WORKING PAPERS

N° 1148

March 2025

“Catastrophes, delays, and learning”

Matti Liski and François Salanié

Toulouse School of Economics

Catastrophes, delays, and learning*

Matti Liski† François Salanié‡

March 25, 2025

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 feebacks 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 and Research Council Finland (project 357065)

‡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 are particularly helpful for our analysis. First, if the planner could learn the outcomes of experiments instantly, there would be no legacy. In that scenario, what would be the long-run stock level, denoted Q^E, at which the planner stops experimenting? Second, suppose instead that the legacy is one, meaning that a catastrophe is bound to occur: which stock level, denoted Q^D, should one aim at? The ordering of Q^E and Q^D then divides possible planning situations into two distinct classes.

Consider for example the management of a pandemic at its start, where the classical trade-off is between economic activity, typically associated with younger 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, as a first step, evaluating and ranking 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 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, as in path I' in the figure. Our first theorem rationalizes such fatalism from a set of well-founded primitives.

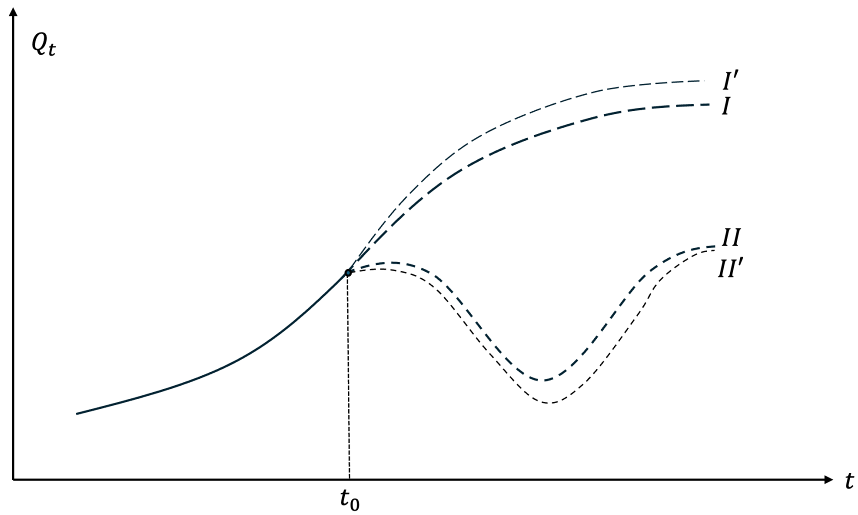

Figure 1

The second theorem applies when $Q^E > Q^D$, which occurs if the planner places a higher value on life compared to economic activity. A possible optimal policy is illustrated by path II in Figure 1. Under intuitive conditions, if the planner faces the same legacy as in the first theorem, she imposes an early, strict lockdown to reduce infections and mitigate the catastrophe's potential impact on the health system. During the lockdown, the legacy diminishes because no new experiments occur, and consequently, the planner becomes more optimistic over time.

We show that this first phase of reducing infections is optimally followed by a second phase that tolerates rising infections, eventually reaching or exceeding the level at which the lockdown began. Hence, the policy is non-monotonic: with the same preferences, a lockdown or higher infections can each be optimal, depending on how the current infection level was reached. Finally, and in contrast to the first theorem, a higher legacy now prompts more caution, leading to a stricter lockdown; however, the optimal asymptotic stock level remains unaffected by the legacy (see path II' in the figure).

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

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

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 our results with two stylized climate-change examples. First, climate-change targets are often expressed as "budgets" for total CO_2 emissions, but the "safe" budget is highly uncertain (van der Ploeg, 2018; IPCC, 2021). We model this unknown budget as a threshold for cumulative emissions beyond which a catastrophe is triggered. From this setup, several policy implications follow. If cutting current emissions does little to mitigate catastrophe damages, then the first theorem applies: policies are monotonic, and a higher total budget is allowed when the legacy is larger (for example, if emissions reached their current level quickly rather than gradually). Otherwise, the second theorem applies. In that case, for a sufficiently high legacy, the policy involves sharp early emissions reductions so that the budget is initially untouched, with actual usage deferred to later. These insights do not emerge from the existing literature.

Second, in a stylized setting, we show that classical climate-economy models balancing consumption against climate damages also fit into our theorems' dichotomy, in a manner similar to the disease-control application.

To put the main analysis into perspective, we consider two extensions that modify some of its key assumptions. First, we introduce a strategic interaction in which two agents share a common legacy, creating information spillovers. We show a "discouragement effect" in experimentation: once one agent stops, the other finds it optimal to stop as well. Second, we allow for "positive catastrophes" (innovations) to demonstrate how delays in observing the outcomes of past innovation efforts can lead to non-monotonic policies. A temporary pause avoids duplicating efforts, and if no results emerge during that pause, resuming later becomes optimal. In this mechanism pausing alters the legacy in ways that eventually encourage restarting. By contrast, when catastrophes are negative, non-monotonicity largely stems from inheriting an unfavourable history—a realistic scenario for pandemics and climate change.

Related literature. Our model ties together two canonical but distinct approaches to modelling catastrophes.[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. Thus, the catastrophe is bound to happen, while action can be taken to delay its occurrence and severity. But there is no memory of the past, and no learning over time. 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) use this approach, that we refer to as the *hazard-rate approach*.

In the second approach—the *unknown threshold 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, but the cost of surpassing it is borne privately by the firm that exceeds it. In contrast, in Diekert (2017) surpassing the threshold imposes a common cost on all agents. 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

[2]Catastrophes, 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.

[3]We discuss these contributions in Section 2.1.

[4]Crépin and Naevdal (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.

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 she 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 negative consequences from triggering a catastrophe are delayed, an effect which trivially supports more experimentation. But delays create a legacy of the past, with an ambiguous impact on experiments. Under the first theorem, a higher legacy encourages to experiment more, because the planner becomes more fatalistic; under the second theorem, the opposite result holds, because lower stocks values reduce the future damages from the catastrophe. This opposition also unifies the literature in a precise sense: the extreme informational assumptions of the literature define two stock-level targets whose comparison tells which one of the theorem applies.

To the best of our knowledge, the only paper that introduces delays in the unknown threshold approach is Guillouët and Martimort (2024), developed simultaneously with this paper. In their model, exceeding an unknown threshold increases the arrival rate of a catastrophe, the occurrence of which ends the game. As in our case, only the occurrence is observed, and beliefs have to be revised accordingly over time; and the optimal path thus depends not only on the present value of the stock, but also on the chronicle of past actions. Their paper provides a mathematical characterization of optimal paths, and then studies their decentralization between different selves of the same planner who may not necessarily observe past actions. They link this scenario to a Precautionary Principle. Our model is more general in several respects, and allows for a unification of previous approaches; notably, catastrophes are no longer inevitable, and we allow for damage mitigation before a catastrophe occurs.

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 an information structure that connects the legacy to past experiments.

Laiho, Murto and Salmi (2025) 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 payoff relevant stock level is reversible. Also, in our model, the chronicle of past actions is essential for revising beliefs; in their setting, the precise timing of past experiments does not matter.[5]

1 Model

The model defines a general framework that encompasses different applications. Consider, for example, the case of greenhouse gases and climate change. In each period, a planner chooses an emission flow, taking into account that emissions accumulate in a stock with harmful effects. Section 1.1 defines and studies this classical stock-flow trade-off. Section 1.2 introduces the possibility of a catastrophe, triggered when the stock exceeds a threshold value but occurs only after a delay, as in the Greenland ice sheet example mentioned at the beginning of the Introduction. Section 1.3 adds uncertainty on both the threshold and the delay. Given these components, 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:

$$\dot{Q}_t = q_t \in [\underline{q}, \overline{q}], \qquad Q_0 \text{ given.} \tag{1}$$

We assume $\underline{q} < 0 < \overline{q}$, so that the stock may increase or decrease over time. In the climate change example, this means that q are net emissions of greenhouse gases (net of decay or absorption by forests), while Q is the stock of CO_2 in the atmosphere. The

[5] 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.

planner's objective function at date zero is the following sum of payoffs, discounted at the rate $\delta > 0$:

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

The instantaneous utility function u thus captures the trade-off between higher emissions (associated to higher production and consumption) and a higher CO_2 stock (impacting utility or production, or both).

The Stock-Flow Problem (SFP) involves maximizing (2) under (1). It can be solved by assuming that the planner has chosen to stabilize the stock at a value Q, setting q permanently to zero. From this stabilized situation, we define the marginal payoff $\nu(Q)$ of additional emissions as the immediate benefit of increasing the flow q, plus the effects of increasing the stock Q, discounted over the entire horizon:[6]

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

Tsur and Zemel (2014) underline the role this function plays in dynamic settings, especially when $\nu(Q)$ is decreasing with the stock. Under this assumption, when the stock is low, the function is positive, encouraging accumulation. Conversely, when the stock is high, the function is negative, encouraging stock reduction. Stabilization at Q thus requires that $\nu(Q)$ be zero. This intuition motivates the following assumption:

Assumption 1 *The function u is twice continuously differentiable, bounded from above, and weakly concave in q. Moreover, for every Q we have:*

$$u_{QQ}(0, Q) \leq 0 \quad \textit{and} \quad u_{qQ}(0, Q) < 0$$

The first part of the assumption is common in the study of dynamic problems. The second part of the assumption ensures that the function ν is strictly decreasing with Q, for every positive value of δ. Hence the definition:

Definition 1 *Q^N (where N stands for "No catastrophe") is the stock level at which $\nu(Q)$ is zero. 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.*

Q^N is thus the long-run target in the absence of catastrophes. In line with the intuition sketched above, we obtain:

[6] Subscripts denote partial derivatives.

Proposition 1 *The Stock-Flow Problem (1)-(2) admits a solution whose path* $(Q_t)_{t\geq 0}$ *is monotonically converging to* Q^N.

The proofs of this result and all other results in this paper can be found in the Appendix. These proofs in fact only rely on the property that ν is weakly decreasing, and on the requirement that different thresholds (Q^N, Q^D, and Q^E – to be defined soon) are uniquely defined. This will allow us to handle other important cases. For example, if the planner cares linearly about consumption and stock, we have $u(q, Q) = u_0 + u_1 q + u_2 Q$, so that $\nu(Q)$ is a constant, and Q^N is uniquely defined as plus or minus infinity, depending on the sign of the constant. Another canonical example considers an agent with a revenue flow y, managing his wealth Q to smooth his consumption c over time. With an interest rate r, the budget constraint writes

$$\dot{Q} = rQ + y - c.$$

Thanks to the change of variable

$$q = rQ + y - c \qquad u(q, Q) = \mathcal{U}(c) = \mathcal{U}(rQ + y - q),$$

we obtain

$$\nu(Q) = \Big(\frac{r}{\delta} - 1\Big)\mathcal{U}'(rQ + y).$$

Then ν is indeed decreasing in Q whenever the utility $\mathcal{U}$ from consumption is concave and $r > \delta$. As in the linear model, Q^N is infinite, meaning that the planner would like to increase the stock without bound in the absence of catastrophes. Other examples include cases in which the level of the stock impacts utility directly, such as fishery management.

1.2 Catastrophes and delays

Catastrophes are irreversible and costly events, that are *triggered* when the stock exceeds a threshold value, but which *occur* only after a delay. To illustrate this key distinction, one may imagine a skater on thin ice. Instantaneous utility flow increases with the speed and/or the distance to the shore at a decreasing rate (Assumption 1), but the ice gets thinner and thinner. When the first crack in the ice appears (the triggering), the skater may turn back as long as the ice is still holding. When the ice finally breaks (the occurrence), the journey finds an abrupt end, and the damage to the skater depends on the remaining distance to the shore.

We assume 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:

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

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:

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

so that $T(S) < t$ if and only if $S < \overline{Q}_t$.

By assumption, the catastrophe itself *occurs* only after a delay $\tau \geq 0$ after the triggering, at date $\mathcal{T} = T(S) + \tau$. Note that, in contrast to the SFP, now the past trajectory of the stock is relevant at time 0, as the catastrophe may have been triggered in the past without occurring yet. After time $\mathcal{T}$, the catastrophe occurs, the game ends, and the planner receives a continuation payoff $V(Q_{\mathcal{T}})$, which depends on the value of the stock at the catastrophe date $\mathcal{T}$.[7] 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 (think to the skater turning back to the shore). Leaving V instead dependent on the threshold S, or on the maximum level tried in the past $\overline{Q}_{\mathcal{T}}$, would eliminate this possibility by assumption.

