



Contents lists available at ScienceDirect

Journal of Economic Behavior and Organization

journal homepage: www.elsevier.com/locate/jebo

Time-consistent renewable resource management with present bias and regime shifts[☆]

Maria Arvaniti^{a,1}, Chandra Kiran B. Krishnamurthy^{c,e,*}, Anne-Sophie Crépin^{b,d}

^a Department of Economics, University of Bologna, Italy

^b Beijer Institute for Ecological Economics, The Royal Swedish Academy of Sciences, Stockholm, Sweden

^c Department of Forest Economics, Swedish University of Agricultural Sciences, Skogsmarksgränd 27, 90183, Umeå, Sweden

^d Stockholm Resilience Centre, Stockholm University, Stockholm, Sweden

^e Center for Environmental and Resource Economics (CERE), Umeå, Sweden

ARTICLE INFO

Article history:

Received 13 April 2021

Revised 6 January 2023

Accepted 18 January 2023

Available online 8 February 2023

JEL classification:

Q20

C61

C73

Keywords:

Renewable resources

Regime shifts

Hyperbolic discounting

Present bias

Uncertainty

Markov equilibrium

ABSTRACT

We investigate the extraction plan of present-biased decision makers managing a renewable resource stock whose growth is uncertain and which could undergo a rapid and significant change when stock falls below a threshold. We show that the Markov-Nash equilibrium extraction policy is unique, time consistent, and increasing in resource stock. An increase in the threshold leads to increased resource extraction, rather than the precautionary reduction in extraction often observed with exponential discounting. An increase in the degree of present bias also leads to an increase in resource extraction. Our analysis suggests that accounting for and appropriately dealing with resource managers' present bias may be important to understand resource use sustainability.

© 2023 The Author(s). Published by Elsevier B.V.
This is an open access article under the CC BY license
(<http://creativecommons.org/licenses/by/4.0/>)

1. Introduction

Managing renewable natural resources over long time horizons raises two special challenges: the first pertains to accounting for potentially complex ecosystem dynamics and their effects upon preferences while the second involves questions related to discounting. As to the first, renewable resource dynamics often involve natural variability in resource growth and potential critical transitions that could trigger abrupt changes in the ecosystems providing the resource. These abrupt

[☆] We thank seminar participants at Umeå University, ETH Zurich and AUEB, and conference participants at ASSET 2018, FSR Climate Conference 2018, International Workshop on The Economics of Climate Change and Sustainability, University of Bologna 2019, EAERE 2019, CRETE 2019, FAERE Conference 2019, the 5th International Workshop on Economic Growth, Environment and Natural Resource at the European University St.Petersburg for many valuable comments and suggestions. We are also grateful to Aart de Zeeuw, Łukasz Woźny and Niko Samuli Jaakkola for valuable discussions and to two anonymous reviewers and the Editor for very helpful comments. The usual disclaimer applies.

* Corresponding author.

E-mail address: chandra.kiran@slu.se (C.K.B. Krishnamurthy).

¹ The project was partly carried out when the author was the holder of a fellowship of the Tore Browaldh Foundation (Handelsbanken foundation). Support from the foundation is gratefully acknowledged.

changes in the structure and function of natural ecosystems – called “regime shifts” – are pervasive. Fisheries can collapse (Pershing et al. (2015)); foragers can hinder forest regrowth (Crépin, 2003); a latent infectious disease or a pest can shift to an outbreak rapidly spreading in the population (Ludwig et al., 1978); a coral reef can exist in multiple regimes (clear or algae-dominated, Norström et al. (2009)); essential elements of the climate system (e.g. sea ice, monsoon) can exist in different regimes (Steffen et al., 2018).

The second aspect, related to discounting, is of considerable importance in view of the long time horizons over which many renewable resources are to be managed. Experimental and empirical evidence suggest that decreasing, not constant, discount rates provide a more plausible representation of individual discounting over time. Policy makers and regulators, who may be more responsive to current voters and their concerns, are likely to over-weight benefits accruing today relative to those in the more distant future, leading to a specific form of declining discount rates. In consequence, the policy framework resembles an inter-generational strategic setting between regulators who make decisions today and those who make decisions tomorrow. We examine the implications for resource extraction policy of decision makers who are “present-biased” (in applying a higher discount rate for the immediate future, relative to the distant future) when resource dynamics are complex, in the sense of involving regime shifts. Our main objective is to understand if regulators exhibit “precautionary behaviour”, in the sense of reducing extraction in anticipation of a regime shift, and, more broadly, how present bias affects resource extraction policy with complex resource dynamics. Current regulators are unable to commit regulators in the distant future to follow a pre-arranged plan. In consequence, we restrict attention to policies that do not involve a specific commitment device, and focus on stationary Markov-Nash equilibria which are time-consistent by construction. In lieu of parametric classes of models, we make use of general properties of resource models with regime shifts and work with general utility functions.

The effects of regime shifts and non-constant discounting for resource and environmental questions have been investigated separately (see e.g., Iverson and Karp (2021); Gerlagh and Liski (2017) for non-constant discounting and Karp and Tsur (2011); Ren and Polasky (2014); Brozović and Schlenker (2011) for regime shifts), but little is known regarding how the two features together influence resource management. While abrupt change can sometimes trigger greater precaution in resource extraction, present bias may be anticipated to incentivise an increase in extraction, meaning that the overall effect on extraction is unclear. Regarding the fundamental question, “Do regime shifts motivate precautionary behaviour”, economic theory remains ambiguous. Many important aspects influence precautionary behaviour, including the ability of resource users to affect the risk of a regime shift, the impact of a regime shift on resource dynamics, and its degree of reversibility (see Polasky et al. (2011); de Zeeuw and He (2017)). Experimental evidence, on the other hand, suggests a tendency for precautionary behaviour amongst groups collectively managing a resource (Lindahl et al., 2016; Schill and Rocha, 2019). Different perceptions of inter-temporal trade-offs could perhaps be one explanation for some of the discrepancy between experimental findings and those from economic theory. A specific pattern of “impatience” is often associated with inter-temporal resource management, with a bias for consumption in the immediate future over the intermediate and long-run, called “present bias”. Nonetheless, the literature on dynamic resource management with a variety of regime shifts has exclusively assumed constant time-discounting, with only one study (Karp and Tsur, 2011) considering a regime shift in a hyperbolic discounting set-up, with significantly reduced action space.

Our objective is to address this gap in the literature: We investigate the renewable resource extraction problem faced by a forward-looking regulator who is present-biased. We consider regime shifts, which are: reversible, meaning that resource growth returns to its original trajectory when resource stock rises beyond a threshold; endogenous, with threshold-crossing depending upon the regulators’ actions; and stochastic, in that random shocks affect resource growth. Resource growth exhibits a rapid and substantial change when entering a fixed region of the state space defined by a known threshold stock level, and this constitutes the regime shift. In some ecosystems, the regime shift occurs immediately after the threshold is crossed, while in others, it occurs once the stock after regeneration falls below the threshold. Our framework accommodates both types of ecosystems.

Our analysis makes use of key features of the stock transition function and highlights the significance of the relationship between stock and the threshold, and stock and reinvestment for equilibrium extraction. Using very general utility functions, and adopting general stochastic production functions whose properties encompass those of special functional forms used in the current literature, we are able to establish the existence of a unique equilibrium-extraction policy. Our approach, which follows the literature using recursive approaches to non-constant discounting (Balbus et al., 2014; 2018), differs from those in the previous theoretical literature, enabling us to derive novel results. It also enables us to relate the regime shift problem to the canonical consumption-savings problems in a quasi-hyperbolic context.

To our knowledge, we are the first to establish key properties of the equilibrium policy for quasi-hyperbolic stochastic resource problems with regime shifts. Our first main finding is that the regulator increases extraction in response to an increase in the threshold. A key aspect driving this result is the Edgeworth complementarity between resource reinvestment and the threshold: reinvestment is more beneficial for thresholds at lower resource stock levels. Indeed the risk of reinvesting in a stock that may grow much more slowly in the future is smaller when the threshold is at a lower stock level. Our second major result is that equilibrium extraction is monotonically increasing in the resource stock. This finding arises from the Edgeworth substitutability between reinvestment and stock: the (marginal) returns to increased reinvestment are decreasing in stock, independent of the threshold, a feature of most resource economic and stochastic growth models. Our third major finding is that equilibrium extraction increases with the degree of present bias. Furthermore, we are also able

to establish the existence of a non-trivial stationary distribution and to characterise it, finding it to be decreasing in the threshold.

All of these results are novel: as to the relationship between extraction and the threshold, findings in the previous literature (using exponential discounting) suggest either a non-monotonic (Brozović and Schlenker, 2011) or a decreasing (Polasky et al., 2011) relationship. No previous study, to our knowledge, has established the direction of the relationship between equilibrium extraction and resource stock or present bias (with or without regime shifts): both have only been conjectured in a previous study without regime shifts (Karp, 2005).

The rest of the paper proceeds as follows: Section 2 provides an overview of the related literature on modelling regime shifts and hyperbolic discounting, while Section 3 details key aspects of our model set up. Section 4 explores the formulation where regime shifts occur with a lag, while Section 5 details the more common case of no lag. Section 6 provides a discussion of how our findings relate to those in the broader literature and Section 7 concludes.

Key proofs are provided in an Appendix while further technical details and proofs are provided in an Online Appendix.

2. Related literature

Our analysis builds on at least two distinct strands of the literature: The first relates to whether precautionary behaviour is an optimal response when managing renewable resources or choosing (carbon) emission mitigation strategies when faced with regime shifts or catastrophes. The second relates to a more general question examining the effects of non-exponential discounting on equilibrium policies, largely without regime shifts. We provide a brief review of both strands, beginning with an overview of different formulations of a regime shift.

2.1. Modelling regime shifts and catastrophes

The most common way of modelling regime shifts in economics is to consider them as penalty functions with an associated hazard rate, triggering reversible or irreversible catastrophic events, as in Clarke and Reed (1994), and Tsur and Zemel (1996, 1998). A common way of solving such models is to transform them into a deterministic control problem with the associated survival probability as the state variable (see Crépin and Nævdal (2019) for an overview). A different approach is to model a regime shift as a change in the system dynamics, often in a deterministic setup using continuous time models with convex-concave growth function generating reinforcing dynamics (Wagener, 2003; Mäler et al., 2003; Brock and Starrett, 2003). Precautionary behaviour is then an optimal strategy for endogenous regime shifts only, with exogenous regime shifts not affecting management (see e.g., Polasky et al. (2011)).

