International Institute for Applied Systems Analysis
IIASA www.iiasa.ac.at

Address: IIASA, Schlossplatz 1, A-2361 Laxenburg, Austria

Email: repository@iiasa.ac.at

Working Paper **WP-24-009**

A model-based measure for the resilience of resource use under the risk of disruption

Michael Kuhn (kuhn@iiasa.ac.at)
Stefan Wrzaczek (wrzaczek@iiasa.ac.at)

Approved by

Hans Joachim Schellnhuber
Director General
13 May 2024

Table of contents

ZVR 524808900

Disclaimer, funding acknowledgment, and copyright information:
IIASA Working Papers report on research carried out at IIASA and have received only limited review. Views or opinions expressed herein do not necessarily represent those of the institute, its National Member Organizations, or other organizations supporting the work.

The authors gratefully acknowledge funding from IIASA and the National Member Organizations that support the institute (The Austrian Academy of Sciences; The Brazilian Federal Agency for Support and Evaluation of Graduate Education (CAPES); The National Natural Science Foundation of China (NSFC); The Academy of Scientific Research and Technology (ASRT), Egypt; The Finnish Committee for IIASA; The Association for the Advancement of IIASA, Germany; The Technology Information, Forecasting and Assessment Council (TIFAC), India; The Indonesian National Committee for IIASA; The Iran National Science Foundation (INSF); The Israel Committee for IIASA; The Japan Committee for IIASA; The National Research Foundation of Korea (NRF); The Mexican National Committee for IIASA; The Research Council of Norway (RCN); The Russian Academy of Sciences (RAS); Ministry of Education, Science, Research and Sport, Slovakia; The National Research Foundation (NRF), South Africa; The Swedish Research Council for Environment, Agricultural Sciences and Spatial Planning (FORMAS); The Ukrainian Academy of Sciences; The Research Councils of the UK; The National Academy of Sciences (NAS), USA; The Vietnam Academy of Science and Technology (VAST).

The authors gratefully acknowledge funding from the Austrian Science Fund (FWF) for the research project 'Life-cycle behaviour in the face of large shocks to health' (No. P 30665-G27).

A model-based measure for the resilience of resource use under the risk of disruption

Michael Kuhn*, Stefan Wrzaczek*

Abstract While useful conceptualizations and measures have been developed that allow to assess the resilience to shocks of eco systems, socioeconomic systems or individuals, these are often context-specific, somewhat ad-hoc, focused on only certain dimensions of resilience, and not always easy to apply to the assessment of concrete policies. In this paper, we build on a simple yet general model of resource use under the risk of potentially catastrophic disruptions, and derive a rigorous, comprehensive, and policy-oriented measure of socioeconomic resilience. By way of numerical analysis, we illustrate how the measure captures the key aspects of resistance, recovery and robustness; how it embraces current resilience and the continuation resilience following future shocks; and how it can be employed to study the resilience of different harvesting policies both from an ex-ante perspective and following a cascade of shocks.

1 Introduction

The notion of resilience has become widely popular if not fashionable (and overused) in the policy and wider public debate. With a rigorous and long established grounding in ecology (e.g., Holling (10)) and psychology (e.g., Bonanno (3)), the concept is increasingly applied for understanding under what conditions socio-economic systems are best able to absorb and survive or even grow on disruptive economic or ecological shocks (e.g., Folke (6); Brown and Westaway (4); Hallegate, 2014 (8); Keating et al. (12); Li et al. (13); Grafton et al. (7)). While there is an increasing number of very useful approaches towards the empirical or simulation-driven measurement of the extent of socio-economic resilience and its determinants (e.g., Hallegatte (8); Angulo (1); Salignac (18); Yonson and Noy (22); Hochrainer-Stigler et al. (9); Hosseini and Ivanov (11); Oliveira et al. (16)) these approaches have mostly been very context-specific in terms of both the conceptualization of resilience and its measurement. Furthermore, in socio-economic and management contexts, the issue of resilience (or its lack) is inherently tied to a set of underlying decisions, which may inadvertently harm resilience or may deliberately chosen to enhance it. Strikingly, and with a few notable exceptions (Li et al. (13); Grafton et al. (7)) the decision- or policy-context of resilience is rarely put at the centre of analysis.

The present work seeks to advance both the conceptualization and decision-orientation of the notion of resilience by proposing a rigorous mathematical model-based, comprehensive, policy-oriented measure of socio-economic resilience. Building on a simple yet general model of resource use under the risk of potentially

*International Institute for Applied Systems Analysis (IIASA), Austria; and Wittgenstein Centre (IIASA, OeAW/VID, University of Vienna), Austria.

catastrophic disruptions, we derive a measure of resilience as an index on the zero-one interval. Our measure is composed of four crucial components:

- a sub-measure of resistance, as the system's capacity to return to a desired path, such as e.g., the return of a diminished resource stock towards a balanced growth path or steady-state, conceptualized as the zero-one indication as to whether or not the system state lies in a region of attraction towards the desired path;
- a sub-measure of recovery, valuing the time for the system it takes to recover towards the desired path, conceptualized as a weighted discount function of the time span up until entry into a target region;
- a sub-measure of robustness against the arrival of future shocks, conceptualized as the life expectancy of the system within its current regime or stage[1];
- While the product of the first three sub-measures makes up a measure of resilience that is referring to the current regime, this does not yet fully capture the resilience of the system to future shocks. We thus add to the the measure of "current resilience" a measure of "continuation resilience" that, broadly speaking, refers to the expected stream of resilience (again based on the current and continuation components) to future regimes as induced by the sequence of probabilistic shocks.

To embrace the policy-relevance of our measure, we include the scope for a decision-maker to choose between different harvesting trajectories that will affect the stock of the resource and consequently its resilience to possible disruptions. In a series of numerical exercises, we then study (i) how the measure of resilience can be decomposed into its components depending on critical systemic features, such as the initial resource stock; (ii) how resilience is affected by different harvesting policies; (iii) how ex-ante resilience is affected by the possibility of multiple shocks; and (iv) how resilience emerges (ex-post) over time for a cascade of shocks and (v) how this depends on the underlying policy. This serves to demonstrate in a sand-box context how our measure of resilience can be usefully applied for policy analysis.

There are a number of innovative aspects about the measure we are suggesting: (a) it is rigorously grounded within a model of intertemporal resource use in a broad sense (and, thus, embracing not only natural resources but also man-made capital, financial assets, human capital (health), social capital, amongst others) and thereby allows for a tight integration of the measure with the processes described in the model; (b) by addressing the three dimensions of resistance, recovery and robustness, it covers both ex-ante (i.e., precaution) and ex-post (adjustment) aspects of resilience; (c) it allows for a forward-looking assessment of resilience, where many approaches are only considering resilience relating to the current stage; (d) by appropriate calibration of the model to specific contexts, it has the potential to be taken to the data; (e) it is flexible in the sense of allowing different policy preferences (relating to the desired path and target region; discount rate on time-to-transition); (f) the measure can, in principle, be applied in the assessment of both given policy scenarios and optimal policies from the perspective of different decision-makers.

[1] Throughout the paper we will speak of "regimes" or "stages" of the system when referring to its positioning towards a shock. This is to reserve the usage of the alternative term "state" for reference to the "state variable" as part of the mathematical description of the system dynamics.

Our work is tying into a small literature that is seeking to conceptualize resilience in the context of (rigorous) modelling of resource management. While aspects of resilience are formalized in Polasky et al. (17) for a similar set-up, as well as in Li et al. (13), these remain somewhat partial and are not included into a comprehensive measure. Asheim et al. (2) provide a rigorous foundation of a resilience measure based on axiomatic theory. However, their approach is oriented towards characterizing the processes of recovery in response to individual health shocks, i.e., resilience at the individual level, with a view to providing a foundation for the empirical study of resilience at the population level. Apart from the difference in context, our approach provides a consistent measure, including sub-measures, for the assessment of the implications for resilience of different policies. Our conceptualization of resilience is echoing Grafton et al. (7) who highlight the roles of resistance, recovery and robustness for a coherent assessment of resilience, but do not fully integrate them into a coherent measure. In further distinction to Grafton et al. (7), our measures takes explicit and rigorous account of the distinction between current resilience and its forward looking continuation.

The remainder of the article is structured as follows. The next Section introduces a simple resource dynamic model, including the scope for disruptive shocks and demonstrates some key aspects of our conceptualization of resilience. We then introduce and discuss the measure and its components before demonstrating its functionality and applicability by way of a number of numerical exercises. The paper concludes with a discussion of applications and extensions.

2 Motivating Model

As a backdrop against which to derive our measure of resilience, we introduce in this section a simple model of resource use in the face of a possible (catastrophic) disruption. We are also employing a simple numerical version of the model to illustrate the logic of the measure.

Resource dynamics subject to disruption

Consider a resource stock $x(t)$ that at each point in time t follows the dynamics $\dot{x}(t) = f(x(t), h(t))$, where $h(t)$ is a quantity (harvest) that is diverted from the resource for some other purpose, e.g., sale or consumption. The notion of "resilience" typically implies the existence of (at least) two regions of attraction, a resilient one, in which the system develops towards a desired state, e.g., a high-level steady state, or a desired (balanced) growth trajectory $x^*(t)$; and a non-resilient one, in which the system develops towards a low steady state, possibly involving the exhaustion of the resource, such that $x(t) \to 0$ within finite time, but more generally also including undesirable outcomes. While per se this is not a prerequisite for the definition of our measure, we note that such settings are generated by specifications of $f(x(t), h(t))$ that are convex-concave in the first-argument, i.e., for which $f_{xx}(0, h(t)) > 0$; $f_{xx}(x(t), h(t)) < 0$ for some $x(t) > \hat{x} > 0$.