To put more structure on payoffs, let us proceed to a natural comparison. At any point in time, if the planner stabilizes the stock Q she obtains $u(0, Q)$ forever, while if instead she experiences the catastrophe, her continuation payoff is $V(Q)$. The following assumption orders these two payoffs:

Assumption 2 *The function $V(Q)$ is twice continuously differentiable and weakly concave in Q. Moreover, for every Q one has*

$$u(0, Q) \geq \delta V(Q) \qquad u_Q(0, Q) \geq \delta V'(Q).$$

Essentially, we assume that a catastrophe reduces utility compared to stabilization and that a catastrophe is more costly when the stock is higher. For further reference, we define the damage function D as follows:

$$D(Q) \equiv \frac{1}{\delta} u(0, Q) - V(Q). \tag{4}$$

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

Assumption 2 thus states that the damage is weakly positive and weakly increasing with respect to the stock value at the date of the catastrophe. Overall, given S, τ, and a path $(Q_t)_{t\in(-\infty,+\infty)}$, one can compute $\mathcal{T} = T(S) + \tau$ from (3), and the planner's payoff from date $t = 0$ onward equals

$$\int_0^{\mathcal{T}} u(q_t, Q_t)\exp(-\delta t)dt + \exp(-\delta\mathcal{T})V(Q_{\mathcal{T}}).$$

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 $[\underline{S}, \overline{S}]$. 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 that it follows an exponential distribution with parameter $\alpha > 0$, with the cumulative distribution function $1 - \exp(-\alpha\tau)$. In particular, τ and S are independent variables. These assumptions are clearly made for tractability, and we will underline their consequences below.

A key assumption is that the planner does not observe the triggering of a catastrophe: she only observes its occurrence. This allows us to capture the idea that a catastrophe might well be underway, although the planner does not know exactly. These uncertainties are often invoked in biology, under the name of the extinction debt (Tilman et al., 1994).

Hence, the only hard information the planner may learn is that a catastrophe has occurred – but at that point the game ends. Thus, the planner's policy concerns actions taken before the game ends.

1.4 The planner's problem

We are now in a position to state the planner's problem. Recall that the prior beliefs characterized by F are given at the beginning of times ($t = -\infty$). By contrast, at the planning date ($t = 0$), the planner inherits the past trajectory of the stock $(Q_t)_{t\leq 0}$, and she also knows that the catastrophe was not triggered in the past, or was triggered but did not happen yet: equivalently, $\mathcal{T} \geq 0$. Therefore, the planner's problem is to find a policy $(q_t, Q_t)_{t>0}$ that maximizes

$$\mathbb{E}\left[\int_0^{\mathcal{T}} u(q_t, Q_t)\exp(-\delta t)dt + \exp(-\delta \mathcal{T})V(Q_{\mathcal{T}}) \,\middle|\, \mathcal{T} \geq 0, (Q_t)_{t\leq 0}\right] \tag{5}$$

subject to (1). While Q_t is continuous by construction, we only require q_t to be piecewise-continuous. We say that a path $(Q_t)_{t\geq 0}$ is monotonic if Q_t is everywhere weakly decreasing, or everywhere weakly increasing, with respect to time. Moreover, we define $\overline{Q}_\infty \leq +\infty$ as the supremum value for the stock. We say that $\overline{Q}_\infty$ is reached in finite time if there exists $T < +\infty$ such that $Q_T = \overline{Q}_\infty$. Otherwise, we say that $\overline{Q}_\infty$ is reached asymptotically, and in this case one has $Q_t \leq \overline{Q}_t < \overline{Q}_\infty$ for all t.

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 $\mathcal{T}$, computed at the beginning of times using the prior beliefs F:

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

To characterize this probability, one may distinguish two possibilities for survival at time t. Either S is above $\overline{Q}_t$, so that no catastrophe could have been triggered before time t, and survival is certain. Or S is below $\overline{Q}_t$, 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

$$p_t = 1 - F(\overline{Q}_t) + \int_{S<\overline{Q}_t} \exp[-\alpha(t - T(S))]dF(S). \tag{6}$$

Hence, the survival probability at time t exceeds $1 - F(\overline{Q}_t)$, 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 2 *For a given path, the legacy of the past at date t is*

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

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

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

The existence of a 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 $1 - F(\overline{Q}_t)$, and π_t is 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 $S < \overline{Q}_t$. 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. By applying (6) at $t = 0$, we get the information content of the data $(Q_t)_{t\le 0}$ relevant for planning:

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

Moreover, by differentiating (6), we obtain a law of motion:

$$\dot{p}_t = \alpha[1 - F(\overline{Q}_t) - p_t] \tag{7}$$

in which $\overline{Q}_t$ is the highest stock on record:

$$\overline{Q}_t = \max(\max_{0\le t'\le t} Q_{t'}, \overline{Q}_0). \tag{8}$$

We are now in a position to rewrite the planner's problem defined in (5). The cumulative distribution function of the occurrence date T is simply $1 - p_{\mathcal{T}}$, and thus the event $\mathcal{T} \ge 0$ has probability p_0. The payoff in (5) becomes

$$\mathbb{E}\left[\int_{t\ge 0} \mathbb{1}_{\mathcal{T}\ge t} u(q_t, Q_t) \exp(-\delta t) dt + \exp(-\delta \mathcal{T}) V(Q_{\mathcal{T}}) \;\middle|\; \mathcal{T} \ge 0, (Q_t)_{t\le 0}\right]$$

$$= \int_{t\ge 0} \frac{p_t}{p_0} u(q_t, Q_t) \exp(-\delta t) dt + \int_{\mathcal{T}\ge 0} \exp(-\delta \mathcal{T}) V(Q_{\mathcal{T}}) \frac{1}{p_0} d(1 - p_{\mathcal{T}}).$$

By leaving out the constant term p_0, and by relabelling $\mathcal{T}$ into t in the second integral, we obtain that the optimal policy maximizes

$$\int_0^\infty [p_t u(q_t, Q_t) - \dot{p}_t V(Q_t)] \exp(-\delta t) dt, \tag{9}$$

subject to (1), (7), (8), and given Q_0, $\overline{Q}_0$ and p_0.

Because time appears only through exponential discounting, the problem is *autonomous*: its evolution depends solely on the state rather than on time. Three variables suffice to describe the state: the stock Q, the maximum historical stock $\overline{Q}$, and the survival probability p. Under the assumption that the waiting time for a catastrophe triggered is exponentially distributed, the initial conditions $\left(Q_0, \overline{Q}_0, p_0\right)$ fully summarize the past trajectory $\left(Q_t\right)_{t\leq 0}$. Equivalently, one may replace p_0 with initial legacy

$$\pi_0 \;=\; 1 \;-\; \frac{1 - F\left(\overline{Q}_0\right)}{p_0},$$

which represents the probability (conditional on survival) that a catastrophe was triggered in the past. Intuitively, $\overline{Q}_0$ reflects the *extent* of past experimentation, while π_0 captures its *timing*. Together, these two variables encapsulate all relevant historical data.

2 Optimal policies

Characterizing the optimal policies is not a simple task, as the problem involves three state variables—one of which is a record process—and allows for nonparametric functions in both the payoffs and the belief distribution. Methods from optimal control theory or the calculus of variations that aim to derive policies from first-order conditions do not readily apply. Standard continuity and convexity requirements for the constraints fail to hold for the general problem, because of the presence of a record process in (8). Consequently, we establish the existence of an optimum and characterize the optimal policies only after identifying qualitative properties of candidate paths and showing that these properties follow from intuitive conditions.

To develop the conditions needed, we next introduce important benchmarks results and their connections to the literature.

2.1 Benchmarks

The case of a past triggering (the hazard-rate approach): Assume that the planner knows the catastrophe was triggered in the past, but has not yet occurred, so that the legacy of the past is permanently set to one. Suppose, moreover, that the stock is stabilized permanently at some value Q, conjectured to be optimal. Thanks to the assumption of an independent Poisson process for the delay, the benefit from a small temporary increase in the flow is easily seen to be equal to:[8]

$$\nu(Q) - \frac{\alpha}{\alpha+\delta} D'(Q). \tag{10}$$

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

Definition 3 *Q^D (where D stands for "Damages") is the stock level at which (10) is zero. By convention, we set $Q^D = +\infty$ if (10) is positive for all Q, and $Q^D = -\infty$ if (10) is negative for all Q.*

This situation refers to the case where $\overline{Q}_0 \ge \overline{S}$ at the planning date 0, so the planner knows that the catastrophe has already been triggered. The law of motion for the survival probability (7) reduces to:

$$p_t = p_0 \exp(-\alpha t). \tag{11}$$

A 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 these 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:

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

Comparing with (11), we see that these 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.

We can now provide a general result illustrating the importance of the threshold value Q^D:

[8] To prove this formula, let 0 be the present date. Recall that by assumption the catastrophe was triggered in the past, at some unknown date $T(S) < 0$, but did not occur yet. Therefore, the additional damages $D'(Q)$ from the future catastrophe must be discounted by

$$E[\exp(-\delta(T(S)+\tau)) | T(S)+\tau > 0].$$

Using the assumption that τ is distributed exponentially with parameter α, we then obtain directly that this value equals $\frac{\alpha}{\alpha+\delta}$. In particular, it does not depend on the distribution of $T(S)$.

Proposition 2 *Suppose $\overline{Q}_0 \geq \overline{S}$. Then for every initial value $Q_0 \leq \overline{Q}_0$, there exists an optimal policy, which moreover is such that the path $(Q_t)_{t\geq 0}$ converges monotonically to the value Q^D.*

Hence, Q^D can be interpreted as the long-run target when one knows that the catastrophe was triggered in the past.

The case of no past triggering (the unknown threshold approach): Tsur and Zemel (1994, 1995, 1996), and more recently Lemoine and Traeger (2014), Diekert (2017), and Chen (2020) all use an unknown threshold approach in which a catastrophe occurs as soon as the threshold is reached, so that there is no delay between triggering and occurrence. In our model, this corresponds to the case when α goes to infinity. Then the law of motion for the survival probability (7) reduces to

$$p_t = 1 - F(\overline{Q}_t),$$

and the legacy of the past is zero at every date. To study this simplified model, assume that the planner has stabilized the stock at some level Q. By experimenting a bit more, the planner would increase her payoff by the following quantity:

$$\nu(Q) - \rho(Q)D(Q). \tag{12}$$

Indeed, the first term is the gain in the absence of catastrophes, while the second term measures the risk that the catastrophe be triggered and occurs immediately. The following result illustrates the role of this expression:

Proposition 3 *Suppose $Q_0 = \overline{Q}_0$. Suppose also that there exists a value $Q^{E0} \in [\underline{S}, \overline{S}]$ such that (12) is zero. In the absence of delay ($\alpha = +\infty$), there exists an optimal path $(Q_t)_{t\geq 0}$, and it is:*

(i) decreasing and converging to Q^N, if $Q_0 > Q^N$;

(ii) constant and equal to Q_0, if $Q_0 \in [Q^{E0}, Q^N]$;

(iii) increasing and converging to Q^{E0}, if $Q_0 < Q^{E0}$.

This 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 $Q_0 = \overline{Q}_0$ is made for simplicity). The striking part is that 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 in the absence of delays.[9]

Our model involves delays; therefore, we reintroduce them by assuming $\alpha < \infty$, while continuing to work under the hypothetical assumption that it is known the catastrophe has not been triggered in the past. In such a situation, one may safely stabilize the stock by playing $q = 0$ forever. One may also experiment by increasing the stock 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:[10]

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

In the case with a past triggering, one was worried about aggravating a catastrophe that was already underway. Now, one is worried about triggering a catastrophe with some probability measured by the hazard rate ρ: hence the difference between (10) and (13). Under our assumptions, expression (13) is weakly decreasing in Q and lies below $\nu(Q)$. Therefore, it may reach zero only at a value $Q^E \leq Q^N$, and once more this value is uniquely defined as follows:

Definition 4 *Q^E (where E stands for "Experimentation") is the stock level at which (13) is zero. By convention, we set $Q^E = \underline{S}$ if (13) is negative at $\underline{S}$, and $Q^E = \overline{S}$ if (13) is positive at $\overline{S}$.*

This threshold value will play an important role in our analysis of the general problem. The above reasoning proves that one should not stabilize the stock below Q^E, and we state it explicitly here for future reference:

Proposition 4 *Suppose that it is optimal to stabilize the stock at some level Q_∞. Then $Q_\infty \geq Q^E$.*

[9] This 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."