We introduce a regime shift into the framework often used in discrete-time stochastic renewable resource models: resource growth is subject to uncertainty, the regime shift threshold is fixed and known, and a regime shift can occur either before or after the extraction decision is made. A few prior studies take this approach to modeling regime shifts (e.g., Peterson et al. (2003); Brozović and Schlenker (2011); Ren and Polasky (2014)), under special conditions including specific functional forms. Departing from the use of specific functional forms for stock transition, we instead represent the “production function” leading to next period stock in an inherently stochastic manner. This implies that regime shifts now occur when entering (exiting) a fixed undesirable region of the state space and lead to a rapid and very substantial change in the distribution function of the resource stock.

2.2. Regime shifts and precautionary behaviour

As already described, it is generally unclear whether potential regime shifts should trigger “precautionary behaviour” (a reduction in extraction induced by the possibility of regime shifts) in view of the inherent trade-offs involved: while the increased possibility of maintaining a resource or pollutant stock at an appropriate level can incentivise a precautionary reduction of extraction, the loss of consumption opportunities associated with this behaviour provides an incentive to increase extraction. In essence, this is the resource economic counterpart (with regime shifts) to the consumption-saving trade-off. The literature (see Li et al. (2018) for an overview) reflects this mixed intuition and sheds some light on specific factors likely to influence outcomes.

Polasky et al. (2011), in a continuous-time regime-shift framework with stochastic timing, find that the optimal policy always turns “precautionary” only in the case of endogenous regimes shifts resulting in changed system dynamics. If instead an endogenous regime shift triggers a stock collapse, the degree of precautionary behaviour depends on the outcome of the trade-off between the desire for precaution and the wish to secure harvest while the resource is still available. No precautionary action is taken for exogenous regime shifts and it is in fact optimal to increase harvesting, when faced with a stock collapse in this case.

Brozović and Schlenker (2011) use instead a discrete-time model of pollution loading, with an additive noise term leading to increased loading once a possibly unknown reversible threshold is crossed. They report a non-monotonic relationship between precautionary activity and uncertainty. An increase in uncertainty may first lead to increased precaution but will always lead to reduced precautionary activities in the long run. Additional precaution here has an added benefit: the expected pollutant stock drops as the threshold is approached. Delays between the triggering event and its impact in the form

of a regime shifts is another factor that can influence the degree of precaution (Crépin and Nævdal, 2019). All of these models assume linearity of benefit from harvest. Using a utility function concave in harvest and a more general resource growth function, Ren and Polasky (2014) and de Zeeuw and He (2017) find that precautionary behaviour is not guaranteed even in the case of endogenous regimes shifts with altered system dynamics.

2.3. Discounting

Our analysis posits that regulators' discounting behaviour exhibits present bias, which may arise for many reasons, with one interpretation closer to many resource extraction and climate change settings: Regulators may discount utility gains within their own generation differently from those later on, incentivising procrastination when taking decisions that impose current costs for future benefits.²

Few studies in the literature address environmental problems with non-constant discounting, and those that do are often unable to characterise the equilibrium or time-consistent policies. Prior studies on resource policy with a stock variable and (quasi-) hyperbolic discounting either restricted attention to a very simplified set of policies (Karp and Tsur, 2011), or focused on numerical characterisation (Karp (2005)), and simulation frameworks (Fujii and Karp, 2008). To our knowledge, Gerlagh and Liski (2017) is the first study that characterises the equilibrium in the context of a stock variable. They focus on optimal climate and capital policy in a multi-sector deterministic growth framework, in a setting with no regime shifts and with very specific functional forms. Iverson and Karp (2021) extend this model to allow for general discounting functions and functional forms. Karp and Tsur (2011) is the only study to embed hyperbolic discounting in a continuous-time model of catastrophic climate related damages. An environmental catastrophe leads to a permanent loss of income, while the level of hazard (and hence the probability of the event) cannot be reduced. For tractability, the planner is restricted to two choices, either abate and stabilise the hazard level or follow business as usual, letting the hazard reach its maximal level. A closed form solution for the resulting equilibrium is obtained but little is known regarding uniqueness or other qualitative properties of the equilibria or time-consistent policies. In any case, all previous studies consider largely deterministic state evolution.

A large literature in economics studies the dynamic problem of “present-biased” regulators who discount utility gains during their own generation differently from those after their time (with resulting discounting also called “quasi-hyperbolic”). Furthermore, the regulator is assumed unable to commit to future actions and cannot dictate the decisions of future regulators. The traditional approach to dynamic decision problems, using recursive decision theory, is challenging with (quasi-) hyperbolic discounting, due to time-inconsistency (see e.g., Krusell et al. (2002); Klein et al. (2008); Karp (2005)): What is optimal for a decision maker at time t is not optimal for the decision maker at time $s \neq t$. Recent studies have attempted to use the approach of stochastic games and apply recursive methods to problems with hyperbolic decision makers, restricting attention to the more practically useful pure strategies (Balbus et al., 2014; 2018). Our analysis builds upon the analytical framework of these studies. In this setting of stochastic dynamic (exponentially) discounted games, the appropriate notion of equilibrium is that of a Stationary Markov Nash Equilibrium (SMNE).

A key advantage of this approach is that, if it exists at all, the SMNE (henceforth simply “equilibrium”) is time-consistent.³ Our analysis will use this notion of equilibrium.

3. Model details

Our set-up consists of a standard regime shift model in discrete-time embedded in a stochastic quasi-hyperbolic capital accumulation setting. It is conventional in the literature to consider a decision maker who views himself as a sequence of “selves” in discrete time, $t \in \{0, 1, \dots\}$. Let $X_t \in \mathcal{X} \subset \mathbb{R}^+$ denote a random variable for the resource stock level at time t , with x the realised stock. Consider a situation when after observing the stock $X_t = x$ at the beginning of period t , the decision maker chooses an extraction level, $q_t := q(X_t) \in [0, x)$, and leaves $a_t := x - q_t \equiv a(X_t)$ as the reinvestment (or residual stock, at end of period t). Reinvestment and initial stock levels, together with the threshold level of stock \underline{X} lead to the next period stock, X_{t+1} via the transition function (or stochastic kernel) $Q(dX_{t+1}|a, x, \underline{X})$. In our setting, the regime shift is such that realised stock levels below \underline{X} trigger a shift to substantially lower stock regrowth compared to that above the threshold. Therefore, instead of specifying X_{t+1} as arising from the addition of an “error term” to a deterministic growth function or objective function, we directly specify the next period stock via the following stochastic specification:

$$X_{t+1} \sim Q(\cdot | a_t = a, X_t = x, \underline{X}). \quad (1)$$

² Many institutions of collective decision making are unable to commit future decision makers. In the case of climate change, for example, Canada withdrew from the Kyoto Protocol a few years after its signing and Japan opted out of its second Phase, while the USA opted out of the Paris Agreement a few years after signing it. Examples drawn from resource management relate to the Atlantic cod collapse in the early 1990's and the Peruvian anchovy collapse in the 1970's. In both cases, time inconsistency, particularly a bias for the immediate present have been suggested to have influenced the outcomes (Heppburn et al. (2010)).

³ The notion of *Stationary Markov Nash Equilibrium* encompasses three distinct concepts. A strategy is Markov if it depends on history only through the current state and it is stationary if it is time invariant. A Nash equilibrium must dominate all potential strategies by definition. Hence a Stationary Markovian strategy that is also a Nash equilibrium is time consistent. See Dutta and Sundaram (1998) for precise definitions and detailed discussion of these notions. An alternative definition of equilibrium, Markov Perfect (or the Markov Nash) Equilibrium, is also used in the literature (e.g., Karp and Tsur (2011)). These equilibria need not be stationary (by definition), and consequently, for hyperbolic discounting agents, not time-consistent.

The transition function, Q , maps the state space to itself and defines a probability distribution over the next period stock.

Consistent with the previous literature, regime shifts are taken to affect the resource stock directly, and utility U , derived from extraction, is continuous, increasing, strictly concave, and bounded, with $U(0) = 0$ and $U(\cdot) \leq \bar{U}$. Finally, parameters $0 < \delta \leq 1$ and $0 < \beta \leq 1$ denote respectively the exponential discount factor and the degree of present bias. Preferences of the decision maker at time t are represented as:

$$U(q_t) + \beta \mathbb{E}_t \left(\sum_{i=t+1}^{\infty} \delta^{i-t} U(q_i) \right), \tag{2}$$

with \mathbb{E}_t the expectation taken w.r.t. the time- t distribution of X_t .

The discounting embodied in Eq. (2) can be explained as follows: starting at time period t , the decision maker uses the discount factor δ to compare pay-offs (consumption) between any two adjacent periods beyond $t + 1$ (e.g., between $t + 2$ and $t + 3$) while using the factor $\beta\delta$ to compare outcomes between period t and $t + 1$. This leads to the following series of discount functions: $1, \beta\delta, \beta\delta^2, \beta\delta^3, \dots, \beta\delta^t, \dots$. For any $\beta < 1$, this discounting framework represents declining discount rates: the “short-run” discount rate, $-\ln \beta\delta$ is larger than the “long-run” discount rate $-\ln \delta$. Lower values of β represent stronger bias for the present, and preferences are time-inconsistent for any value of $\beta < 1$. Preferences exhibiting these characteristics are termed quasi-hyperbolic or present biased. The interpretation of ‘present bias’ follows from the observation that all future pay-offs are discontinuously discounted by the amount β , relative to the case of exponential discounting (see e.g., Benhabib et al. (2010); Angeletos et al. (2001)). While discount rates that decline over time may appear desirable from a long-run decision making perspective, this form of discounting in fact strengthens the preference for immediate pay-off, relative to exponential discounting.

Following the previous literature, we assume that the regulator shares the consumer’s preferences,⁴ and cannot commit to future actions. In this case, the best a regulator can do is to play a dynamic game with future regulators and manipulate their decisions through his own policies for the present. Of course, if the regulator could commit to future policies, she could achieve a higher “ex-ante” welfare i.e., from the perspective of the current self’s preferences. In that case, the “commitment” plan would involve two different extraction policies, one for the current self and one for all future selves. The extraction policy for the current self would be the same as for a “naive” regulator i.e., the extraction level chosen expecting that future selves will follow through. For the future selves, the extraction policy is the same as with standard exponential discounting, i.e., $\beta = 1$.⁵ However, it is not at all evident how the regulator could provide a commitment mechanism, in essence restricting the actions of future regulators whose preferences differ from those of the current regulator. As Iverson and Karp (2021) show, even if such a mechanism existed it would have very little value unless it was permanent. In any case, given the lack of a commitment mechanism for regulators, we opt for a time consistent equilibrium, which is the only sustainable policy for the management of the resource stock (since future regulators would deviate from any other plan).⁶ We return to this aspect briefly in Section 5.