The resource is subject to a disruptive shock that shifts the resource into a new regime. The shock arrives at the hazard rate $\eta = \eta(t, x(t))$, which may depend on the resource stock itself (for instance, it has been shown that local and regional warming decreases with the size of the remaining intact forest, with implications for the risk of (devastating) droughts and forest fires (5)). A disruptive shock at time τ may entail (i) a jump in the resource stock at τ that leaves $\lim_{t\to\tau^+} x(t) = \lim_{t\to\tau^-} \varphi(\tau, x(\tau)) \leq x(\tau)$ and/or (ii) a change in the growth function to $\tilde{f}(x(t), h(t), x(\tau), \tau) \leq f(x(t), h(t))$ for $t \geq \tau$. Both the extent of instantaneous change and the extent to which future resource growth is changed may depend on the size of the stock at the point of disruption, $x(\tau)$ as well as the time τ of disruption itself. Note that here we are deliberately describing the jump of the resource stock and the change in the growth function in neutral terms, as certain shocks, such as new discoveries of a natural resource or the arrival of resource-enhancing technologies, may be beneficial. Nevertheless, we frame the remainder of the analysis in the context of negative shocks, involving the resource destruction and/or an inhibition of resource renewal.

Also note at this stage that while for fixing ideas we stick to a narrative of natural resource exploitation in the presence of random resource collapse, the model can also be read in (i) a stock pollution sense, where $x(t)$ represents the remaining absorptive capacity and $h(t)$ a flow of pollutants; or in (ii) an entrepreneurial sense, where $x(t)$ represents capital that generates a certain return, which can then either be reinvested or withdrawn for different purposes at rate $h(t)$. The model is easily generalized to include further aspects of disruptive shocks, e.g., the notion that the impact of disruptions may increase in the number of prior shocks or, indeed, vice versa.

A simple example

In the following, we introduce a simple example that will help to illustrate the logic of our measure and to demonstrate how it can be applied in assessing the impact of policy options on resilience.

Thus, assume the resource dynamics are given by

$$\dot{x}(t) = f(x(t), h(t)) = \frac{ax(t)^2}{b + x(t)^2} - h(t), \tag{1}$$

where resource renewal (as given by the first term) follows a convex-concave pattern and where we assume that the harvest policy

$$h(t) = hx(t) \tag{2}$$

follows a feedback rule at rate $h \geq 0$. For h sufficiently low, this system typically yields three steady states with respect to the resource stock: a low one at $x = 0$ which is stable, a high one at $x_H > 0$ and an intermediate one which is unstable. For a full characterization of the steady state and stability properties see Appendix A.

We assume the shock to arrive at a non-negative rate

$$\eta(t, x(t)) = \eta(x(t)) = \max\left\{0, \bar{\eta}\frac{\bar{x} - x(t)}{\bar{x}}\right\} \tag{3}$$

and to destroy a share $1 - \varepsilon \in [0, 1]$ of the resource stock, such that a system hit by a shock at τ enters the post-shock stage with a remaining stock of $\varphi(\tau, x(\tau)) = \varepsilon x(\tau) \in [0, x(\tau)]$. The hazard rate depends negatively on the value of the stock, implying a (linear) decrease from a maximal value $\bar{x} > 0$ for $x(t) = 0$ to a zero level if the stock reaches a 'safe' level of $\bar{x}$.

We begin by studying how the shock and more specifically the timing of its arrival bear on the scope for policies to generate sustained trajectories of resource use, i.e., be resistant to the shock as one key feature of resilience. Figure 1 provides an illustration for a harvesting policy at rate $h = 0.1$ and for two different

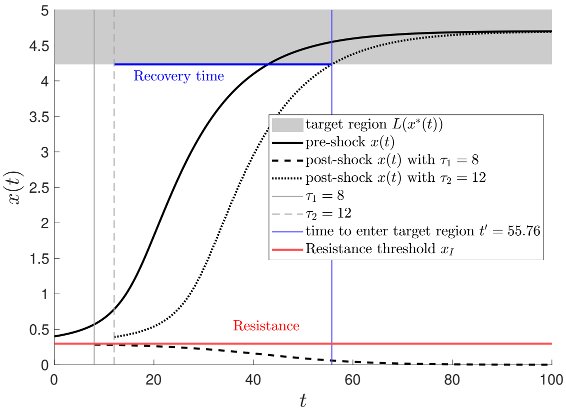

Figure 1: Resistance and recovery depending on the arrival of the shock. Parameters: $a = 0.5$, $b = 1.4$, $\varepsilon = 0.5$, $h = 0.1$.

arrival dates of the shock ($\tau_1 = 8$ and $\tau_2 = 12$).[2] The shock is assumed to destroy half of the stock $\varepsilon = 0.5$. The (resistance) threshold x_I (which coincides with the unstable steady state, see Appendix A) describes the resource level below which the resource stock $x(t)$ will be asymptotically driven to extinction for the harvesting policy $h = 0.1$. We assume that the decision-maker targets the high and stable steady state $x^* = x^H$ for which, in the present context, the system is also resilient in the long-run. As the steady-state is attained only asymptotically, we assume that the decision-maker considers it sufficient if the state falls into a neighbourhood of x^* (corresponding to the grey-shaded area in Figure 1), which for the purpose of the present paper we call a target region $L(x^*)$.[3] When starting from a low initial level of the stock (this could be the stock level immediately following a past shock, reflecting that the time line may have started earlier and within a different regime), the policy allows the gradual accumulation of the stock, asymptotically approaching the high steady state and, thus, meeting the threshold for the target region in finite time. For an early shock at $\tau_1 = 8$ not enough stock has been accumulated, such that, following the shock, the remaining stock falls below the resistance threshold and the resource system expires. For a later shock at $\tau_2 = 12$ the resource level proves to be sufficient for the stock to recover.

The figure allows us to illustrate two crucial elements of resilience, namely resistance and recovery, as discussed in the introduction. First, by allowing the arrival time of the shock τ to increase from low levels, it is easy for this simple model to infer a cut-off time $\tau_I \in (8, 12)$ such that the system is resistant if and only if $\tau > \tau_I$. As will be seen below, our measure of resilience will assign a value of 0 to all arrival times of the shock that are below τ_I and a value of 1 otherwise, such that by counting up, the measure will increase with the instances in which the system proves to be resistant. Second, we arrive at a measure of recovery by counting the spell of time $t' - \tau$ it takes to attain the threshold for the target region following a shock (or, indeed, from the starting point). In the present example, and assuming a shock at $\tau_2 = 12$ to which the system is resistant, the recovery time equals $t' - \tau_2 = 55.76 - 12 = 43.76$. Later shocks for which a greater resource stock has been built would then be associated with lower recovery times. Following an appropriate transformation, "recovery"

[2]We consider distinct exogenous arrival dates only for the purpose of this first illustrative exercise. In all subsequent analyses they will be generated according to the endogenous hazard rate.

[3]For the purpose of the present analysis, we define the target region to include all $x(t) > 0.9x^H$.

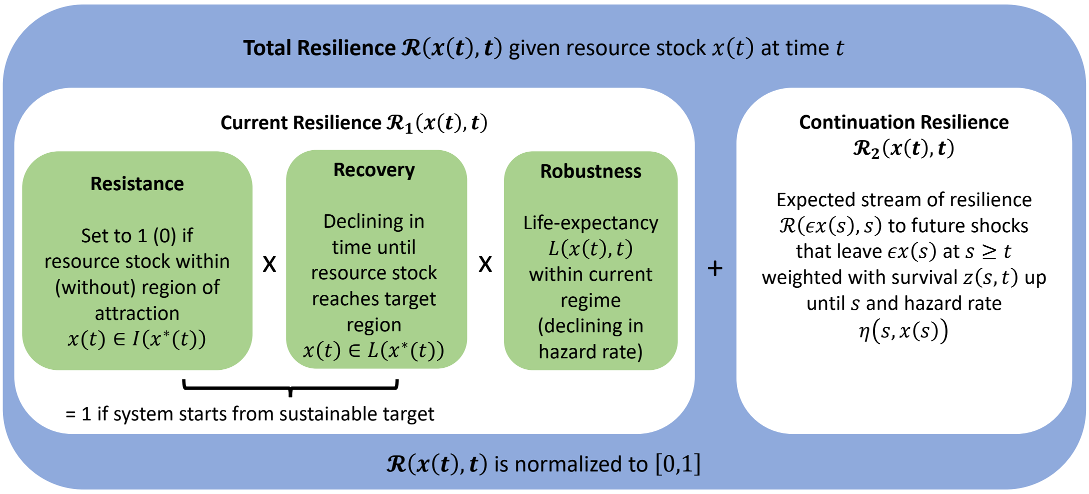

Figure 2: Schematic representation of the resilience measure and its components.

and "resistance" constitute two out of three components of the resilience of a system from the perspective of the regime it is currently in - "robustness" against future shocks being the third one. In the following section, we proceed to develop a more complete and formal measure of resilience.

3 Deriving the Measure

To capture the three elements of resistance, robustness, and recovery as well as the possibility of future (sequences of) system shocks in a coherent and rigorous way, we introduce in the following a mathematical measure of resilience. For ease of exposition, we relegate a number of the more rigorous and general mathematical definitions and descriptions to Appendix B.

Specifically, we define total resilience

$$\mathcal{R}(x(t),t) = \mathcal{R}_1(x(t),t) + \mathcal{R}_2(x(t),t) \tag{4}$$

at time t and for a system state $x(t)$ as the sum of *current resilience* $\mathcal{R}_1(x(t),t)$ and *continuation resilience* $\mathcal{R}_2(x(t),t)$. Here, current resilience refers to the regime the system is currently in. The measure of current resilience we propose follows the suggestion by Grafton et al. (7) in embracing the three components of (i) *resistance*, (ii) *recovery* and (iii) *robustness*. Consider the system to be in its current regime, as described by $x(t)$ and $f(x(t),h(t))$ at t. We then define $\mathcal{R}_1(x(t),t)$ by the product of three sub-measures of resistance, recovery and robustness, each of which is normalised to lie on the $[0,1]$ interval. Our measure of current resilience can reflect (i) the aftermath of a shock, allowing an assessment on the extent to which the system is resistant, can recover, and is robust against *further* shocks, or (ii) the system's "original" status, allowing an assessment of the robustness against a (first-time) shock. Notably, the measure of current resilience needs to be complemented by a forward-looking component of "continuation resilience", reflecting the resilience of the system following each of a sequence of future random shocks. Figure 2 summarizes and illustrates how the overall resilience measure is composed. With this backdrop in mind, we now discuss in more detail the various components of the measure.