[10] See footnote 8.

To summarize: So far, we have defined three unique long-run targets:

- Q^N: target in the absence of a catastrophe;
- Q^D: target when it is known that the catastrophe was triggered in the past;
- Q^E: stock level below which stabilization should not occur.

We also know that the last two values lie below Q^N. As we will see in the next section, the ranking between Q^D and Q^E is key to our main theorems. The symmetry in equations (10)–(13), together with our monotonicity assumptions, makes it straightforward to find conditions for the ranking. For example, we have:

Lemma 1 *If the function* $(\frac{1}{\delta}u(0,Q) - V(Q))(1 - F(Q))$ *increases (resp. decreases) with* Q *at* $Q = Q^D$*, then* $Q^D < Q^E$ *(resp.* $Q^D > Q^E$*).*

Two polar cases come to mind. If a catastrophe reduces the stabilization value of the stock, $u(0,Q)$, by a fixed amount, then the damage $D(Q) = \frac{1}{\delta}u(0,Q) - V(Q)$ is constant. In such a case, modifying the value of the stock is of no help if one wants to reduce damages, and from (3) we obtain $Q^D = Q^N \geq Q^E$. Conversely, if the damage increases sharply with the stock level at the time the catastrophe occurs, then it becomes highly valuable to reduce the stock; in this case, Q^D is small, and lies below Q^E.

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

The first theorem applies when

$$\overline{Q}_0 < Q^E < Q^D < \min(Q^N, \overline{S}). \tag{14}$$

In words, at the initial date, experimentation has barely begun, so the initial stock is low. The situation is one in which the damage does not depend too much on the value of the stock when the catastrophe occurs: mitigation strategies are not very effective. The inequalities in (14) lead to our first theorem, with the following sequence of arguments.

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 still optimally increase the stock toward Q^D. Lemma C.1 in the Appendix extends this result 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 small, and then one should not experiment any further if one is already above Q^E, or the legacy is high, and then one should optimally come close to the long-run target Q^D (see 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 $Q_T \le Q^D$ at some date $T \le +\infty$. Because the path is monotonic, the record-process plays no role, and existence of optimal paths is 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?[11]

Theorem 1 (Case $Q^E < Q^D$) *Suppose* (14) *holds. Then there exists an optimal policy. Under this policy, the path $\left(Q_t\right)_{t\ge 0}$ is weakly increasing and converges to $\overline{Q}_\infty \in [Q^E,\, Q^D]$, reached at some (possibly infinite) time T. Moreover:*

1. *If $\overline{Q}_\infty$ is reached only asymptotically (i.e. $T = +\infty$), then necessarily $\overline{Q}_\infty = Q^E$.*

2. *In every case (whether T is finite or infinite), one has*

$$\nu(Q_T) \; = \; \frac{\alpha}{\alpha+\delta}\Big[\big(1-\pi_T\big)\,\rho(Q_T)\,D(Q_T) \; + \; \pi_T\,D'(Q_T)\Big]. \tag{15}$$

Finally, condition (15) *implies that a higher Q_T is associated with a higher π_T.*

Condition (15) nicely consolidates the conditions supporting the definitions of Q^D and Q^E, with weights determined by the legacy at the time when the experimentation stops. Notice that when delays are infinite (α vanishes), we are back to the no-catastrophe case, for which $Q^E = Q^D = Q^N$, and to the optimal path characterized in Proposition 1. Similarly, in the absence of delays ($\alpha = +\infty$) the legacy is identically zero, and therefore we confirm the result in Proposition 3 that the stock converges to the value Q^{E0} from a low initial level.

It is also remarkable that a higher legacy π_T is associated with a higher long-run value of the stock, that is, more experimentation in total. Indeed, immediate consumption becomes more of a priority when it is more likely that a catastrophe was triggered in

[11] The proof of Theorem 1 is in Appendix F, and the dynamic programming interpretation in Appendix I.1.

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 Q^E, towards Q^D.

However, whether a higher initial legacy π_0 leads to more experimentation in total requires global comparative statics, involving variations in the entire policy path from the initial date to the conclusion of experimentation. We provide such an analysis in Proposition 5 of Section 3.2 for a simple model of climate policies, confirming that a higher initial π_0 indeed leads to more experimentation in total for an explicit optimal policy solution. The disease control example further supports this result through simulations.[12]

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.

Theorem 1 assumes that $\overline{Q}_0$ is low enough for condition (14) to hold. In particular, this condition ensures the monotonicity of Q_t and guarantees that condition (15) holds at Q_T. These properties are essential for the characterization.[13]

2.3 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 enough to the stock level to imply

$$Q^D < Q^E < \min(Q^N, \overline{S}). \tag{16}$$

In this situation, a striking result is that the long-run target for the stock can be easily computed. Indeed, if the stock remains below $\overline{S}$, then the legacy of the past must vanish in the long run. This implies that stabilizing below Q^E is suboptimal, since further experimentation would still be valuable. Conversely, additional experimentation above Q^E is suboptimal: when the legacy is zero, this follows directly from the definition of

[12] Comparative statics with respect to α also require global analysis for the same reasons. Additionally, the variations of α and π_0 are linked: the value of α impacts the computation of π_0 from historical data $(Q_t)_{t\leq 0}$.

[13] In contrast, if $\overline{Q}_0 > Q^D > Q^E$, a number of cases can arise. For example, the catastrophe may be triggered with certainty, but establishing such a result would require a more detailed specification of the model.

Q^E, and when the legacy is high, one should instead aim for a lower target, closer to Q^D. Building on these intuitions, we obtain:

Theorem 2 (Case $Q^D < Q^E$) *Suppose* (16) *holds, and let* $\left(Q_t\right)_{t\geq 0}$ *be an optimal path. If*

$$\overline{Q}_\infty \; < \; \min\left(Q^N, \, \overline{S}\right),$$

then:

1. $\left(Q_t\right)_{t\geq 0}$ *converges to* $\overline{Q}_\infty$.
2. $\overline{Q}_\infty = \max\left(\overline{Q}_0, \, Q^E\right)$.
3. *If* $\overline{Q}_0 > Q^E$, *then convergence occurs in finite time.*

A key implication of the theorem is that the optimal path converges to a steady state, a result that is generally not guaranteed when more than two state variables are involved (see, e.g., Benhabib and Nishimura, 1979). The interpretation is especially clear if one starts at a low level of the stock: in that situation, any optimal path that remains below the no-catastrophe target Q^N—and that does not trigger a catastrophe with certainty—must converge to the unique value of the stock for which further experimentation has no marginal value. Notably, in contrast with the scenario in the previous theorem, this long-run target does not depend on the initial legacy.

On the other hand, despite this convergence, the path need not be monotonic in the short run. Indeed, Lemma C.1 (in the Appendix) shows that an optimal path can decrease at some date if the legacy of the past lies above a certain threshold at that date. The applications below demonstrate that such non-monotonic but transitory paths can arise—even though ultimately, under the conditions considered, the trajectory still settles to its long-run target. Intuitively, in these illustrations the planner inherits a "bad history" (i.e., a large Q_0); economically, such transitory non-monotonicity is easy to understand from a damage-mitigation perspective.[14]

[14] By contrast, our extension to positive catastrophes shows that, under a different set of assumptions, permanently ceasing experimentation conflicts with the long-run evolution of the legacy.

3 Applications

3.1 Disease control and social distancing

We now present a simple model of a pandemic that incorporates the trade-off between social distancing and economic activity—a common theme in the literature (see, e.g., Bloom, Kuhn and Prettner (2022) for a review). Additionally, our model accounts for the possibility of a breakdown in the health system or even the entire economy. This catastrophic risk, which is novel in the literature, generates a rich set of predictions, as we now demonstrate.

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:

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

The recovery rate r and the death rate d are positive parameters. Variable $R_t \in [0, \overline{R}]$ measures new infections, with maximum value $\overline{R} > r + d$ 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 with 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

$$Y(R) - wdI.$$

This model is a special case of our general framework, with the formulas applying directly under the transformation $Q = \log(I)$ (see Appendix G.1). When we assume catastrophes are ruled out, balancing the benefits and costs from increasing the stock of infected agents leads to the long-run target

$$I^N = \frac{\delta\, Y'(r+d)}{w\, d}. \tag{17}$$

The policy target I^N varies intuitively with the model's parameters and can be reached over time by a social distancing policy satisfying $R > r + d$ if and only if $I_0 < I^N$.

However, planning in a pandemic 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. When a catastrophe occurs, the planner loses control: the matching rate takes an exogenous value R^*, and output remains fixed at a low level $Y^* < Y(r+d)$. The death rate increases to $d^* > d$, the recovery rate becomes r^*, and the resulting rate of increase q^* of infected agents is assumed to satisfy the following inequalities:[15]

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

If the catastrophe occurs at time $\mathcal{T}$, at infections level $I_{\mathcal{T}}$, the discounted value of the continuation payoff can be readily obtained as $\frac{Y^*}{\delta} - \frac{wd^*}{\delta - q^*} I_{\mathcal{T}}$. The damage from the event is then the discounted sum of production losses and the value of the mortality increases:

$$\frac{Y(r+d) - Y^*}{\delta} + w\mu^* \frac{d}{\delta} I, \qquad\qquad \mu^* \equiv \frac{\frac{d^*}{\delta - q^*} - \frac{d}{\delta}}{\frac{d}{\delta}} > 0. \tag{18}$$

where the parameter μ^* measures the increase in mortality. Having established this damage, we can now see how it affects the planning target in comparison to the no-catastrophe target, I^N. When the planner is certain the catastrophe will arrive but it has not yet done so, the trade-offs familiar from the general model lead to

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

which is the infection level targeted under certainty of a future catastrophe that has not yet occurred. One rationally braces for the catastrophe by reducing infections below the no-catastrophe target I^N, and does so more drastically the larger the change in mortality measured by μ^*.

The target I^E applies when the catastrophe is not deemed inevitable. Its expression is more involved and is therefore omitted here. However, we obtain a conclusive result:

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

$$\frac{1}{1 + \frac{Y(r+d) - Y^*}{w\mu^* d I^D}} < \rho(I^D), \tag{19}$$

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

[15] The last inequality avoids infinite values for the discounted welfare cost of deaths. Alternatively, one could assume 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 using a full S-I-R model instead of a simple exponential.

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

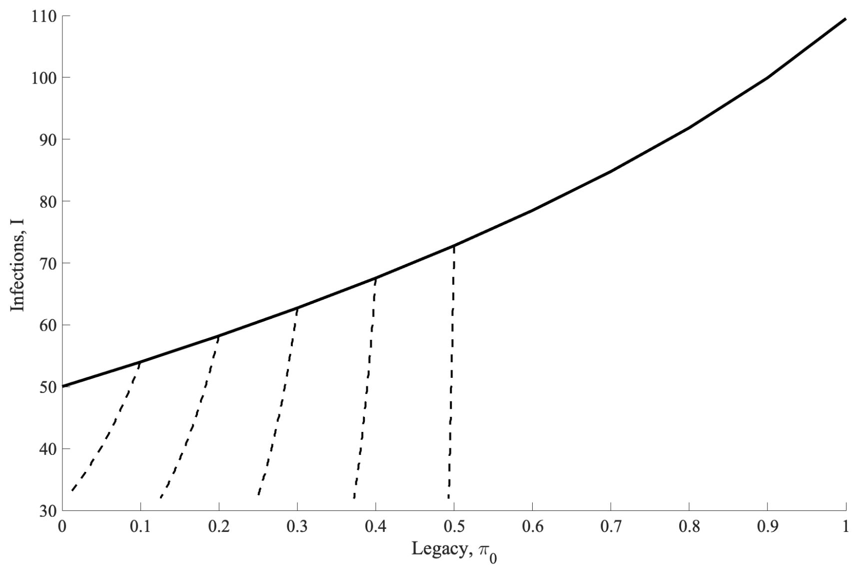

Optimal paths in the plane (π, I) for a linear production function $Y(R) = Y_0 R$. Parameters are: $\delta = 0.03, q^* = 0.01, w = 1, d = 0.1, r = 0.9, d^* = 0.2, \alpha = 0.2, Y_0 = 1000, Y^* = 950$, $I_0 = 32$, $\bar{q} = 1$. Distribution F for $\log(I)$ is uniform: $f = 1/6$. Benchmark values are $I^N = 300$, $I^D = 110$, and $I^E = 50$.

