshold S , with a cumulative distribution function $F(S)$ and the associated hazard rate $\rho(S)$; and it occurs after an exponential delay τ , with parameter $\alpha > 0$. After the occurrence, the planner gets a continuation payoff that we allow to depend on cumulative consumption: we set $V(Q) = -v_0 Q$, where Q is the cumulative consumption at the occurrence time, with $v_0 \geq 0$.

In terms of the general model, the primitives of the cake-eating problem are now:

$$u(q, Q) = u_0 + u_1 q \quad V(Q) = -v_0 Q \quad u_1 > 0 \quad u_0, v_0 \geq 0.$$

We obtain:

$$\nu(Q) = u_1 > 0 \quad D(Q) = \frac{u_0}{\delta} + v_0 Q.$$

Therefore, Q^N equals plus infinity: there is no reason to limit consumption if catastrophes are excluded by assumption. On the other hand, if the catastrophe was triggered with certainty in the past, the relevant long-run target is now Q^D , and it is easily seen that this target is plus or minus infinity, according to whether $u_1 - \frac{\alpha}{\alpha + \delta} v_0$ is positive or negative.¹⁰ Intuitively, both the marginal gains and expected losses from consumption are constant, and thus the planner optimally either reaps consumption gains or mitigates damages

¹⁰For simplicity, we ignore here the natural constraint $Q \geq 0$.

as much as possible before the catastrophe occurs. Finally, when it is known that the catastrophe has not been triggered at all, the experimentation threshold is Q^E , implicitly defined by the following equality:

$$u_1 = \frac{\alpha}{\alpha + \delta} \rho(Q^E) \left(\frac{u_0}{\delta} + v_0 Q^E \right),$$

provided such a value belongs to the support of S (see Definition 3). We now distinguish two cases.

Theorem 1 applied: Assume $u_1 > \frac{\alpha}{\alpha + \delta} v_0$, so that the ranking is $Q^E < Q^D = +\infty$. From Theorem 1, the optimal policies are weakly increasing. Consequently, the problem in (10)-(13) simplifies because \bar{Q} equals Q everywhere, and constraint (12) has disappeared. Because the utility function and the constraints all are linear in q , the Pontryagin principle is easily applied: we show in Appendix that the optimal policy consist of setting q at its maximum level \bar{q} in some time interval $[0, T]$, and then, after time T , stabilizing the stock forever by setting $q = 0$. The following result characterizes the optimal values of this stabilization time, which may be finite or infinite:

Proposition 3 *For the cake-eating problem, let $u_1 > \frac{\alpha}{\alpha + \delta} v_0$ and $Q_0 = \bar{Q}_0 < Q^E$. Then there exists a value π^* such that:*

(i) *If the initial legacy π_0 is below π^* , the optimal policy is to set the control variable at its maximum level until some finite date T : $q_t = \bar{q}, t \in [0, T]$, and to stabilize the stock forever after this date: $q_t = 0, t \geq T$.*

(ii) *If the initial legacy π_0 is above π^* , the optimal policy consists in triggering the catastrophe with certainty, by setting q_t at the maximum level \bar{q} forever.*

(iii) *The stabilization date ($T \in [0, +\infty]$) and the final stock level Q_T are nondecreasing functions of the initial legacy π_0 .*

This result nicely formalizes the main intuition in this case. With a low consumption in the past, one is confident that the cake will not disappear soon, and this makes it worth being cautious and to avoid experimentation. Conversely, after a high past consumption, one expects the cake to disappear anyway, and therefore it becomes optimal to allow for even more consumption while this is possible. The key result is the third one, proving that higher legacies lead to more experimentation.

Theorem 2 applied: Conversely, assume $u_1 < \frac{\alpha}{\alpha+\delta}v_0$, so that the ranking now is $Q^D < Q^E$. For a stark illustration, suppose further that we start planning after intensive experimentation in the recent past: the level of the stock is equal to the highest level on record, the level itself is quite high in the following sense:

$$Q^E < Q_0 = \bar{Q}_0 < \min(Q^N, \bar{S}) \quad u_1 < \frac{\alpha}{\alpha + \delta}v_0. \quad (19)$$

Proposition 4 *Consider the cake-eating example, and suppose (19) holds. If the legacy π_0 is small enough ($u_1 > \pi_0 \frac{\alpha}{\alpha+\delta}v_0$), then there exists an optimal path, which consists in stabilizing the stock forever: $q_t = 0$ for all t . Otherwise, there exists a unique optimal path, characterized by two dates t_1 and t_2 such that $0 < t_1 < t_2 < +\infty$, and which are increasing with π_0 , such that:*

- $q_t = \underline{q} < 0$ for $t < t_1$;
- $q_t = \bar{q} > 0$ for $t_1 < t < t_2$;
- $q_t = 0$ and $Q_t = \bar{Q}_0$ for $t > t_2$.

This result thus proves formally that optimal policies can be non-monotonic. It is interesting also to compare to proposition 3: now, a higher legacy of the past makes the planner more cautious in the short-run, since the threat of pending catastrophes leads him to reduce the stock more. In the long-run, the legacy vanishes, and convergence to the initial value \bar{Q}_0 follows.

5 Disease control and social distancing

We now provide a simple model of a pandemic that includes a trade-off between social distancing and economic activity. This trade-off is common in the literature; see, e.g., Bloom, Kuhn and Prettner (2022) for a review. Additionally, we include the possibility that the health system, or even the entire economy, breaks down when the situation becomes severe enough. The threat of such a catastrophe is new in the literature, and allows to derive a rich set of predictions, as we now show.

Consider a population of agents whose mass is normalized to one. During the early stages of the pandemic, the population I_t of infected agents at time t follows a simple law of motion:

$$\dot{I}_t = (R_t - (r + d))I_t, \quad I_0 > 0 \text{ given.}$$

The recovery rate r and the death rate d are positive parameters. Variable $R_t \in [0, \bar{R}]$ measures new infections, with maximum value $\bar{R} > r + d$ attained when people behave as in the absence of the pandemic. By mandating social distancing, the social planner can reduce the value of R_t , so that stabilization occurs when $R = r + d$, and complete isolation is associated to the value $R = 0$. The benefit from social distancing is to eventually reduce the number of deaths, with a value of statistical life $w > 0$. But this reduction comes at an economic cost: the value of production at time t is an increasing and concave function $Y(R)$ of R . Therefore, the instantaneous payoff is

$$Y(R) - wdI.$$

If we now define

$$Q \equiv \log I, \quad q \equiv R - (r + d),$$

we are back to our model, with

$$u(q, Q) = Y(q + r + d) - wd \exp(Q), \quad \dot{Q} = q \in [\underline{q} \equiv -r - d, \bar{q} \equiv \bar{R} - r - d],$$

and an initial value $Q_0 = \log I_0$.¹¹ Variable q is thus the rate of increase of the population of infected agents. Function ν is defined as in Assumption 1:

$$\nu(Q) = u_q(0, Q) + \frac{u_Q(0, Q)}{\delta} = Y'(r + d) - \frac{wd}{\delta} \exp(Q), \quad (20)$$

and it is indeed decreasing with Q . In the absence of catastrophes, the long-run target for the stock of infected agents is defined by the equality $\nu(Q^N) = 0$, or equivalently:

$$I^N = \exp(Q^N) = \frac{\delta Y'(r + d)}{wd}. \quad (21)$$

The policy target I^N varies intuitively with parameters, and reaching it over time follows from a social distancing policy implementing $R > r + d$ on a path starting at $I_0 < I^N$. However, planning in a pandemics may not be such a smooth operation. One may worry that society, or the health system, breaks down if the number of infected agents is too high, or that the pathogen mutates into something much more dangerous. We now add the possibility of such concerns to the simple stock-flow problem above.

Assume that a catastrophe is triggered when the logarithm of the number of infected agents exceeds a threshold S whose value is unknown. With this interpretation, the

¹¹In epidemiology, R is the matching rate. The reproduction rate, as commonly defined, is the matching rate R times the time spent in the infected state, $\frac{1}{r+d}$. Therefore, $q = 0$ corresponds to a matching rate equal to $r + d$ and a reproduction rate equal to 1.

distribution F of S on the support $[0, \bar{S}]$ and the associated hazard rate ρ are as defined in the general model.

When the catastrophe occurs, the planner loses control: the matching rate takes an exogenous value R^* , and the output remains fixed at a low level $Y^* < Y(r + d)$. The proportion of deaths increases to $d^* > d$, the recovery rate becomes r^* , and the resulting rate of increase q^* of the number of infected is assumed to satisfy the following inequalities:¹²

$$0 < q^* \equiv R^* - (r^* + d^*) < \delta.$$

After the catastrophe has occurred at time κ , we therefore have $I_t = I_\kappa \exp[q^*(t - \kappa)]$, and the continuation payoff can be computed explicitly:

$$V(Q_\kappa) = \int_\kappa^{+\infty} [Y^* - wd^*I_t] \exp[-\delta(t - \kappa)] dt = \frac{Y^*}{\delta} - \frac{wd^*}{\delta - q^*} \exp(Q_\kappa).$$

Consequently, the damage function can be written as the sum of a production loss, and of the value of the mortality increase:

$$D(Q) = \frac{u(0, Q)}{\delta} - V(Q) = \frac{Y(r + d) - Y^*}{\delta} + w\mu^* \frac{d}{\delta} \exp(Q), \quad (22)$$

where the parameter μ^* measures the increase in mortality:

$$\mu^* \equiv \frac{\frac{d^*}{\delta - q^*} - \frac{d}{\delta}}{\frac{d}{\delta}} > 0.$$

After some manipulations, we obtain the long-run target I^D :

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

meaning that one rationally braces for the catastrophe by reducing infections below the no-catastrophe target I^N , and the more so the bigger is the change in mortality measured by μ^* . The expression for I^E is more complicated. Some calculus leads to the following result:

Lemma 2 *In the disease control model, if one has*


$$\frac{1}{1 + \frac{Y(r+d) - Y^*}{w\mu^* d I^D}} < \rho(I^D), \quad (23)$$

then $I^E < I^D$, and Theorem 1 applies. Otherwise, $I^E > I^D$, and Theorem 2 applies.