To do so, let us define $x^*(t)$ as a target path of the resource state that is defined by the decision-maker or the modeller. Typically, $x^*(t)$ could relate to a desired steady state or balanced growth path, which is the way we will define it in our example. However, in some circumstances $x^*(t)$ could also relate to a subsistence level of the resource state that ensures the survival of the system. In either case, $x^*(t)$ would typically also imply full resilience or a specified (high) level of resilience, again to be determined by the decision-maker or modeller. For many dynamic models, a steady state is reached only asymptotically. For this reason, we define a neighbourhood $L(x^*(t))$ of $x^*(t)$ as a target region for which the decision-maker is satisfied of having moved sufficiently close to $x^*(t)$. For some applications, the target region can be bounded only from one side, as for instance in the case in which $x^*(t)$ relates to a steady-state of the resource stock that is feasible in the long-run even if the decision-maker would prefer $x(t) > x^*(t)$. Finally, we define $\mathcal{I}(L(x^*(t))) = \mathcal{I}(x^*(t))$ as the region of attraction, out of which the system reaches the target region within finite time. Here, we denote by t' the point in time at which the system state permanently enters the target region.[4]

Resistance is then satisfied if and only if $x(t)$ lies in the attractive region $\mathcal{I}(x^*(t))$. This is measured by an indicator function $\mathbb{I}_{[x(t)\in\mathcal{I}(x^*(t))]}$ set to one, while the indicator is set to zero if $x(t)$ lies in the attractive region of the (undesired) low steady state.

Recovery is measured by a function $g(t', t)$ on the $[0, 1]$ interval that decreases in the length of time $t' - t$ required for the state to enter the target region $L(x^*(t))$. A general definition and discussion of the function $g(\cdot)$ is is provided in Appendix B.

Robustness is measured by the remaining life expectancy $\mathcal{L}(x(t), t)$ within the current regime, i.e., the expected time until the arrival of a (new) shock. Note that life expectancy $\mathcal{L}(x(t), t) = \int_t^\infty z(s, t)\mathrm{d}s$ is, through the survival function $z(s, t) = \exp\left[-\int_t^s \eta(\hat{s}, x(\hat{s}))\mathrm{d}\hat{s}\right]$, dependent on the hazard rate of future shocks $\eta = \eta(s, x(s))$ for $s \geq t$. Normalisation to the $[0, 1]$ interval of this component is attained by division through $\mathcal{L}(x(t), t) + 1$.

Altogether we can then write

$$\mathcal{R}_1(x(t), t) = \frac{\mathcal{L}(x(t), t)}{\mathcal{L}(x(t), t) + 1}\mathbb{I}_{[x(t)\in\mathcal{I}(x^*(t))]}\, g(t', t) \tag{5}$$

as our measure of *current resilience*, where $\frac{\mathcal{L}(x(t),t)}{\mathcal{L}(x(t),t)+1} \in [0, 1]$, $\mathbb{I}_{[x(t)\in\mathcal{I}(x^*(t))]} \in [0, 1]$, and $g(t', t) \in [0, 1]$ measure robustness, resistance and recovery, respectively. With this measure capturing the system's current resilience following a shock just prior to time t or, alternatively, the system's resilience within a regime before any shock has occurred,[5] we are not yet accounting for the future pathway of resilience should a shock arrive and take the system into a new regime. This part is captured by the measure

$$\mathcal{R}_2(x(t), t) = \frac{1}{\mathcal{L}(x(t), t) + 1}\int_t^\infty z(s, t)\eta(s, x(s))\mathcal{R}(\varphi(s, x(s)), s)\,\mathrm{d}s \tag{6}$$

of *continuation resilience*, which captures the expected stream of resilience $\mathcal{R}(\varphi(s, x(s)), s)$ to a sequence of possible future shocks at time $s \geq t$ that reduce the resource stock to $\varphi(s, x(s)) \leq x(s)$ in case of a damaging

[4] Additional details and a number of generalizations of these concepts can be found in Appendix B.

[5] In case of the system setting out from a regime in which no shock has occurred yet, it would seem reasonable to set $\mathbb{I}_{[x(t)\in\mathcal{I}(x^*(t))]} = 1$ on the presumption that the pre-shock regime is "sustainable", and $g(t', t) = 1$ on the presumption that the system is on a desired path (or, indeed, in equilibrium).

shock.[6] As argued above, the shock may also shift the dynamics of resource renewal. Resilience at each future time $s \geq t$ is weighted with the arrival rate of a shock $\eta(s, x(s))$ for that particular time and with the survival $z(s,t)$ of the system from the current period t until s. Note that the survival function implies the discounting of high order shocks that do not arise until far into the future.

The measure can be used as a flexible tool (cf., descriptive and controlled models, calibrated theoretical and data driven models, etc.) to assess the resilience of a system, as described by the relevant functions and parameters; to disentangle the various sub-components and channels; and to analyze how resilience varies with the specific policy under consideration. To demonstrate the scope for such analysis, we revert to our motivating model and calculate resilience for an illustrative parameter constellation.

4 Applying the Measure

We now revert to our numerical example, as outlined earlier, to demonstrate the working of the measure as well as some potential applications. For the purpose of this exercise we thus specify the target state to coincide with the high steady-state of the system, $x^*(t) = x^H(t)$ and the target region to include all states that are not below 90% of the high steady state, i.e., $L(x^*(t)) = [0.9x^H, \infty)$. Furthermore, we note that the region of attraction corresponds to the set of all states that are above the unstable steady state, i.e. $\mathcal{I}(x^*(t)) = [x_I, \infty)$. Finally, we specify the recovery function $g(t', s) = e^{-r(t'-s)}$ with $r \in [0, \infty)$.

In the state dynamics (7) we set the values $a = 0.5$ and $b = 1.4$ throughout. In the specifiction of the hazard rate (3), we assume $\bar{x}$ to equal the high steady state x^H in most examples. In those cases in which we compare different harvest rates h leading to different steady states we assume $\bar{x}$ to equal the lowest steady state. Furthermore, we set $\bar{\eta} = 2$ in (3), $\varepsilon = 0.5$ in the damage function, and $r = 0.02$ in the recovery function. Further specifications are adjusted to the specific examples and presented in footnote 7.

Specifically, we perform the following exercises:[7] (i) We demonstrate how for a given harvesting policy, the measure allows for a decomposition into the three dimensions of (current) resilience; as well as into current and continuation resilience (Figure 3). (ii) We show how the measure allows us to assess the resilience of different harvesting policies in a comparative way (Figure 4). Exercises (i) and (ii) assume that only a singular shock may arrive in the future. We subsequently lift this assumption and illustrate (iii) how the scope for cascading shocks bears on resilience from an ex-ante perspective (Figure 5) and (iv) how resilience emerges over time for a cascade of shocks (Figure 6).

Figure 3 depicts resilience as a function of the initial resource stock x_0. For the sake of the argument assume here, that the system has undergone a shock in the past, leaving the initial resource stock, and is subject

[6] More generally, our measure of resilience may also relate to positive shocks such as the arrival of a technological breakthrough that shifts the technology from $x(t)$ to $\varphi(x(t)) \geq x(t)$.

[7] We set $x_0 = 0.5$ with the exception of Figure 3, which depicts resilience as a function of x_0, and Figure 6, where we set $x_0 = 1$ for a better visualization. Furthermore we set $h = 0.1$ in all but Figures 4 and 6, in which h is varied for comparison. For Figures 3 and 5 we (arbitrarily) assume that $\bar{x} = x^H|_{h=0.1}$, which ensures that the system reaches and maintains full resilience in its steady state; As multiple steady states are feasible in Figures 4 and 6 we (arbitrarily) assume that $\bar{x} = x^H|_{h=0.13}$, which ensures that the system reaches and maintains full resilience when reaching the steady state for the highest harvest rate (and thus the lowest steady state level) under consideration. Likewise, we assume that $x^* = 0.9x^H|_{h=0.1}$ in Figures 3 and 5 and $x^* = 0.9x^H|_{h=0.13}$ in Figures 4 and 6.

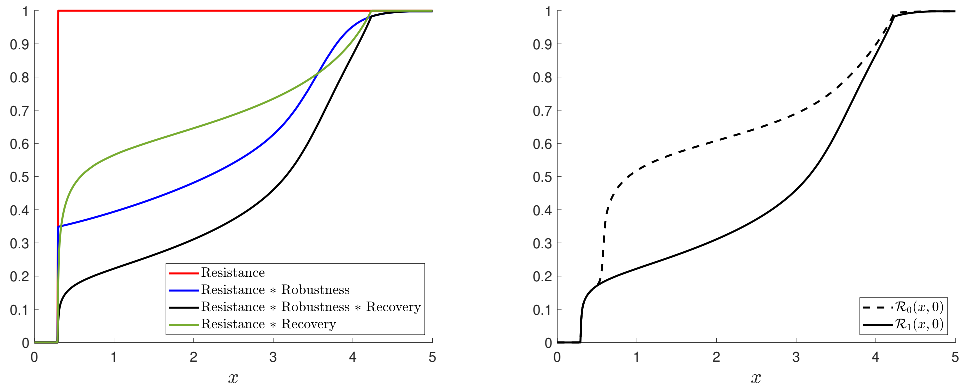

2

Figure 3: Decomposing current resilience (left panel) and overall resilience (right panel). Parameters: $a = 0.5$, $b = 1.4$, $\varepsilon = 0.5$, $\bar{x} = x^H$, $h = 0.1$, $\bar{\eta} = 2$, $r = 0.02$, $L(x^*(t)) = [0.9x^H, \infty)$.