3.1. Uncertainty and regime shifts in resource dynamics

It will be useful to first understand how uncertainty and regime shifts in resource growth models are commonly modelled. Reversible regime shifts are often represented as altered system dynamics once a threshold is crossed. Let \underline{X} denote a generic parameter that governs a change in system regime, e.g., a change in some parametric specification of the production function such as carrying capacity (Ren and Polasky, 2014) or pollutant loading (Brozović and Schlenker, 2011). Regime shifts can then be formalised as the effect of this parameter on a generic resource production function, $L(a, x, \epsilon; \underline{X})$, which transforms current stock and reinvestment (a, x) into next period stock with a random growth term ϵ following some distribution. In our context, this is a threshold level of stock that divides the state space, $\mathcal{X} \subset \mathcal{R}_+$, into two regions, a desirable one, (\underline{X}, ∞) , and an undesirable one, $(0, \underline{X}]$, with $\underline{X} = 0$ corresponding to the case of no regime shift. Then, for every stationary policy, $a(X_t)$, the resulting stock series, $\{X_t\}_{t=1}^{\infty}$, can be shown to lead to a Markov chain with transition function Q as in Eq. (1).⁷ In the context of economic dynamics with a production interpretation (Amir, 1997), the direct specification

⁴ This approach abstracts away from the question of which set of preferences the planner would like to maximise. If one believes that present bias represents a behavioural failure that an agent would like to correct, then a planner would act paternalistically and opt to maximise the “long-run” or normative preferences: the preferences an agent has without present bias (see e.g., Kang (2019)).

⁵ Krusell et al. (2002) and Gerlagh and Liski (2017) show that the full commitment solution would consist of one policy for today -the same as in the equilibrium- and a different policy for all future periods, the optimal policy with exponential discounting ((footnote 23 characterises these policies for our case). This is because the “naive” and “sophisticated” (SMNE) solutions coincide when utility is logarithmic as it is assumed in their framework. Karp and Tsur (2011) explore the idea of “restricted” commitment where the decision maker cannot switch between policies as a more realistic alternative.

⁶ Hepburn et al. (2010) show how following any other policy could result in an unforeseen resource collapse. Note that the inability to commit cannot be alleviated by using policy instruments e.g., an extraction tax: in the absence of an externality, any corrective policy undertaken by the planner will not alter the chosen (time-consistent) equilibrium extraction policy.

⁷ More formally: fix an \underline{X} , let ϵ be defined on \mathcal{Z} (a compact subset of \mathbb{R}_+), ψ defined on $\mathcal{P}_{\mathcal{Z}}$, the set of all probability distributions on \mathcal{Z} , with $\epsilon \sim \psi$, denote by \mathcal{X} the state space and by \mathcal{B} any Borel set on \mathcal{X} (i.e., $\mathcal{B} \in \mathcal{B}_{\mathcal{X}}$) and let $L : \mathcal{X} \times \mathcal{X} \times \mathcal{Z} \rightarrow \mathcal{X}$. Then $Q(\mathcal{B}|a, x, \underline{X}) := \mathbb{P}(L(a, x, \epsilon; \underline{X}) \in \mathcal{B}) = \int \mathbb{1}_{\mathcal{B}}[L(a, x, \epsilon; \underline{X})] \psi(d\epsilon)$.

of the production function in this form is often more convenient than the indirect representation through $L(a, x, \epsilon; \underline{X})$ and most of the previously cited literature in stochastic games (which we follow), uses this type of production function.⁸

In our model, the transition structure has the following additive form:⁹

$$Q(\cdot|a, x, \underline{X}) = (1 - P(\mathcal{X}|a, x, \underline{X}))\delta_0(\cdot) + P(\cdot|a, x, \underline{X}), \tag{3}$$

where $P(\cdot|a, x, \underline{X})$, a function of (a, x) parameterised by \underline{X} , is a probability measure and δ_0 is the Dirac measure concentrated at 0 stock. The transition function, Q is thus a mixture of the measures P and δ_0 , meaning that next period stock is either drawn from P or is zero (with measure $(1 - P) > 0$).¹⁰ A regime shift is encapsulated in the dependence of the measure P upon the threshold, \underline{X} . Intuitively, the transition function in Eq. (3) has to satisfy the following two properties: First, the shift of current stock from the desirable to the undesirable region of the state space has a very large effect on next period stock, consistent with the idea of significant worsening of stock regeneration associated with the undesirable region. We label this property S1. Second, a higher threshold (meaning an enlarged undesirable region of the state space) is associated with an increased adverse effect on stock regeneration. We label this property S2.

While we will discuss these properties later on, two points regarding the structure of the transition are worth mentioning: First, from the structure of the transition in Eq. (3), it is evident that the distribution of X_{t+1} has a strictly positive probability of reaching state 0 from any other stock value. Given that utility will be bounded below (by 0), this structure implies that reaching 0 stock leads to the collapse of this resource, and, more importantly, that the probability of this event is strictly positive for all t .¹¹ This feature is very similar to many models in resource economics, including those of catastrophic regime shifts, resource extinction, and to models of non-convex stochastic growth with renewable resources.¹² Second, we note that all the major findings in this study hinge upon assumptions on these model primitives (properties of Q), as we discuss further in Section 4.

Example 1: Consider the following definition of a measure P ,

$$F_p(\cdot|X_t = x, a_t = a, \underline{X}) = M_1(\cdot|x)\theta(a, x, \underline{X}) + M_2(\cdot|x)(1 - \theta(a, x, \underline{X})) \tag{4}$$

where F_p is the cdf associated with P . Here M_1, M_2 represent two stock regeneration distributions with M_2 being associated with a higher growth regime than M_1 , while $\theta \in [0, 1]$ is a function that satisfies two properties: it is increasing in the threshold \underline{X} and in the stock x ; it changes rapidly and substantially around the threshold. In particular, this choice ensure that θ puts more weight on M_1 , the lower growth regime, when $x \leq \underline{X}$, meaning that $1 - \theta(a, x, \underline{X}) < \theta(a, x, \underline{X})$ whenever $x \leq \underline{X}$.

We now verify that the form of F_p in Eq. (4) satisfies properties S1 and S2. For S1, it follows that if M_2 dominates M_1 in the sense of first order stochastic dominance,¹³ then by construction lower next period stock occurs whenever current stock is in the undesirable region. First order stochastic dominance is arguably the most natural extension of the deterministic notion for the most common deterministic transition, and has been used before in (irreversible) regime shift contexts (Diekert, 2017).

We also note that the degree to which the undesirable region affects next period stock depends upon both the distributions M_1, M_2 and the function θ : if θ is very steep, then the shift approximates an “abrupt” one.¹⁴ As to property S2, since θ is increasing in the threshold, F_p is “stochastically decreasing” in the threshold. This implies that \underline{X} shifts F_p in the sense

⁸ It is more general and often more convenient to work with Eq. (1). Otherwise, ad hoc assumptions would be needed for example on ϵ (e.g., finiteness) to ensure that X_{t+1} is drawn from a compact set. Such assumptions are often difficult to justify, they require bounds on ϵ (which may depend upon X), and they may have little to do with the structure of the problem at hand.

⁹ The set-up here is for a regime shift with a short lag (of one period), meaning that current stock and reinvested stock together, determines the shift. The case of regime shifts occurring immediately when reinvested stock falls below the threshold is considered in Section 5.

¹⁰ A more detailed discussion of this form of the transition structure is provided in Balbus et al. (2014, 2018) (and references therein). It turns out that this mixing condition is fundamental for uniqueness of equilibrium, in whose absence only existence can be established.

¹¹ The regime shift we model is “reversible” in the sense that if resource stock at time t falls below the threshold i.e., if $0 < x_t \leq \underline{X}$, then the probability that stock will rise to being above the threshold the very next (or any future) period is strictly positive. The only exception is at 0 stock: the resource stock can never recover (“grow”) once it falls to 0, which can be considered an “irreversible” collapse (see footnote 12).

¹² For instance, Leizarowitz and Tsur (2012) consider an irreversible catastrophe occurring at some random point in time; once it occurs, the stock is driven to some fixed level (possibly zero) forever while Mitra and Roy (2007) consider resource extinction. In both cases, conditional on the irreversible event (catastrophe and extinction, respectively) not occurring, stock evolution is Markovian. Furthermore, the probability of the event occurring is strictly positive $\forall t$ and is a function of the extraction policy, similarly to our case (where $1 - P$, the probability of having 0 stock next period, is a function of a, x). Finally, in the framework of a stochastic growth model with non-convex production, Kamihigashi and Stachurski (2014) allow for arbitrarily bad shocks to stock regeneration (which is Markovian), implying that the probability of having arbitrarily low stock is again strictly positive.

¹³ Let $\mathcal{B}_{\mathcal{X}}$ denote the Borel sigma algebra for the state space, \mathcal{X} , with $\mathcal{M}(\mathcal{X})$ the set of distributions defined on $(\mathcal{X}, \mathcal{B}_{\mathcal{X}})$. For two distributions, M_1, M_2 in $\mathcal{M}(\mathcal{X})$, we note that if M_2 (first order) stochastically dominates M_1 , denoted $M_2 \geq M_1$, then for every increasing, continuous and bounded function h on \mathcal{X} , $\int hM_2(dy) \geq \int hM_1(dy)$.