¹²The last inequality allows to avoid infinite values for the discounted welfare cost of deaths. Alternatively, one could assume that a vaccine is discovered after some (exogenous but possibly stochastic) date T ; or one could endogenize the value of R^* after the catastrophe by allowing the planner to control it; or one could impose that the number of infected agents cannot exceed the population size, by switching from a simple exponential model to a full S-I-R model.

This result underlines the role played by the ratio $\frac{Y(r+d)-Y^*}{w\mu^*d}$, which measures the relative importance of economic losses vis-à-vis mortality increases. It is remarkable that this simple parameter determines important characteristics of optimal paths, as we now explain by ways of simulations.

Figure 2



Optimal paths in the plane (π, I) for a linear production function $Y(R) = Y_0 R$. Parameters are: $\delta = 0.03, q^* = 0.01, w = 1, d = 0.01, r = 0.99, d^* = 0.2, \alpha = 0.2, Y_0 = 1000, Y^* = 900, I_0 = 32, \bar{q} = 1$. Distribution F for $\log(I)$ is uniform: $F = 1/6$. Benchmark values are $I^N = 300, I^D = 110$, and $I^E = 50$.

Theorem 1 applied: Consider first the case of a planner who cares more about economic activity than about deaths, so that condition (23) is satisfied. Assume that the production function is linear, and initial beliefs are uniform (all parameters are specified in Figure 2). The solid curve Figure 2 depicts the infection level I satisfying the stopping condition of theorem 1, eq. (17), as a function of legacy π . When there is no legacy ($\pi_0 = 0$), the infection level is $I^E = 50$, and similarly, when $\pi_0 = 1$ we get $I^D = 110$. Under the conditions in Theorem 1, optimal policy paths are monotonic and must stabilize the infection levels at a point (π, I) from this solid curve.


Let us then see how the legacy and the infection level jointly evolve before the stabilization.¹³ Each dotted curve depicts this relationship, for varied initial legacies, but

¹³The problem is linear in q , and, by standard dynamic programming arguments, the optimal control takes the maximum value \bar{q} under the conditions in theorem 1 until the stopping condition holds. This gives a differential equation for the legacy. We solved the differential equation and the two-point

with the same initial infection level set at $I_0 = 32$. As in the cake-eating example, one observes that a higher initial legacy leads to a higher long-run value for the stock.¹⁴ The intuition is the same: if the stock of infected agents has increased very rapidly before time zero, then the probability that the catastrophe was triggered is high, and the planner chooses to privilege high production levels before the event occurs, at the price of additional deaths. Another noteworthy remark is that along each optimal path the legacy π_t is increasing with t : this means that the planner allows the stock of infected to increase quite fast, thereby increasing the probability that a catastrophe is triggered. This fatalistic behavior is at odds with what prudence would recommend; but it is the rational consequence of an emphasis on production, relative to deaths.

Theorem 2 applied: Let us now enter the domain of Theorem 2, by assuming that the planner mainly aims at reducing the number of deaths, so that inequality (23) is reversed. For this illustration, assume that planning starts so late that the infected population I_0 is close to the long-run target in the absence of catastrophes I^N . By Theorem 2, optimal paths must converge to this initial level in the long-run. We compute the solution path in Appendix G.2.

Figure 3




The population of infected agents over time, for a linear production function. Parameters are: $\delta = 0.025$, $q^* = 0.02$, $w = 5$, $d = 0.02$, $r = 0.98$, $d^* = 0.2$, $\alpha = 0.2$, $Y_0 = 1000$, $I_0 = 230$.

Figure 3 depicts the optimal time path of the stock of infected agents for different boundary value problem numerically to reach the stopping condition from given (π_0, I_0) .

¹⁴Note that $\pi_0 = 1 - (1 - F(Q_0))/p_0$ cannot exceed $F(Q_0)$. This is why π_0 only takes values below 0.5 in the graph.

values of the legacy at the initial date. A complete lockdown turns out to be optimal in a first phase, as soon as the legacy is strictly positive. After a while, if the catastrophe does not occur the planner becomes more and more convinced that the catastrophe was not triggered in the past, and chooses to gradually relax the lockdown. In the long-run, it is optimal to increase the stock up to the initial value, because the probability that the threshold lies below it has become negligible.

Figure 4



The optimal control, for a linear production function. Parameters are: $\delta = 0.025, q^* = 0.02, w = 5, d = 0.02, r = 0.98, d^* = 0.2, \alpha = 0.2, Y_0 = 1000, I_0 = 230$.

Figure 4 depicts the optimal time path of the control variable R_t , corresponding to the paths in Figure 3. The Figure confirms that with a higher initial value for the legacy the lockdown lasts longer, and the recovery is slower, though in the long-run all paths converge to the same level. We conclude that contrary to what happened before, a higher legacy makes the planner initially more cautious. Finally, the optimality of early containment followed by a relaxation and increasing infections resembles the so-called hammer-and-dance policies for Covid-19. This learning-based rationale for the hammer-and-dance policy differs from those surveyed in Assenza et al. (2020).

6 Remarks on climate change

Consider a pollution stock Q_t that follows a simple law of motion:

$$\dot{Q}_t = E_t - \gamma Q_t, \quad (24)$$

where E_t is the pollution flow, and $\gamma > 0$ is the constant decay rate of the stock. The output, denoted by Y_t , is

$$Y_t = \exp(-\theta Q) K^{1-\beta} E_t^\beta \quad (25)$$

where K stands for capital, which we will set to 1 in this illustration, E_t measures the fossil-fuel energy use, and $\beta \in (0, 1)$ is the factor share. With $\theta > 0$, the first term corresponds to the production losses due to the accumulation of the pollution stock. Production is entirely consumed at each date, so that $C_t = Y_t$. Instantaneous utility of consumption is $U(C) = \log C$.

We are back to our model if we set $q = E - \gamma Q$. Then,

$$u(q, Q) = \beta \log(q + \gamma Q) - \theta Q, \quad \nu(Q) = \frac{\beta \gamma + \delta}{Q} - \frac{\theta}{\delta}.$$

and solving $\nu(Q^N) = 0$ yields

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

These stock-flow tradeoffs present a toy model for climate change, inspired by Golosov et al. (2014). The target Q^N increases in the abatement cost β , in the rate at which CO₂ disappears from the atmosphere γ , and declines in the percentage of output lost per unit increase in the stock θ .

It is a common concern that such smooth stock-flow tradeoffs may not well describe the climate change problem (e.g., Pindyck, 2014). 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). The 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. We may capture this impact by making θ to increase by a factor $k > 1$, and we assume that this shock is important enough:

$$k > 1 + \frac{\gamma}{\delta}.$$

After the catastrophe has occurred, the planning goes on, and the continuation value $V(Q)$ becomes:¹⁵

$$V(Q) = \frac{-k\theta}{\delta + \gamma}Q + \frac{\beta}{\delta}[\log(\gamma Q^N) - \log k - 1]. \quad (26)$$

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

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

This simple setting allows highlighting the basic conceptual differences in two the approaches to modelling a catastrophe in the literature.

In the first approach, the catastrophe is pending. For example, van der Ploeg and de Zeeuw (2017) is explicit about the idea that the ultimate arrival of the catastrophe is evident, and the focus is on how to prepare for such an event. In our toy model, the corresponding target is Q^D , and from Definition 2 we get:

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

which indeed is less than Q^N .

In the second approach, there is no legacy from the past because there is no delay between triggering and occurrence. Then the relevant target Q^E can be expressed, thanks to Definition 3, as a solution to equation

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

where function g is defined as

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

Function g is strictly increasing and concave with $g(Q^N) > 0$, so it holds that $Q^E < Q^N$. When there is no delay we have $\frac{\alpha}{\alpha + \delta} = 1$, and then the information structure is no different from that, for instance, in Lemoine and Traeger (2014).

By comparing the above equations, one immediately obtains that Q^E is below QD if and only if the hazard rate is high enough, as already observed in Lemma 2 for the pandemic case. In light of this one-parameter variation, we observe that both theorems

¹⁵The planning problem is to maximize $V(Q_0) = \max \int_0^\infty \ln C_t \exp(-\delta t) dt$, 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 $V(Q)$.

are potentially relevant for the optimal policies. But this is only the first step of planning, as the planner should also evaluate the legacy, the information content of the past experiments. In fact, our results suggest an agenda for the applied quantitative research evaluating the optimal climate policies with detailed climate-economy descriptions: 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 (2023), offer frameworks for exploring the question.

7 Connections to the literature

The hazard rate approach. The comparison with the approach used in Clarke and Reed (1994), Polasky, de Zeeuw and Wagener (2011), Sakamoto (2014), van der Ploeg and de Zeeuw (2017), or Besley and Dixit (2019) is instructive. 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 reads as:

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

Comparing with (14), we see that those works can be interpreted to assume that a catastrophe was triggered in the past. They then focus on how to best manage two distinct elements. First, the delay before the catastrophe occurs can be controlled by reducing the stock since they assume that h is an increasing function of Q . We do not allow for this possibility in our model, as our delay follows a process with a constant hazard rate α . Second, the damage from the catastrophe can be controlled by varying the stock, as in our model; this effect is stronger if the damage varies more with the stock, which makes Q^D lower compared to Q^N . Overall, by assuming exogenous delays our setting is somewhat less general but, on the other hand, it allows to deal with the question of whether to trigger a catastrophe in the first place.

The uncertain threshold approach. Tsur and Zemel (1994, 1995, 1996), and more recently Lemoine and Traeger (2014), Diekert (2017), and Chen (2020) all use an uncertain threshold approach in which a catastrophe occurs as soon as the threshold is reached, so that there is no delay between triggering and occurrence. Consequently, there is no legacy of the past. In our model, this corresponds to the case when α goes to

infinity. The definition for Q^E in (6) is now modified since $\alpha/(\alpha + \delta)$ goes to 1, and the target stock Q^{E0} (superscript 0 stands for the absence of delays) is such that

$$\nu(Q^{E0}) = \rho(Q^{E0})D(Q^{E0}).$$

The following result was first obtained in Tsur and Zemel (1994). Since our assumptions are weaker than theirs, we offer a general proof in the Appendix. The statement $Q_0 = \bar{Q}_0$ is made for simplicity.