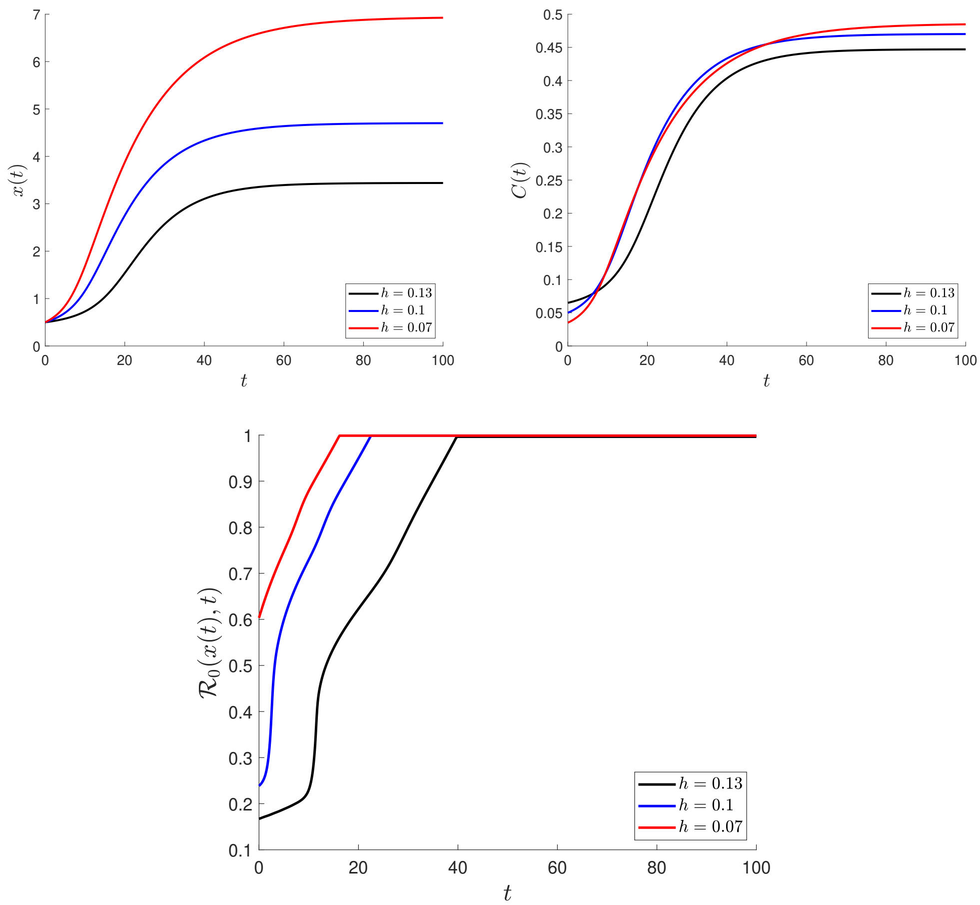

Figure 4: Resilience for different harvesting policies. Parameters: $a = 0.5$, $b = 1.4$, $\varepsilon = 0.5$, $\bar{x} = x^H|_{h=0.13} = 3.4391$, $\bar{\eta} = 2$, $r = 0.02$, $L(x^*(t)) = [0.9x^H|_{h=0.13}, \infty) = [3.0952, \infty)$, $x_0 = 0.5$.

to the risk of a single future shock. The left panel then illustrates how the three constituents – resistance, robustness and recovery – contribute to current resilience (black line). Altogether, we see that for the model we are considering current resilience is increasing in the level of the initial resource stock and converges to one for

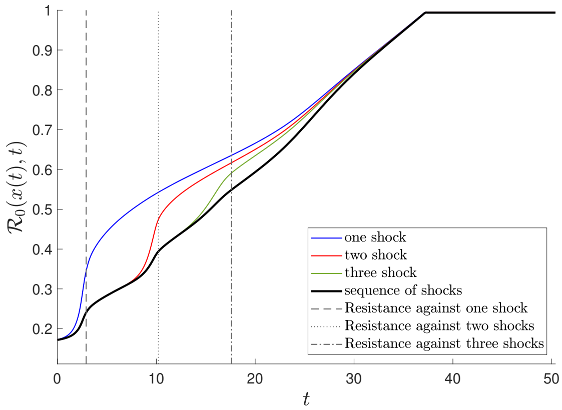

Figure 5: (Ex-ante) resilience for multiple shocks. Parameters: $a = 0.5$, $b = 1.4$, $\varepsilon = 0.5$, $\bar{x} = x^H$, $h = 0.1$, $\bar{\eta} = 2$, $r = 0.02$, $L(x^*(t)) = [0.9x^H, \infty)$, $x_0 = 0.5$.

large enough levels of the initial stock.[8] The graphical representation allows to infer for each level of the initial resource stock the contribution of each constituent component. This is achieved by starting from resistance, i.e., the step function $\mathbb{I}_{[x(t)\in L(x^*(t))]} \in 0, 1$ (red plot), then adding either the robustness weight $\frac{\mathcal{L}(x(t),t)}{\mathcal{L}(x(t),t)+1} \in [0, 1]$ (blue plot) or the recovery weight $g\,(t', t) \in [0, 1]$ (green plot), and finally adding the respective remaining third weight to arrive at the measure of current resilience (black plot). We see that resistance is foundational, as the system persists into the long-run only if the resource stock exceeds a certain threshold level. But even if resistance is guaranteed, the hazard of an additional shock taking the system out of its current regime leads to a sizeable reduction in the measure (difference between red and blue plot). A downward adjustment (difference between red and green plot) also follows when the time it takes for the system for recovering and reaching the target region is taken into account. By comparison, we see that for the specification we have chosen recovery contributes less to resilience than robustness for relatively low initial levels of the resource stock but that this relationship reverses for high levels of the initial stock.[9] We also note that for the multiplicative nature of the resilience measure, recovery and robustness do not contribute to current resilience in an additive way. For the context under consideration, both robustness and recovery improve in the level of the initial stock, implying the convergence of the current resilience measure to one.

The right panel of Figure 3 decomposes total resilience $\mathcal{R}(x(t), t)$ (solid plot) into its current component $\mathcal{R}_1(x(t), t)$ (dashed plot) and continuation component $\mathcal{R}_2(x(t), t)$ (difference between the two plots). While unsurprisingly both total and current resilience increase with the level of the initial resource stock, notably this expansion follows a (smoothed) step-wise form for total resilience. Here, the first step coincides with the step-wise increase in current resilience and reflects the system's resistance in its current stage (i.e., resistance to a past shock). As long as current resistance is not achieved, the system fails and for that reason, continuation resilience does not matter. The second step in total resilience is associated with an initial resource stock for

[8]This may of course differ for other contexts, e.g., for a setting with an exogenous hazard rate, where the robustness part will imply a constant "drag".

[9]The valuation of "recovery" crucially depends on the discount factor applied to time interval that is required to reach the target region. The higher this factor the stronger the focus on the present - at which time the target region is still distant - and thus the lower is the value of recovery.

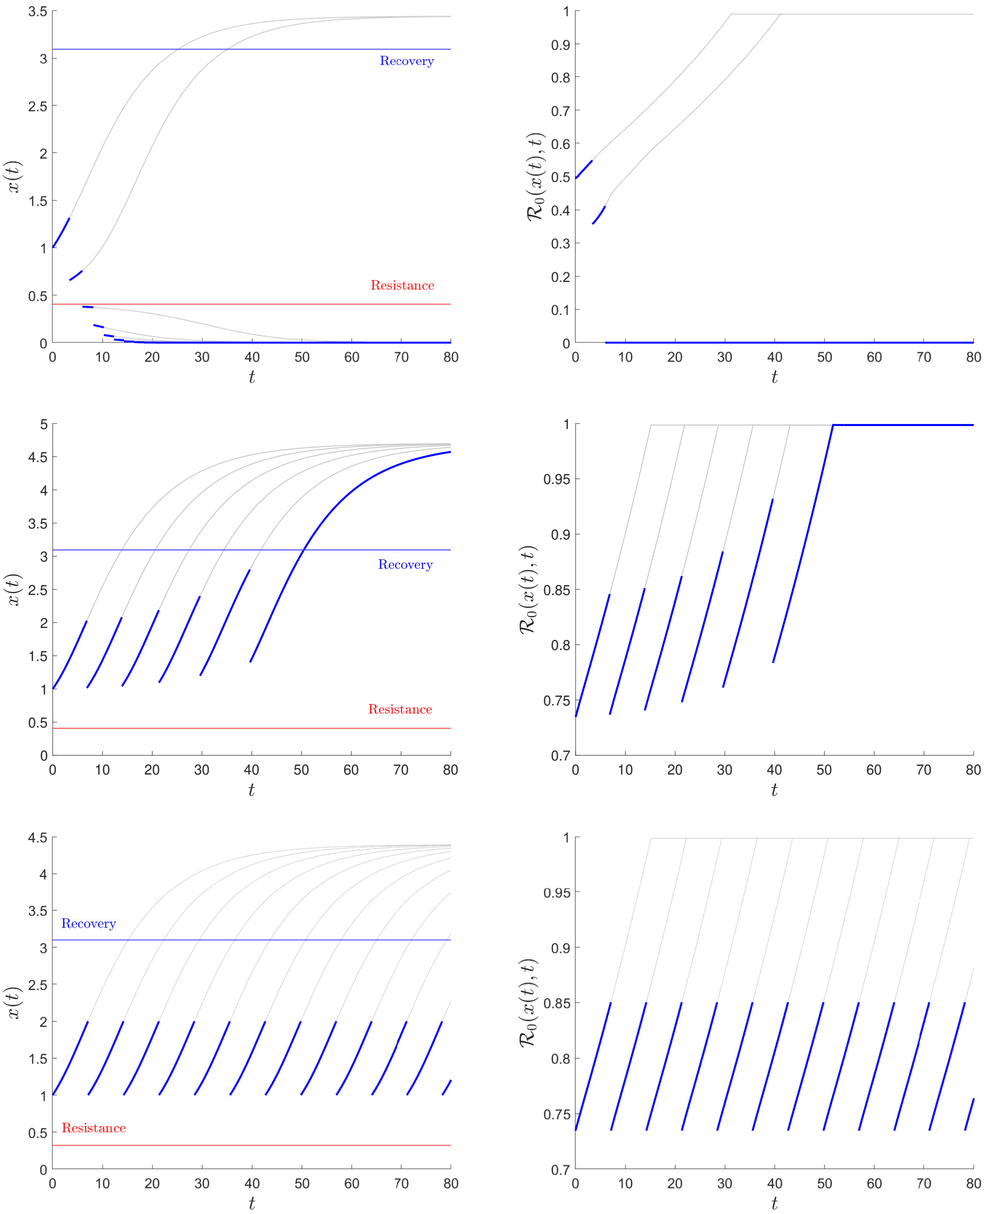

Figure 6: Emergence of resilience for cascading shocks. Parameters: $a = 0.5$, $b = 1.4$, $\varepsilon = 0.5$, $\bar{x} = x^H|_{h=0.13} = 3.4391$, $h = 0.13$ (upper panels) $h = 0.1$ (middle panels), and $h = 0.1061$ (lower panels), $\bar{\eta} = 0.5$, $r = 0.02$, $L(x^*(t)) = [0.9x^H|_{h=0.13}, \infty) = [3.0952, \infty)$, $x_0 = 0.5$.