Theorem 1 applied: Consider first the case of a planner who prioritizes economic activity over deaths, in the sense that condition (19) 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. (15), 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.[16] Each dotted curve depicts this relationship, for varied initial legacies, but with

[16]The problem is linear in q, and, by standard 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 boundary value problem

the same initial infection level set at $I_0 = 32$. One observes that a higher initial legacy leads to a higher long-run value for the stock.[17] 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 (19) is reversed. For this illustration, assume that planning starts so late that 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.[18]

Figure 3

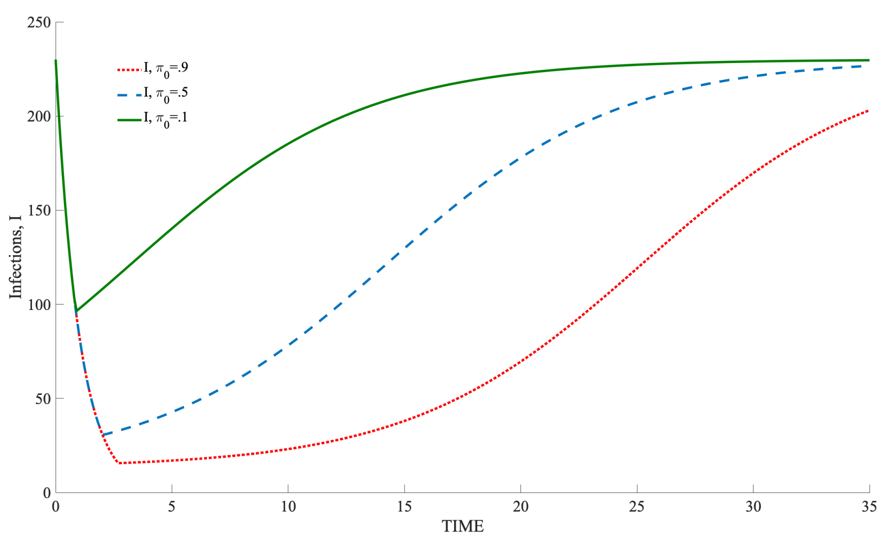

The population of infected agents over time, for a linear production function. Parameters are: $\delta = 0.023, q^* = 0.02, w = 1, d = 0.1, r = 0.98, d^* = 0.25, \alpha = 0.2, Y_0 = 1000, I_0 = 230$.

numerically to reach the stopping condition from given (π_0, I_0).

[17] Note 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.

[18] We compute the solution path in Appendix G.1.

Figure 3 depicts the optimal time path of the stock of infected agents for different 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

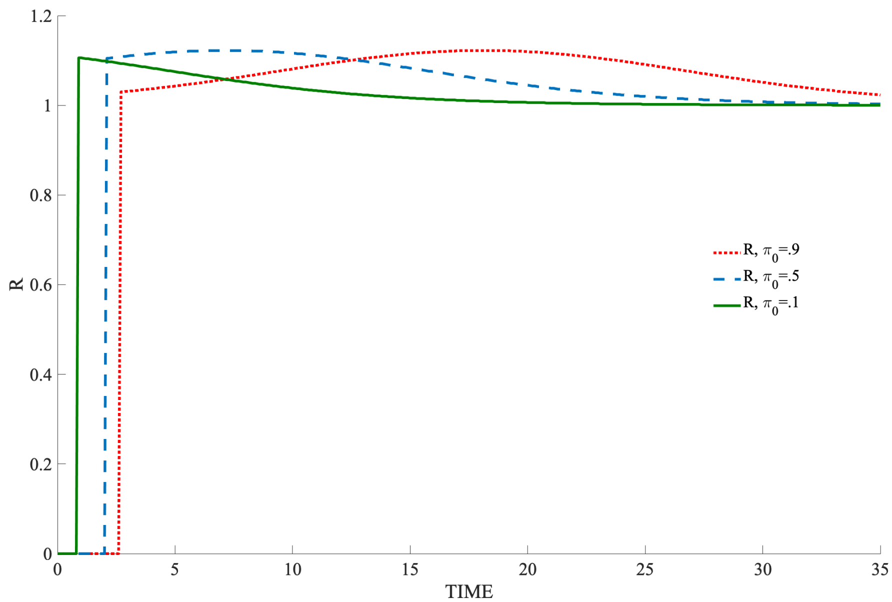

The optimal control, for a linear production function. Parameters are: $\delta = 0.023, q^* = 0.02, w = 1, d = 0.1, r = 0.98, d^* = 0.25, \alpha = 0.2, Y_0 = 1000, I_0 = 230$.

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 in the previous case, 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).

3.2 Climate change

3.2.1 Optimal carbon budget

Studies of climate change often mentions a safe carbon budget whose value is uncertain (van der Ploeg, 2018) and should not be exceeded, lest a catastrophe be triggered. Formally, this problem relates to a seminal work by Kemp (1976) who studies a cake-eating problem in which the size of the cake is initially unknown. We extend this model by incorporating a delay between the triggering and the occurrence of a catastrophe, during which it is not known if the safe budget has been exceeded. 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 chooses a net consumption $q_t \in [\underline{q}, \overline{q}]$ and receives an instantaneous payoff $u_0 + u_1 q_t$, where u_0 is the non-use value of the climate as a resource and u_1 is the value of a unit of consumption. The catastrophe is triggered when the cumulative consumption Q_t exceeds an unknown threshold. After the catastrophe occurs, the planner obtains a continuation payoff $-v_0 Q$, where Q is the cumulative consumption at the occurrence time. In terms of our general framework, the primitives of this problem are

$$u(q, Q) = u_0 + u_1 q, \quad V(Q) = -v_0 Q, \quad \nu(Q) = u_1, \quad D(Q) = \frac{u_0}{\delta} + v_0 Q,$$

with $u_1 > 0, u_0, v_0 \geq 0$. In contrast to the disease control model, an infinite carbon budget is optimal in the absence of catastrophes: $Q^N = +\infty$. On the other hand, if the catastrophe was triggered with certainty in the past, the relevant target budget is Q^D. It is straightforward to see that Q^D equals $+\infty$ or $-\infty$, depending on whether

$$u_1 - \frac{\alpha}{\alpha + \delta}\, v_0$$

is positive or negative.[19] Intuitively, both the marginal gains and expected losses from consumption are constant, so the planner either reaps as much consumption as possible or mitigates damages as much as possible before the catastrophe occurs. Finally, when it is known that the catastrophe has not been triggered at all, then the optimal carbon budget is Q^E, implicitly defined by

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

[19]For simplicity, we ignore the natural constraint $Q \geq 0$.

provided such a value belongs to the support of S (see Definition 4). We now distinguish two cases.

Theorem 1 applied: Assume $u_1 > \frac{\alpha}{\alpha+\delta} v_0$. This implies $Q^E < Q^D = +\infty$, so that Theorem 1 applies. The optimal policies are weakly increasing, as stated in the theorem, and strongly depend on the legacy π_0 at which the planning begins.

Proposition 5 ***(Carbon budget I)*** *Let $u_1 > \frac{\alpha}{\alpha+\delta} v_0$ and $Q_0 = \overline{Q}_0 < Q^E$. Then there exists a critical legacy π^* such that:*

(i) If the initial legacy π_0 is below π^*, the optimal policy is to consume maximally, $q_t = \overline{q}$, until reaching a finite date T, and stop thereafter, $q_t = 0$. The optimal carbon budget is such that $Q^E < Q_T < Q^D$.

(ii) If the initial legacy π_0 is above π^*, the optimal carbon budget is unbounded: the period of maximal consumption T extends to infinity, triggering the catastrophe with certainty.

(iii) The stopping date ($T \in [0, +\infty]$) and the optimal budget Q_T are nondecreasing functions of the initial legacy π_0.

With a low consumption in the past, one is confident that the budget has not been exceeded yet, and this makes it worth being cautious and to avoid experimentation. Conversely, after a high past consumption, one expects the consumption opportunities 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: Assume now $u_1 < \frac{\alpha}{\alpha+\delta} v_0$, so that the benchmark carbon budgets are ranked as $Q^D < Q^E$. For a stark illustration, suppose further that we start planning after intensive experimentation in the past: the stock has already exceeded the benchmark budgets when the planning starts at $t = 0$. Formally, we assume:

$$Q^E < Q_0 = \overline{Q}_0 < \min(Q^N, \overline{S}) \qquad u_1 < \frac{\alpha}{\alpha + \delta} v_0. \tag{20}$$

Proposition 6 ***(Carbon budget II)*** *Assume (20) 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* $t < t_1$*;*
- $q_t = \overline{q} > 0$ *for* $t_1 < t < t_2$*;*
- $q_t = 0$ *and* $Q_t = \overline{Q}_0$ *for* $t > t_2$*.*

Thus, in both situations the optimal carbon budget is Q_0.

This result thus proves formally that optimal policies can be non-monotonic. It is interesting also to compare to Proposition 5: 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 Q_0 follows.

3.2.2 Stock-flow trade-offs in climate change

Considering the carbon budget as a resource with uncertain size provides new insights, yet this perspective does not neatly align with climate-economy models that analyse the trade-offs between consumption and gradually accruing damages, as well as potential tipping points.[20] To study these elements under delays, we consider a toy model for climate change, inspired by Golosov et al. (2014). Consider a pollution stock Q_t that follows a simple law of motion:

$$\dot{Q}_t = E_t - \gamma Q_t, \tag{21}$$

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

$$Y_t = \exp(-\theta Q) K^{1-\beta} E_t^{\beta} \tag{22}$$

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 $U(C) = \log C$.

[20] It is noteworthy that Nordhaus' seminal contributions initially focused on setting a carbon budget; only later developments incorporated damages from climate-economy interactions (Nordhaus, 1975).

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

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

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

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

The target Q^N increases in the abatement cost β, in the decay rate 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, 2013). 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 in Cai and Lontzek (2019), when occurring the catastrophe irreversibly changes the production possibility frontier. We may capture this impact by increasing θ by a factor $k > 1$, and we assume that this shock is important enough:

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

This simple setting highlights the basic conceptual differences in the main two approaches in the literature. In the hazard rate 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. By contrast, in the unknown threshold approach, there is no legacy from the past because there is no delay between triggering and occurrence. Without 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). Then the relevant target is Q^E.

Proposition 7 *In the toy model of climate change, it holds that $Q^E < Q^D$ if and only if*

$$\frac{\alpha+\delta}{\alpha+\delta \ + \ \alpha\,\rho(Q^E)\,g\big(Q^E\big)} \quad < \quad \frac{\gamma+\delta+\alpha}{\gamma+\delta+k\,\alpha}.$$

where function g is defined as

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

By comparing the above equations, one obtains that Q^E is below Q^D if and only if the hazard rate ρ 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 are relevant for the optimal policies. However, this is only the first step in planning. The planner must also assess the legacy of past experiments and their information content, which suggests an agenda for applied quantitative research on optimal climate policies within detailed climate-economy models. These models can quantify the information content of past (unplanned) experiments to provide a structural interpretation of 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.

4 Extensions

In this Section, we extend the model in two directions. The first extension studies a climate change game with two players, and shows how commitment by one player can influence the other to reach efficient outcomes. The second extension proposes to see research innovations as positive catastrophes, and discusses the optimality of research cycles.