Proposition 5 *Suppose $Q_0 = \bar{Q}_0$. In the absence of delay ($\alpha = +\infty$), there exists an optimal path, and it is:*

- (i) *decreasing and converging to Q^N , if $Q_0 > Q^N$;*
- (ii) *constant, if $Q_0 \in [Q^{E0}, Q^N]$;*
- (iii) *increasing and converging to Q^{E0} , if $Q_0 < Q^{E0}$.*

In particular, the optimal path is a constant in case (ii): one does not want to experiment further because the stock is already above Q^{E0} , and reducing the stock is also useless, as the current situation is safe.¹⁶

8 Concluding remarks

Inferences about catastrophes are difficult before they actually happen. This paper developed a novel approach for optimal experimentation with catastrophes that have delayed observable impacts and severity depending on past actions. The model interprets historical data for obtaining 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 optimal actions forward. For crises such as Covid-19, the model predicts that similar planners can take very different optimal courses of actions depending on the legacy. 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

¹⁶This confirms the findings in the literature, as summarized in the following citation (Tsur and Zemel, 1996, page 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.”

if the current level was approached slowly. The lesson for climate change would be: The “lockdown of emissions” may be optimal until unknowns can be ruled out.

The information structure developed in this paper seems broadly applicable, including also good events such as breakthroughs in basic science and technology development. In fact, the gestation times in basic research are measured in decades (e.g., Adams, 1990), and therefore the delay between the cause and the impact seems essential in assessments of past investments in research. Should basic research, private or government sponsored, be conducted steadily over time or as intensive bursts? The delayed learning from knowledge stocks could offer a new avenue studying such questions and implications for policies that seek to internalize the spillovers between research and commercially oriented R&D (see, e.g., Akcigit, Hanley and Serrano-Velarde, 2020).

References

- Adams, James D.** 1990. “Fundamental Stocks of Knowledge and Productivity Growth.” *Journal of Political Economy*, 98(4): 673–702.
- Akcigit, Ufuk, Douglas Hanley, and Nicolas Serrano-Velarde.** 2020. “Back to Basics: Basic Research Spillovers, Innovation Policy, and Growth.” *The Review of Economic Studies*, 88(1): 1–43.
- 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.
- 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.
- Bloom, David E., Michael Kuhn, and Klaus Prettnner.** 2022. “Modern Infectious Diseases: Macroeconomic Impacts and Policy Responses.” *Journal of Economic Literature*, 60(1): 85–131.
- 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.

- Chen, Yi.** 2020. “A revision game of experimentation on a common threshold.” *Journal of Economic Theory*, 186: 104997.
- 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*, 122(4): 1259–1285.
- Diekert, Florian K.** 2017. “Threatening thresholds? The effect of disastrous regime shifts on the non-cooperative use of environmental goods and services.” *Journal of Public Economics*, 47(1): 30–49.
- 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.** 2018. “Carbon Prices for The Next Hundred Years.” *The Economic Journal*, 128(609): 728–757.
- 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.** 2020. “Precaution, Information and Time- Inconsistency: On The Value of the Precautionary Principle.” Centre for Economic Policy Research.
- IPCC.** 2021. *The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change [Masson-Delmotte, V., P. Zhai, A. Pirani, S.L. Connors, C. Péan, S. Berger, N. Caud, Y. Chen, L. Goldfarb, M.I. Gomis, M. Huang, K. Leitzell, E. Lonnoy, J.B.R. Matthews,*

T.K. Maycock, T. Waterfield, O. Yelekçi, R. Yu, and B. Zhou (eds.). Cambridge University Press.

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.

Laiho, T., P. Murto, and J. Salmi. 2023. “Gradual Learning from Incremental Actions.” Working paper.

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 de Zeeuw, and Florian Wagener. 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.
- 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 P.** 2023. “ACE—Analytic Climate Economy.” *American Economic Journal: Economic Policy*, 15(3): 372–406.
- 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, Frederick.** 2018. “The safe carbon budget.” *Climatic Change*, 47(1): 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.

Online Appendix

Table of Contents

A Preliminaries	37
B A useful inequality	37
C Consequences for monotonicity	40
D Consequences for convergence	41
E Consequences for benchmarks	49
F Optimal policies: main theorems	50
G Applications	53
G.1 Cake-eating	53
G.2 Social distancing illustration	56
H Additional results	56
H.1 Dynamic programming and optimal stopping	56
H.2 The model without delay	57

A Preliminaries

Note: to alleviate notations, we often omit arguments when there is no ambiguity, and we write e for $\exp(-\delta t)$, LHS for left-hand side, and RHS for right-hand side. We also use the convention that the hazard rate ρ is zero outside the support of S .

For a given admissible path, we can easily compute the unique solution to the differential equation (13):

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:

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

Here are also some simple facts about the survival probability and the legacy of the past.

The survival probability p_t is by definition strictly positive before the catastrophe occurs. From (13), p_t is weakly decreasing and thus must converge. Therefore \dot{p}_t goes to zero, and p_t goes to $1 - F(\bar{Q}_\infty)$, where \bar{Q}_∞ is the supremum of stock values for a given path.

The legacy of the past is $\pi_t = 1 - (1 - F(\bar{Q}_t))/p_t \in [0, F(\bar{Q}_t)]$. π_T is zero only when there has been no experimentation at all in the past, i.e., $\bar{Q}_T \leq \underline{S}$. Finally, one useful property is $\dot{p}_t = -\alpha p_t \pi_t$.

B A useful inequality

The next result will be used repeatedly to study the monotonicity and convergence of optimal paths. It follows from replacing, on an interval $[t_1, t_2]$, the candidate optimal path by a constant path.

Lemma B.1 *Let $(Q_t)_{t \geq 0}$ be an optimal path. Then*

$$\int_{t_1}^{t_2} (Q_t - Q_{t_1}) \left(\dot{p}_t (D'(Q_{t_1}) - \frac{\alpha + \delta}{\alpha} \nu(Q_{t_1})) + \delta (1 - F(\bar{Q}_t)) \nu(Q_{t_1}) \right) \exp(-\delta t) dt$$

$$\geq \frac{\alpha\delta}{\alpha + \delta} \rho(Q_{t_1}) D(Q_{t_1}) \int_{t_1}^{t_2} (1 - F(\bar{Q}_t)) (\bar{Q}_t - \bar{Q}_{t_1}) \exp(-\delta t) dt, \quad (\text{B.1})$$

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

- *Case (i)*: $0 \leq t_1 < t_2 = +\infty$.
- *Case (ii)*: $0 \leq t_1 < t_2 < +\infty$, $Q_{t_1} = Q_{t_2}$, $\bar{Q}_{t_1} = \bar{Q}_{t_2}$.

Proof of Lemma B.1: First, let us compute the payoff W from the optimal path on an interval $[t_1, t_2]$. Recall that by definition $V(Q) = \frac{u(0, Q)}{\delta} - D(Q)$, so that

$$W \equiv \int_{t_1}^{t_2} [pu - \dot{p}V] edt = \int_{t_1}^{t_2} [pu - \dot{p} \frac{u(0, Q)}{\delta} + \dot{p}D] edt.$$

Integrate by parts $\dot{p}u(0, Q)/\delta$ to get:

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

The concavity of u in q implies:

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

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

$$W \leq \underbrace{-[p \frac{u(0, Q)}{\delta} e]_{t=t_1}^{t=t_2}}_{=A} + \underbrace{\int_{t_1}^{t_2} [pq\nu + \dot{p}(D - D(Q_{t_1}))] edt}_{=B} + \underbrace{D(Q_{t_1}) \int_{t_1}^{t_2} \dot{p} edt}_{=C}. \quad (\text{B.2})$$

Now, consider an alternative path $(q'_t, Q'_t)_{t \geq t_1}$ that consists in setting $q'_t = 0$ on the same interval $[t_1, t_2]$, so that the stock remains set at Q_{t_1} on this interval; complete this path by setting $q'_t = q_t$ after time t_2 . Proceeding as above, we obtain that the payoff for this new path on the interval $[t_1, t_2]$ equals

$$W_0 = \underbrace{-[P \frac{u(0, Q_{t_1})}{\delta} e]_{t=t_1}^{t=t_2}}_{=A_0} + \underbrace{D(Q_{t_1}) \int_{t_1}^{t_2} \dot{P} edt}_{=C_0}. \quad (\text{B.3})$$

where the survival probability P is now given by (A.2). In case (i) of the Lemma, the optimality of the initial path implies the inequality $W \geq W_0$. In case (ii), the condition $(Q_{t_1} = Q_{t_2}, \bar{Q}_{t_1} = \bar{Q}_{t_2})$ ensures that the survival probability is the same under both paths on the interval $[t_1, t_2]$, and that the payoff from both paths is the same after t_2 . Therefore, once more the inequality $W \geq W_0$ must hold. We now compare the different terms in this inequality.

Observe first that A in W equals A_0 in W_0 . Indeed, in case (i) the bracketed terms are equal at $t = t_1$, and also at $t_2 = +\infty$ because the exponential is zero. In case (ii), this follows because $Q_{t_1} = Q_{t_2}$ and $\bar{Q}_{t_1} = \bar{Q}_{t_2}$, so that P and p are everywhere equal in this interval.

Secondly, in case (ii) C equals C_0 , because as already observed we have $p = P$. In case (i), we compute the difference between these terms by integrating by parts:

$$C - C_0 = \int_{t_1}^{+\infty} (\dot{p} - \dot{P})e dt = \delta \int_{t_1}^{+\infty} (p - P)e dt.$$

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

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

Another integration by parts yields:

$$C - C_0 = \frac{\alpha \delta}{\alpha + \delta} \int_{t_1}^{+\infty} (F(\bar{Q}_{t_1}) - F(\bar{Q}_t)) \exp(-\delta t) dt.$$

Now, we have:

$$F(\bar{Q}_t) - F(\bar{Q}_{t_1}) = \int_{\bar{Q}_{t_1}}^{\bar{Q}_t} f(S) dS = \int_{\bar{Q}_{t_1}}^{\bar{Q}_t} (1 - F(S)) \rho(S) dS \geq (1 - F(\bar{Q}_t)) \rho(\bar{Q}_{t_1}) (\bar{Q}_t - \bar{Q}_{t_1})$$

because ρ is increasing and $1 - F$ is decreasing.¹⁷ This yields the RHS in (B.1).