which the system turns out to be resistant in expected terms to the arrival of a (singular) future shock. This step is quantitatively large and (by construction) only affects continuation resilience. Notably, from this point onwards the lack of robustness that continues to depress current resilience turns out to be less relevant for total resilience, with the system being resistant to future shocks and the remaining drag only being down to incomplete recovery. To the extent that further increases in the initial resource stock also render the system robust against future shocks, current and total resilience converge. We note from this exercise that our resilience measure allows for an intuitive understanding of the relative impact of robustness, resistance and recovery on total resilience for different configurations of the system (in this case in regard to the initial resource stock).

Figure 4 presents the development of resilience over time for a set of different harvesting policies ($h = 0.07, 0.1, 0.13$), again for a set up with a single possible future shock. We consider this jointly with the development of the resource stock and the aggregate harvest over time. Focusing on the two policies with comparatively high harvesting levels, we see that the most ambitious one ($h = 0.13$) allows high levels of the harvest early on but stifles the growth of the resource stock. In comparison to a more conservative policy ($h = 0.1$), this implies not only a lower aggregate harvest in the long-term but also leads to a marked drag on resilience. A comparison between the two conservative policies ($h = 0.07$ and $h = 0.1$) shows that while the harvesting policies can be unambiguously ranked in terms of the resource trajectories and in terms of the resilience they imply, this is not true for the aggregate harvest. Here, very conservative policies may trigger lower levels than less conservative policies not only in the short-term but also in the mid- or long-term.[10]

Our measure allows to study a system's resilience to a cascade of multiple future shocks that can be distinct in their arrival rate and the damage they cause. Figure 5 illustrates this from an ex-ante perspective. Specifically, the figure plots how total resilience $\mathcal{R}(x(t), t)$ emerges over time, conditional on the assumption that no (new) shock has happened up until t. We distinguish scenarios with a single, two and three future shocks, respectively, as well as a scenario with a whole sequence of future shocks.[11] In all settings, total resilience grows in line with the resource stock in a smoothed step-wise fashion. Each of these steps occurs at a point in time at which the resource stock reaches a level at which the system becomes resistant against a sequence of n consecutive shocks out of a sequence of $N \geq 1$ shocks. For a setting with only a single shock ($N = 1$; blue plot) resilience increases by a large amount at around $t = 2.5$ when the resource stock crosses a threshold for which the system is fully resistant against this first (and only) shock $n = 1 = N$ and any further increase in resilience relates only to the minor impact of a further increase in robustness and recovery. Note that the threshold level of the resource stock is directly presented in the left panel in Figure 3. Contrasting this against a setting with two possible shocks ($N = 2$; red plot), we see that while total resilience is also taking a step-wise increase at $t = 2.5$ when the system becomes resilient to the first shock ($n = 1 < 2 = N$), the increase in resilience is less pronounced. This is because full resistance is attained after further accumulation of the resource stock up to a level that renders the system resistant to a sequence of two consecutive shocks ($n = 2 = N$). This only happens around time $t = 10$ where a further step-wise increase is observed and from whereon resilience asymptotically approaches the level

[10] It is easy to see that the long-term (steady-state) aggregate harvest may be lower under conservative policies when considering a harvest rate close to zero. While such a rate leads to the highest possible – but finite – steady state level of the resource stock, it nevertheless sustains a close to zero aggregate harvest.

[11] Strictly speaking the sequence consists of eight shocks, but as we are arguing in the text, due to discounting this already provides a good approximation to a scenario with an infinite number of possible shocks.

of one. The argument naturally extends to the setting with $N = 3$ shocks (green plot) where full resistance is attained only for $n = 3 = N$. More generally, for a setting with $N > 1$ shocks, resilience is built up in line with the system's resistance to an increasing number of n consecutive shocks; or, conversely, total resilience continues to experience a drag in line with the number of $N - n$ consecutive shocks to which the system is not yet resistant. Note, however, that both the step increases and the drag diminish in magnitude over time due to the declining survival probability that is assigned to the system prevailing in its current regime over time. This discounting with the hazard rate also implies an asymptotic time path of resilience for cascades of a large enough (possibly infinite) number of shocks. Altogether, the drag on ex-ante resilience of an additional shock $N+1$ is relevant over an intermediate spell of time between the expected points in time for which full resistance is attained for $n = N$ and $n = N+1$, respectively, but largely diminishes in magnitude with increasing N. This is important from a practical policy perspective, as for the assessment of resilience, the policy maker can focus on a limited number of future shocks.

We conclude our analysis by considering in Figure 6 the emergence of (ex-post) resilience for cascading shocks over time. The figure plots the development of the resource stock (left panel) and total resilience (right panel) under the assumption that shocks reoccur at intervals, the length of which corresponds to the hazard rate $\eta(x(t))$ (as defined by (3)) and evaluated at the time of a shock. Specifically, for a system suffering the n-th shock at time t with a (post-shock) resource stock $x(t)$, the $n+1$-st shock arrives at (precisely) $\mathcal{L}(x(t), t)$ as long as long as the resource stock is below the level at which the hazard rate turns to nil, i.e., as long as $x(t) < \bar{x}$. The two upper and middle row panels depict settings with a relatively high ($h = 0.13$) and low ($h = 0.1$) harvest rate, respectively, whereas the two bottom panels consider an intermediate harvest rate ($h = 0.1061$). Our assumption that the hazard rate itself depends on the resource stock implies that the intervals at which shocks will actually occur vary with the resource stock. Thus, if the resource stock that is left after each shock is diminishing with each additional shock, as it does for the high harvest rate (upper panel), then the hazard of a renewed shock increases over time, leaving less and less time to rebuild the resource stock up until the next shock. As we see from the left upper panel, although resilience is recuperated, following each shock, the fact that the hazard of a renewed shock is escalating with the number of past shocks, this leads to the gradual erosion of resilience over time and to its ultimate collapse at the point at which the system loses resistance to the next shock. Such a process is well compatible with the climate mechanics described in Wunderling et al. (21), where the risk of tipping cascades increases in the duration of temperature overshoot, and with the processes documented by Machado-Silva et al. (14) for the Amazon rain forest, where increasing spells of drought lead to a reduced capacity for rain forest recovery and, thus, a decline in resilience. For the low harvest rate at $h = 0.1$ (middle panel), in contrast, there is enough time, following each shock, to rebuild the resource stock to an extent that the hazard of a renewed shock and, thus, the (expected) frequency of shocks decline. This process allows a long-run, if staggered, build-up of the resource stock that is reflected in the gradual increase in resilience.

These exercises illustrate the scope for policies not only to affect resilience in the short-term, where indeed effects may be limited, but also their scope to erode or nurture resilience in the long run if policies remain unaltered. For the system under consideration, there is a knife-edge harvesting policy at $h = 0.1061$ for which the sequencing of shocks is such that following each shock there is just enough time to rebuild the resource stock to its pre-shock level and resilience is following a stable seesaw pattern, as can be glanced from the bottom panel. It is then easy to infer intuitively that all harvest policies with $h > 0.1061$ embrace a tendency towards the ultimate collapse of the system, whereas policies with $h < 0.1061$ tend to be compatible with a resilient steady-state (for a formal proof see Appendix C). We should caution that the random arrival process of the shock may lead to situation, where a (lucky) long delay of the shock beyond its expected arrival date, may leave enough time to build up a (more) resilient resource stock even under a harvest policy that would lead to the erosion of resilience in expected terms. Conversely, the (unlucky) early advent of a shock may trigger a process towards system collapse even for low harvest rates. The ranking of policies nevertheless applies in expected terms: when re-running the system many times under a given harvest rate, the count of runs at which the system emerges to be resilient in the long-run exceeds (falls short of) the count of runs at which the system collapses when $h < 0.1061$ ($h > 0.1061$).

We conclude the discussion by noting that while the examples illustrate how resilience behaves in respect to certain interesting dimensions, i.e., in terms of time t, initial resource stock $x(0)$, or the harvest rate h our measure allows a representation of resilience along many other dimensions.

5 Discussion and Conclusions

Against a background in which an understanding of resilient policy making continues to be much needed, this note proposes a model-derived measure of resilience that is mathematically rigorous, comprehensive (in terms of including all relevant dimensions of resilience), policy-relevant in the sense that it can be used for the assessment of different policy pathways. While our measure provides a single index of resilience, centred on the zero-one interval, this index (i) can be decomposed into the sub-components resistance, recovery and robustness, allowing for a clear identification of their contribution, and (ii) is forward looking in a way that embraces the system's anticipated response to (endogenous) sequences of future shocks.