4.1 Strategic interactions

Consider a game in which two countries $n = i, j$ face the risk of a catastrophic climate change triggered when their aggregate net emissions exceeds an unknown threshold, and occurring after a stochastic delay. Players share common prior beliefs on the threshold S (with a cumulative distribution function $F(S) = 1 - \exp(-\rho S)$) and on the delay τ (Poisson distribution with parameter $\alpha > 0$). Both players observe the full history of actions. For simplicity, we endow players with similar preferences as in the carbon budget application of Section 3.2, and we only allow for two actions, either positive business-as-usual net emissions, or zero net emissions:

$$u(q, Q) = u_1 q,\ u_1 > 0,\ \text{with } q = 0 \text{ or } q = \overline{q} > 0.$$

Moreover, the switch from $\overline{q} > 0$ to zero is irreversible and must be decided at date 0, as it involves significant investments. In this simplified model of climate change, each player only has to decide when to stop. The occurrence of a catastrophe ends the game, yielding each player a continuation payoff of $-v_0 < 0$. Therefore, damages are independent of the stock, making players fatalistic if it is likely that a catastrophe was triggered in the past.

Consider the case when aggregate emissions q are fixed forever. Then it is easily shown that the rate of change of the legacy of the past is given by

$$\dot{\pi} = (1-\pi)(\rho q - \alpha\pi).$$

We assume $\rho\overline{q} < \alpha < 2\rho\overline{q}$. Business-as-usual emissions by both countries are thus high enough to make the past triggering of a catastrophe more and more likely. By contrast, if only one player emits, then the legacy of the past converges to a value $\pi^{\infty} \equiv \rho\overline{q}/\alpha$ strictly between 0 and 1. We also assume that the damage v_0 is high enough.[21]

Now, suppose that player j has decided to emit forever: $T_j = +\infty$. We show in Appendix H.1 that player i's best response depends on the initial legacy: if it is high enough, precisely if

$$\pi_0 > \pi_i^* \equiv 1 - \frac{u_1}{v_0}\frac{\alpha+\delta}{\alpha\rho}\frac{\rho\overline{q}+\delta}{\delta},$$

then the player becomes fatalistic, and its best response is also to pollute forever, i.e. $T_i = +\infty$.

Conversely, suppose that player j has decided to stop immediately: $T_j = 0$. Once more, the best response of player i is characterized by a threshold, and the Appendix shows that player i chooses $T_i = 0$ if

$$\pi_0 < \pi_i^{**} \equiv 1 - \frac{u_1}{v_0}\left(\frac{\alpha+\delta}{\alpha\rho}\right).$$

Note that this second threshold is higher than the first one, which itself is below one under our assumptions. We obtain:

Proposition 8 *Suppose the initial legacy takes an intermediate value: $\pi^* < \pi_0 < \pi^{**}$. Then both $(T_i = 0, T_j = 0)$ and $(T_i = +\infty, T_j = +\infty)$ are Nash Equilibria of the game.*

Thus, if player i commits to T_i, it follows that $T_j = 0$. By committing to stop, a player shapes the legacy path faced by the non-committing party, thereby making immediate stopping the best-response strategy. While the precise mechanism differs, the

[21] Precisely, we impose the condition $\frac{u_1}{v_0} < \rho\frac{\alpha-\rho\overline{q}}{\alpha+\delta}$.

outcome resembles the "encouragement effect" discussed in the strategic experimentation literature (Bolton and Harris, 1999), as player i's action influences player j's beliefs about the rewards from continuation. This contrasts with commitments in common-pool resource problems, where commitment can yield a strategic advantage by enabling one party to over-exploit the resource at others' expense.[22]

4.2 Positive catastrophes: innovations

For broader implications of our information structure, we next consider positive catastrophes, such as breakthroughs in basic science and technology. The gestation periods in basic research are often measured in decades (e.g., Adams, 1990); thus, the delay between cause and innovation seems crucial when assessing past research investments. Should basic research—whether privately or government sponsored—be conducted steadily over time or in intensive bursts?

In innovation contests, as in Halac, Kartik and Liu (2017), learning involves updating beliefs based on the absence of success after attempts to solve the problem. Setting aside strategic interactions in this illustration, the planner increasingly assigns more weight to the possibility that the problem is unsolvable as accumulated effort grows. The problem is solvable only if an unknown state is good; the prior for this is $\hat{p}_0$. Success occurs when the state is good and the cumulative innovation effort Q_t exceeds an unknown threshold S, which is assumed to follow an exponential distribution with parameter ρ. The hazard rate for success is then

$$\frac{\rho}{1+\frac{1-\hat{p}_0}{\hat{p}_0}\exp(\rho Q_t)} = \rho\hat{p}_t$$

where $\hat{p}_t$ represents the belief that the state is good, conditional on no success by time t.

Suppose that the planner launches an R&D program in which innovation intensity remains fixed at $\overline{q} > 0$ for as long as the program continues. A fixed cost associated with launching the program may prevent adjustments to innovation intensity during its duration. Consequently, the policy decision reduces to selecting the time interval $[0, T]$, determining the program's length. We may write the planner's payoff as in Halac, Kartik and Liu (2017):

$$\int_0^T \big(\hat{p}_t\rho v - c\big)p_t\overline{q}\exp(-\delta t)dt,$$

[22]This mechanism drives the results in, for example, Harstad (2012), where parties have incentives to secure commitments prior to negotiations to strategically exploit others.

in which v is the prize from the success, c is the unit cost of effort, $\overline{q}$ is the innovation intensity, and p_t is the survival probability, i.e the probability that no success is achieved by time t. Since the belief that the problem is solvable declines with cumulative effort, it is optimal to abandon the project after a finite time. This happens when the belief reaches the threshold

$$\hat{p}^* = \frac{c}{\rho v}.$$

By contrast, our information structure with delays leads to significantly different predictions in this model. Assume that the state is good with probability one, $\hat{p}_0 = 1$, but there is a delay between exceeding the unknown threshold S and learning about it. The interpretation of parameters remains unchanged, except that we now introduce a delay captured by α. Our key result on payoffs (Lemma F.3) establishes that the planner evaluates the following integral:

$$\int_0^T \big((1-\pi_t)\frac{\alpha}{\alpha+\delta}\rho v - c\big)\overline{q}p_t \exp(-\delta t)dt.$$

Assuming $0 < \frac{\alpha+\delta}{\alpha\rho}\frac{c}{v} < 1$, we can define a critical legacy

$$\pi^* = 1 - \frac{\alpha+\delta}{\alpha\rho}\frac{c}{v}$$

at which the integrand vanishes. The interpretation of the legacy is the same as before: it captures the probability that the innovation has already been triggered by past actions. We see that if the initial legacy π_0 is below π^*, then it is optimal to launch the program.

To know when to stop it, we have to study the evolution of π_t through time. It is straightforward to show that when T goes to infinity π_t goes to a limit $\pi^\infty = \min(1, \frac{\rho}{\alpha}\overline{q})$. Appendix H.2 shows the following result:

Proposition 9 *Assume $\pi_0 < \pi^* < \pi^\infty$. Conditional on no success, the research program optimally stops at a finite T^* at which π_t has increased to π^*. After stopping, π_t decreases and, in finite time, returns to π_0, the level at which the program started.*

The conditions in the Proposition imply that π is strictly increasing during the R&D program – intuitively, the planner is more and more optimistic that the success has been triggered. It becomes optimal to stop to avoid *duplication of own effort*, a concept that to our knowledge has not been discussed in the literature.[23] However, after ceasing the

[23]The duplication of effort typically refers to efforts among innovators, not within one innovation program.

program, no news is bad news and π declines monotonically towards zero and therefore it must reach value π_0 in finite time. This suggests that resuming the program is optimal. We leave the details of the optimal innovation cycle for future research.

5 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 highlights the importance of timing of past actions. The planner has different information about the consequences of past actions depending on whether the same stock level was reached gradually or rapidly; therefore, the forward-looking plan also differs accordingly. For crises such as Covid-19, the model predicts that similar planners can take very different optimal courses of actions depending on the legacy of the past. Late planning starting after an explosion of infections can justify the optimality of a lockdown, but the same infection level can justify further steps forward if the current level was approached slowly. The lesson for climate change would be: The "lockdown of emissions" may be optimal until unknowns can be ruled out.

We see that the setting offers several openings for future work. Following our extensions, it is natural to study a game between multiple players who share a common legacy of the past. What are the implications for climate contracts when parties can influence the legacy before negotiations? What types of contracts best induce participation in climate treaties? For innovation contests, it would be interesting to revisit the canonical model when duplication of effort can occur both within a player's own actions and between innovators.

References

Adams, James D. 1990. "Fundamental Stocks of Knowledge and Productivity Growth." *Journal of Political Economy*, 98(4): 673–702.

Assenza, T., F. Collard, M. Dupaigne, P. Fève, C. Hellwig, S. Kankanamge, and N. Werquin. 2020. "The Hammer and the Dance: Equilibrium and Optimal Policy during a Pandemic Crisis." TSE Working Paper, n. 20-1099, May 2020.

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

Benhabib, Jess, and Kazuo Nishimura. 1979. "The Hopf bifurcation and the existence and stability of closed orbits in multisector models of optimal economic growth." *Journal of Economic Theory*, 21(3): 421 – 444.

Besley, Timothy, and Avinash Dixit. 2019. "Environmental catastrophes and mitigation policies in a multiregion world." *Proceedings of the National Academy of Sciences*, 116(12): 5270–5276.

Bloom, David E., Michael Kuhn, and Klaus Prettner. 2022. "Modern Infectious Diseases: Macroeconomic Impacts and Policy Responses." *Journal of Economic Literature*, 60(1): 85–131.

Bolton, Patrick, and Christopher Harris. 1999. "Strategic Experimentation." *Econometrica*, 67(2): 349–374.

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.

Crépin, Anne-Sophie, and Eric Naevdal. 2019. “Inertia Risk: Improving Economic Models of Catastrophes.” *The Scandinavian Journal of Economics*, 122(4): 1259–1285.

Cropper, M.L. 1976. “Regulating activities with catastrophic environmental effects.” *Journal of Environmental Economics and Management*, 3(1): 1 – 15.

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. 2024. “Precaution, Information and Time- Inconsistency: On The Value of the Precautionary Principle.” Toulouse School of Economics WP 1411.

Halac, Marina, Navin Kartik, and Qingmin Liu. 2017. “Contests for Experimentation.” *Journal of Political Economy*, 125(5): 1523–1569.

Harstad, Bård. 2012. "Climate Contracts: A Game of Emissions, Investments, Negotiations, and Renegotiations." *The Review of Economic Studies*, 79(4): 1527–1557.

Hartl, Richard F. 1987. "A Simple Proof of the Monotonicity of the State Trajectories in Autonomous Control Problems." *Journal of Economic Theory*, 41: 211–215.

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, Tuomas, Pauli Murto, and Julia Salmi. 2025. "Gradual learning from incremental actions." *Theoretical Economics*, 20(1): 93–130.

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.

Nordhaus, William. 1975. "Can We Control Carbon Dioxide?" IIASA Working Paper. IIASA, Laxenburg, Austria: WP-75-063.

Pindyck, Robert S. 2013. "Climate Change Policy: What Do the Models Tell Us?" *Journal of Economic Literature*, 51(3): 860–72.

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.

Tsur, Yacov, and Amos Zemel. 2014. "Steady-state properties in a class of dynamic models." *Journal of Economic Dynamics and Control*, 39: 165–177.

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

For the sake of brevity, 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.

Let $(Q_t)_{-\infty<t<+\infty}$ be an admissible path. The maximum stock on record at date t is $\overline{Q}_t \equiv \max_{t' \leq t} Q_{t'}$, and $\overline{Q}_\infty \equiv \sup_{t'} Q_{t'}$ is the supremum of stock values.

The survival probability p_t is by assumption strictly positive at date 0, and above $1 - F(\overline{Q}_0)$. From (7), one easily shows that p_t remains strictly positive. It is weakly decreasing and thus converges. Therefore $\dot{p}_t$ goes to zero, and p_t goes to $1 - F(\overline{Q}_\infty)$.