There only remains to study B in (B.2). To do so, define the function $N(Q) \equiv \int_{Q_{t_1}}^Q \nu(x) dx$. We have:

$$\int_{t_1}^{t_2} pqv e dt = [pN(Q)e]_{t=t_1}^{t=t_2} - \int_{t_1}^{t_2} N(Q)(\dot{p} - \delta p) e dt.$$

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:

$$B = \int_{t_1}^{t_2} [\delta p N + \dot{p}(D - D(Q_{t_1}) - N)] e dt.$$

Assumption 1 implies that N is concave, so that

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

¹⁷Recall that by convention ρ is zero outside the support of S . Thus this inequality also holds when \bar{Q}_{t_1} is below \underline{S} or above \bar{S} .

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

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

Since $\dot{p} \leq 0$, we obtain:

$$\delta p N + \dot{p}(D - D(Q_{t_1}) - N) \leq (Q_t - Q_{t_1})[\delta p \nu(Q_{t_1}) + \dot{p}(D'(Q_{t_1}) - \nu(Q_{t_1}))].$$

Finally, we use (13) to replace p by $1 - F(\bar{Q}) - \frac{1}{\alpha}\dot{p}$:

$$\delta p N + \dot{p}(D - D(Q_{t_1}) - N) \leq (Q_t - Q_{t_1})[\delta(1 - F(\bar{Q}))\nu(Q_{t_1}) + \dot{p}(D'(Q_{t_1}) - \frac{\alpha + \delta}{\alpha}\nu(Q_{t_1}))].$$

This yields the remaining terms in (B.1), and concludes the proof. \blacksquare

C Consequences for monotonicity

Corollary C.1 *Every optimal path is weakly increasing when Q_t is below Q^D .*

Proof of Corollary C.1: Let us proceed by contradiction. Consider an optimal path such that $Q^D > Q_T > Q_{T'}$ at some dates $T < T'$. One 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 (B.1) at $t_1 = T$ and $t_2 = +\infty$ (case (i)): $Q_t - Q_T$ is negative, and sometimes strictly negative; the first product is negative because $Q^D > Q_T$ and $\dot{p}_t \leq 0$, the second product is negative as $Q^N > Q_T$, and the RHS is zero as \bar{Q} is a constant.¹⁸

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 we are able to obtain a similar contradiction by checking (B.1) in case (ii). \blacksquare

Corollary C.2 *Suppose $Q^D < Q_{t_1}$, and that the stock decreases at the right of t_1 . Then $\pi_{t_1} > \frac{\delta}{\alpha} \frac{\nu(Q_{t_1})}{D'(Q_{t_1}) - \nu(Q_{t_1})}$, and this bound is strictly below 1, and is strictly positive if and only if $Q_{t_1} < Q^N$.*

¹⁸Strictly speaking, this only shows that the LHS is weakly negative, while the RHS is zero. But for the LHS to be exactly zero one would need to find points at which $\dot{p}_t = 0 = 1 - F(\bar{Q}_t)$, and from (13) we would obtain $p_t = 0$, in contradiction with the fact that the catastrophe did not occur yet. In other proofs, we shall skip such arguments, for the sake of brevity.

Proof of Corollary C.2: because the stock decreases at the right of t_1 , it is possible to apply Lemma B.1 in case (i) or (ii), on an interval $[t_1, t_2 \leq +\infty[$ on which $Q_t < Q_{t_1}$, so that $\bar{Q}_t = \bar{Q}_{t_1}$. Then there must exist t such that the function in the integral in (B.1) is negative, so that:

$$\dot{p}_t(D'_1 - \frac{\alpha + \delta}{\alpha}\nu_1) + \delta(1 - F(\bar{Q}_{t_1}))\nu_1 < 0.$$

Since $\pi_t = 1 - \frac{1-F(\bar{Q}_t)}{p_t}$, we obtain $\dot{p} = \alpha(1 - F(\bar{Q}_t) - p_t) = -\alpha p\pi$, and thus:

$$(D'_1 - \frac{\alpha + \delta}{\alpha}\nu_1)\alpha p_t \pi_t > \delta p_t(1 - \pi_t)\nu_1,$$

or equivalently $\pi_t > \frac{\delta}{\alpha} \frac{\nu(Q_{t_1})}{D'(Q_{t_1}) - \nu(Q_{t_1})}$. Moreover, because \bar{Q}_t is a constant, π_t is decreasing, and therefore $\pi_{t_1} \geq \pi_t$. This establishes the announced inequality. The bound is below one because $Q_{t_1} > Q^D$ (check the definition in (5)); it is above zero if and only if $\nu(Q_{t_1}) > 0$, or equivalently $Q_{t_1} < Q^N$, as announced. \blacksquare

D Consequences for convergence

The next result establishes the convergence of optimal paths for which there is a positive probability of not triggering the catastrophe:

Lemma D.1 *If an optimal path is such that $\bar{Q}_\infty < \bar{S}$ and $\bar{Q}_\infty \leq Q^N$, then it converges to some value Q_∞ as time goes to infinity. Moreover, one of the three following cases must hold:*

- (i) $Q_\infty = Q^N$;
- (ii) $Q_\infty > Q^N$, and the stock value Q_t is weakly decreasing for t high enough;
- (iii) $Q_\infty < Q^N$, and the stock value Q_t is weakly increasing for t high enough.

Proof of Lemma D.1: For t, t_1 such that $t \geq t_1$, define the function

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

Step 1: we first show that for every $\varepsilon > 0$, there exists a date $\Gamma(\varepsilon) < +\infty$ such that, for all t and t_1 such that $t \geq t_1 > \Gamma(\varepsilon)$, one has

$$|\dot{p}_t(D'(Q_{t_1}) - \frac{\alpha + \delta}{\alpha}\nu(Q_{t_1}))| < \delta(1 - F(\bar{Q}_t))\varepsilon. \quad (\text{D.1})$$

Indeed, the right-hand side is at least $\delta(1 - F(\overline{Q}_\infty))\varepsilon$, which is strictly positive. Moreover, recall that $D'(Q) - \frac{\alpha + \delta}{\alpha}\nu(Q)$ is weakly increasing from Assumption 2, and let us distinguish two cases:

- Either $Q_{t_1} < Q^D$, and therefore the path is increasing from date 0 to date t_1 , from Corollary C.1. Then we have

$$D'(Q_0) - \frac{\alpha + \delta}{\alpha}\nu(Q_0) \leq D'(Q_{t_1}) - \frac{\alpha + \delta}{\alpha}\nu(Q_{t_1}) \leq 0.$$

- Either $Q_{t_1} \geq Q^D$, and because $Q_{t_1} \leq \overline{Q}_\infty < \overline{S}$, which is finite, we have

$$0 \leq D'(Q_{t_1}) - \frac{\alpha + \delta}{\alpha}\nu(Q_{t_1}) \leq D'(\overline{S}) - \frac{\alpha + \delta}{\alpha}\nu(\overline{S}).$$

This shows that in any case the factor of \dot{p}_t in (D.1) is bounded. Since \dot{p}_t goes to zero, the result follows.

Step 2: suppose that there exists t_1 such that $Q_{t_1} < Q^N$ and $t_1 > \Gamma(\nu(Q_{t_1}))$. We show that Q_t must be weakly increasing at the right of Q_{t_1} .

Indeed, under our assumption in this step, from Step 1 (D.1) must hold at $\varepsilon = \nu(Q_{t_1}) > 0$, for all $t \geq t_1$. This implies that $B(t_1, t)$ is strictly positive for every $t \geq t_1$, as the second term in B is strictly positive as $Q_{t_1} < Q^N$, and this term is strictly above the absolute value of the first term.

Now, notice that the expression inside the integral in (B.1) is at most equal to $(Q_t - Q_{t_1})B(t_1, t)$. Therefore, if Q_t lies below Q_{t_1} for all $t \geq t_1$, and is sometimes strictly below Q_{t_1} , we reach a contradiction with inequality in (B.1) in case (i). And if there exists $t_2 \geq t_1$ such that $Q_{t_1} = Q_{t_2} \geq Q_t$ for all $t \in [t_1, t_2]$, with sometimes a strict inequality, we once more reach a contradiction with (B.1) in case (ii). Therefore, Q_t must be weakly increasing at the right of Q_{t_1} , as announced.

Step 3: suppose that for all t_1 above a threshold, we have either $(Q_{t_1} < Q^N$ and $t_1 > \Gamma(\nu(Q_{t_1}))$), or $(t_1 \leq \Gamma(|\nu(Q_{t_1})|))$. If the second domain is bounded, then after a threshold date the path must fully belong to the first domain, so that the path is weakly increasing after this threshold date, from Step 2. Since the path is bounded by the finite value \overline{Q}_∞ , it must converge. In particular, if it converges to a value strictly below Q^N , then it must be weakly increasing for t high enough, as announced in case (iii) of the Lemma.

Alternatively, if the second domain is unbounded, as t_1 grows without bounds in this domain the inequality $t_1 \leq \Gamma(|\nu(Q_{t_1})|)$ implies that $\nu(Q_{t_1})$ must get closer and closer to zero, so that Q_{t_1} must get arbitrarily close to Q^N ; and whenever t_1 belongs to the first domain, then Q_t must be weakly increasing at the right of t_1 from Step 2, thus getting closer to Q^N . This shows that the path converges to Q^N , as in case (i) of the Lemma.

Step 4: Otherwise, for every T there exists $t_1 \geq T$ such that $Q_{t_1} > Q^N$ and $t_1 > \Gamma(-\nu(Q_{t_1}))$.

A first possibility is that, after some threshold date, the path is weakly decreasing whenever it is above Q^N . From our assumption in this step, the path must therefore remain above Q^N . Being weakly decreasing and bounded from below, the path must converge. In particular, if it converges to a value strictly above Q^N , then it must be weakly decreasing for t high enough, as announced in case (ii) of the Lemma.