¹⁴ We note that the most commonly used deterministic equation for the general regime shift case can be written as

$$X_{t+1} = \begin{cases} L_1(a_t, X_t), & X_t \leq \underline{X} \\ L_2(a_t, X_t), & \text{else} \end{cases} \tag{5}$$

where L_1 and L_2 are deterministic stock growth functions that satisfy the condition $L_2 \geq L_1$. In this case, θ is an indicator function for $X_t \leq \underline{X}$.

of (first order) stochastic dominance, meaning that increases in \underline{X} shift probability mass to the left, putting more weight on lower values of resource stock.¹⁵

3.2. Recursive approach to Markov equilibrium

At time t , the forward looking regulator inherits the state variable, X_t , knows the threshold, \underline{X} , and makes the extraction and investment decisions, $q(x), a(x)$.¹⁶ These choices together with X_t determine X_{t+1} via the transition kernel, Q . Regulators at time t play an inter-generational game with all regulators at time $t + 1, t + 2, \dots$. We use a recursive approach to characterise the Markov equilibrium of this inter-generational game. Consequently, our set-up parallels the more common Bellman recursion for exponential discounting, with only key differences highlighted here (see the Appendix for further details). In this spirit, we denote by w a (pure strategy) stationary Markov Nash equilibrium (“equilibrium”, henceforth) for a quasi-hyperbolic agent that satisfies the following functional equation:¹⁷

$$w(x) \in \arg \max_{q \in [0, \underline{x}]} \left[U(q) + \beta \delta \int_{\mathcal{X}} V_w(x') Q(dx' | x - q, x, \underline{X}) \right], \tag{6}$$

where $V_w : \mathcal{X} \rightarrow \mathbb{R}^+$, the continuation value function for the household of future selves, satisfies

$$V_w(x) = U(w(x)) + \delta \int_{\mathcal{X}} V_w(x') Q(dx' | x - w(x), x, \underline{X}). \tag{7}$$

Now, defining the value function for the self at time period t to be

$$W_w(x) := U(w(x)) + \beta \delta \int_{\mathcal{X}} V_w(x') Q(dx' | x - w(x), x, \underline{X}), \tag{8}$$

one obtains the relation

$$V_w(x) = \frac{1}{\beta} W_w(x) - \frac{1 - \beta}{\beta} U(w(x)). \tag{9}$$

Equation (9) is the so-called generalised Bellman equation, which “adjusts” the conventional one (by the term $\frac{1 - \beta}{\beta} U(w(x))$) to account for preferences that are present-biased. The generalised Bellman equation captures interactions between different “generations” of the regulator (the different “selves”) and thus represents a strategic version of the conventional Bellman equation. A fixed point, V^* , of a suitably defined operator upon V_w corresponds to the value function of a time-consistent equilibrium. If V^* is a unique fixed point, then there is a unique, pure strategy, time-consistent equilibrium extraction, w^* , representing the equilibrium outcome of the inter-generational game between selves. Our results pertain, as already mentioned, to the case of “no commitment” of future generations to current policies, which is the most commonly used equilibrium notion.

We emphasise that our approach to analysing equilibrium policy in hyperbolic settings differs significantly from that used in most previous studies, particularly addressing environmental questions (Gerlagh and Liski, 2017; Karp, 2005; Karp and Tsur, 2011). The differences lie in many aspects: the explicit focus on uncertainty, in using a fully recursive approach to equilibrium, and in not relying on specific functional forms for utility and the transition. More specifically, our model setting differs from the commonly used one in Krusell and Smith Jr (2003) and in Harris and Laibson (2001), two canonical yet different approaches to the problem. Our approach is instead closely related to stochastic games and the recursive approaches to time-inconsistent problems applied in Balbus et al. (2014, 2018).

4. Characterising the equilibrium extraction policy

We examine the nature of the equilibrium extraction policy for the resource extraction problem characterised by the generalised Bellman equation in Eq. (9) and the stochastic production function in Eq. (3), where we consider the case of regime shifts depending upon pre-extraction stock, X_t .

4.1. Uniqueness

We seek a stationary, Markov Nash equilibrium extraction policy w^* , which, by definition, is the policy followed by all regulators i.e., $w_t(X_t) = w^*(x) := w^*, \forall t$. Without assuming either any specific parametric function for utility and transition or the existence of a differentiable equilibrium, we establish the existence of a unique equilibrium extraction policy w^* .

¹⁵ More formally, this means that $1 - F_p(\cdot | a, x, \underline{X}) := 1 - \mathbb{P}(X_{t+1} \leq b | a, x, \underline{X})$ is decreasing in \underline{X} (see Online Appendix for a formal definition). For simplicity and consistency with standard definitions, we say “ F_p is stochastically decreasing”, whenever we mean to say that X_{t+1} is.

¹⁶ The extraction and investment decisions $q(x), a(x)$, as well as the value functions, V_w and W_w , all depend upon the threshold \underline{X} . For notational simplicity, we henceforth suppress the explicit dependence of functions upon \underline{X} whenever its effects are not explicitly discussed.

¹⁷ We note that w denotes a stationary policy followed by future selves while q denotes the extraction decision of the current self.

Our results hereon build upon:

Assumption 1. We assume

- a. $U : \mathcal{X} \rightarrow \mathbb{R}^+$ is positive, increasing and strictly concave, with $U(0) = 0$ and $U(\cdot) \leq \bar{U}$.
- b. for $a, x, \underline{X} \in \mathcal{X}$, the probability measure $P(\cdot|a, x, \underline{X})$ (and its cdf F_P) satisfy the following properties,
 - i. for $x \in \mathcal{X} \setminus \{0\}$, $a \in [0, x]$, $P(\mathcal{X}|a, x, \underline{X}) < 1$, $P(\mathcal{X}|0, 0, \underline{X}) = 0$;
 - ii. for every bounded function $V : \mathcal{X} \rightarrow \mathbb{R}^+$, the function $\int_{\mathcal{X}} V(x') dF_P(x'|a, x, \underline{X})$ is continuous in (a, x) , and increasing and concave in a (for every given x).

A brief discussion of these conditions will be relevant before we proceed to discuss how these relate to the specifics of our problem. Assumption 1(a) is standard, specifying that utility is bounded above. Assumption 1(b) deals with the structure of Q : (ii) states the requirement that the expected continuation value function, $\int V(x') dF_P(x'|a, x, \underline{X})$, is increasing and concave in a . It can be shown that this is true whenever increases in reinvestment lead to a larger measure $P(\cdot|a, x, \underline{X})$ and hence larger next period stock but at a decreasing rate (see Remark 4 in Online Appendix). This is only meaningful if $P(\cdot|a, x, \underline{X}) < 1$, $\forall a, x > 0$ (if $P(\cdot|a, x, \underline{X}) = 1$ for some a, x , then (ii) trivial), which is the content of (i). We discuss next what this property means and why it follows from intuitive properties of resource growth functions.

As mentioned earlier, we use only key properties of the transition kernel, which can be directly linked to commonly used properties of production functions in models of stochastic growth. Following Amir (1997), we assume that F_P , the stochastic production function, is stochastically increasing and concave in a , a common property in stochastic growth theory and in models of resource extraction and regime shifts.¹⁸ F_P stochastically increasing in a means (see footnote 15) that increases in reinvestment, a , increase the probability of obtaining higher stock levels next period, i.e., $P(\cdot|a, x, \underline{X})$ is increasing in a . Similarly, F_P stochastically concave in a means that increases in reinvestment, a , yield diminishing marginal returns for X_{t+1} , hence $P(\cdot|a, x, \underline{X})$ is concave in a and Assumption 1(b)(ii) is satisfied. From now on, we will refer to F_P , the associated cdf of $P(\cdot|a, x, \underline{X})$, whenever we need to specify properties of the distribution of X_{t+1} .

We now turn to our main result for this section (which holds under Assumption 1):

Theorem 1. For the renewable resource extraction problem considered here, there is a unique, bounded value function, V^* and correspondingly, a unique time-consistent equilibrium extraction policy w^* .

A proof can be constructed following an application of the Banach Fixed Point Theorem (see e.g., Balbus et al. (2018), Theorem 1). As already mentioned, uniqueness of equilibrium is unusual in the non-constant discounting literature and in the regime shifts literature. Here it follows from the concavity of the expected continuation value function, which, in turn, arises from the structure of the stochastic transition (see footnote 10). Conditions similar to the strong mixing condition we use here have been used in a variety of previous literature (Amir, 1996; Chassang, 2010), including stochastic growth model contexts (Kamihigashi and Stachurski (2014)). It is worth noting that, compared to deterministic studies of either regime shifts or non-constant discounting, uncertainty in stock growth can be seen to “convexify” our problem. We note that the convexifying effect of uncertainty has also been observed before in related literature on hyperbolic discounting (Harris and Laibson, 2001; Chatterjee and Eyiungor, 2016) and regime shifts (Brozović and Schlenker, 2011; Barrett, 2013).

4.2. The relationship between extraction, stock level and the threshold

We next turn to understanding two aspects of the equilibrium extraction policy: do increases in beginning-of-period stock and the threshold level lead to unconditional increased (or decreased) reinvestment? These properties hinge on those of the stochastic production function, $F_P(\cdot|a, x, \underline{X})$ and in particular on the substitutability between reinvestment, a , and either the state x , or the threshold \underline{X} . The notion of substitutability relevant here relates to the effect of these inputs on the marginal returns (in terms of next period stock), which can be formalised by Edgeworth substitutability: two inputs to production are Edgeworth substitutes if having more of one reduces the marginal returns to having more of the other. Edgeworth substitutability is formalised by the notion of increasing (decreasing) differences: A deterministic production function H exhibits increasing (decreasing) differences in its inputs a, x (or a, \underline{X}) if marginal returns to a are increasing (decreasing) in x (and vice-versa). Our analysis uses stochastic equivalents of this definition (defined formally in the Online Appendix): the production function, F_P , satisfies stochastically decreasing differences in (a, x) and in (a, \underline{X}) . The interpretation is intuitive: if $F_P(\cdot|a, x, \underline{X})$ satisfies stochastic decreasing differences in (a, x) , then the marginal return to increased stock (i.e., the probability of reaching a higher stock in the next period) is decreasing in a , and vice-versa. Similarly, if $F_P(\cdot|a, x, \underline{X})$ satisfies stochastic decreasing differences in (a, \underline{X}) then the marginal return to increased a is decreasing in the threshold, \underline{X} . In the

¹⁸ This corresponds to the standard assumption in deterministic growth theory, that the production function $f(a; \underline{X})$ is increasing and concave in a , and is defined in the Online Appendix. For stochastic models, in all the state transition equations in the cited literature increases in reinvestment generate higher next-period stock. Furthermore, concavity also follows, either trivially – as a result of linearity – or explicitly, e.g., Ren and Polasky (2014).

previous literature, special functional form assumptions ensure that these conditions are always fulfilled¹⁹ : X_{t+1} is usually separable both in (a, x) and in (a, \underline{X}) .²⁰

4.2.1. Is equilibrium extraction increasing in resource stock?

With exponential discounting, larger resource stocks often lead to greater extraction, a conclusion that does not always hold with non-convexities induced by stock-dependence in the growth function even in the absence of regime shifts (see e.g., Krishnamurthy (2017)). With regime shifts, competing aspects are at play: higher stocks can enhance incentives to extract but this may depend upon the stock and threshold levels (since the threshold can alter incentives to extract). Consequently, even without present bias, monotonic optimal extraction policy with complex regeneration functions is rare and has not been established for regime shifts.