The differential equation (7) admits a unique solution: for $t \geq T$, one has

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

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

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

Recall also the definitions of the legacy of the past, and of the damage function:

$$\pi_t = 1 - \frac{1 - F(\overline{Q}_T)}{p_t} \qquad D(Q) = \frac{u(0,Q)}{\delta} - V(Q).$$

Finally, for a given path we define the payoff associated with playing $q = 0$ forever from date T onward:

$$Z_0(T) = \int_T^{+\infty} (P_t u(0, Q_T) - \dot{P}_t V(Q_T)) \exp(-\delta(t-T)) dt,$$

where P is defined in (A.2). We obtain

$$Z_0(T) = (1 - F(\overline{Q}_T)) \frac{u(0,Q_T)}{\delta} + (p_T - 1 + F(\overline{Q}_T)) \frac{u(0,Q_T) + \alpha V(Q_T)}{\alpha + \delta},$$

which, thanks to the definitions of π and D, reduces to:

$$Z_0(T) = p_T \left(\frac{u(0,Q_T)}{\delta} - \frac{\alpha}{\alpha+\delta} \pi_T D(Q_T) \right). \tag{A.3}$$

B A useful inequality

The following result will be used repeatedly to study 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*

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

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

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):* $0 \leq t_1 < t_2 < +\infty$, $Q_{t_1} = Q_{t_2}$, $\overline{Q}_{t_1} = \overline{Q}_{t_2}$.

Proof of Lemma B.1: First, let us compute the payoff W from the optimal path on the interval $[t_1, t_2]$. We have

$$W = \int_{t_1}^{t_2}(pu - \dot{p}V)edt,$$

and by integrating by parts the second term we get

$$W = -[pVe]_{t_1}^{t_2} + \int_{t_1}^{t_2} p(u - \delta V + qV')edt.$$

Since by definition $V(Q) = \frac{u(0,Q)}{\delta} - D(Q)$, we get:

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

The concavity of u in q implies:

$$W \leq -[pVe]_{t_1}^{t_2} + \int_{t_1}^{t_2} p\left(q(u_q(0,Q) + \frac{u_Q(0,Q)}{\delta}) + \delta D - qD'\right)edt.$$

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

$$W \leq W^+ \equiv \underbrace{-[pVe]_{t_1}^{t_2}}_{=A} + \int_{t_1}^{t_2} p\left(q(\nu(Q) - D') + \delta D\right)edt.$$

Now, consider an alternative path (q'_t, Q'_t) that consists in setting $q'_t = q_t$ before t_1 and after t_2, and $q'_t = 0$ on $[t_1, t_2]$, so that the stock remains set at Q_{t_1} on this interval. Proceeding as above, the payoff for this new path on the interval $[t_1, t_2]$ equals

$$W_0 = \underbrace{-[P_tV(Q_t)e]_{t_1}^{t_2}}_{=A_0} + \int_{t_1}^{t_2} P_t\delta D(Q_{t_1})edt,$$

where the survival probability P is now given by (A.2) with $T = t_1$. In case (i) of the Lemma, the optimality of the initial path implies the inequality $W^+ \geq W_0$. In case (ii), the condition ($Q_{t_1} = Q_{t_2}$, $\overline{Q}_{t_1} = \overline{Q}_{t_2}$) 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 $\overline{Q}_t$ is the same constant for both path, so that $P = p$ on $[t_1, t_2]$. The inequality $W^+ \geq W_0$ thus reduces to

$$\underbrace{\int_{t_1}^{t_2} p[q(\nu(Q) - D') + \delta(D - D(Q_{t_1}))]edt}_{=B} \geq \underbrace{\delta D(Q_{t_1}) \int_{t_1}^{t_2} (P - p)edt}_{=B_0}.$$

In case (ii), B_0 is zero, and is trivially above the RHS in (B.1), which is also zero. Let us show the same result in case (i). To evaluate B_0, we apply (A.1) at $T = t_1$ to both p and P:

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

Another integration by parts yields:

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

Now, we have:

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

because ρ is increasing and $1 - F$ is decreasing.[24] This implies that B_0 is above the RHS in (B.1), as announced.

Finally, to evaluate B we define the function $N(Q) \equiv \int_{Q_{t_1}}^{Q} \nu(x)dx$. We have:

$$\int_{t_1}^{t_2} pq(\nu - D')edt = [p(N - D)e]_{t_1}^{t_2} - \int_{t_1}^{t_2} (N - D)(\dot{p} - \delta p)edt,$$

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

so that

$$B=[p(N-D)e]_{t_1}^{t_2}+\int_{t_1}^{t_2}(\dot{p}(D-N)+\delta pN-\delta pD(Q_{t_1}))\,edt$$

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

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

And Assumption 1 implies that N is concave, so that

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

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

$$B\leq[p(N-D)e]_{t_1}^{t_2}+\int_{t_1}^{t_2}(Q_t-Q_{t_1})[\dot{p}(D'(Q_{t_1})-\nu(Q_{t_1}))+\delta p\nu(Q_{t_1})]edt+D(Q_{t_1})\int_{t_1}^{t_2}(\dot{p}-\delta p)edt.$$

The first term and the last term on the RHS cancel each other under both case (i) and case (ii). Finally, we use (7) to replace p by $1-F(\overline{Q})-\frac{1}{\alpha}\dot{p}$. This implies that B is below the LHS in (B.1), and concludes the proof. ■

C Consequences for monotonicity

The following result is derived from Lemma B.1, and shows that every optimal path is weakly increasing when Q_t is below Q^D:

Lemma C.1 *Suppose that the stock decreases at the right of t_1. Then $Q_{t_1}\geq Q^D$, and $\pi_{t_1}>\frac{\delta}{\alpha}\frac{\nu(Q_{t_1})}{D'(Q_{t_1})-\nu(Q_{t_1})}$.*

Proof of Lemma C.1: 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 $\overline{Q}_t=\overline{Q}_{t_1}$. Then there must exist t such that the function in the integral in (B.1) is negative, so that:

$$\dot{p}_t(D_1'-\frac{\alpha+\delta}{\alpha}\nu_1)+\delta(1-F(\overline{Q}_{t_1}))\nu_1<0.$$

Since $\pi_t=1-\frac{1-F(\overline{Q}_t)}{p_t}$, we obtain $\dot{p}=\alpha(1-F(\overline{Q}_t)-p_t)=-\alpha p\pi$, 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 $\overline{Q}_t$ is a constant, π_t is decreasing, and therefore $\pi_{t_1}\geq\pi_t$. This establishes the announced inequality. Finally, notice that the right-hand side of this inequality is below 1 if and only if $Q_{t_1}\geq Q^D$. ■

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:

Proposition D.1 *Suppose an optimal path is such that $\overline{Q}_\infty < \overline{S}$. Then the stock Q_t converges to a value $Q_\infty \in [Q^E, Q^N]$ when t goes to infinity. Moreover, if $Q_\infty < Q^N$, then $Q_\infty = \overline{Q}_\infty$, and the path is weakly increasing for t high enough.*

To prove this result, we begin by the following Lemma:

Lemma D.1 *If an optimal path is such that $\overline{Q}_\infty < \overline{S}$, 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

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

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

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

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 Lemma 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 $Q_{t_1} \leq \overline{Q}_\infty$, we have

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

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) equals $(Q_t - Q_{t_1})B(t_1, t)$. 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 every t_1 above a threshold belong either to the domain $\{t_1 : Q_{t_1} < Q^N and t_1 > \Gamma(\nu(Q_{t_1}))\}$, or to the domain $\{t_1 : t_1 \leq \Gamma(|\nu(Q_{t_1})|)\}$. 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 $\overline{Q}_\infty$, it must converge to a limit below Q^N. 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 $\overline{Q}_{t_1} = \overline{Q}_{t_2}$, 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 $\overline{Q}_{t_1}$, 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 nth 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 from Step 2. This implies that the limit of Q_{τ_n} must be at least Q^N, and since Q_{τ_n} decreases the stock must remain above Q^N forever. Symmetrically, $Q_{\tau'_n}$ must be increasing with n, and because $Q_{\tau'_n} = \overline{Q}_{\tau'_n}$ this sequence must converge to $\overline{Q}_\infty$.

In the long-run, $\overline{Q}_t$ becomes arbitrarily close to $\overline{Q}_\infty$, 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. ■

We can now prove the Proposition:

Proof of Proposition D.1: We know from the previous Lemma that the stock converges to a value Q_∞. The proof consists of six steps.

Step 1: 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

$$z \equiv (1 - F(\overline{Q}_\infty))\frac{u(0, Q_\infty)}{\delta}$$

as T goes to infinity. A first remark is that since the path is optimal, then $W(T)$ is at least the payoff $Z_0(T)$ from stabilizing the stock forever at its level Q_T. Thanks to (A.3), we have

$$Z_0(T) = p_T \left(\frac{u(0, Q_T)}{\delta} - \frac{\alpha}{\alpha+\delta}\pi_T D(Q_T)\right).$$

Since p_T converges to $1 - F(\overline{Q}_\infty)$, and π_T goes to zero, we have shown:

$$W(T) \geq Z_0(T) \qquad \text{and} \qquad \lim_{T\to+\infty} Z_0(T) = z. \tag{D.2}$$

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

$$(p_t - 1 + F(\overline{Q}_\infty))u(q_t, Q_t)$$

$$+(1 - F(\overline{Q}_\infty))(u(q_t, Q_t) - u(q_t, Q_\infty))$$

$$+(1 - F(\overline{Q}_\infty))(u(q_t, Q_\infty) - u(0, Q_\infty))$$

$$+(1 - F(\overline{Q}_\infty))u(0, Q_\infty)$$

$$-\dot{p}_t V(Q_t).$$

Because u is concave in q, the third line is less than $(1 - F(\overline{Q}_\infty))q_t u_q(0, Q_\infty)$; 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 $[\underline{q}, \overline{q}] \times A$, 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 $W(T)$ converges to z, as announced.

Step 3: in this step, given $(T,\ \hat{q} \in [\underline{q}, \overline{q}],\ a > 0)$ we define an alternative path: play $\hat{q}$ on $[T, T+a]$, and 0 afterwards. We obtain new trajectories for the variables $(\hat{q}, \hat{Q}, \overline{\hat{Q}}, \hat{p})$:

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

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

and from (A.1):

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

- For $t \geq T + a$:

$$\hat{q}_t = 0 \qquad \hat{Q}_t = \hat{Q}_{T+a} \qquad \overline{\hat{Q}}_t = \overline{\hat{Q}}_{T+a}$$

and from (A.2):

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

Overall, this alternative path yields the following payoff:

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

Consider now the following condition:

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

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 $W_1(T, \hat{q}, 0) = Z_0(T)$, and we know from Step 2 that $Z_0(T)$ and W(T) have the same limit z when T goes to infinity. Therefore:

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

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 $\underline{q} \leq \hat{q} < 0$. 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 $\overline{\hat{Q}}_t$ is a constant, equal to $\overline{Q}_T$. 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 $t \geq T + a$:

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

and, since $\dot{\hat{p}} = \alpha(1 - F(\overline{Q}_T) - \hat{p}_t)$:

$$\frac{\partial \dot{\hat{p}}_t}{\partial a} = -\alpha \frac{\partial \hat{p}_t}{\partial a}. \tag{D.6}$$

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

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

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

$$(1 - F(\overline{Q}_\infty)) \exp(-\delta a)\hat{q} \Big[\frac{u(\hat{q}, Q_\infty + a\hat{q}) - u(0, Q_\infty + a\hat{q}))}{\hat{q}} + \frac{u_Q(0, Q_\infty + a\hat{q})}{\delta}\Big]. \tag{D.8}$$

Finally, recall that we chose $\hat{q}$ to be strictly negative, and notice that the bracketed term is also strictly negative for $\hat{q}$ close enough to zero, as its limit when $\hat{q}$ goes to zero 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 $Q_\infty < \overline{Q}_\infty$ and $Q_\infty < Q^N$. Choose $\bar{a}$ small enough so that $Q_\infty + \bar{a}\hat{q} < \overline{Q}_\infty$. Choose $(T, \hat{q}, a)$ such that $0 < \hat{q} \leq \overline{q}$ and

$0 < a \leq \overline{a}$. Consider the alternative path $(T,\, \hat{q},\, a)$. It is is such that that the highest stock on record $\overline{\hat{Q}}_t$ is a constant, equal to $\overline{Q}_T$. 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, $Q_\infty < Q^N$ implies $Q_\infty = \overline{Q}_\infty$, as announced in the Lemma.

