Catastrophes, delays, and learning*

Matti Liski† François Salanié‡

March 15, 2026

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 thank the editor, Bård Harstad, and four anonymous referees for their valuable feedback. We have benefited from comments received 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 their 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.

1 Introduction

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 value of 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 results in a higher value of the legacy 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 has a value equal to 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 value for the 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 value for the 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.

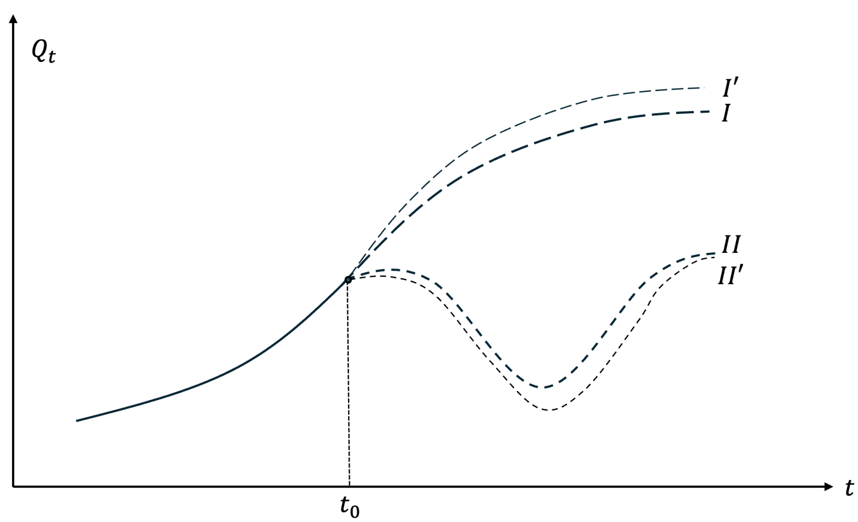

Figure 1: Some possible paths for the stock: I and I' for the case $Q^E < Q^D$, II and II' for the case $Q^E > Q^D$.

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 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 that trade off consumption against climate damages also fit into our theorems' dichotomy, in a manner similar to the disease-control application.

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

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

[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.

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 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

[3] We discuss these contributions in Section 3.1.

[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.

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]

[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;

2 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 2.1 defines and studies this classical stock-flow trade-off. Section 2.2 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 2.3 adds uncertainty on both the threshold and the delay. Given these components, Section 2.4 formulates the complete planning problem.

2.1 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 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). To understand the forces shaping optimal policies, it is useful to begin with one admissible policy in which the planner holds the stock constant at some value Q by choosing $q = 0$. Such a stabilized

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.

path need not be optimal—in particular, the optimal long-run stock level Q^N may lie at infinity. Nevertheless, this stationary situation provides a convenient benchmark for identifying the marginal trade-offs faced by the planner: starting from $q = 0$, a small increase in q both raises current utility and increases the stock, with future consequences discounted over time. These considerations define the marginal payoff,[6]

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

which measures the net gain from slightly increasing the flow when the stock is held at 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:

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

[6] Subscripts denote partial derivatives.

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 $u(q,Q) = u_0+u_1q+u_2Q$, 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.

2.2 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:

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

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

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

$(Q_t)_{t\in(-\infty,+\infty)}$, one can compute $\mathcal{T} = T(S) + \tau$ from (3), and the planner's payoff from date $t = 0$ onward equals

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

2.3 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.

2.4 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 $(Q_t)_{t\leq 0}$ relevant for planning:

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

Similarly, the maximum stock on record can be computed as a function of $\overline{Q}_0$:

$$\overline{Q}_t = \max(\max_{0\leq t'\leq t} Q_{t'}, \overline{Q}_0). \tag{7}$$

Moreover, differentiating (6) with respect to t yields

$$\dot{p}_t = -f(\overline{Q}_t)\dot{\overline{Q}}_t \; + \int_{S<\overline{Q}_t} (-\alpha)\, e^{-\alpha(t-T(S))}\, dF(S) \; + \; e^{-\alpha(t-T(\overline{Q}_t))} f(\overline{Q}_t)\dot{\overline{Q}}_t. \tag{6'}$$

The first and third terms cancel in both relevant situations: (i) record not updated, $\dot{\overline{Q}}_t = 0$, so both terms are zero; (ii) record updated, $\overline{Q}_t = Q_t$ and $T(\overline{Q}_t) = t$, implying $e^{-\alpha(t-T(\overline{Q}_t))} = 1$, so the two boundary terms are equal in magnitude and opposite in sign. Thus the law of motion reduces to

$$\dot{p}_t = -\alpha \int_{S<\overline{Q}_t} e^{-\alpha(t-T(S))}\, dF(S) = \alpha\big[1 - F(\overline{Q}_t) - p_t\big]. \tag{8}$$

Current actions q_t affect $\dot{p}_t$ only through their occasional effect on the running maximum $\overline{Q}_t$, but the expression above remains valid even at those instants.