Similar to the case of uniqueness of equilibrium extraction in Theorem 1, the key consideration determining the relationship between extraction and stock is whether the expected continuation value function satisfies an important property, here stochastic decreasing differences in (a, x) . Our arguments for this case parallel those for Theorem 1: if F_p satisfies stochastic decreasing differences in (a, x) (as it does for our case), then so does $\mathbb{E}[V](a, x, \underline{X})$ (see Online Appendix for a proof). Assumption 2 formalises the discussion:

Assumption 2. The function $\int V(x')dF_p(x'|a, x, \underline{X})$ satisfies decreasing differences in (a, x) .

Assumption 2 enables us to establish that resource extraction is in fact increasing in stock, which is the content of

Theorem 2. Under Assumptions 1 and 2, equilibrium extraction, w^* , is increasing in the resource stock, x .

This result, while intuitive in a general resource extraction or growth model setting, may seem surprising with a regime shift at a first glance, since higher extraction (for any given threshold) increases the probability of a regime shift. Recall however from our discussion above that marginal increases in reinvestment are taken to be less beneficial when stocks are higher, irrespective of the distance from the threshold. This aspect (which is common in the literature), together with the decreasing returns to reinvestment, can explain why extraction is increasing in stock. In the terminology of Karp (2005); Iverson and Karp (2021), action (extraction) and stock in our case are “strategic complements”.

To our knowledge, we are the first to characterise this relationship for cases with regime shifts, with or without hyperbolic discounting, and that without using special functional forms. The only comparable results involve either special functional forms (e.g., Brozović and Schlenker (2011)) or hypothesised findings (Karp (2005)).²¹

Remark 1. Note that Theorem 2 does not imply that next period stock, X_{t+1} is increasing in current stock X_t . Karp (2005) provides conditions under which, for his deterministic framework, the trajectory of the stock is a monotonic function of time but cannot exclude the less natural case of an oscillating stock trajectory. In our case, the equivalent (to “ X_{t+1} is increasing in X_t ”) notion is that the distribution function F_p is increasing in x . A sufficient condition for this typically involves a Lipschitz condition on extraction along with the transition being stock independent, which is not fulfilled in our case. It has been noted before that in resource extraction problems with exponential discounting, stock dependence can lead to Lipschitz continuity of extraction not holding (Krishnamurthy (2017)). Our findings suggest that stock dependence also leads to essentially the same complexity in reinvestment decisions in hyperbolic discounting settings.

Remark 2. We note that $\int V(x')Q(dx'|x - w, x, \underline{X})$ is continuous in x , for each w (from Assumption 1). Furthermore, it is also known that the SMNE, $w^*(x) := \arg \max_{q \in A(x)} [u(q) + \beta \delta \int V^*(x')Q(dx'|x - q, x, \underline{X})]$ is unique. It therefore follows, from the maximum theorem, that $w^*(x)$ is continuous in x if $A(x)$ is non-empty and compact (Stachurski (2009, Theorem B.1.3)).

¹⁹ Studies considering the effects of a threshold (e.g., Brozović and Schlenker (2011), Peterson et al. (2003)) have often opted for a linear and separable relationship between a, \underline{X} (by adopting a piece-wise specification for the transition equation), see Eq. (5) for a generic example. In contrast Ren and Polasky (2014) use a parameterised deterministic transition equation, which assumes that reductions in “environmental quality” lead to reduced stock growth. This can be interpreted in terms of Edgeworth substitution, meaning that having “more of” \underline{X} reduces the marginal returns to a . In other words, they assume that a regime shift would not only reduce resource production but also the gross rate of return to reinvestment. In any case, as in footnote 20, linearity and separability ensure that the condition of decreasing differences is (trivially) satisfied.

²⁰ For the discrete-time setting, the literature considered three types of relationships between (a, x) : without an explicit state-based regime shift (e.g., Leizarowitz and Tsur (2012)), with parameterised regime shifts (Ren and Polasky, 2014), and with linear and separable functions in (a, x) (Brozović and Schlenker, 2011; Karp, 2005). In all cases, the transition equation can be decomposed to a linear or separable function in a, x . The key benefit of separability is in ensuring that the marginal effect of a change in current stock, x , upon future stock, X_{t+1} , does not depend upon reinvestment, a . This means that the condition of decreasing differences is satisfied (the boundary case of 0 cross-derivative).

²¹ The study of Krusell and Smith Jr (2003) suggests that equilibrium in a canonical deterministic quasi-hyperbolic setting may be non-monotonic. However, this result is not applicable to many settings using alternative approaches to equilibrium in a quasi-hyperbolic framework, including those of Harris and Laibson (2001) (as stated in footnote 3 of Krusell and Smith Jr (2003)), Chatterjee and Eyigungor (2016) (see footnote 2 of that study), our study and for Karp (2005).

4.2.2. Equilibrium extraction and the threshold

Perhaps the most interesting question related to regime shifts is how the threshold affects the extraction decision. One may anticipate, for instance, that increases in the threshold, leading to enlarged adverse region of the state space, will lead to reduced extraction. We investigate whether this seemingly intuitive result holds. Before doing so, however, we first articulate two key properties of a production function encapsulating a regime shift. The first is that increases in the threshold reduce the probability of reaching a larger next-period stock, which implies that the production function, F_P , is stochastically decreasing in \underline{X} . The second pertains to the substitutability between a and \underline{X} on the margin: following previous discussion, the marginal reinvestment rate is considered to exhibit decreasing returns, meaning that marginal return to reinvestment is non-increasing in the threshold (irrespective of the current stock level). We emphasise that these two properties do not, by themselves, preclude either monotonic or non-monotonic extraction.

As before, the crucial question is whether these two properties (decreasing differences in (a, \underline{X}) , and decreasing in \underline{X}) hold for the expected continuation value function. Paralleling the argument for [Theorems 1,2](#), $\mathbb{E}[V(a, x, \underline{X})]$ can be shown to satisfy decreasing differences in (a, \underline{X}) , and to be decreasing in \underline{X} if F_P satisfies stochastic decreasing differences in (a, \underline{X}) and is stochastically decreasing in \underline{X} . The interpretation of these properties for $\mathbb{E}[V(a, x, \underline{X})]$ is also intuitive: $\mathbb{E}[V(a, x, \underline{X})]$ decreasing in \underline{X} means that increases in the threshold lead to a reduction of the expected value function by increasing the size of the undesirable region of state space. $\mathbb{E}[V(a, x, \underline{X})]$ satisfying decreasing differences in (a, \underline{X}) implies that for the expected value function, increases in reinvestment are more valuable when the threshold is smaller than when it is larger.

Assumption 3 formalises these two conditions on F_P ²² :

Assumption 3. We assume:

- a. u does not depend on \underline{X} and satisfies [Assumption 1](#).
- b. for $a, x \in \mathcal{X}$ and $\underline{X} \in \underline{\mathcal{X}} \subset \mathbb{R}_+$, Q has the structure in [Eq. \(3\)](#) with $P(\cdot|a, x, \underline{X})$ satisfying [Assumption 1](#).
- c. for every bounded function $V : \mathcal{X} \rightarrow \mathbb{R}^+$, the function $\int V(x')dF_P(x'|a, x, \underline{X})$ is decreasing on $\underline{\mathcal{X}}$, and submodular (decreasing differences) in (a, \underline{X})

Then, the following result holds:

Theorem 3. Under Assumptions 1 and 3, equilibrium extraction, $w_{\underline{X}}^*$, is increasing in the threshold, \underline{X} and the continuation value $V_{\underline{X}}^*(x)$ is decreasing in \underline{X} .

While the proof is provided in the Online Appendix, we provide an intuitive discussion of this finding here. Our result that increases in the threshold (i.e., an enlarged adverse region of the state space) result in increased extraction is due to the interplay of three forces: first, an increase in the threshold reduces future stock growth, incentivising reducing (or at least not increasing) extraction in order to ensure adequate stock in subsequent periods; second, increased extraction yields a utility benefit today; and third, given that reinvestment a is less “productive the higher is the threshold, extracting today is more beneficial than reinvesting for extraction in the future. On balance, increasing extraction is more beneficial.

This unambiguous finding of a monotonic relationship between the threshold and extraction is unique in the literature. For instance, [Brozović and Schlenker \(2011\)](#) find a region of non-monotonic behaviour—with an initial decrease as the threshold increases and a subsequent increase. Continuous time frameworks with a time-based regime shift (such as in [Polasky et al. \(2011\)](#)) do not feature this behaviour at all. Our results suggest in fact that the presence of a threshold does not lead to “precautionary behaviour ” i.e., a reduction in extraction with a hope of avoiding the regime shift in question. This is also the finding of the simulation-based study by [Peterson et al. \(2003\)](#), which investigates optimal emission loading with exponential discounting.

Remark 3. Note that for the case of exponential discounting (obtained by setting $\beta = 1$) for the stochastic regime shift set-up used in [Sections 4 and 5](#), many of our results, including [Theorem 3](#), hold. This is because the crucial aspect for [Theorem 3](#) (and in fact for [Theorems 2 and 4](#)) is the structure imposed on the stochastic production function F_P .

5. Regime shifts without lags

The literature often considers ecosystems with rapid dynamics ([Brozović and Schlenker, 2011](#); [Ren and Polasky, 2014](#); [Diekert, 2017](#)), in which case current rather than past action determines the state of the system and threshold crossing. Therefore, regime shifts are conditioned on reinvestment (or post-extraction stock), a_t , and the transition function from [Eq. \(3\)](#) is analogously modified as

$$Q(\cdot|a, \underline{X}) = (1 - P(\cdot|a, \underline{X}))\delta_0(\cdot) + P(\cdot|a, \underline{X}). \tag{10}$$

²² Following the general structure of F_P from [Example 1](#) (for which assumption 3a and b are satisfied), consider $F_P(\cdot|a, \underline{X}) = M_1(\cdot|x)\theta(a, \underline{X}) + M_2(\cdot|x)(1 - \theta(a, \underline{X}))$ with $\theta(a, \underline{X}) = e^{-a}e^{-\underline{X}} + (1 - e^{-\underline{X}})$, and M_1 first order stochastically dominating M_2 . It is not difficult to verify that this specific $F_P(\cdot|a, \underline{X})$ exhibits increasing differences in (a, \underline{X}) and is increasing in \underline{X} , satisfying [Assumption 3\(c\)](#).