Step 6: once more proceeding by contradiction, we now assume $Q_\infty < Q^E$, so that $Q_\infty = \overline{Q}_\infty$ from Step 5. Choose $(T, \hat{q}, a)$ such that $0 < \hat{q} \leq \overline{q}$ and $a > 0$. Consider the alternative path $(T,\, \hat{q},\, a)$. A new feature is that the 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 $\max(Q_T + a\hat{q}, \overline{Q}_T)$. We therefore define the indicator function $1_{Q_T + a\hat{q} \geq \overline{Q}_T}$, 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, since $\dot{\hat{p}} = \alpha(1 - F(\overline{\hat{Q}}_{T+a}) - \hat{p}_t)$:

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

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

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

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 sum of the last two integrals goes to

$$f(Q_\infty + a\hat{q})\hat{q} \exp(-\delta a) \left(u(0, Q_\infty + a\hat{q})(\frac{1}{\alpha + \delta} - \frac{1}{\delta}) + V(Q_\infty + a\hat{q})\frac{\alpha}{\alpha + \delta}\right).$$

By choosing $\hat{q}$ strictly positive but small enough, and using the definitions of D and ρ, we can make this expression arbitrarily close to

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

We add the limit of (D.8) when $\hat{q}$ goes to zero, to get

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

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. ■

E Consequences for benchmarks

In the main text we have defined two benchmarks: the Stock-Flow Problem without catastrophes, and the case when the catastrophe was triggered with certainty in the past. Both problems are autonomous, with only one state variable. From Hartl (1987), we deduce that if there exists a solution, then there exists a monotonic solution. We show here both existence, and monotonicity of all solutions. Under assumptions that are different from ours, Tsur and Zemel (2014) give a number of results about the long-run stability of these solutions; our function ν is in fact what they call a L-function.

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 4.) 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_∞. Proposition D.1 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 (G.1) 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. ■

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 $Q^E < Q^D$ and $\overline{Q}_0 \leq Q^D$. Then optimal paths cannot exceed Q^D.*

Proof of Lemma F.1: Suppose first $Q^D \geq \overline{S}$. Proposition 2 implies that as soon as the stock exceeds $\overline{S}$, it must converge to Q^D in a monotonic way. This shows the result.

Suppose now $Q^D < \overline{S}$. Suppose an optimal path sometimes exceeds Q^D. Because $\overline{Q}_0 \leq Q^D$, there exists $t_1 \geq 0$ such that Q_t crosses Q^D from below for the first time at time t_1. Moreover, from Lemma C.1 we know that after t_1 the path must remain above Q^D. Therefore, for $t \geq t_1$ we have

$$\overline{Q}_t \geq Q_t > Q_{t_1} = \overline{Q}_{t_1} = Q^D \geq Q^E.$$

We can then apply the inequality (B.1) in case (i): on the LHS, the first term in the parenthesis is zero by definition of Q^D. On the RHS, we have $\overline{Q}_t - \overline{Q}_{t_1} = \overline{Q}_t - Q_{t_1} \geq Q_t - Q_{t_1}$, so that (B.1) implies

$$\int_{t_1}^{t_2} (Q_t - Q_{t_1})(1 - F(\overline{Q}_t)) \exp(-\delta t) dt \left(\nu(Q_{t_1}) - \frac{\alpha\delta}{\alpha+\delta} \rho(Q_{t_1}) D(Q_{t_1}) \right) \geq 0.$$

The first term is strictly positive because the difference $(Q_t - Q_{t_1})$ is positive, and sometimes strictly so, and because $1 - F(\overline{Q}_t)$ is strictly positive (at least for t close to t_1). But the second term is negative, because $Q^D > Q^E$ (see Definition (4)). ■

Lemma F.2 *Suppose an optimal path is such that $Q^E < \overline{Q}_\infty < \min(Q^N, \overline{S})$. Then $\overline{Q}_\infty$ is reached in finite time.*

Proof of Lemma F.2: let us proceed by contradiction. Suppose that $\overline{Q}_\infty$ is reached only asymptotically, necessarily from below. From Proposition D.1, the path converges, to a value Q_∞ which is at most $\overline{Q}_\infty$, and thus strictly below Q^N. Using once more Proposition D.1, we obtain $Q_\infty = \overline{Q}_\infty$, and also that there exists T such that Q_t is weakly increasing for $t \geq T$. Because convergence is asymptotic, we get $\overline{Q}_t = Q_t$ for $t \geq T$. Moreover, because $\overline{Q}_\infty > Q^E$, we can choose T such that $Q_T > Q^E$. We therefore have, for every $t \geq T$, $\overline{Q}_t = Q_t \geq Q_T > Q^E$.

Referring to Lemma B.1, for $t \geq t_1 \geq T$ consider the function

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

Because $Q_{t_1} > Q^E$, the second term is strictly negative; in fact, because $\overline{S} > \overline{Q}_\infty > Q_{t_1} \geq Q_T$ this second term is strictly less than

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

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 $Q_t = \overline{Q}_t$, and therefore the function which is summed 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 $\overline{Q}_\infty$. So its integral in case (i) cannot be weakly positive, and we have a contradiction with (B.1). This shows that $\overline{Q}_\infty$ must be reached in finite time. ■

Lemma F.3 *Consider an admissible path that is constant after a finite date $T \geq 0$. Then:*

(i) The planner's payoff at time 0 *equals*

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

where

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

$$C_t \equiv \nu(Q_t) - \frac{\alpha}{\alpha + \delta} \left[(1 - \pi_t) \rho(Q_t) D(Q_t) 1_{Q_t = \overline{Q}_t \text{ and } q_t \geq 0} + \pi_t D'(Q_t) \right]. \tag{F.1}$$

(ii) If moreover this path is optimal, then

$$\frac{\alpha}{\alpha + \delta} \pi_T D'(Q_T) \leq \nu(Q_T) \leq \frac{\alpha}{\alpha + \delta} \left[(1 - \pi_T) \rho(Q_T) D(Q_T) + \pi_T D'(Q_T) \right]. \tag{F.2}$$

Moreover, in this expression the first inequality is an equality if $Q_T < \overline{Q}_T$, and the second inequality is an equality if $Q_t = \overline{Q}_t$ at the left of T.

Proof of Lemma F.3: (i) The planner's payoff is

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

where the survival probabilities p and P are given in (A.1) and (A.2). The second integral is in fact $\exp(-\delta T) Z_0(T)$, where Z_0 was defined in (A.3). From the identity

$$\exp(-\delta T) Z_0(T) = Z_0(0) + \int_0^T \frac{d}{dt} [\exp(-\delta t) Z_0(t)] dt,$$

we get

$$W = Z_0(0) + \int_0^T [p_t u(q_t, Q_t) - \dot{p}_t V(Q_t) + Z_0'(t) - \delta Z_0(t)] e dt.$$

From (A.3), the first term is the one given in the Lemma. There only remains to compute the term below the sum sign. To do so, we use the following identities:

$$\dot{p}_t = -\alpha p_t \pi_t$$

$$\dot{\pi}_t = (1-\pi_t)(\rho(Q_t)q_t 1_{Q_t=\overline{Q}_t \text{ and } q_t \geq 0} - \alpha\pi_t)$$

$$Z_0(t) = p_t\Big[\frac{u(0,Q_t)}{\delta} - \frac{\alpha}{\alpha+\delta}\pi_t D(Q_t)\Big].$$

Then $pu - \dot{p}V + Z_0' - \delta Z_0$ equals p, times

$$u + \alpha\pi(\frac{u_0}{\delta} - D) - \alpha\pi(\frac{u_0}{\delta} - \frac{\alpha}{\alpha+\delta}\pi D)$$

$$+\frac{u_Q(0,Q_t)q_t}{\delta} - \frac{\alpha}{\alpha+\delta}(1-\pi_t)(\rho(Q_t)q_t 1_{Q_t=\overline{Q}_t \text{ and } q_t \geq 0} - \alpha\pi)D$$

$$-\frac{\alpha}{\alpha+\delta}\pi_t D'(Q_t)q_t - u(0,Q_t) + \frac{\alpha\delta}{\alpha+\delta}\pi_t D(Q_t).$$

which simplifies to the expression $B_t + q_t C_t$ given in the Lemma.

(ii) Consider a deviation that consists in playing a small quantity q on a small interval just after T. Then step (i) applies. By continuity of π and Q in t, for the deviation to be unprofitable it must be that $b(q) + qc(q)$ is at most zero for q close to zero, where

$$b(q) = u(q,Q_T) - u(0,Q_T) - qu_q(0,Q_T)$$

$$c(q) = \nu(Q_T) - \frac{\alpha}{\alpha+\delta}\left[(1-\pi_T)\rho(Q_T)D(Q_T)1_{Q_T=\overline{Q}_T \text{ and } q\geq 0} + \pi_T D'(Q_T)\right].$$

But $b(q)$ is second-order when q is small. Therefore, only the sign of $c(q)$ matters, but this sign may depend on whether q is positive or negative. We can only impose that the sign of $c(q)$ is weakly positive for $q < 0$, and weakly negative for $q > 0$: hence, we only obtain the inequalities in (F.2).

Consider now the case when $Q_T < \overline{Q}_T$. Now, whatever q $c(q)$ is the same constant:

$$c(q) = \nu(Q_T) - \frac{\alpha}{\alpha+\delta}\pi_T D'(Q_T)$$

and we can apply the same reasoning: c must be zero, as announced in the Lemma.

Finally, consider the case when $Q_t = \overline{Q}_t$ on an interval $[T', T]$, with $T' < T$. Then $q_t \geq 0$ on this interval. Suppose $C_t < 0$. Then $C_t < 0$ for t close to T, by continuity. Because $B_t \leq 0$ in any case, one would be better off deviating to q_t small but strictly negative, a contradiction. Therefore $C_T = 0$, as announced in the Lemma. ■

Proof of Theorem 1: from Lemma F.1, optimal paths cannot exceed Q^D; from Lemma C.1, optimal paths are weakly increasing, so that $Q_t = \overline{Q}_t$ 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 $\overline{Q}_\infty \in [Q^E, Q^D]$, from Proposition D.1, and that the stock level is constant after $\overline{Q}_\infty$ is reached.

If $\overline{Q}_\infty$ is reached asymptotically, then a contrapositive to Lemma F.2 implies that the stock converges to $\overline{Q}_\infty = Q^E$. Since the legacy π_T vanishes when T goes to infinity, we obtain that (15) indeed holds at $T = +\infty$, $\pi_T = 0$, and $Q_T = Q^E$, as announced.

If $\overline{Q}_\infty$ is reached at time $T < +\infty$, then (ii) in Lemma F.3 implies (15). ■

Proof of Theorem 2: Because $\overline{Q}_\infty < \overline{S}$, Proposition D.1 applies: the path converges to a value $Q_\infty \in [Q^E, Q^N]$. Since $Q_\infty \leq \overline{Q}_\infty$, and $\overline{Q}_\infty < Q^N$ by assumption, from the same Proposition one must have $Q_\infty = \overline{Q}_\infty$, and the path is weakly increasing for t high enough. This shows $\overline{Q}_\infty \geq \max(\overline{Q}_0, Q^E)$.

Now, let us proceed by contradiction, and suppose $\overline{Q}_\infty > \max(\overline{Q}_0, Q^E)$. Then Lemma F.2 implies that $\overline{Q}_\infty$ is reached in finite time, say T, and $T > 0$ since $\overline{Q}_\infty > \overline{Q}_0$. Then $Q_t = \overline{Q}_t$ at the left of T. For simplicity, focus on the case when after T Q remains constant. Then we can apply Lemma F.3 to obtain that the second inequality in (F.2) holds as an equality. But this equality implies that $\overline{Q}_\infty$ lies between Q^E and Q^D, in contradiction with our assumption $Q^D < Q^E < \overline{Q}_\infty$.