Otherwise, for every T there exists $t_1 \geq T$ such that $Q_{t_1} > Q^N$, $t_1 > \Gamma(-\nu(Q_{t_1}))$, and the path is increasing at the right of Q_{t_1} . From Step 1, $B(t_1, t)$ is strictly negative for all $t \geq t_1$. If the path remains weakly above Q_{t_1} for $t \geq t_1$, we obtain a contradiction with (B.1) in case (i); so that the path must at some point go strictly below Q_{t_1} . Therefore, there exists $t_2 > t_1$ such that $Q_{t_1} = Q_{t_2} \leq Q_t$ for $t \in [t_1, t_2]$. If moreover $\bar{Q}_{t_1} = \bar{Q}_{t_2}$, we obtain a contradiction with (B.1) in case (ii). We have thus shown that for each such t_1 , after t_1 the path strictly exceeds the maximum stock on record \bar{Q}_{t_1} , before going strictly below Q_{t_1} . Therefore the path fluctuates an infinite number of times, and the amplitude of each fluctuation is increasing as time goes by.

Let us number these fluctuations using an integer index n . What we have shown is that there exists two increasing and unbounded sequences $(\tau_n)_{n \geq 0}$ and $(\tau'_n)_{n \geq 0}$ such that $\tau_n < \tau'_n < \tau_{n+1}$, and τ_n is the date at which the n^{th} fluctuation reaches a minimum, and τ'_n is the date at which this fluctuation reaches a maximum. Moreover, Q_{τ_n} must be decreasing with n ; but if it goes down below Q^N , then it must be that $\tau_n \leq \Gamma(\nu(Q_{\tau_n}))$, because otherwise the stock is increasing. This inequality implies that the limit of Q_{τ_n} must be at least Q^N , and therefore that the stock remains above Q^N forever. Symmetrically, $Q_{\tau'_n}$ must be increasing with n , and because $Q_{\tau'_n} = \bar{Q}_{\tau'_n}$ this sequence must converge to \bar{Q}_∞ .

In the long-run, \bar{Q}_t becomes arbitrarily close to \bar{Q}_∞ , and the planner learns almost nothing new from each fluctuation. Therefore π_t must go to zero, and the planner's problem becomes identical to the Stock Flow Problem without catastrophes, for which

all solutions are weakly decreasing paths that converge to Q^N . We thus have reached a contradiction. \blacksquare

The following result characterizes the possible limit values for the stock:

Lemma D.2 *Suppose an optimal path is such that $\bar{Q}_\infty < \bar{S}$. Suppose moreover that the stock Q_t converges to a value Q_∞ when t goes to infinity. Then $Q^E \leq Q_\infty \leq Q^N$, and $Q_\infty = Q^N$ if $Q_\infty < \bar{Q}_\infty$.*

Proof of Lemma D.2: The proof consists of six steps.

Step 1: we first state several simple results. By assumption, Q_t converges to Q_∞ , and \bar{Q}_t converges to \bar{Q}_∞ , and $Q_\infty \leq \bar{Q}_\infty < \bar{S}$.

From (13), p_t is weakly decreasing and thus must converge. Therefore \dot{p}_t goes to zero, and p_t goes to $1 - F(\bar{Q}_\infty)$.

Because Q_T goes to Q_∞ , the difference $Q_\infty - Q_T = \int_{t \geq T} q_t dt$ goes to zero. Moreover, an integration by parts yields

$$\int_{t \geq T} q_t \exp(-\delta(t - T)) dt = \int_{t \geq T} q_t dt - \delta \int_{t \geq T} \left(\int_{\tau \geq t} q_\tau d\tau \right) \exp(-\delta(t - T)) dt$$

so that the left-hand side goes to zero when T goes to infinity.

Step 2: we now show that the planner's payoff:

$$W(T) \equiv \int_{t \geq T} (p_t u(q_t, Q_t) - \dot{p}_t V(Q_t)) \exp(-\delta(t - T)) dt$$

converges to the value $Z \equiv (1 - F(\bar{Q}_\infty)) \frac{1}{\delta} u(0, Q_\infty)$.

A first remark is that since the path is optimal, then $W(T)$ is at least the payoff $W_0(T)$ from stabilizing the stock forever at its level Q_T . Using (A.2), we compute

$$\begin{aligned} W_0(T) &\equiv \int_{t \geq T} (P_t u(0, Q_T) - \dot{P}_t V(Q_T)) \exp(-\delta(t - T)) dt \\ &= u(0, Q_T) \left(\frac{1 - F(\bar{Q}_T)}{\delta} + \frac{p_T - 1 + F(\bar{Q}_T)}{\alpha + \delta} \right) + \frac{\alpha}{\alpha + \delta} V(Q_T) (p_T - 1 + F(\bar{Q}_T)). \end{aligned}$$

Since p_T converges to $1 - F(\bar{Q}_\infty)$, we have shown:

$$W(T) \geq W_0(T) \quad \text{and} \quad \lim_{T \rightarrow +\infty} W_0(T) = Z. \quad (\text{D.2})$$

A second remark is that one can decompose $p_t u(q_t, Q_t) - \dot{p}_t V(Q_t)$ into

$$\begin{aligned}
& (p_t - 1 + F(\bar{Q}_\infty))u(q_t, Q_t) \\
& + (1 - F(\bar{Q}_\infty))(u(q_t, Q_t) - u(q_t, Q_\infty)) \\
& + (1 - F(\bar{Q}_\infty))(u(q_t, Q_\infty) - u(0, Q_\infty)) \\
& + (1 - F(\bar{Q}_\infty))u(0, Q_\infty) \\
& - \dot{p}_t V(Q_t).
\end{aligned}$$

Because u is concave in q , the third line is less than $(1 - F(\bar{Q}_\infty))q_t u_q(0, Q_\infty)$; and the last result in Step 1 implies that the integral on $t \geq T$ of this last expression, weighted by $\exp(-\delta(t - T))$, goes to zero as T goes to infinity.

Now, for T high enough one can restrict attention to Q taking values in a bounded neighborhood A of Q_∞ . Because u and u_Q are bounded on $[q, \bar{q}] \times A$, and V is bounded on A , the first, second, and last terms go to zero as t goes to infinity, and so do their integrals when weighted by $\exp(-\delta(t - T))$.

Overall, by integrating on $t \geq T$ we obtain that $W(T)$ is below the weighted integral of the fourth term, which is Z , plus some terms that go to zero as T goes to infinity. Together with (D.2), this establishes that both $W(T)$ and $W_0(T)$ converge to Z , as announced.

Step 3: in this step, we define an alternative, feasible path with payoff W_1 , and we exhibit a simple condition on W_1 ensuring that this path dominates the candidate path – something that should not happen since the candidate path is supposed to be a solution to the planner’s problem. Each of the next three steps then uses this result in different cases to prove different parts of the Lemma by contradiction.

Given $(T, \hat{q} \in [q, \bar{q}], a > 0)$, we define an alternative path: play \hat{q} on $[T, T + a]$, and 0 afterwards. This path is feasible, and we obtain new trajectories for the variables $(\hat{q}, \hat{Q}, \bar{Q}, \hat{p})$:

- For $t \in [T, T + a]$:

$$\hat{q}_t = \hat{q} \quad \hat{Q}_t = Q_T + \hat{q}(t - T) \quad \bar{Q}_t = \max(\bar{Q}_T, \hat{Q}_t)$$

and from (A.1):

$$\hat{p}_t = p_T \exp(-\alpha(t - T)) + \alpha \exp(-\alpha t) \int_T^t (1 - F(\bar{Q}_\tau)) \exp(\alpha \tau) d\tau.$$

- For $t \geq T + a$:

$$\hat{q}_t = 0 \quad \hat{Q}_t = \hat{Q}_{T+a} \quad \bar{Q}_t = \bar{Q}_{T+a}$$

and from (A.2):

$$\hat{p}_t = 1 - F(\bar{Q}_{T+a}) + (\hat{p}_{T+a} - 1 + F(\bar{Q}_{T+a})) \exp(-\alpha(t - T)). \quad (\text{D.3})$$

Overall, this alternative path yields a payoff W_1 , as follows:

$$\begin{aligned} W_1(T, \hat{q}, a) &\equiv \int_T^{T+a} (\hat{p}_t u(\hat{q}, \hat{Q}_t) - \dot{\hat{p}}_t V(\hat{Q}_t)) \exp(-\delta(t - T)) dt \\ &+ \int_{t \geq T+a} (\hat{p}_t u(0, \hat{Q}_{T+a}) - \dot{\hat{p}}_t V(\hat{Q}_{T+a})) \exp(-\delta(t - T)) dt. \end{aligned}$$

Consider now the following condition:

$$\exists \hat{q}, \bar{a} > 0, k > 0, \quad \forall a \in]0, \bar{a}[, \quad \lim_{T \rightarrow +\infty} \frac{\partial W_1}{\partial a}(T, \hat{q}, 0) > k. \quad (\text{D.4})$$

Suppose it holds. Then we have

$$W_1(T, \hat{q}, \bar{a}) = W_1(T, \hat{q}, 0) + \int_0^{\bar{a}} \frac{\partial W_1}{\partial a}(T, \hat{q}, a) da.$$

Moreover, we have $W_1(T, \hat{q}, a) = W_0(T)$, and we know from Step 2 that $W_0(T)$ and $W(T)$ have the same limit when T goes to infinity. Therefore:

$$\lim_{T \rightarrow +\infty} [W_1(T, \hat{q}, \bar{a}) - W(T)] = \int_0^{\bar{a}} \lim_{T \rightarrow +\infty} \frac{\partial W_1}{\partial a}(T, \hat{q}, a) da > \bar{a}k > 0,$$

which means that for T high enough the alternative path (T, \hat{q}, \bar{a}) dominates the initial path. This is impossible, as the initial path was assumed to be a solution. In each of the last three steps, we thus only have to show that (D.4) holds to reach a contradiction.

Step 4: in this step, we proceed by contradiction, by assuming $Q_\infty > Q^N$. Choose \hat{q} such that $q \leq \hat{q} < 0$. For some T and $a > 0$, consider the alternative path (T, \hat{q}, a) .