Taken together, these two forces lead to the following objective:

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

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

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

[8]That this objective follows from the planner's problem defined in (5) can be formally proven. Note first that the cumulative distribution function of the occurrence date $\mathcal{T}$ is $1 - p_{\mathcal{T}}$, so the event $\mathcal{T} \geq 0$

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.

3 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 (7). 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 and their connections to the literature.

3.1 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

---

has probability p_0. The payoff in (5) can therefore be written as

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

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

By omitting the constant factor p_0 and relabelling $\mathcal{T}$ as t in the second integral, we obtain (9).

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:[9]

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

From Assumptions 1-2, this expression is decreasing in Q and lies below $\nu(Q)$. Therefore, it may reach zero only at a value $Q^D \leq Q^N$, 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 $\overline{Q}_0 \geq \overline{S}$ at the planning date 0, so the planner knows that the catastrophe has already been triggered. The law of motion for the survival probability (8) 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

[9] 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 $T(S)$.

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:

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 (8) 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:

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

Recall that $\rho(Q) = f(Q)/[1 - F(Q)]$ is the hazard rate associated with the threshold distribution F. Here, 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.[10]

Proposition 3 also highlights the differences between the unknown-threshold and hazard-rate models. In the no-delay limit ($\alpha = +\infty$), the effects of an uncertain threshold with a continuous distribution seem equivalent to a change in the payoff function: an increase in the running maximum $\overline{Q}_t$ indeed entails a smooth expected marginal cost $\rho(Q)D(Q)$, as in (12). This does *not*, however, imply that the unknown-threshold mechanism can be represented by a stylized model in which the flow payoff is modified by some function of the current stock; nor does it make it equivalent to the hazard-rate model studied above. First, the two benchmarks feature different damage formulas: when triggering occurred in the past, marginal incentives are governed by $D'(Q)$ (cf. (10)), whereas under an unknown threshold without delay they are governed by the expected loss from immediate triggering, $\rho(Q)D(Q)$ (cf. (12)). Second, even abstracting from this payoff difference, the hazard-rate intuition applies only when policy increases the running maximum (so that $\overline{Q}_t$ moves). When the stock is reduced after an earlier increase, $\overline{Q}_t$ is unchanged and the triggering risk is already determined; policy then concerns the management of damages conditional on the past maximum. With lags ($\alpha < \infty$), this history dependence becomes central through the legacy, which is precisely what our general analysis captures.

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

[10] 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."

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:[11]

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

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 $Q_\infty \geq Q^E$.*

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:

[11] See footnote 9.

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 we obtain $Q^D = Q^N \geq Q^E$ from Definitions 1-3. 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.

3.2 First theorem: experimentation first ($Q^E < Q^D$)

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 $Q_T \leq Q^D$ at some date $T \leq +\infty$. 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?[12]

Theorem 1 (Case $Q^E < Q^D$) *Suppose* (14) *holds. Then there exists an optimal policy. Under this policy, the path* $\left(Q_t\right)_{t\geq 0}$ *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) makes explicit how our model *unifies* the two approaches developed in the literature, the hazard-rate approach and the unknown-threshold approach. It shows that both marginal conditions used in these literatures are in fact limiting cases of the *same* optimality condition, with the legacy at the stopping time determining the relative weight placed on each. In this sense, (15) consolidates the definitions of Q^D and Q^E within a single framework. Notice that when delays are infinite (α vanishes), we return to the no-catastrophe case, for which $Q^E = Q^D = Q^N$, and to the optimal path characterized in Proposition 1. Conversely, in the absence of delays ($\alpha = +\infty$) the legacy is identically zero, which reproduces Proposition 3 and the convergence of the stock to 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 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 4.2 for a simple model of climate policies, confirming that a higher

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

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.[13]

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.[14]

3.3 Second theorem: mitigation first ($Q^E > Q^D$)

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 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* $\big(Q_t\big)_{t\geq 0}$ *be an optimal path. If*

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

then:

[13] 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 $(Q_t)_{t\leq 0}$.