This formulation, which allow us to establish the existence of an Euler equation, is helpful not only i connecting our analysis to the recent literature on hyperbolic discounting and on regime shifts but also in deriving intuition regarding certain aspects of the problem.

5.1. Main results

Theorem 1, guaranteeing a unique equilibrium extraction policy and **Theorem 3** regarding the effect of a regime shift continue to hold. **Theorem 4**, (whose proof is presented in the Online Appendix (§1.2)), is an extended version of **Theorem 2**.

Theorem 4. For the renewable resource problem considered here (with transition as in Eq. (10)), under Assumption 1, equilibrium extraction w^* is increasing in stock, x , and Lipschitz with modulus 1.

Theorem 4 is the basis for all further results in this section and states that optimal extraction and reinvestment policy are: increasing in the current stock; differentiable; and the absolute value of their derivatives is bounded above by 1. An important consequence of **Theorem 4** is that F_p is stochastically increasing in current stock, x , from which it follows that future stock is increasing in current stock (in the sense mentioned in **Remark 1**). Our study is the first to report this finding. As discussed in **Remark 1**, the only directly comparable study examining this aspect but without regime shifts, **Karp (2005)**, while providing sufficient conditions for this property, is unable to definitively establish it.

Remark 4. The case of stock dependent utility in stochastic growth models with renewable resources has been explored in some of the older literature (e.g., **Mitra and Roy (2007)**), which finds that complementarity between state and policy variables in the reward function is sufficient for optimal policies to be non-decreasing in the resource stock. Similarly, **Karp (2005)** shows that the future stock of emissions is increasing in the current stock when the stock and flow of emissions are complements in the utility function. In our case, stock dependent utility does not alter our main results. In particular **Theorems 1** and **3** hold without any additional conditions. As for **Theorem 2**, equilibrium extraction continues to be increasing in the resource stock, as long as $U(q, x)$ has increasing differences in (q, x) i.e., extraction and stock are strategic complements, as has been suggested in the literature. **Theorem 4**, however, survives only partly: equilibrium extraction w^* is increasing in stock, but equilibrium reinvestment a^* no longer is.

5.2. The euler equation

In the prior literature focused on examining environmental questions with present bias, the Euler equation was the only way to establish existence of an equilibrium policy. The existence of an Euler equation in these studies follows either from the assumption that the policy function is continuously differentiable (**Karp, 2005**) or from special functional forms and settings leading to the existence of smooth Markovian equilibria (**Gerlagh and Liski, 2017**). Unlike these studies, we are able to establish that the Markovian equilibrium is differentiable without resorting to such special assumptions.

Theorem 4 established Lipschitz continuity of the equilibrium extraction policy. Using this, we show that the equilibrium policy and value function are differentiable. Subsequently, we present a version of the Generalised Euler Equation (‘Euler equation’, henceforth) for our problem, which is used to explore the nature of the relationship between present bias and resource extraction. A set of technical conditions, Assumption 4 (stated in the Online Appendix (§1.2)), is required for our result. Beyond standard conditions that allow for an application of the implicit function theorem, the only noteworthy aspect is an Inada-like condition ensuring that e.g., the last fish left unharvested has a very high value (a standard assumption in the literature). We can then establish the following result, whose proof is discussed in the Online Appendix (§1.2):

Theorem 5. For the renewable resource problem considered here (with transition as in Eq. (10)), under Assumptions 1 and 4, equilibrium extraction policy w^* and the value function V^* are differentiable on $(0, \infty)$.

Relying on **Theorem 5**, we can now write our version of the Euler equation characterising the equilibrium reinvestment as (see the Online Appendix for derivation)

$$u'(x - a^*(x)) = - \int \left[u'(x' - a^*(y)) \left(\beta \delta (1 - a^{*'}(y)) + \delta a^{*'}(y) \right) F'_p \right] dy \tag{11}$$

where x, y denote respectively current and next period stock, F'_p denotes $\frac{\partial F_p(x_{t+1} | a_t, X)}{\partial a_t}$ and $a^{*'}$, which lies in $[0,1]$ (by **Theorem 4**) can be termed the “Marginal Propensity to Reinvest” (MPR). The difference with its counterpart when discounting is exponential is that the constant exponential discount factor δ of the latter is replaced with the “effective discount factor”, $EDF := \beta \delta (1 - a^{*'}(y)) + \delta a^{*'}(y)$. This effective discount factor is endogenous and stochastic (**Harris and Laibson, 2001**) as it depends on the MPR. Since $\beta < 1$, the effective discount factor is positively related to the future MPR, $a^{*'}(y)$. Put more plainly, the EDF increasing in $a^{*'}(y)$ means that the valuation of the future increases as $a^{*'}(y)$ increases (recall that increases in the discount factor imply the future is discounted less).

The intuition behind this result is as follows: Since the regulator at time t values marginal reinvestment at $t + 1$ more than the $t + 1$ regulator does, the regulator at time t acts strategically in an intergenerational game: the lower the expected

marginal propensity to reinvest at $t + 1$, the less the regulator values the future and the more he will extract in period t . In summary, the fact that a present-biased regulator discounts the short-term more than the long term causes him to value future reinvestment higher than the future regulator will, leading to a strategic reduction in current reinvestment (relative to when he would have discounted exponentially).²³

5.3. The effect of present bias

An important question that arises with quasi-hyperbolic discounting relates to how time-consistent optimal policies compare to those derived using standard exponential discounting. More specifically, it is interesting to know whether present bias leads to an increase in optimal extraction (as is often surmised) and a reduction in optimal reinvestment. The key aspect that allows us to provide precise answers to this question is that the short term discount factor, β , can be treated simply as a parameter of our key operators (defined in the Appendix). Consequently, the approach we previously used for comparative statics (in Section 4.2.2) is applicable.²⁴

While details of this approach are provided in the Appendix, we present here our main result (which is also applicable for the case of regime shift with lags) and provide some intuition for it:

Theorem 6. *Under Assumption 1, equilibrium extraction w^* is decreasing in β .*

Proof. See Section A.3. \square

In other words, a decrease in β from $\beta = 1$ to $\beta < 1$ implies an increase in equilibrium extraction. Recalling that a reduction in β implies increased present bias, it follows that an increase in (the degree of) present bias leads to an increase in extraction. Thus, a regulator who is present biased always extracts more in an equilibrium than if he were instead an exponential discounter. Since $a^* = x - w^*$, equilibrium reinvestment is therefore decreasing in the degree of present bias.

While this result is independent of the existence of the Euler equation (and of Theorems 4, 5), it is nonetheless instructive to connect this result to our discussion of the Euler Equation and the strategic reasoning involved in arriving at this outcome. Our previous discussion of the Euler equation suggests that a reduction in β has opposing effects: on the one hand, it implies higher discounting of the future by the time- t regulator, suggesting that increased extraction at time t is optimal. On the other hand, since the time- $(t + 1)$ regulator discounts the short-run future with $\beta\delta$, he values reinvestment even less than before. In view of the divergence between the time- t and time- $(t + 1)$ regulators' valuation of the future (in particular, of time- $(t + 2)$), the time- t regulator will wish to manipulate the time- $(t + 1)$ regulator to reduce his extraction at time- $(t + 1)$. He can do so by reducing his own time- t extraction thereby increasing time- $t + 1$ stock, in response to which (with a positive future MPR), the regulator at $(t + 1)$ will increase reinvestment. The fact that $0 < a^*(y) < 1$ (from Theorem 4), however means that the former effect is larger, leading to an increase in extraction at time t . The argument holds for any t and, in view of Theorem 6, must constitute an equilibrium.

We are unaware of any previous study that explicitly focused on this question, and established a formal result of the form we do in Theorem 6.²⁵ In the environmental literature, the only study considering this aspect is Karp (2005); this study hypothesizes, but is unable to establish conclusively, that emissions would be higher with a present biased regulator.

Remark 5. Note that Theorems 4, 6 and the Euler equation (see discussion in footnote 23), all continue to hold for the case of $\underline{X} = 0$. In other words, whenever a stochastic renewable resource extraction problem with present-biased agents has a transition function that only depends upon reinvestment (not upon stock) and can be expressed as in Eq. (10), these results hold (along with Theorem 1). Since a transition function independent of stock is the most common one in the literature, our results suggest that for most renewable resource problems with simple transition functions, and a strictly positive probability of collapse (e.g., Leizarowitz and Tsur (2012)), equilibrium extraction is increasing in the resource stock and the degree of present bias, and an Euler equation formulation can be written down. We are unaware of results of this degree of generality regarding renewable resource problems with random growth for a present biased decision maker.

²³ Note that if the regulator had a commitment device at her disposal then the extraction plan would involve two different extraction policies, one for the current regulator and one for all future regulators. These plans may be derived from the GEE, with the plan for the current regulator being $u'(x - a^N(x)) = -\beta\delta \int u'(y - a^N(y))F_p^y dy$, which corresponds to the extraction policy chosen by a "naive" agent who wrongly believes that future selves will follow through with her plan. Instead, the extraction policy for all future regulators would be given by $u'(x - a^E(x)) = -\delta \int u'(y - a^E(y))F_p^y dy$, which is the optimal extraction policy for the problem with standard exponential discounting.

²⁴ We are very grateful to Prof. Łukasz Woźny for suggesting this approach.

²⁵ We are unaware of any formal analysis of the effect of the quasi-hyperbolic discounting factor upon equilibrium policy in the standard literature (Krusell et al., 2002; Krusell and Smith Jr, 2003; Harris and Laibson, 2001; Chatterjee and Eyigungor, 2016; Balbus et al., 2014; 2018).

5.4. Stationary distribution

We next take a brief look at the structure of the Stationary Markov Equilibrium²⁶ resulting from following the SMNE, $w^*(x)$. For this discussion, we consider a more specific form for the transition function detailed in Eq. (3):

$$Q(\cdot|w(x), x, \underline{X}) = g_0(w(x), x)\delta_0(\cdot) + \sum_i^J g_i(w(x), x)\lambda_i(\cdot|x, \underline{X}), \tag{12}$$

for any SMNE, $w(x)$. Here, $g_0^w(x) := g_0(w(x), x)$, $g_i^w(x) := g_i(w(x), x)$ are positive functions, indexed by an SMNE, $w(x)$, that are defined on the unit interval (and sum to unity), and depend only upon the state. The $\{\lambda_i(\cdot|x, \underline{X})\}$ are continuous measures (with no atoms) that may depend upon the state but not on the action (i.e., not on the extraction policy) and are parameterised by the threshold. The interpretation of Eq. (12) is straightforward: under SMNE $w(x)$, with probability g_i^w , the next-period state is a draw from the measure $\lambda_i(\cdot|x, \underline{X})$ and with probability $g_0^w \left(:= 1 - \sum_i^N g_i^w \right)$ it is instead 0.