Because $\hat{q} < 0$, this alternative path is such that the highest stock on record \bar{Q}_t is a constant, equal to \bar{Q}_T . Let us compute the derivative of $W_1(T, \hat{q}, a)$ with respect to a . Using Step 3, we compute the following expressions, for $t \geq T + a$:

$$\frac{\partial \hat{p}_t}{\partial a} = \frac{\partial \hat{p}_{T+a}}{\partial a} \exp(-\alpha(t - T)) = \alpha(1 - F(\bar{Q}_T) - \hat{p}_{T+a}) \exp(-\alpha(t - T)) \quad (\text{D.5})$$

$$\frac{\partial \dot{\hat{p}}_t}{\partial a} = -\alpha \frac{\partial \hat{p}_t}{\partial a}. \quad (\text{D.6})$$

Therefore, $\frac{\partial W_1}{\partial a}(T, \hat{q}, a)$ equals

$$\begin{aligned} & \hat{p}_{T+a}u(\hat{q}, \hat{Q}_{T+a}) \exp(-\delta a) - \hat{p}_{T+a}u(0, \hat{Q}_{T+a}) \exp(-\delta a) \\ & + \int_{t \geq T+a} \alpha(1 - F(\bar{Q}_T) - \hat{p}_{T+a})(u(0, \hat{Q}_{T+a}) + \alpha V(\hat{Q}_{T+a})) \exp(-(\alpha + \delta)(t - T)) dt \\ & + \int_{t \geq T+a} (\hat{p}_t u_Q(0, \hat{Q}_{T+a}) \hat{q} - \dot{\hat{p}}_t V'(\hat{Q}_{T+a}) \hat{q}) \exp(-\delta(t - T)) dt. \end{aligned} \quad (\text{D.7})$$

From Step 1, as T and $t \geq T + a$ go to infinity, \hat{Q}_{T+a} goes to $Q_\infty + a\hat{q}$, $\overline{\hat{Q}_{T+a}}$ goes to \bar{Q}_∞ , \hat{p}_{T+a} and \hat{p}_t both go to $1 - F(\bar{Q}_\infty)$, and $\dot{\hat{p}}_t$ goes to zero; recall also that V' is bounded on a neighborhood of \bar{Q}_∞ from Assumption 2. Therefore, $\lim_{T \rightarrow +\infty} \frac{\partial W_1}{\partial a}(T, \hat{q}, a)$ equals

$$(1 - F(\bar{Q}_\infty)) \exp(-\delta a) \hat{q} \left[\frac{u(\hat{q}, Q_\infty + a\hat{q}) - u(0, Q_\infty + a\hat{q})}{\hat{q}} + \frac{u_Q(0, Q_\infty + a\hat{q})}{\delta} \right]. \quad (\text{D.8})$$

Finally, recall that we chose \hat{q} to be strictly negative, and notice that the bracketed term is strictly negative for \hat{q} close enough to zero, as its limit is $\nu(Q_\infty) < 0$. This shows (D.4), and we obtain a contradiction thanks to the reasoning at the end of Step 3. This shows that Q_∞ cannot exceed Q^N .

Step 5: in this step, we proceed by contradiction, by assuming $Q_\infty < \bar{Q}_\infty$ and $Q_\infty < Q^N$. Choose \hat{q} such that $0 < \hat{q} \leq \bar{q}$. For some T and $a > 0$, consider the alternative path (T, \hat{q}, a) .

Because $Q_\infty < \bar{Q}_\infty$ one can choose \bar{a} small enough so that $Q_\infty + \bar{a}\hat{q} < \bar{Q}_\infty$, and therefore the alternative path is such that the highest stock on record $\overline{\hat{Q}_t}$ is a constant, equal to \bar{Q}_T .

We then proceed exactly as in Step 4, to get a contradiction: the final expression for the limit is unchanged, and it is strictly positive because now \hat{q} and $\nu(Q_\infty)$ are both strictly positive. Hence, $Q_\infty = Q^N$ if $Q_\infty < \bar{Q}_\infty$, as announced in the Lemma.

Step 6: once more proceeding by contradiction, we now assume $Q_\infty < Q^E$, so that $Q_\infty = \bar{Q}_\infty$ from Step 5. Choose \hat{q} such that $0 < \hat{q} \leq \bar{q}$. For some T and $a > 0$, consider the alternative path (T, \hat{q}, a) .

A new feature is that the new highest stock on record may now depend on a , since $\hat{q} > 0$. Referring to (D.3), we note that we only have to care about the value of $\overline{\hat{Q}_{T+a}}$, which now equals $\max(Q_T + a\hat{q}, \bar{Q}_T)$. We therefore define the indicator function $1_{Q_T + a\hat{q} \geq \bar{Q}_T}$,

and the only changes to our computations are in (D.5) and (D.6), which we rewrite into:
for $t \geq T + a$,

$$\frac{\partial \hat{p}_t}{\partial a} = \alpha(1 - F(\overline{\hat{Q}_{T+a}}) - \hat{p}_{T+a}) \exp(-\alpha(t-T)) + 1_{Q_T + a\hat{q} \geq \overline{Q}_T} f(Q_T + a\hat{q}) \hat{q} (\exp(-\alpha(t-T)) - 1)$$

and

$$\frac{\partial \dot{\hat{p}}_t}{\partial a} = -\alpha^2(1 - F(\overline{\hat{Q}_{T+a}}) - \hat{p}_{T+a}) \exp(-\alpha(t-T)) - \alpha 1_{Q_T + a\hat{q} \geq \overline{Q}_T} f(Q_T + a\hat{q}) \hat{q} \exp(-\alpha(t-T)).$$

The derivative $\frac{\partial W_1}{\partial a}(T, \hat{q}, a)$ now becomes

$$\begin{aligned} & \hat{p}_{T+a} u(\hat{q}, \hat{Q}_{T+a}) \exp(-\delta a) - \hat{p}_{T+a} u(0, \hat{Q}_{T+a}) \exp(-\delta a) \\ & + \int_{t \geq T+a} \left(\alpha(1 - F(\overline{\hat{Q}_{T+a}}) - \hat{p}_{T+a}) (u(0, \hat{Q}_{T+a}) + \alpha V(\hat{Q}_{T+a})) \right) \exp(-(\alpha + \delta)(t - T)) dt \\ & \quad + \int_{t \geq T+a} \left(\hat{p}_t u_Q(0, \hat{Q}_{T+a}) \hat{q} - \dot{\hat{p}}_t V'(\hat{Q}_{T+a}) \hat{q} \right) \exp(-\delta(t - T)) dt \\ & + 1_{Q_T + a\hat{q} \geq \overline{Q}_T} f(Q_T + a\hat{q}) \hat{q} u(0, \hat{Q}_{T+a}) \int_{t \geq T+a} (\exp(-\alpha(t - T)) - 1) \exp(-\delta(t - T)) dt \\ & \quad + \alpha 1_{Q_T + a\hat{q} \geq \overline{Q}_T} f(Q_T + a\hat{q}) \hat{q} V(\hat{Q}_{T+a}) \int_{t \geq T+a} \exp(-\alpha(t - T)) \exp(-\delta(t - T)) dt. \end{aligned}$$

We now compute the limit of this derivative when T goes to infinity. Since $Q_\infty = \overline{Q}_\infty$ and $\hat{q} > 0$, \hat{Q}_{T+a} and $\overline{\hat{Q}_{T+a}}$ both go to $Q_\infty + a\hat{q}$, and the first three lines converge as before to (D.8). Also, for T high enough $1_{Q_T + a\hat{q} \geq \overline{Q}_T}$ is 1. Finally, using the definition of D , the last two integrals go to

$$-(1 - F(Q_\infty + a\hat{q})) \exp(-\delta a) \hat{q} \frac{\alpha}{\alpha + \delta} \rho(Q_\infty + a\hat{q}) D(Q_\infty + a\hat{q}).$$

By choosing \hat{q} strictly positive but small enough, we can make this limit arbitrarily close to

$$(1 - F(\overline{Q}_\infty)) \exp(-\delta a) \hat{q} [\nu(Q_\infty) - \frac{\alpha}{\alpha + \delta} \rho(Q_\infty) D(Q_\infty)],$$

which is strictly positive because $\hat{q} > 0$ and $Q_\infty < Q^E$. Using the same reasoning as at the end of Step 3, we obtain a contradiction. Therefore, Q_∞ has to be at least Q^E , as announced in the Lemma. \blacksquare

E Consequences for benchmarks

We can now apply our useful inequality (B.1) to two benchmarks: the Stock-Flow Problem without catastrophes, and the case when the catastrophe was triggered with certainty in the past.

Proof of Proposition 1: Existence of a solution to the SFP 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 3.) Then (B.1) 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 B.1. 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 sometimes strict. 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.

This also implies that the path converges to some value Q_∞ . Lemma D.2 then implies that this value is Q^N , since $Q^E = Q^N$. ■

Proof of Proposition 2: The proof follows exactly the proof of Proposition 1, since in the two problems the constraint sets are identical; and the objectives (2) and (15) are formally identical; and u and $u + \alpha V$ share the same properties. In particular, recall how ν is built from u and δ , and proceed similarly with the new objective function $\varphi \equiv u + \alpha V$ and $\alpha + \delta$: we have

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

and using the definition of V in (4) this expression reduces to $\nu(Q) - \frac{\alpha}{\alpha+\delta}D'(Q)$, which is decreasing in Q from our assumptions. This is the only property we need to apply the proof of Proposition 1. \blacksquare

F Optimal policies: main theorems

We begin by a few intermediate results that we will use repeatedly in the proofs of our main theorems.

Lemma F.1 *Suppose $Q^E \leq Q^D$ and $\bar{Q}_0 \leq Q^D$. Then optimal paths cannot exceed Q^D .*

Proof of Lemma F.1: If $Q^D \geq \bar{S}$, Proposition 2 implies that if the stock exceeds \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 Corollary C.1 we know that after t_1 the path must remain above Q^D . Therefore, we have $\bar{Q}_t \geq Q_t > Q_{t_1} = \bar{Q}_{t_1} = Q^D \geq Q^E$ for all $t \geq t_1$, and we can apply Lemma (B.1) in case (i): The first term is zero by definition of Q^D , and the difference between the second term and the RHS is strictly negative, as $1 - F(\bar{Q}_t)$ is strictly positive (at least for t close to t_1), $\bar{Q}_t - \bar{Q}_{t_1} \geq Q_t - Q_{t_1} > 0$, and $Q_{t_1} > Q^E$. This contradicts inequality (B.1). \blacksquare

Lemma F.2 *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 F.2: let us proceed by contradiction. Suppose that \bar{Q}_∞ is reached only asymptotically, necessarily from below. From Lemma D.1, there exists T such that Q_t is weakly increasing for $t \geq T$, and converges asymptotically to \bar{Q}_∞ , so that $\bar{Q}_t = Q_t$ for $t \geq T$. Moreover, because $\bar{Q}_\infty > Q^E$, we can choose T such that $Q_T > Q^E$. We therefore have, for every $t \geq T$, $\bar{Q}_t = Q_t \geq Q_T > Q^E$.