Our measure is also general in its conceptualization. While for the purpose of demonstration we apply it to a parsimonious setting of the exploitation of a single resource under the risk of collapse, it is operationalizable for much richer and very different settings. Relevant generalizations include the following. First, by assuming a constant rate of harvesting, we allow only mechanical policy adjustment to changes in the size of the resource stock. The model easily extends to time variable policies, i.e., harvesting rates, or policy rules, specifying, e.g., that the harvesting rate should be adjusted in response e.g., to shocks. The framework can also be applied to optimal policy making, i.e., in the present setting the choice of an optimal control path of harvesting with a view to maximizing an intertemporal profit or utility function.

Second, by reading the state and policy variables as vectors and the resource dynamics as as system of differential equations, our model generalizes to multiple resource and multiple policy settings. This would include e.g., the analysis of common pool problems, where a single resource is exploited by multiple actors; and setting in which a whole portfolio or resources is managed, up to the point of an earth system perspective, in which the vector resources includes renewable and non-renewable resources to be extracted as well as absorbing capacities. We should also stress here that the model would extend to the business economic setting of the management of a portfolio of corporate investments and activities. Depending on the specific context some of these settings may raise computational challenges for the analysis but there is no principle caveat on measuring resilience in such settings. A choice needs to be made, however, as to what aspect of the system would we apply the notion of resilience, too. One way of proceeding is to measure the resilience of each individual state and then aggregate in an appropriate way into an overall index of resilience, where the resilience of states that are essential for the resilience of the system enters in a multiplicative way, whereas the resilience of expendable states enters in a (weighted) additive way. Another way would lie into applying the resilience measure to an overarching outcome, e.g., human welfare or the long-run profit of a firm.

To this end, we note that, third, our measure can be applied to system outcomes rather than states. Thus, in our simple model we could, for instance, specify a utility function $u[hx(t)]$ and define the measure of resilience on the value of this function rather than on the resource state $x(t)$. In settings in which multiple resources contribute to utility, i.e., where $u[h_1x_1(t), h_2x_2(t), ...]$, this allows us to derive at a compact measure of resilience that endogenously takes account of possible substitutability or complementarity across the different resources.

Based on these generalizations, we see much scope for applying our measure of resilience for the assessment of real-world policy contexts. Indeed, there are two possible ways of approaching the calculation of our measure for a given policy context. One would involve the elicitation of expert or stakeholder opinions on the various components of the measure, including (i) the assessment whether or not the current (or an expected) level of the resource stock following a shock lies in the set $L(x^*(t))$ for which a long-term recovery of the system is expected; (ii) an estimate of the expected time up until the arrival of a shock – or equivalently the hazard rate – and an estimate of the damage sustained; (iii) an estimation as to how long the system is taking to recover to reach a desired path, depending on the resource stock and the extent of damage; (iv) an assessment of how a given policy, e.g., the harvesting policy discussed, bears on the renewal of the resource. Based on these pieces of information, the measure can be calculated in an approximate way. Alternatively, based on empirical insight into the key functional relationships, namely the resource renewal function, including the impact of the policy, as well as the hazard and damage functions, one can calibrate and run the model for relevant policy scenarios. Based on the model output one can then calculate the resilience measure. Where not enough data is available, e.g., on the hazard or damage functions, one resort to scenario analysis to examine the robustness of policy resilience against variations in the assumptions on these functions.

Our research agenda foresees a number of the generalizations discussed as well as the application of the measure to assess how different policies bear on systemic resilience in areas of resource management, climate change, and risky economic activity.

References

[1] Angulo, A. M., Mur, J., and Trívez, F. J. Measuring resilience to economic shocks: an application to Spain. The Annals of Regional Science 60, 2 (Mar. 2018), 349–373.

[2] Asheim, G. B., Bossert, W., D'Ambrosio, C., and Vögele, C. The measurement of resilience. Journal of Economic Theory 189 (Sept. 2020), 105104.

[3] Bonanno, G. A. Loss, trauma, and human resilience: have we underestimated the human capacity to thrive after extremely aversive events? American Psychologist 59, 4 (2004), 20–28.

[4] Brown, K., and Westaway, E. Agency, Capacity, and Resilience to Environmental Change: Lessons from Human Development, Well-Being, and Disasters. Annual Review of Environment and Resources 36, 1 (Nov. 2011), 321–342.

[5] Butt, E. W., Baker, J. C. A., Bezerra, F. G. S., Von Randow, C., Aguiar, A. P. D., and Spracklen, D. V. Amazon deforestation causes strong regional warming. Proceedings of the National Academy of Sciences 120, 45 (Nov. 2023), e2309123120.

[6] Folke, C. Resilience: The emergence of a perspective for social–ecological systems analyses. Global Environmental Change 16, 3 (Aug. 2006), 253–267.

[7] Grafton, R. Q., Doyen, L., Béné, C., Borgomeo, E., Brooks, K., Chu, L., Cumming, G. S., Dixon, J., Dovers, S., Garrick, D., Helfgott, A., Jiang, Q., Katic, P., Kompas, T., Little, L. R., Matthews, N., Ringler, C., Squires, D., Steinshamn, S. I., Villasante, S., Wheeler, S., Williams, J., and Wyrwoll, P. R. Realizing resilience for decision-making. Nature Sustainability 2, 10 (Oct. 2019), 907–913.

[8] Hallegatte, S. Economic Resilience: Definition and Measurement. Policy Research Working Papers. The World Bank, May 2014.

[9] Hochrainer-Stigler, S., Velev, S., Laurien, F., Campbell, K., Czajkowski, J., Keating, A., and Mechler, R. Differences in the dynamics of community disaster resilience across the globe. Scientific Reports 11, 1 (Sept. 2021), 17625.

[10] Holling, C. S. Resilience and Stability of Ecological Systems. Annual Review of Ecology and Systematics 4 (1973), 1–23. Publisher: Annual Reviews.

[11] Hosseini, S., and Ivanov, D. A new resilience measure for supply networks with the ripple effect considerations: a Bayesian network approach. Annals of Operations Research 319, 1 (Dec. 2022), 581–607.

[12] Keating, A., Campbell, K., Mechler, R., Magnuszewski, P., Mochizuki, J., Liu, W., Szoenyi, M., and McQuistan, C. Disaster resilience: what it is and how it can engender a meaningful change in development policy. Development Policy Review 35, 1 (Jan. 2017), 65–91.

[13] Li, C.-Z., Crépin, A.-S., and Folke, C. The Economics of Resilience. International Review of Environmental and Resource Economics 11, 4 (Aug. 2018), 309–353.

[14] Machado-Silva, F., Peres, L. F., Gouveia, C. M., Enrich-Prast, A., Peixoto, R. B., Pereira, J. M. C., Marotta, H., Fernandes, P. J. F., and Libonati, R. Drought resilience debt drives npp decline in the amazon forest. Global Biogeochemical Cycles 35, 9 (2021), e2021GB007004. e2021GB007004 2021GB007004.

[15] Mäler, K.-G., Xepapadeas, A., and de Zeeuw, A. The Economics of Shallow Lakes. Environmental and Resource Economics 26, 4 (Dec. 2003), 603–624.

[16] Oliveira, B. M., Boumans, R., Fath, B. D., and Harari, J. Socio-ecological systems modelling of coastal urban area under a changing climate – Case study for Ubatuba, Brazil. Ecological Modelling 468 (June 2022), 109953.

[17] Polasky, S., de Zeeuw, A., and Wagener, F. Optimal management with potential regime shifts. Journal of Environmental Economics and Management 62, 2 (Sept. 2011), 229–240.

[18] Salignac, F., Marjolin, A., Reeve, R., and Muir, K. Conceptualizing and Measuring Financial Resilience: A Multidimensional Framework. Social Indicators Research 145, 1 (Aug. 2019), 17–38.

[19] Skiba, A. K. Optimal Growth with a Convex-Concave Production Function. Econometrica 46, 3 (May 1978), 527.

[20] Wrzaczek, S., Kuhn, M., and Frankovic, I. Using age structure for a multi-stage optimal control model with random switching time. Journal of Optimization Theory and Applications 184, 3 (Mar. 2020), 1065–1082.

[21] Wunderling, N., Winkelmann, R., Rockström, J., Loriani, S., Armstrong McKay, D. I., Ritchie, P. D. L., Sakschewski, B., and Donges, J. F. Global warming overshoots increase risks of climate tipping cascades in a network model. Nature Climate Change 13, 1 (Jan. 2023), 75–82.

[22] Yonson, R., and Noy, I. Disaster risk management policies and the measurement of resilience for philippine regions. Risk Analysis 40, 2 (2020), 254–275.

A Properties of an illustrative model of resource use

In this section we present and investigate the simple resource use model as used in the main body of the paper to motivate the measure on resilience.

The resource stock $x(t)$ at time t increases by a natural rate of growth $g(x(t))$ and decreases by a quantity $H(x(t))$ that is harvested, implying the dynamics:

$$\dot{x}(t) = g(x(t)) - H(x(t)). \tag{7}$$

For the natural growth function we use the classic Holling type III functional response (see, for instance, Mäler et al. (15)) for an ecological context or Skiba (19) for the context of economic growth), which implies convex-concave growth along the stock x. For the harvesting we employ a simple linear feedback rule on the resource stock. Thus,

$$g(x) = \frac{ax^2}{b+x^2} \tag{8a}$$

$$H(x) = hx, \tag{8b}$$

where a, b and h are model parameters. a and b in (8a) refer to g's maximum value (realized as x tends to ∞) and the slope, which becomes flatter for increasing b (implied by $g_b(x) < 0$). The linear parameter h in (8b) depicts the harvesting rate of the resource at t.

Solving $0 = x\left(ax - h\left(b + x^2\right)\right)$ the following three steady states can be characterized analytically:

$$x_1 = 0, \tag{9a}$$

$$x_2 = \frac{a}{2h} - \sqrt{\frac{a^2}{4h^2} - b}, \tag{9b}$$

$$x_3 = \frac{a}{2h} + \sqrt{\frac{a^2}{4h^2} - b}. \tag{9c}$$

The stability of x_i $(i = 1, 2, 3)$ is determined by the Jacobian $J = \frac{2abx}{(b+x^2)^2} - h$ associated with (7).