Denote the set of all probability measures on \mathcal{X} by $\mathcal{P}_{\mathcal{X}}$ and by $\mathcal{B}_{\mathcal{X}}$ the Borel sets on \mathcal{X} , and suppose that the distribution of stock at time t , when following the unique SMNE $w^*(x)$, represented by the Markov chain X_t , is ϕ i.e., $X_t \sim \phi \in \mathcal{P}_{\mathcal{X}}$. Then, the distribution of X_{t+1} can be obtained from Eq. (12) using the operator $T_{w^*} : \mathcal{P}_{\mathcal{X}} \rightarrow \mathcal{P}_{\mathcal{X}}$ defined as

$$T_{w^*}(\phi)(B) = \int_{\mathcal{S}} Q(B|w^*(x), x, \underline{X})\phi(dx), \quad B \in \mathcal{B}_{\mathcal{X}}. \tag{13}$$

In essence, the operator T_{w^*} shifts the distribution of the state forward by one period, meaning that for $X_t \sim \phi$, $X_{t+1} \sim T_{w^*}(\phi)$ (with some abuse of notation). A non-trivial invariant distribution $\Psi \in \mathcal{P}_{\mathcal{X}}$ (if it exists) is the fixed point of this operator i.e., it satisfies $T_{w^*}(\Psi) \equiv \Psi$. In other words, it is the case that for $X_t \sim \Psi$, $X_{t+1} \sim T_{w^*}(\Psi) \equiv \Psi$. For stating our main results in this section, Theorems 7 and 8, we need a set of technical conditions, Assumption 5, which are stated and discussed in the Online Appendix (§1.3). Proofs of these two results and further discussion are also provided there. The only noteworthy condition in Assumption 5 pertains to the existence of at least one non-trivial value of the state at which the probability of attaining 0 stock vanishes, which is clearly the case in almost all comparable regime shift models.

We now state our main result for this section, which is

Theorem 7. For a continuous and unique SMNE $w^*(x)$, under Assumption 5, there exists a non-trivial steady state distribution for the resource stock s , of the following form:

$$\Psi(\cdot) = \eta\Psi_1(\cdot) + (1 - \eta)\delta_0, \tag{14}$$

with $\eta \in (0, 1)$ and Ψ_1 a non-atomic distribution.

In other words, despite the assumption that for large enough values of x , the probability of stock collapse is exactly 0, we find that the invariant distribution, like the transition function, is a combination of a continuous distribution with an atom at 0 (with strictly positive probability).

Remark 6. We note that continuity of $w^*(x)$ for the problem in Section 4 was discussed in Remark 2 while for the set-up in Section 5, it is implied by the stronger result in Theorem 4. Consequently, Theorem 7 holds for both cases.

Given that increases in the threshold lead to the stock distribution shifting left, as discussed in §4.2.2, we anticipate that such a property will also hold for the invariant distribution. In other words, we expect that if we evaluate the stock evolution for two regime shift settings differing only in the threshold, then a (first order) stochastic dominance result will hold: the invariant distribution corresponding to the lower threshold should stochastically dominate the invariant distribution corresponding to the higher threshold. Consequently, the mean “long-run” stock level is higher for lower threshold values.

We first formalise this intuitive result before we discuss why the needed conditions are satisfied in our case.

Theorem 8. Consider a regime shift model with a transition kernel that is: continuous and increasing in stock, x , and stochastically ordered in the threshold stock \underline{X} . For this setting, under the SMNE $w^*(x)$, every invariant distribution is stochastically ordered in the threshold, \underline{X} . Clearly, it is the case that the mean long-term stock level is larger at lower thresholds.

Remark 7. Given that continuity of $w^*(x)$ (and therefore of the stochastic kernel Q in x) holds for both regime shift settings in Sections 4 and 5, that a stationary distribution exists for the regime shift in both these settings (see Remark 6), and that Q is stochastically ordered for both settings, the key question relates to T_{w^*} increasing in x . An important consequence of Theorem 4 is that for the regime shift setting in Section 5, the transition kernel Q is increasing in stock, x . Consequently, Theorem 8 is directly applicable for this case. However, as stated in Remark 1, it is not a given that the transition kernel Q is increasing in stock for the problem in Section 4. Consequently, the applicability of Theorem 8 to the set-up in Section 4 is not direct.

²⁶ We are thankful to an anonymous reviewer for suggesting this line of investigation.

6. Discussion

We model reversible regime shifts as changes in the distribution of the future stock level of a renewable resource in a discrete-time stochastic framework. We study two types of regime shifts, those for which the shift to a new regime lags behind the time of threshold crossing and those without such a lag, and consider the perspective of forward-looking present-biased regulators. Our results suggest that for many renewable resource problems, equilibrium extraction is increasing in resource stock and that there is no “precautionary” reduction in resource extraction. For regime shifts occurring immediately after threshold crossing, we find that reinvestment is increasing in the stock level. For both types of regime shifts, we find that extraction increases in the degree of present bias. These findings are in contrast to those of much of the large literature on regime shifts, which use exponential discounting and find some evidence for precautionary behaviour.

Regarding extraction and the critical threshold, we show that equilibrium extraction is in fact increasing in the threshold: the larger the undesirable region becomes, the more the regulator would like to extract, implying a higher risk of actually crossing the threshold. In contrast, [Brozović and Schlenker \(2011\)](#) found a non-monotonic relationship between optimal pollutant loading and the threshold-stock level. As to the effect of discounting upon extraction for present biased preferences, we are unaware of any exploration for resource extraction problems, with or without regime shifts. Our findings here suggest that increased present bias leads to an increase in resource extraction. Some of the related literature using non-exponential discounting but without regime shifts either cannot establish uniqueness, monotonicity or related properties of the equilibrium ([Karp, 2005](#)) or use specific functional forms to do so ([Gerlagh and Liski, 2017](#)). The only directly comparable study to ours, which considers irreversible regime shifts and hyperbolic discounting, is [Karp and Tsur \(2011\)](#), whose simplified model differs from ours in many ways, in particular allowing for only two actions (business-as-usual and full stabilization), and is also unable to establish any qualitative properties of the equilibrium. In our analysis, which allows an unrestricted action space, we are able to characterise the equilibrium policy and establish many qualitative properties, including monotonicity in stock and threshold.

Our results may seem to provide gloomy perspective on “sustainable management” as an optimal strategy. Indeed, if the optimal strategy of present-biased resource users facing a threshold is to increase resource exploitation, then the risk of crossing the threshold increases, jeopardising sustainable resource use. Two important sets of dynamics not explicitly considered here however may yet influence long-run resource use decisions. First, heterogeneous discounting of multiple actors involved in resource management suggests that bargaining over discounting could influence long-run outcomes (see e.g., [Denant-Boemont et al. \(2017\)](#)). Second a greater understanding of preference formation, which has been suggested to include aspects such as uncertainty, goal proximity, opportunity costs, resource scarcity, and emotions ([Dasgupta and Maskin, 2005](#)) can provide channels for policy to influence preferences in ways complementary to more challenging long-term commitment mechanisms (whose value may be far less than initially assumed, see [Iverson and Karp \(2021\)](#)).

Finally, most of the results discussed so far build on modelling evidence, with little grounding in observed behaviour. Two types of empirical observations provide suggestive evidence regarding the degree to which our findings may reflect reality: general observations showing global trends in resource use and more specific experimental evidence investigating the collective behaviour of resource users. The general empirical observation pertains to the “great acceleration” in the use of natural resources, particularly since the second world war ([Costello et al., 2016](#)) with associated regime shifts either already having occurred or soon predicted to occur. This observation is suggestive of increasing resource use intensity, rather than precaution, when faced with regime shifts, lending a measure of support to our findings. In contrast, emerging experimental evidence suggests that users who manage a common pool resource can adapt their behaviour if they become aware of a potential regime shift. Lab-based experiments reveal that resource users can exhibit precautionary behaviour when the resource can undergo a regime shift ([Lindahl et al., 2016](#)) and that uncertainty influences decisions in that context ([Schill et al., 2015](#)). Framed-field experiments with local resource users in many contexts lend support to these experimental findings ([Lindahl and Jarungrattanapong, 2018](#)). Given that decision makers’ discounting characteristics are embedded in their decisions in experiments, it is difficult to state whether non-exponential discounting actually influenced their decisions. Our results suggest more empirical and/or experimental investigation of discounting are required to better understand their effects on resource use.

7. Conclusions

We investigated the decision process that forward-looking present-biased resource users face when managing a renewable resource stock subject to the threat of a reversible regime shift. Our results indicate that potential regime shifts matter for management in ways that are substantial and complex. While our analysis focused on the case of the planner without a credible commitment mechanism, consideration of regulation and commitment is clearly central in many resource extraction and pollution control settings. In this context, an extension of our framework to address the following two problems appears most interesting: the first problem pertains to regulation and commitment in resource use contexts with regime shifts involving “inertia”, as in e.g., [Crépin and Nævdal \(2019\)](#), leading to a potential mechanism for some form of commitment (as in [Gerlagh and Liski \(2017\)](#)). The second one is an exploration of equilibrium carbon emissions with a fossil (i.e., scarce) and a man-made capital stock when faced with a regime shift and present-biased decision makers, necessitating an extension of our framework to multiple state and control variables.

To the extent that features such as present bias are important in long-run resource management problems, an understanding of the underlying incentives and motivations facing decision makers is likely an important step in addressing some of the most pressing environmental problems today.

Declaration of Competing Interest

All authors approve the resubmission and declare no conflicts of interest.

Data availability

No data was used for the research described in the article.

Appendix A

A1. Dynamic program and operators

The objective of the decision maker at time t can be represented by

$$U(q_t) + \beta \mathbb{E}_t \left(\sum_{i=t+1}^{\infty} \delta^{i-t} U(q_i) \right), \tag{15}$$

where $0 < \beta \leq 1$ and $0 < \delta < 1$, u is an instantaneous utility function and \mathbb{E}_t is the expectation taken w.r.t. the time- t distribution of X_t . The dynamic program associated with the optimisation problem with quasi-hyperbolic discounting is defined analogous to the exponential discounting case. We denote by $w \in \mathcal{A} := \{w : \mathcal{X} \rightarrow A; w \text{ bounded, } w(x) \in A(x)\}$ a (pure strategy) stationary Markov Nash Equilibrium for a quasi-hyperbolic agent that satisfies the functional equations detailed in §3.2.