Referring to Lemma B.1, consider the function

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

defined for $t \geq t_1 \geq T$. Because $Q_{t_1} > Q_T > Q^E$, the second term is strictly negative; in fact, it is strictly less than

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

Because \dot{p} goes to zero, the first term becomes negligible compared to k^- when t is high enough, so that we can choose t_1 high enough so that $B(t_1, t) < 0$ for all $t \geq t_1$. Finally, because the stock is weakly increasing, we have $Q_t = \bar{Q}_t$, and therefore the function in (B.1) 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 with (B.1). This shows that \bar{Q}_∞ must be reached in finite time. \blacksquare

Lemma F.3 *Let $T \in]0, +\infty[$. Suppose an optimal path Q_t is weakly increasing on $[0, T]$ and constant afterwards, with $Q_t = Q_T = \bar{Q}_\infty$ for $t \geq T$. Then the planner's payoff at time 0 equals*

$$p_0 \left(\frac{u(0, Q_0)}{\delta} - \frac{\alpha}{\alpha + \delta} \pi_0 D(Q_0) \right) + \int_0^T p_t [B_t + q_t C_t] \exp(-\delta t) dt$$

and moreover $C_T = 0$, where

$$B_t \equiv u(q_t, Q_t) - u(0, Q_t) - q_t u_q(0, Q_t),$$

and

$$C_t \equiv \nu(Q_t) - \frac{\alpha}{\alpha + \delta} [(1 - \pi_t)\rho(Q_t)D(Q_t) + \pi_t D'(Q_t)]. \quad (\text{F.1})$$

Proof of Lemma F.3: the planner's payoff is

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

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 (13), and then we use (A.2) to compute \dot{P}_t/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 $V(Q) = \frac{u(0, Q)}{\delta} - D(Q)$, this can be simplified into

$$Z(T) \equiv p_T \frac{u(0, Q_T)}{\delta} - \frac{\alpha}{\alpha + \delta} (p_T - 1 + F(Q_T)) D(Q_T). \quad (\text{F.2})$$

In particular, we have $W(0) = Z(0)$, and thus $W(T) = Z(0) + \int_0^T W'(t) dt$, where $W'(t)$ stands for the left-derivative of W at t . There remains to compute $W'(T)$, which is $\exp(-\delta T)$, times

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

We have $\dot{p}_T = \alpha(1 - F(Q_T) - p_T)$. Let u_0 denote $u(0, Q_T)$, and u_{0Q} denote $u_Q(0, Q_T)$. The above expression becomes:

$$\begin{aligned} & p_T u(q_T, Q_T) - \alpha(1 - F - p) \left(\frac{u_0}{\delta} - D \right) + \alpha(1 - F - p) u_0 / \delta + p q \frac{u_{0Q}}{\delta} \\ & - \frac{\alpha}{\alpha + \delta} (\alpha(1 - F - p) + qf) D - \frac{\alpha}{\alpha + \delta} (p - 1 + F) q D' - p u_0 + \delta \frac{\alpha}{\alpha + \delta} (p - 1 + F) D. \end{aligned}$$

Almost all terms in u_0 and D cancel each other. We obtain that the left-derivative of $W(T)$ is $p_T \exp(-\delta T)$, times

$$u(q_T, Q_T) - u_0 + q \frac{u_{0Q}}{\delta} - \frac{\alpha}{\alpha + \delta} \frac{qfD}{p} - \frac{\alpha}{\alpha + \delta} \left(1 - \frac{1 - F}{p} \right) q D'.$$

We finally use the definitions $\max \pi = 1 - (1 - F)/P$, $\rho = f/(1 - F)$, and $\nu = u_{0q} + u_{0Q}/\delta$ to get the announced expression for the planner's payoff.

There remains to show that $C_T = 0$. Because the path is optimal, it must be that $T > 0$ maximizes the planner's payoff $W(T)$, as computed above. Two optimality conditions must hold. Firstly, the first-order condition $W'(T) = 0$ yields $B_T + q_T C_T = 0$ at the optimal date T , so that:

$$u(q_T, Q_T) - u(0, Q_T) - q_T u_q(0, Q_T) + q_T C_T = 0.$$

Secondly, $W(T)$ is also maximized with respect to the path before T , and the value of q_T is free; therefore, q_T must maximize $B_T + q_T C_T$, so that

$$q_T \in \arg \max_{q \in [\underline{q}, \bar{q}]} u(q, Q_T) - q u_q(0, Q_T) + q C_T.$$

The first condition expresses that both q_T and 0 are solution to the program in the second condition. Because 0 is interior to the interval $[\underline{q}, \bar{q}]$, this implies that the derivative of the objective of the objective function is zero at $q = 0$. This shows $C_T = 0$, as

announced. ■

Proof of Theorem 1: from Lemma F.1, optimal paths cannot exceed Q^D ; from Corollary C.1, optimal paths are weakly increasing, 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). We also obtain that the optimal path converges toward the maximum value $\bar{Q}_\infty \in [Q^E, Q^D]$, from Lemma D.2, and that the stock level is constant after \bar{Q}_∞ is reached. Lemma F.3 then gives the planner's payoff.

If \bar{Q}_∞ is reached asymptotically, then a contrapositive to Lemma F.2 implies that the stock converges to $\bar{Q}_\infty = Q^E$. Since the legacy π_T vanishes when T goes to infinity, we obtain that (17) indeed holds at $T = +\infty$, $\pi_T = 0$, and $Q_T = Q^E$, as announced.

If \bar{Q}_∞ is reached at time $T < +\infty$, 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 $W(T)$, and Lemma F.3 implies $C_T = 0$, as announced. ■

Proof of Theorem 2: Because $\bar{Q}_\infty < \min(Q^N, \bar{S})$, case (iii) in Lemma D.1 is the only possibility. Therefore the path converges to a value Q_∞ strictly below Q^N , and Lemma D.2 implies $\bar{Q}_\infty = Q_\infty$. And \bar{Q}_∞ is at least \bar{Q}_0 by definition of the maximum stock on record, and at least Q^E from (iii) in Lemma D.1.

Now, let us proceed by contradiction, and suppose $\bar{Q}_\infty > \max(\bar{Q}_0, Q^E)$. Then Lemma F.2 implies that \bar{Q}_∞ is reached in finite time. (iii) in Lemma D.1 implies that Q_t is weakly increasing for t high enough. Together, these properties imply that there exists $T < +\infty$ such that Q_t is weakly increasing just before T , and equals \bar{Q}_∞ after T . Then we can apply Lemma F.3 to obtain that (17) holds. But this equality implies that Q_T lies between Q^E and Q^D , and this contradicts the inequalities $Q^D \leq Q^E < Q_T = \bar{Q}_\infty$. ■

G Applications

G.1 Cake-eating

Proof of Proposition 3: let us first consider the case when after some time t_0 the optimal policy exceeds the upper value \bar{S} , so that the catastrophe is triggered with certainty. Then we know that the optimal policy maximizes (15), which in this simple

case is

$$\int_{t \geq t_0} (u_0 + u_1 q_t - \alpha v_0 Q_t) \exp(-(\alpha + \delta)t) dt.$$

Thanks to a simple integration by parts, this objective can be transformed into

$$\int_{t \geq t_0} (u_0 + (u_1 - \frac{\alpha}{\alpha + \delta} v_0) q_t) \exp(-(\alpha + \delta)t) dt.$$

Because we have assume $u_1 > \frac{\alpha}{\alpha + \delta} v_0$, the solution consists in setting $q_t = \bar{q}$ forever.

We can now examine the optimal policy when the stock lies below \bar{S} . We can focus on “bang-bang” policies that set q_t either to zero or to \bar{q} . Recall that from Lemma F.3 we can consider that the control is set to zero after some date $T \leq +\infty$. This Lemma also gives a useful expression for the payoff. Since by linearity of function u the B_t terms are identically zero, this payoff reduces to

$$\int_0^T q_t p_t C_t \exp(-\delta t) dt,$$

where we can use the definition of π to replace in (F.1):

$$\begin{aligned} p_t C_t &= p_t u_1 - \frac{\alpha}{\alpha + \delta} [(1 - F(Q_t)) \rho(Q_t) D(Q_t) + (p_t - 1 + F(Q_t)) D'(Q_t)] \\ &= p_t (u_1 - \frac{\alpha}{\alpha + \delta} v_0) - \frac{\alpha}{\alpha + \delta} (1 - F(Q_t)) (\rho(Q_t) D(Q_t) - D'(Q_t)). \end{aligned}$$

Now, suppose that q_t is zero between two dates t_0 and t_1 , and is \bar{q} just after t_1 . This implies that the integral after time t_1 is positive; otherwise, one would play $q_t = 0$ forever. But then the idle time between t_0 and t_1 is wasted: it would be better to instead play at $t \geq t_0$ what is scheduled for $t \geq t_1$. Indeed, not only one would follow the same path for the stock at a earlier date, but in addition one would also benefit from a a higher survival probability; and in the expression above this higher probability is beneficial because it is multiplied by a positive coefficient.