[14] 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.

1. *If $\overline{Q}_0 \geq Q^E$, then the path $\left(Q_t\right)_{t\geq 0}$ converges to $\overline{Q}_0$ in finite time, and for every t we have $Q_t \leq \overline{Q}_0$.*

2. *If $\overline{Q}_0 < Q^E$, then the path $\left(Q_t\right)_{t\geq 0}$ converges to Q^E, and for every t we have $Q_t \leq Q^E$.*

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. The assumption in the Theorem is typically satisfied when damages are sufficiently convex in the stock, so that the planner prefers not to trigger the event with certainty.

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. A particularly transparent illustration occurs when $\overline{Q}_0 > Q^E$. If the value of legacy is high enough, then it is optimal to reduce the stock in the short run in order to mitigate damages. After a while, if, contrary to the initial expectation, the event does not occur, the legacy gradually declines, and eventually the planner is content to return to $\overline{Q}_0$—but not above it, since $\overline{Q}_0 > Q^E$. Intuitively, by doing so the planner offsets the impact of the past trajectory; economically, this non-monotonicity follows naturally from the desire to mitigate damages.

4 Applications

4.1 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:

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

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.

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.

[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.

Figure 2

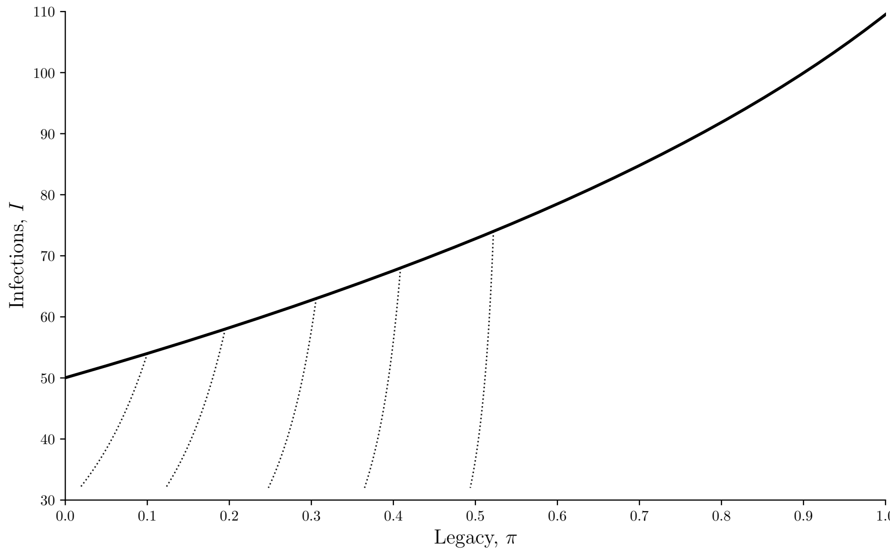

Optimal paths in the plane (π, I) for a linear production function $Y(R) = Y_0 R$. Parameters are: $\delta = 0.03, q^* = 0.01, w = 1, d = 0.1, r = 0.9, d^* = 0.2, \alpha = 0.2, Y_0 = 1000, Y^* = 950$, $I_0 = 32$, $\bar{q} = 1$. Distribution F for $\log(I)$ is uniform: $f = 1/6$. Benchmark values are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

[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 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.