The low steady state x_1 is always stable. The middle and high steady states only exist for $\frac{a^2}{4h^2} > b$ (positive discriminant) and are strictly positive for $b > 0$. The steady state x_2 is unstable, while x_3 is stable. For $\frac{a^2}{4h^2} = b$ both steady states collide at $x_{2,3} = \frac{a}{2h}$ with 'semi-stability' realizing at a harvesting rate $h^* = \frac{a}{2\sqrt{b}}$, implying stability from the right (i.e., for $x > x_{2,3}$) and instability from the left (i.e., $x < x_{2,3}$).

The stability properties of the three steady states (9) are shown in Figure 7 in the form of a phase diagram in the $(x, \dot{x})$-space for different harvesting rates. For $h = 0.02$ (black), $h = 0.1$ (blue), and $h = 0.2$ (red) all three steady states exist and x_3 is stable. For $h = h^* \approx 0.2113$ the discriminant of (9) is nil and x_2 collides with x_3. The zoom of the phase diagram (right panel) shows the semi-stability. Harvesting rates that exceed h^* converge to the (stable) steady state at the origin x_1. Additionally, Figure 7 shows that (for $h > h^*$) the level of x_2 (which serves as the resistance threshold in our model) increases and x_3 (long-run steady state that is reached, following a resilient resource trajectory) decreases in the harvesting rate.

B General definition of the resilience measure

Within this section we provide a general definition of the resilience measure. As a preparatory step, we present and discuss in the following a number of functions that play a crucial role as part of the measure:

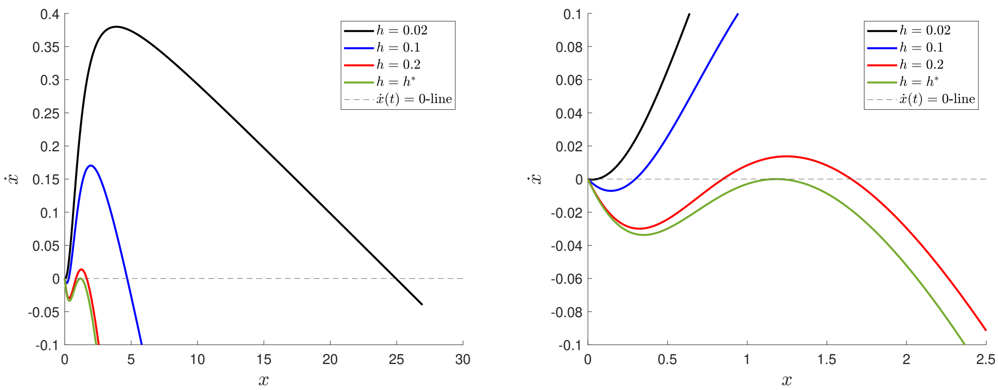

Figure 7: Phase diagram of resource dynamics (7). Trajectories correspond to different values of the harvesting rate ($a = 0.5$, $b = 1.4$). The right panel zooms in.

- Let $y^*(t)$ denote a desired (target) time path for a target variable $y(t)$ (or, indeed, a vector of target variables). What is "desired" here is up to the preferences of the decision-maker. Plausibly, it could involve the steady-state or a balanced growth trajectory of the target variable $y(t)$, but it could also refer e.g., to a safe or long-run resilient outcome. Similarly, while in many plausible cases, the target $y(t)$ can be stated in respect to the underlying state $x(t)$, it could also refer to different outcomes, such as profit or utility. (i) Acknowledging that full attainment of a target path may not always be feasible, e.g., if $y^*(t)$ is defined in terms of an asymptotic steady-state, and (ii) taking account of a possible need to transform units of the target variable $y^*(t)$ into units of the state $x(t)$ for the purpose of keeping track of the dynamics in the state-space, we define $L(t, y^*(t))$ as an acceptable or satisficing target region (short: target region), where the path $x(t)$ is said to lie in the target region at t if $x(t) \in L(t, y^*(t))$. For concreteness, consider the following two examples.

  Example 1: In line with the model sketched in the main body of the paper, consider $x(t)$ as a one-dimensional state and assume that this state also constitutes the target such that $y^*(t) = x^*(t)$. We can then define the target region as $L(t) = L(t, x^*(t)) := [x^*(t) - \delta_1(t), x^*(t) + \delta_2(t)]$ for $t \in [0, \infty)$. The deviation parameters $\delta_i(t)$ ($i = 1, 2$), which may be time specific, define the maximal deviation from $x^*(t)$ for every t. A target region that is bounded only by one side is obtained by appropriate choice of the deviation parameters. E.g., if the resilience measure should count the periods in which the resource stock is not below $\delta \cdot 100\%$ of the target path (as is the assumption in Figures 3-6), $\delta_1(t) = (1 - \delta)x^*(t)$ and $\delta_2(t) = \infty$ implies $L(t) = [\delta x^*(t), \infty)$.

  Example 2: Assume that the target variable is the utility $u(t) = \hat{u}(x(t))$ associated with the state $x(t)$ at t and consider a desired utility level $y^*(t) = u^*(t)$ as the relevant target. We would then say that $x(t)$ lies in the target region if $x(t) \in L(t, u^*(t)) = \left[\hat{u}^{-1}(u^*(t)) - \delta_1(t), u^{-1}(u^*(t)) + \delta_2(t)\right]$ where $\hat{u}^{-1}(\cdot)$ denotes the inverse of the utility function. Note that this is equivalent to $u(x(t)) \in L^u(t, u^*(t)) = [u^*(t) - \delta_1(t), u^*(t) + \delta_2(t)]$ in utility space.

- Based on the previous notions of a target path and target region, we can say that a system is resistant if and only if the state (or vector of states) $x(t)$ follow a path that ultimately leads into the target region.

This is equivalent to requiring the system state(s) $x(t)$ to lie in a region of attraction into the target region. Thus, we define $\mathcal{I}(L(s, y^*(s))) = \mathcal{I}(s, y^*(s))$ as the region of attraction at time s, where $x(s) \in \mathcal{I}(s, y^*(s))$ at time s implies $x(t) \in L(t, y^*(t)) \, \forall t \geq t' \geq s$, with t' denoting the time at which the path enters the target region without leaving it again. We can then employ the indicator function

$$\mathbb{I}_{[x(s) \in \mathcal{I}(s, y^*(s))]} = \begin{cases} 1 & \text{if } x(s) \in \mathcal{I}(s, y^*(s)) \\ 0 & \text{else} \end{cases} \tag{10}$$

to determine as to whether or not the system is resistant. Note that the indicator function can be written in more general form, using the logical 'true' and 'false'. This is important if a set of different metrics are combined for defining the target region.[12]

- We define the hazard rate $\eta(s, x(s))$ as the density of the shock at time s divided by the probability that it has not arrived by s, assuming that it (possibly) depends on time s as well as on the state $x(s)$.

- The function $z(s, t)$ denotes the probability that the disruptive shock does not arrive by s, conditional on it not yet having arrived by $t < s$, i.e., $z(s, t) = \mathbb{P}(\tau \geq s | \tau > t)$. The dynamics of $z(s, t)$ are determined by the hazard rate:

$$\dot{z}(s, t) := -\eta(s, x(s)) z(s, t), \quad z(t, t) = 1, \forall s \geq t. \tag{11}$$

- The life-expectancy measures the time in which the system is expected to remain in the current regime, as of time t, or equivalently the expected time until the next shock, and is defined as $\mathcal{L}(x(t), t) := \int_t^\infty z(s, t) \mathrm{d}s$. Note that $\mathcal{L}(\cdot)$ may be equal to ∞ or take finite values in case the condition $\int_t^\infty \eta(s, x(s)) \mathrm{d}s = \infty$ holds. In case of $\mathcal{L}(\cdot) = \infty$, our expression for current resilience in (16b) implies an abuse of notation, as in this case ∞ would be divided by ∞. For such a case, we therefore define

$$\frac{\mathcal{L}(x(t), t)}{\mathcal{L}(x(t), t) + 1} = \lim_{T \to \infty} \frac{\int_t^T z(s, t) \mathrm{d}s}{\int_t^T z(s, t) \mathrm{d}s + 1}. \tag{12}$$

- The function $g(t', s)$ discounts to the point in time s the time it takes for the trajectory $x(t)$ to enter the target region $L(t, y^*(t))$ at time t'. Here, g has to fulfill the following properties: $g \in [0, 1]$, $g(t', s) = 1$ if $(t' = s)$, $\lim_{t' \to \infty} g(t', s) = 0$ if $s < \infty$, $g_{t'} < 0$. Different forms are possible. A convenient form satisfying the properties is the exponential $(g(t', s) = e^{-r(t'-s)}$, with discount rate $r \in [0, \infty])$. The time at which $x(t)$ enters $L(t, y^*(t))$ can formally be defined as

$$t'(x(s), s) := \inf \{t | x(\bar{t}) \in L(\bar{t}, y^*(\bar{t})) \text{ for } \forall \bar{t} \geq t; \text{ and } \dot{x}(t) = f(x(t), t), t \geq s\}. \tag{13}$$

- Following a shock at τ the measure of resilience is evaluated at the time of the shock τ and the corresponding value of the state $\varphi(\tau, x(\tau))$. The function φ includes the possibility that the state $x(t)$ just prior to the shock suffers a disruptive change (i.e., a drop or a jump) at τ. Formally this means (cf., (20))

$$\lim_{t \to \tau+} x(t) = \lim_{t \to \tau-} \varphi(t, x(t)). \tag{14}$$

[12]Consider, for instance, a setting, where the target variable is an output or utility that is derived from two "input" states $x_1(t)$ and $x_2(t)$. Depending on whether they are complements or substitutes in generating the target variable. In this case resistance would require that $x_1(t)$ *and* $x_2(t)$ lie in their state-specific target region / region of attraction it the states are complements, and that $x_1(t)$ *or* $x_2(t)$ lie in their state-specific target region / region of attraction if they are substitutes.

Based on these features, the resilience measure can be defined as follows.