This shows that in any case the control variable must be equal to \bar{q} until some date $T \leq +\infty$, and be zero afterwards. We thus have to maximize on T the objective

$$\bar{q} \int_0^T p_t C_t \exp(-\delta t) dt.$$

Moreover, from (A.1) we have

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

so that the cross-derivative in (p_0, T) of the objective above is strictly positive. By supermodularity, this implies that a higher p_0 leads to a higher choice of the stopping time T , and therefore to a higher value for the final stock. Since we have

$$\pi_0 = 1 - \frac{1 - F(Q_0)}{p_0},$$

a higher p_0 is equivalent to a higher initial legacy π_0 . This shows the existence of a threshold π^* , and the results in the Proposition follow. ■

Proof of Proposition 4: 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):

$$p_t = 1 - F(Q_0) + (p_0 - 1 + F(Q_0)) \exp(-\alpha t),$$

and the objective function

$$W = \int_0^{+\infty} [p_t(u_0 + u_1 q_t) + \dot{p}_t v_0 Q_t] \exp(-\delta t) dt$$

becomes

$$W = p_0 \int_0^\infty q_t a_t \exp(-\delta t) dt + C,$$

where C is a constant, and

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

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 , which is easily found to be increasing in π_0 . 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 $\underline{q}t_1 + \bar{q}t_2 = 0$, so that t_2 is also increasing in π_0 . The optimal policy is thus as stated in the claim. ■

G.2 Social distancing illustration

Theorem 2 applied: Under constraint $I_t \leq I_0$, we can solve for p explicitly and, by using the same arguments as in the proof of Proposition 4, we write the general payoff as

$$\int_0^\infty b_t q_t \exp(-\delta t) dt + B,$$

where B is a constant and

$$b_t \equiv (1 - \pi_0) \left(Y_0 - \frac{wdI_t}{\delta} \right) + \pi_0 \left(Y_0 - \left(\frac{\alpha wd^*}{\delta - q^*} + wd \right) \frac{I_t}{\alpha + \delta} \right) \exp(-\alpha t).$$

The flow payoff is thus proportional to the distancing measure q_t with $b_0 < 0$: the first term of b_0 is zero by the definition of I^N , and for the second term note that $\left(Y_0 - \left(\frac{\alpha wd^*}{\delta - q^*} + wd \right) \frac{I^N}{\alpha + \delta} \right) = \left(\frac{wdI^N}{\delta} - \left(\frac{\alpha wd^*}{\delta - q^*} + wd \right) \frac{I^N}{\alpha + \delta} \right) < \left(\frac{1}{\delta} - \frac{1}{(\delta - q^*)} \right) \frac{\alpha}{\alpha + \delta} wd^* I^N < 0$. A complete lockdown, $q = -(r + d)$ implying $R = 0$, is thus optimal at $t = 0$ and, in fact, for all t with $b_t < 0$. But the lockdown must end: b_t turns positive at some finite $t' > 0$ when the lockdown policy is followed at all times t prior to t' . The optimal policy after t' is to relax social distancing so that I grows back to I_0 . When infections grow we must have $b_t q_t \geq 0$ which holds with $b_t = 0$ unless with choice set for q binds. The numerical recovery paths in the Figures satisfy $b_t = 0$. ■

H Additional results

H.1 Dynamic programming and optimal stopping

We develop the stopping condition by variational methods, after several intermediate steps needed for the validity of the approach (see the proof of theorem 1). Taking these steps as given, for intuition, we now invoke a dynamic programming argument to describe the tradeoff at T .

Consider the part of the overall welfare that accrues after stopping in $[T, \infty)$, as defined by the objective (10). Noting that in $[T, \infty)$ the stock is stabilized $q = 0$, and then the survival probability p_t follows a formula (lemma A.1) that allows us to express the said welfare as a product of discount factor $\exp(-\delta T)$ and¹⁹

$$p_T \frac{u_T^0}{\delta} - \frac{\alpha}{\alpha + \delta} (p_T + F_T - 1) D_T,$$

¹⁹This expression is line (F.2) above.

where we use shorthands $u_T^0 = u(0, Q_T)$, $F_T = F(\bar{Q}_T)$, and $D_T = D(Q_T)$. Throughout this paper, the planner stands at $t = 0$ but think, momentarily, that the planner has survived to T . Multiply the welfare expression above by $1/p_T$ to condition on survival and use $\pi = 1 - (1 - F)/p$ to see that the planner's welfare, standing at the stopping time T , takes the following intuitive form: $z_T \equiv u_T^0/\delta - \frac{\alpha}{\alpha + \delta}\pi_T D_T$. Alternatively, the survivor could continue experimenting for a short interval of time $[T, T + \Delta]$ with $q_T > 0$, and after this time stop with $q_{T+\Delta} = 0$. By the above logic, the welfare at $T + \Delta$ is

$$z_{T+\Delta} = \frac{1}{p_T} \left[p_{T+\Delta} \frac{u_{T+\Delta}^0}{\delta} - \frac{\alpha}{\alpha + \delta} (p_{T+\Delta} + F_{T+\Delta} - 1) D_{T+\Delta} \right].$$

The flow gain from this one-shot experiment follows from the objective (10) that, together with the discounted $z_{T+\Delta}$, leads to the full welfare at T

$$\frac{1}{p_T} [p_T u_T - \dot{p}_T V_T] \Delta + \exp(-\delta \Delta) z_{T+\Delta}.$$

This one-shot experimentation welfare can be better grasped by rewriting with $\pi = 1 - (1 - F)/p$, $D = u^0/\delta - V$, and the first-order approximation of $\exp(-\delta \Delta) z_{T+\Delta}$ with respect to Δ ,

$$[u_T + \alpha \pi_T (\frac{u_T^0}{\delta} - D_T)] \Delta + z_T - \delta z_T \Delta + z'_T \Delta,$$

where $z'_T = \frac{\partial z_{T+\Delta}}{\partial \Delta} |_{\Delta=0}$. Now, at optimal T , the planner cannot strictly prefer one of the two options. Using this indifference and choosing the optimal experimentation intensity q_T gives the condition:

$$0 = \max_{q_T} \left\{ u_T + \alpha \pi_T (\frac{u_T^0}{\delta} - D_T) - \delta z_T + z'_T \right\}.$$

After careful evaluation of terms, this condition becomes

$$0 = \max_{q_T} \left\{ u(q_T, Q_T) - u(0, Q_T) - q_T u_q(0, Q_T) + q_T C(T) \right\}$$

where

$$C_T \equiv \nu(Q_T) - \frac{\alpha}{\alpha + \delta} [(1 - \pi_T) \rho(Q_T) D(Q_T) + \pi_T D'(Q_T)]. \quad (\text{H.1})$$

H.2 The model without delay

Proof of Proposition 5: since α is infinite, the problem under study is to maximize

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

under the constraints $\dot{Q}_t = q_t \in [\underline{q}, \bar{q}]$, $p_t = 1 - F(\bar{Q}_t)$, $\bar{Q}_t = \max_{0 \leq t' \leq t} Q_{t'}$, Q_0 being given. Consider a candidate path, and let us proceed by necessary conditions.

Step 1: we first show that one may focus on monotonic paths. Suppose there exist two arbitrary dates 0 and $T > 0$, such that $Q_0 = Q_T \geq Q_t$ for $t \in [0, T]$. In such a case, the maximum stock on record 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 all periods of time 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 $\dot{Q}_t = q_t$ and $Q_t \leq Q_T$. If $Q_T \leq Q^N$, the best thing to do is to make the last constraint binding everywhere,²⁰ 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 optimal policy after time T is to behave as in the SFP, and to adopt a path that is decreasing (see Proposition 1) for t above T . For $t < T$, because the stock level 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 SFP program when the initial stock value is Q . The left-derivative with respect to T of this expression is $\exp(-\delta T)$, times

$$Z_T \equiv (1 - F(Q_T)) \underbrace{(u(q_T, Q_T) - \delta W(Q_T) + q_T W'(Q_T))}_{=A} + f(q_T)q_T \underbrace{(V(Q_T) - W(Q_T))}_{=B}.$$

Now, by definition of W we have, for $0 < T' < T$,

$$\exp(-\delta T')W(Q_{T'}) \geq \int_{T'}^T u(q_t, Q_t) \exp(-\delta t) dt + \exp(-\delta T)W(Q_T),$$

and since the difference is zero at $T' = T$, its derivative wrt T' at the left of T is weakly negative, and we exactly obtain $A \leq 0$. Similarly, Assumption 2 states that $V(Q)$ is at

²⁰This is easily shown: this problem is autonomous, and consequently it admits a monotonic solution.

most $\frac{u(0,Q)}{\delta}$, which is the payoff from stabilizing the stock forever, and is thus below $W(Q)$. This shows $B \leq 0$. Finally, because Q is increasing at the left of T , we get $q_T \geq 0$, and therefore Z_T is weakly negative for every $T > 0$. This shows that one may as well apply the SFP solution from date zero onwards, so that once more we obtain a monotonic path.

Step 2: From Step 1, we easily obtain that a solution exists. Indeed, either the candidate path is weakly decreasing: then catastrophes cannot occur, p_t is a constant $(1 - F(Q_0))$ forever, and we are back to the SFP case with the additional constraint $q_t \leq 0$, for which existence of a solution is easily proven. Or the candidate path is weakly increasing, so that $\bar{Q}_t = Q_t$ everywhere, and $p_t = 1 - F(Q_t)$. The objective function becomes

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

to be maximized under the constraint $\dot{Q}_t = q_t \geq 0$, Q_0 given. This problem is autonomous, and once more our assumptions ensure the existence of a solution.²¹ Overall, a solution follows from the comparison of these two candidates.

In case (i) of the Proposition, suppose that the path Q_t is weakly increasing, so that $Q_t = \bar{Q}_t$ and $p_t = 1 - F(Q_t)$. We can then study the inequality (B.1) at $(t_1 = 0, t_2 = +\infty)$. The expression under the integral is $Q_t - Q_0$, which is positive, times

$$\dot{p}_t(D'(Q_0) - \nu(Q_0)) + \delta(1 - F(Q_0))(\nu(Q_0) - \rho(Q_0)D(Q_0)),$$

and both terms are negative, a contradiction. Therefore, in case (i) the path must be weakly decreasing, and by construction such a path involves no experimentation. The best path is thus the SFP path, and it converges to Q^N , as announced.

In cases (ii) and (iii), a weakly decreasing path would involve no experiment, and therefore would maximize $\int ued t$, with the additional constraint $q_t \leq 0$. But because $Q_0 < Q^N$, the solution to the SFP 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. The problem now consists of maximizing

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

²¹See Theorem 15, p. 237, in Seierstad and Sydsaeter (1987).

under the constraints $\dot{Q}_t = q_t \geq 0$, with an initial value $Q_0 < Q^N$. As explained above, a solution exists. The problem is autonomous, and we can proceed as in Proposition 1 to show that the optimal stock level 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^{E0} . This is possible if $Q^{E0} \geq \bar{Q}_0$ (case (ii.c)). Otherwise, the constraint $q \geq 0$ binds, and the stock remains forever set at Q_0 . ■