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

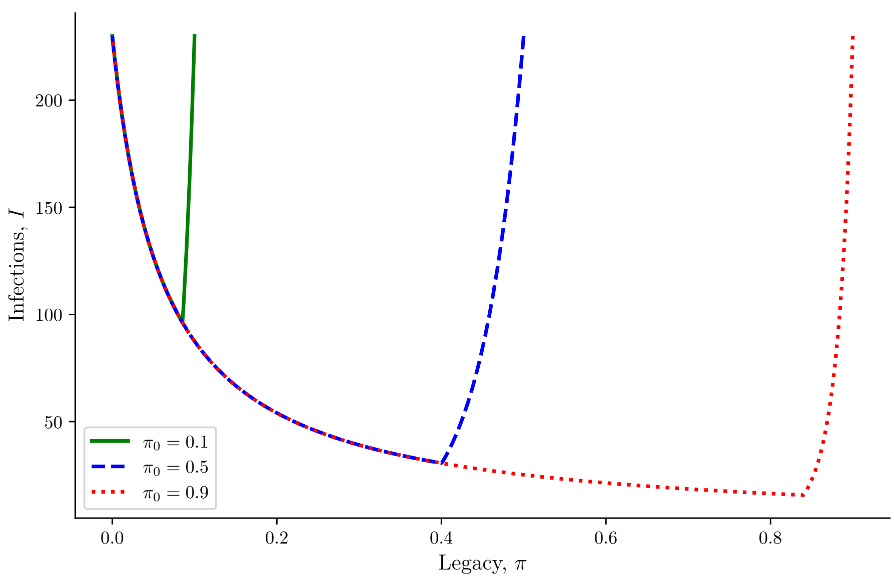

The population of infected agents over time, for a linear production function. Parameters are: $\delta = 0.023, q^* = 0.02, w = 1, d = 0.1, r = 0.98, d^* = 0.25, \alpha = 0.2, Y_0 = 1000, I_0 = 230$.

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 increasingly 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

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

to its initial value, because the probability that the threshold lies below it has become negligible.

Figures 3–4 illustrate this optimal policy from two complementary angles. Figure 3 (phase plot) shows, for three initial legacies, how the legacy declines during the lockdown, while infections initially fall but eventually return to their starting level, since the planner becomes convinced that no catastrophe is coming. This is confirmed in Figure 4, which shows the paths associated to these dynamics.

Figure 4

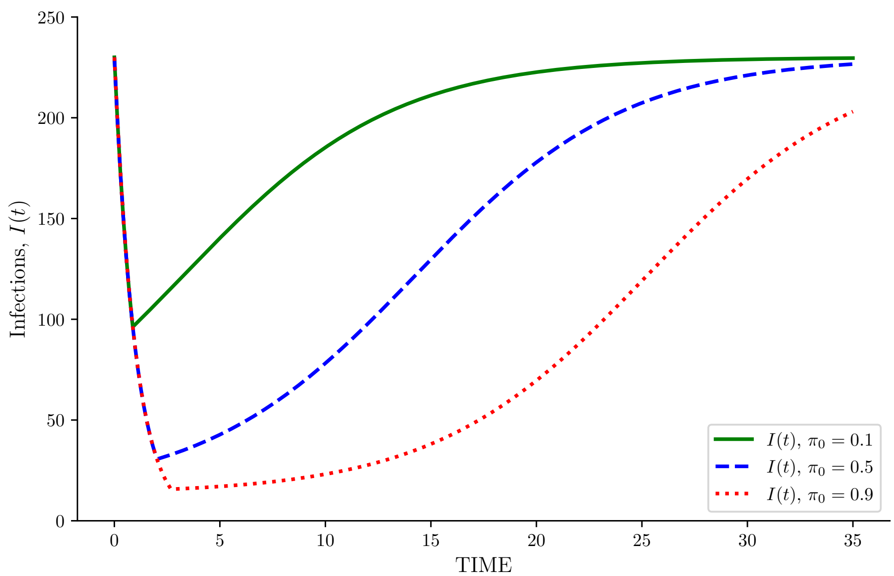

The optimal control, for a linear production function. Parameters are: $\delta = 0.023, q^* = 0.02, w = 1, d = 0.1, r = 0.98, d^* = 0.25, \alpha = 0.2, Y_0 = 1000, I_0 = 230$.

The figures confirm 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).

4.2 Climate change

4.2.1 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.

4.2.2 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

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

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.

5 Conclusion

We conclude by discussing the policy context and potential extensions.

Policy context. The theoretical results shed light on the diversity of real-world policy responses during Covid-19. The Nordic countries share relatively similar primitives facing the pandemic; in particular, they have similar healthcare capacity and similar values regarding economy-life trade-offs. Yet their policies diverged sharply. Sweden adopted relatively lenient policies, while Finland and Norway pursued stricter social distancing and early suppression (Yarmol-Matusiak, Cipriano and Stranges, 2021; Irfan et al., 2022; Larkin et al., 2022). Sweden experienced a sharp early surge in infections, whereas Finland and Norway delayed and flattened the first waves. Nevertheless, evidence shows that by late 2022 population exposure and hybrid immunity had converged across these countries, and excess mortality over the three pandemic years was of comparable magnitude (Forthun et al., 2024; Hallberg et al., 2025).