Definition 1 ***Resilience measure.*** *Let $x(t)$ denote the path of a relevant variable (or vector of variables) of a dynamic process that is subject to a stochastic shock at some time τ that arrives according to the exogenous or endogenous hazard function $\eta(t, x(t))$. The resilience of $x(t)$ at time t is then measured by*

$$\mathcal{R}(x(t),t) = \frac{1}{\mathbb{E}_{\tau\in[t,\infty)} \int_t^\tau 1 ds + 1} \mathbb{E}_{\tau\in[t,\infty)} \left[\int_t^\tau \mathbb{I}_{[x(t)\in\mathcal{I}(t,y^*(t))]} g\left(t'(x(t),t),t\right) ds + \mathcal{R}\left(\varphi(\tau,x(\tau)),\tau\right) \right]. \quad (15)$$

Using basic probability theory $\mathcal{R}(x(t),t)$ as defined in (15) can be transformed to the following deterministic representation.

Corollary 2 *The resilience measure as defined in Definition 1 can be equivalently represented by*

$$\mathcal{R}(x(t),t) = \mathcal{R}_1(x(t),t) + \mathcal{R}_2(x(t),t) \quad (16a)$$

where

$$\mathcal{R}_1(x(t),t) = \frac{\mathcal{L}(x(t),t)}{\mathcal{L}(x(t),t)+1} \mathbb{I}_{[x(t)\in\mathcal{I}(t,y^*(t))]} g\left(t'(x(t),t),t\right), \quad (16b)$$

$$\mathcal{R}_2(x(t),t) = \frac{1}{\mathcal{L}(x(t),t)+1} \int_t^\infty z(s,t)\eta(s,x(s))\mathcal{R}(\varphi(s,x(s)),s) ds, \quad (16c)$$

define the current (expr. (16b)) and continuation (expr. (16c)) resilience of $x(t)$ at t.

Proof: Let us first discuss the denominator of the fraction of (15) which can be written as

$$\mathbb{E}_{\tau\in[t,\infty)} \int_t^\tau 1\mathrm{d}s = \int_t^\infty \left\{ \int_t^\tau 1\mathrm{d}s \right\} \eta(\tau,x(\tau)) e^{-\int_t^\tau \eta(t,x(t))\mathrm{d}t} \mathrm{d}\tau. \quad (17)$$

Integrating by parts and using the assumption

$$\lim_{\tau\to\infty} \int_t^\tau 1\mathrm{d}s \cdot e^{-\int_t^\tau \eta(t,x(t))\mathrm{d}t} = 0 \quad (18)$$

yields

$$\begin{aligned} \mathbb{E}_{\tau\in[t,\infty)} \int_t^\tau 1\mathrm{d}s &= \int_t^\infty e^{-\int_t^\tau \eta(t,x(t))\mathrm{d}t} \mathrm{d}\tau \\ &= \int_t^\infty z(\tau,t)\mathrm{d}\tau \\ &= \mathcal{L}\left(x(t),t\right). \end{aligned} \quad (19)$$

Plugging in the fraction proves the for of (16b) and (16c). Secondly, we consider the second term of (15) for which we can apply the same method (already using the definition of $z(s,t)$):

$$\begin{aligned} &\mathbb{E}_{\tau\in[t,\infty)} \left[\int_t^\tau \mathbb{I}_{[x(t)\in\mathcal{I}(t,y^*(t))]} g\left(t'(x(t),t),t\right) \mathrm{d}s + \mathcal{R}\left(\varphi(\tau,x(\tau)),\tau\right) \right] \\ =\ & \int_t^\infty \left[\int_t^\tau \mathbb{I}_{[x(t)\in\mathcal{I}(t,y^*(t))]} g\left(t'(x(t),t),t\right) \mathrm{d}s \right] z(\tau,t)\eta(\tau,x(\tau))\mathrm{d}\tau \\ & + \int_t^\infty z(\tau,t)\eta(\tau,x(\tau))\mathcal{R}\left(\varphi(\tau,x(\tau)),\tau\right) \mathrm{d}\tau \\ =\ & \mathbb{I}_{[x(t)\in\mathcal{I}(t,y^*(t))]} g\left(t'(x(t),t),t\right) \int_t^\infty z(\tau,t)\mathrm{d}\tau \\ & + \int_t^\infty z(\tau,t)\eta(\tau,x(\tau))\mathcal{R}\left(\varphi(\tau,x(\tau)),\tau\right) \mathrm{d}\tau. \end{aligned} \quad (20)$$

Using the definition of the life expectancy and dividing by $\mathcal{L}(x(t),t)+1$ proves the corollary. □

C Proof of stable harvest rate: Existence and stability

This section proves the existence and stability of a "stable" harvest rate that is defined by (a) it generating a balanced sequence of (equally sized) shocks, each of which leaves exactly the time it takes to rebuild the resource stock $x(t)$ to a level such that the next shock reduces it to the level it had assumed following the last shock; and thus by (b) generating stable see-saw trajectory of the state.

Lemma 3 ***Existence.*** *For a fixed set of model parameters (a, b, ε, x_0) and hazard rate function η there exists a unique* ***stable harvest rate*** $\bar{h}\,(a, b, \varepsilon, \eta, x_0) < h^* = \frac{a}{2\sqrt{b}}$ *such that the resource stock drops to the initial stock x_0 if the shock arrives exactly in line with the life expectancy $\mathcal{L}(x_0, 0)$, i.e. the expected time of the system in its current regime. Hence,*

$$x(\tau+) = \varepsilon x(\tau-) = x(0) = x_0, \qquad \textit{for } \tau = \mathcal{L}(x_0, 0) = \int_0^\tau z(s, 0)ds, \ \textit{and } h = \bar{h}\,(a, b, \varepsilon, \eta, x_0)\,. \tag{21}$$

Proof: Let us denote the dependence of the life-expectancy on the harvest rate as $\mathcal{L}(x_0, 0; h)$. From the derivative $\mathcal{L}_h = -\int_0^\tau z(s, 0)\eta_x(s, x(s))\frac{\partial x(s)}{\partial h}\mathrm{d}s$ we have $\mathcal{L}_h < 0$, as $\eta_x \leq 0$ and $\frac{\partial x(s)}{\partial h} < 0$. For $h = 0$ the life-expectancy reaches a maximum level $\bar{\mathcal{L}}$, and for h tending to ∞ we have $\mathcal{L}(x_0, 0; \infty) = 0$. Furthermore, $\mathcal{L}(x_0, 0; h)$ is continuous in h.

The time $\tau(x_0, h)$ at which the resource level reaches $\frac{x_0}{\varepsilon}$ is implicitly defined by

$$\frac{x_0}{\varepsilon} \quad = \quad x_0 + \int_0^\tau \dot{x}(s)\mathrm{d}s \tag{22}$$

and also a function of h, where $\tau_h > 0$. Here, $h = \infty$ implies $\tau = \infty$, and $h = 0$ implies $\underline{\tau}$ (minimal time to reach the $\frac{x_0}{\varepsilon}$).[13] Again $\tau(x_0, h)$ is continuous in h.

The above conditions for $\mathcal{L}\,(\cdot)$ and $\tau\,(\cdot)$, both functions of h, allow for the application of the mean value theorem, which proves a unique intersection point denoted by $\bar{h}\,(a, b, \varepsilon, \eta, x_0)$. $\square$

Lemma 4 ***Stability.*** *For a fixed set of model parameters (a, b, ε, x_0), hazard rate function η and a stable harvest rate $\bar{h}\,(a, b, \varepsilon, \eta, \bar{x})$ the corresponding resource stock $\bar{x}$ is stable, i.e., for $x_0 = \bar{x} + \kappa$ ($\kappa > 0$ small)*

$$\bar{x} \leq x\left(\mathcal{L}(x_0, 0; \bar{h})+\right) = \varepsilon x\left(\mathcal{L}(x_0, 0; \bar{h})-\right) < x_0, \qquad \textit{for } \tau \textit{ defined by } \frac{x_0}{\varepsilon} = x_0 + \int_0^\tau \dot{x}(s)ds,$$
$$\textit{and } h = \bar{h}\,(a, b, \varepsilon, \eta, \bar{x})\,, \tag{23}$$

and opposite for $\bar{x} > x_0 = \bar{x} - \kappa$, if $\tau_x(\bar{x}, \bar{h}) > \mathcal{L}_x(\bar{x}, 0; \bar{h})$ holds.

Proof: For the proof we consider a small deviation from $\bar{x}$ by a small positive value κ, i.e., $x_0 = \bar{x} + \kappa$, and recall from Lemma 3 that $\bar{h}$ is defined by the condition $\tau(\bar{x}; \bar{h}) = \mathcal{L}(\bar{x}, 0; \bar{h})$ (see (21)). The resource stock follows a stable see-saw pattern on the interval $x(t) \in [x_0, x_0/\varepsilon]$, where $\bar{x} = x_0$, if (i) the expected life-expectancy of a non-disrupted path from x_0 is lower than the time that would be needed to reach the resource level $\frac{x_0}{\varepsilon}$, and if (ii) the resource stock after the shock at $\tau(x_0, \bar{h})$ is lower than x_0, but bounded from below by $\bar{x}$. Condition (i) basically means that the next shock on a path that is starting out from x_0) sets in before the resource stock reaches the level x_0/ε at which the shock is exactly compensated. This is implied by $\tau_x(x_0) > \mathcal{L}_x(x_0, 0)$.

[13] Note here that $h = \infty$ implies the instantaneous and complete depletion of the stock at $t = 0$ such that $x(s) = \dot{x}(s) = 0 \forall s \geq 0$ and by implication $\tau = \infty$.

Condition (ii) means that the sequence of starting values of the resource sub-trajectories following each of a sequence of shocks converges to $\bar{x}$. This is again implied by the same condition $\tau_x(x_0) > \mathcal{L}_x(x_0, 0)$, which in turn implies $\tau\left(x_0, \bar{h}\right) > \mathcal{L}\left(x_0, 0; \bar{h}\right)$. $\square$