Therefore we have $\overline{Q}_\infty = \max(\overline{Q}_0, Q^E)$. Finally, if $\overline{Q}_0 > Q^E$, then Lemma F.2 implies that the limit $\overline{Q}_\infty = \overline{Q}_0$ is reached in finite time. ■

G Applications

G.1 Social distancing illustration

Linking to the general framework: If we define

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

we are back to our general framework, with

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

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

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

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:

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

This is the expression for I^N in the main text.

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 distribution F of S on the support $[0, \overline{S}]$ and the associated hazard rate ρ are as defined in the general model. After the catastrophe has occurred at time $\mathcal{T}$, we therefore have $I_t = I_{\mathcal{T}} \exp[q^*(t - \mathcal{T})]$, and the continuation payoff can be computed explicitly:

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

Consequently, the damage function becomes

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

where the parameter μ^* is as in the main text

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

Condition $\nu(Q) - \frac{\alpha}{\alpha+\delta} D'(Q) = 0$ from (10) gives, after some manipulations, the long-run target I^D:

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

We can similarly use the condition $\nu(Q) - \frac{\alpha}{\alpha+\delta}\rho(Q)D(Q) = 0$ from (13) for characterizing the target $I^E = \exp(Q^E)$. This and the comparison to I^D leads to Lemma 2 in the main text.

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 6, we write the general payoff as

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

where B is a constant and

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

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 $\left(Y_0 - (\frac{\alpha w d^*}{\delta - q^*} + wd)\frac{I^N}{\alpha+\delta}\right) = \left(\frac{wdI^N}{\delta} - (\frac{\alpha w d^*}{\delta - q^*} + wd)\frac{I^N}{\alpha+\delta}\right) < \left(\frac{1}{\delta} - \frac{1}{(\delta - q^*)}\right)\frac{\alpha}{\alpha+\delta} w d^* I^N < 0$. 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 $b_t = 0$. ■

G.2 Optimal carbon budget

Proof of Proposition 5: Let us first consider the case when after some time t_0 the optimal policy exceeds the upper value $\overline{S}$, so that the catastrophe is triggered with certainty. Then we know that the optimal policy maximizes

$$\int_0^{+\infty} [u(q_t, Q_t) + \alpha V(Q_t)] \exp(-(\alpha + \delta)t) dt \tag{G.1}$$

which in this simple case is

$$\int_0^{+\infty} (u_0 + u_1 q_t - \alpha v_0 Q_t) \exp(-(\alpha + \delta)t) dt.$$

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

$$\int_0^{+\infty} (u_0 + (u_1 - \frac{\alpha}{\alpha + \delta} v_0) q_t) \exp(-(\alpha + \delta)t) dt.$$

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

We can now examine the optimal policy when the stock lies below $\overline{S}$. We can focus on "bang-bang" policies that set q_t either to zero or to $\overline{q}$. 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):

$$p_t C_t = p_t u_1 - \frac{\alpha}{\alpha + \delta} [(1 - F(Q_t))\rho(Q_t) D(Q_t) + (p_t - 1 + F(Q_t)) D'(Q_t)]$$

$$= p_t(u_1 - \frac{\alpha}{\alpha+\delta}v_0) - \frac{\alpha}{\alpha+\delta}(1-F(Q_t))(\rho(Q_t)D(Q_t) - D'(Q_t)).$$

Now, suppose that q_t is zero between two dates t_0 and t_1, and is $\overline{q}$ 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 $\overline{q}$ until some date $T \leq +\infty$, and be zero afterwards. We thus have to maximize on T the objective

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

Moreover, from (A.1) we have

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

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 6: From Theorem 2, we can focus on paths that converge to $\overline{Q}_\infty = \overline{Q}_0 = Q_0$. 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

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

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

Otherwise, if $u_1 < \pi_0 \frac{\alpha}{\alpha+\delta} v_0$, then a_t is initially negative, before becoming positive at some strictly positive time t_1, 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 $q = \overline{q}$ until the stock is back to Q_0, at time t_2 such that $\underline{q}t_1 + \overline{q}t_2 = 0$, so that t_2 is also increasing in π_0. The optimal policy is thus as stated in the claim. ■

G.3 Stock-flow trade-offs in climate change

After the catastrophe has occurred, the planning goes on, and the continuation value $V(Q)$ becomes:[25]

$$V(Q) = \frac{-k\theta}{\delta+\gamma}Q + \frac{\beta}{\delta}[\log(\gamma Q^N) - \log k - 1]. \tag{G.2}$$

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

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

Consider first the target Q^D. We obtain from Definition 3 that

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

which indeed is less than Q^N.

Consider then the target Q^E can be expressed. Thanks to Definition 4, it can be found as a solution to equation

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

[25]The 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 $V(Q)$.

giving

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

where function g is defined as

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

Proposition 7 follows directly from these definitions.

H Extensions

H.1 Strategic interactions

Proof of Proposition 8: Computations are simple, but tedious and cumbersome. We only give here the main intermediate results.

Step 1: In this preliminary step, suppose that the aggregate emissions equal a constant q_0 until time $T \geq 0$, and equal a constant q_1 after T. Our aim here is to compute

$$Z(T) = \int_T^{+\infty} p_t \exp(-\delta t) dt.$$

To do so, we first compute p_t at date $t > T$, using (A.1) and $1 - F(Q) = \exp(-\rho Q)$:

$$p_t = p_T \exp(-\alpha(t - T)) + \frac{\alpha}{\alpha - \rho q_1} \exp(-\rho Q_T)\Big(\exp(-\rho q_1(t - T)) - \exp(-\alpha(t - T))\Big).$$

Then we sum over $t \geq T$ to obtain

$$Z(T) = \frac{\exp(-\delta T)}{\alpha + \delta}\Big(p_T + \exp(-\rho Q_T)\frac{\alpha}{\rho q_1 + \delta}\Big).$$

As in the proof of Lemma F.3, we now write

$$Z(T) = Z(0) + \int_0^T Z'(t) dt$$

and we use the identity $\exp(-\rho Q_t) = p_t(1 - \pi_t)$ to end up with

$$Z(T) = \frac{1}{\alpha + \delta} p_0 \Big(1 + (1 - \pi_0)\frac{\alpha}{\rho q + \delta}\Big) - \int_0^T p_t \Big(1 + \frac{\alpha}{\alpha + \delta} \frac{q_0 - q_1}{\delta + \rho q} \rho(1 - \pi_t)\Big) e dt. \quad \text{(H.1)}$$

Step 2: Suppose player j chooses $T_j = +\infty$. If player i chooses to stop at T, then aggregate emissions are $2\overline{q}$ until T, and then $\overline{q}$. From (9), player's i payoff is

$$W(T) = \int_0^T (p_t u_1 \overline{q} + \dot{p} v_0) e + \int_T^{+\infty} \dot{p}_t v_0 e dt.$$

An integration by parts yields

$$\int_0^{+\infty} \dot{p}_t e dt = -p_0 + \delta \int_0^{+\infty} p_t e dt,$$

so that

$$W(T) = -p_0 v_0 + (u_1 \bar{q} + \delta v_0) \int_0^T p_t e dt + \delta v_0 \int_T^{+\infty} p_t e dt.$$

We use (H.1) to compute the second integral, and we obtain

$$W(T) = \frac{\alpha}{\alpha+\delta} v_0 p_0 \Big(\frac{\delta}{\rho\bar{q}+\delta}(1-\pi_0) - 1 \Big) + \bar{q} \int_0^T p_t \Big(u_1 - \frac{\alpha\delta}{\alpha+\delta} \frac{\rho v_0}{\delta+\rho\bar{q}}(1-\pi_t) \Big) e dt.$$

Because $\pi_0 > \pi^*$, the derivative with respect to T at $T = 0$ is positive. Therefore, i must choose $T > 0$, but this makes the derivative with respect to T even more positive. This shows that the best response of player i is to emit forever: $T_i = +\infty$.

Step 3: The case when $T_j = 0$ now obtains by setting $\bar{q}$ to zero at the right places in the formulas in Step 2. We obtain that player i's payoff is a constant, plus

$$\bar{q} \int_0^T p_t \Big(u_1 - \frac{\alpha}{\alpha+\delta} \rho v_0 (1-\pi_t) \Big) e dt.$$

Because $\pi_0 < \pi^{**}$, the derivative with respect to T at $T = 0$ is negative. Under our assumptions, if i emits, then π_t goes to $\pi^\infty < \pi^{**}$; if i does not emit, then π_t goes to zero. Thus, the condition $\pi_0 < \pi^{**}$ ensures that in any case π_t remains below π^{**}. Therefore, the above payoff is decreasing in T, and the best response is not to emit at all. This concludes the proof. ■

H.2 Positive catastrophes

Proof of Proposition 9: In terms of the general model, we have

$$u(q, Q) = -c\bar{q}, \quad V(Q) = v, \quad \nu(Q) = -c, \quad D(Q) = -v.$$

From Lemma F.3, we have

$$B_t = 0, \quad C_t = -c - \frac{\alpha}{\alpha+\delta} \Big[(1-\pi)\rho(-v) \Big]$$

and thus the planner's payoff is

$$\frac{\alpha}{\alpha+\delta} \pi_0 v + \int_0^T \big((1-\pi_t)(\frac{\alpha}{\alpha+\delta})\rho v - c \big) \bar{q} p_t \exp(-\delta t) dt. \tag{H.2}$$

The main text reports the second term in this expression, as only this part can be affected by the policy. To determine if T is finite, we consider the evolution of π_t during the program. To do so, we use the now-familiar law of motion

$$\dot{\pi}_t = (1 - \pi_t)(\rho\overline{q} - \alpha\pi),$$

implying that π_t is increasing and goes to $\pi^\infty = \min(1, \frac{\rho}{\alpha}\overline{q})$ if the program never stops. Consider then the statement in Proposition 9. Assumptions $\pi_0 < \pi^* < \pi^\infty$ imply that π is strictly increasing for a constant $\overline{q} > 0$. Inspecting $(H.2)$ shows that the payoff is strictly increasing in T for all $\pi \in [\pi_0, \pi^*)$, and strictly decreasing for π outside this set. Thus, stopping is optimal when π_t reaches π^*. Because π_t declines monotonically to zero, it must reach value π_0 in finite time.

I Additional results

I.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 (9). Noting that in $[T, \infty)$ the stock is stabilized $q = 0$, and then the survival probability p_t follows a formula (A.1) that allows us to express the said welfare as a product of discount factor $\exp(-\delta T)$ and[26]

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

where we use shorthands $u_T^0 = u(0, Q_T)$, $F_T = F(\overline{Q}_T)$, 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

[26] This expression comes from (A.3).

after this time stop with $q_{T+\Delta}=0$. By the above logic, the welfare at $T+\Delta$ is

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

The flow gain from this one-shot experiment follows from the objective (9) 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,

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

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:

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

After careful evaluation of terms, this condition becomes

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

where

$$C_T\equiv\nu(Q_T)-\frac{\alpha}{\alpha+\delta}\left[(1-\pi_T)\rho(Q_T)D(Q_T)+\pi_T D'(Q_T)\right]. \tag{I.1}$$

I.2 The model without delay

Proof of Proposition 3: 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 $\dot{Q}_t=q_t\in[\underline{q},\overline{q}]$, $p_t=1-F(\overline{Q}_t)$, $\overline{Q}_t=\max_{0\le t'\le t}Q_{t'}$, 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 ($\overline{Q}_0=\overline{Q}_T$), and therefore the problem at time zero and the problem at time T are identical. This proves that at time zero the planner could as well adopt the strategy she 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,[27] 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 $\overline{Q}_t = Q_t$. Therefore $p_t = 1 - F(Q_t)$, and the complete payoff from the candidate path is:

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

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 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 $(1 - F(Q_0))$

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

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 $\overline{Q}_t = Q_t$ 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.[28] 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 $Q_t = \overline{Q}_t$ 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

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

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 $\int uedt$, 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

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

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 $Q^{E0} \geq \overline{Q}_0$ (case (ii.c)). Otherwise, the constraint $q \geq 0$ binds, and the stock remains

[28] See Theorem 15, p. 237, in Seierstad and Sydsaeter (1987).

forever set at Q_0. ■