From the perspective of our model, such divergent transition paths among otherwise similar societies could be explained by different legacy beliefs. Sweden's state epidemiologist, Anders Tegnell, argued that the virus was essentially uncontainable and that strict lockdowns would only postpone the inevitable. This corresponds to a high-legacy view, where the catastrophe is perceived as already seeded, making a "fatalistic" strategy of preserving economic and social activity rational. The policies in Finland and Norway instead seemed to be based on low-legacy beliefs, justifying strict early lockdowns in the hope of preventing the catastrophe. Thus, the Nordic comparison illustrates how legacy beliefs, rather than underlying preferences, can rationalize starkly different policy choices that nevertheless converge to similar long-run outcomes—a pattern captured by the non-monotonic policy paths, as our opening Figure 1.

A parallel interpretation to climate policy would be: societies with broadly similar economic capacity and values may nonetheless adopt sharply different mitigation strategies depending on their legacy beliefs. A high-legacy view—that critical thresholds for irreversible damage have already been crossed—can rationalize an *economy first* strategy of continued emissions, putting priority on growth over costly abatement. By contrast, a low-legacy view supports a *life first* strategy of stringent early reductions, "a lockdown of emissions", until catastrophic risks are ruled out. The model highlights how divergent policy paths can emerge endogenously from different beliefs about legacy, even when underlying trade-offs are otherwise aligned.

Extensions. There are at least two natural extensions. First, in a common-pool setting, a strategic extension could introduce players who share a common "legacy of the past" when deciding on a resource use. While the general analysis is open for future research, our working paper version considers a two-player example in which players must decide whether to commit to emissions reductions. The analysis shows that commitments can act as strategic complements: when one country stops emitting, this can induce the other to stop as well, sometimes resulting in the first-best outcome. This mechanism resembles the "encouragement effect" from the strategic experimentation literature (Bolton and Harris, 1999), but it contrasts with common-pool resource problems where commitments typically create incentives to over-exploit. This opens a broader research agenda on the implications of legacy-dependent risks for climate negotiations and contract design.

Second, the information structure developed in this paper seems broadly applicable, for example to beneficial events such as breakthroughs in basic science and technology development. In fact, the gestation times in basic research are measured in decades (e.g., Adams, 1990), and therefore the delay between the cause and the impact seems essential in assessments of past investments in research. Should basic research, private or government sponsored, be conducted steadily over time or as intensive bursts? Building on canonical models of innovation races (e.g. Halac, Kartik and Liu, 2017), our working paper version shows that delays between discovery and recognition fundamentally alter the planner's problem: rather than a one-shot stopping rule based on declining beliefs, optimal R&D may follow a cycle of launching and suspending programs. A novel feature is that past efforts exert an incentive to stop and wait: one would like to avoid duplicating own effort – a different concern from the usual fear of duplicating effort across innovators. While our analysis is highly stylized, it opens a promising line of research on the timing and design of innovation policies when research inputs and observable breakthroughs are separated by long and uncertain delays.

The replication package for figures in this research is available on Zenodo at https://doi.org/10.5281/zenodo.18760575.

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.

Forthun, Ingvild, Kjersti Reiakvam, Rickard Ljung, Anders Engeland, and Torkild H. Lyngstad. 2024. "Excess mortality in Denmark, Finland, Norway and Sweden during the COVID-19 pandemic 2020–2022: A comparative study of weekly mortality." *European Journal of Public Health*, 34(4): 737–743.

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. "Acting in the Darkness: Towards some Foundations for the Precautionary Principle." Toulouse School of Economics WP 23-1411.

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

Hallberg, Petter, Lars C. Stene, Eythor Eythorsson, Asko Järvinen, Camilla H. Hansen, and Pär-Daniel Sundvall. 2025. "Epidemiological outcomes and policy implementation in the Nordic countries during the COVID-19 pandemic: A comparative review." *Archives of Public Health*, 83(1): 15.

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.

Irfan, Muhammad, Muhammad Ahmad, Shahzad Hussain, and Syed Ali Raza Gilani. 2022. "Coronavirus pandemic in the Nordic countries, Sweden, Denmark, Norway and Finland: A comparative review." *Journal of Global Health*, 12: 05017.

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.

Larkin, Harlan, Kjetil Andersen, Matthew Hall, Olle Lindahl, Katrine Madsen, and Maria Rönkkö. 2022. "Health Policies and Economies in Nordic Countries During the COVID-19 Pandemic." *JAMA*, 328(6): 575–576.

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.

Yarmol-Matusiak, Elena A., Lauren E. Cipriano, and Saverio Stranges. 2021. "A comparison of COVID-19 epidemiological indicators in Sweden, Norway, Denmark, and Finland." *Scandinavian Journal of Public Health*, 49(1): 69–77.