Next, consider a vector space, \mathcal{V} , of bounded (in the sup-norm), positive functions defined on the state space, \mathcal{X} . For $V \in \mathcal{V}$, $\beta \in (0, 1]$ and $x \in \mathcal{X}$, we define an operator T as

$$TV(x) = \frac{1}{\beta} AV(x) - \frac{1-\beta}{\beta} u(BV(x)), \tag{16a}$$

$$\text{with the operators } A \text{ and } B \text{ defined as } AV(x) = \max_{q \in A(x)} \left[u(q) + \beta \delta \int_{\mathcal{X}} V(x') Q(dx' | x - q, x, \underline{X}) \right], \tag{16b}$$

$$BV(x) = \operatorname{argmax}_{q \in A(x)} \left[u(q) + \beta \delta \int_{\mathcal{X}} V(x') Q(dx' | x - q, x, \underline{X}) \right], \tag{16c}$$

with

$$Q(\cdot | a, x, \underline{X}) = (1 - P(\mathcal{X}' | a, x, \underline{X})) \delta_0 + P(\cdot | a, x, \underline{X}). \tag{17}$$

Remark 8. It follows from its definition in Eq. (17), and the fact that $U(0) = 0$, that integration w.r.t. Q is identical to integration w.r.t the measure P , a fact that will be repeatedly used.

A2. Regime shifts with lags

Proof of Theorem 2: Let $w^* = BV^*$ and $V^* = TV^*$. Consider the function

$$G(q, x, V^*) = u(q) + \beta \delta \int_{\mathcal{X}} V^*(x') dF_p(x' | x - q, x, \underline{X}). \tag{18}$$

Observe that G is supermodular in q on a lattice $[0, x]$, the feasible action set, which is also increasing in Veinott's strong set order. Observe that submodularity of $\int_{\mathcal{X}} V^*(x') dF_p(x' | a, x, \underline{X})$ in (a, x) implies supermodularity in (q, x) (since $a = x - q$), which combined with $\int_{\mathcal{X}} V^*(x') dF_p(x' | a, x, \underline{X})$ being concave in a suffices for $G(q, x, V^*)$ to be supermodular in (q, x) . Then, from Topkis (2011, Theorem 2.8.1)), $w^*(x, V^*) := \operatorname{argmax}_{q \in A(x)} G(q, x, V^*)$ is increasing in x on \mathcal{X} .

A3. Regime shifts without lags

Proof of Theorem 6: Consider a modified operator \hat{T} , defined, as before, on \mathcal{V} such that

$$\hat{T}V(x) = \hat{A}V(x) - (1 - \beta)u(\hat{B}V(x)), \tag{19a}$$

$$\text{with the operators } \hat{A} \text{ and } \hat{B} \text{ defined as } \hat{A}V(x) = \max_{q \in A(x)} \left[u(q) + \delta \int_{\mathcal{X}} V(x') Q(dx' | x - q, x) \right], \quad (19b)$$

$$\hat{B}V(x) = \arg \max_{q \in A(x)} \left[u(q) + \delta \int_{\mathcal{X}} V(x') Q(dx' | x - q, x) \right]. \quad (19c)$$

We note that $\hat{T}, \hat{A}, \hat{B}$ are variants of our main operators in Eq. (16). It is also evident that $\hat{A}(\beta V(x)) = A(V(x))$ and $\hat{B}(\beta V(x)) = B(V(x))$. Our goal is now to establish exactly the same properties for these operators as we did for those in Eq. (16). This is easily established (proof available upon request) from the following observation: \hat{V}^* is a fixed point of \hat{T} iff $V^* := \frac{\hat{V}^*}{\beta}$ is a fixed point of T . Next, to establish parametric monotonicity in β , we follow the same approach as for establishing Theorem 3, by considering a parameterised fixed point problem that explicitly depends upon β : $\hat{T}_\beta V(x) = \hat{A}V(x) - (1 - \beta)u(\hat{B}V(x))$. Notice that for a given V , \hat{T} is increasing in β (recall that \hat{A} and \hat{B} are independent of β). It can now be established (proof available upon request) that \hat{T} is a monotone contraction. Under Assumption 1, it therefore has a unique fixed point. Next, following similar arguments as in the proof of Theorem 3, \hat{V}^* is increasing in β . Combining with the fact that $\hat{B}V(x)$ is decreasing in V , it follows that $w_\beta^* = \hat{B}\hat{V}_\beta^* = \beta V_\beta^*$ is decreasing in β .

Supplementary material

Supplementary material associated with this article can be found, in the online version, at [doi:10.1016/j.jebo.2023.01.016](https://doi.org/10.1016/j.jebo.2023.01.016)

References

- Amir, R., 1996. Strategic intergenerational bequests with stochastic convex production. *Economic Theory* 8 (2), 367–376.
- Amir, R., 1997. A new look at optimal growth under uncertainty. *Journal of Economic Dynamics and Control* 22 (1), 67–86.
- Angeletos, G.-M., Laibson, D., Repetto, A., Tobacman, J., Weinberg, S., 2001. The hyperbolic consumption model: calibration, simulation, and empirical evaluation. *Journal of Economic Perspectives* 15 (3), 47–68.
- Balbus, Ł., Reffett, K., Woźny, Ł., 2014. A constructive study of Markov equilibria in stochastic games with strategic complementarities. *J Econ Theory* 150, 815–840.
- Balbus, Ł., Reffett, K., Woźny, Ł., 2018. On uniqueness of time-consistent Markov policies for quasi-hyperbolic consumers under uncertainty. *J Econ Theory* 176, 293–310.
- Barrett, S., 2013. Climate treaties and approaching catastrophes. *J Environ Econ Manage* 66 (2), 235–250.
- Benhabib, J., Bisin, A., Schotter, A., 2010. Present-bias, quasi-hyperbolic discounting, and fixed costs. *Games Econ Behav* 69 (2), 205–223.
- Brock, W.A., Starrett, D., 2003. Managing systems with non-convex positive feedback. *Environ. Resour. Econ.* 26 (4), 575–602.
- Brozović, N., Schlenker, W., 2011. Optimal management of an ecosystem with an unknown threshold. *Ecol. Econ.* 70 (4), 627–640.
- Chassang, S., 2010. Fear of miscoordination and the robustness of cooperation in dynamic global games with exit. *Econometrica* 78 (3), 973–1006.
- Chatterjee, S., Eyigungor, B., 2016. Continuous Markov equilibria with quasi-geometric discounting. *J Econ Theory* 163, 467–494.
- Clarke, H.R., Reed, W.J., 1994. Consumption/pollution tradeoffs in an environment vulnerable to pollution-related catastrophic collapse. *Journal of Economic Dynamics and Control* 18 (5), 991–1010.
- Costello, C., Ovando, D., Clavelle, T., Strauss, C.K., Hilborn, R., Melnychuk, M.C., Branch, T.A., Gaines, S.D., Szuwalski, C.S., Cabral, R.B., et al., 2016. Global fishery prospects under contrasting management regimes. *Proceedings of the national academy of sciences* 113 (18), 5125–5129.
- Crépin, A.-S., 2003. Multiple species boreal forests—what Faustmann missed. *Environ. Resour. Econ.* 26 (4), 625–646.
- Crépin, A.-S., Nævdal, E., 2019. Inertia risk: improving economic models of catastrophes. *Scand J Econ.*
- Dasgupta, P., Maskin, E., 2005. Uncertainty and hyperbolic discounting. *American Economic Review* 95 (4), 1290–1299.
- Denant-Boemont, L., Diecidue, E., L'haridon, O., 2017. Patience and time consistency in collective decisions. *Experimental economics* 20 (1), 181–208.
- Diekert, F.K., 2017. Threatening thresholds? the effect of disastrous regime shifts on the non-cooperative use of environmental goods and services. *J Public Econ* 147, 30–49.
- Dutta, P.K., Sundaram, R.K., 1998. The equilibrium existence problem in general Markovian games. *Organizations with incomplete information* 5, 159–207.
- Fujii, T., Karp, L., 2008. Numerical analysis of non-constant pure rate of time preference: a model of climate policy. *J Environ Econ Manage* 56 (1), 83–101.
- Gerlagh, R., Liski, M., 2017. Consistent climate policies. *J Eur Econ Assoc* 16 (1), 1–44.
- Harris, C., Laibson, D., 2001. Dynamic choices of hyperbolic consumers. *Econometrica* 69 (4), 935–957.
- Hepburn, C., Duncan, S., Papachristodoulou, A., 2010. Behavioural economics, hyperbolic discounting and environmental policy. *Environ. Resour. Econ.* 46 (2), 189–206.
- Iverson, T., Karp, L., 2021. Carbon taxes and climate commitment with non-constant time preference. *Rev Econ Stud* 88 (2), 764–799.
- Kamihigashi, T., Stachurski, J., 2014. Stochastic stability in monotone economies. *Theoretical Economics* 9 (2), 383–407.
- Kang, M., 2019. Pareto-improving tax policies under hyperbolic discounting. *International Tax and Public Finance* 26 (3), 618–660.
- Karp, L., 2005. Global warming and hyperbolic discounting. *J Public Econ* 89 (2), 261–282.
- Karp, L., Tsur, Y., 2011. Time perspective and climate change policy. *J Environ Econ Manage* 62 (1), 1–14.
- Klein, P., Krusell, P., Rios-Rull, J.-V., 2008. Time-consistent public policy. *Rev Econ Stud* 75 (3), 789–808.
- Krishnamurthy, C.K.B., 2017. Optimal management of groundwater under uncertainty: a unified approach. *Environ. Resour. Econ.* 67 (2), 351–377.
- Krusell, P., Kuruşçu, B., Smith, A.A., 2002. Equilibrium welfare and government policy with quasi-geometric discounting. *J Econ Theory* 105 (1), 42–72.
- Krusell, P., Smith Jr, A.A., 2003. Consumption-savings decisions with quasi-geometric discounting. *Econometrica* 71 (1), 365–375.
- Leizarowitz, A., Tsur, Y., 2012. Renewable resource management with stochastic recharge and environmental threats. *Journal of Economic Dynamics and Control* 36 (5), 736–753.
- Li, C.-Z., Crépin, A.-S., Folke, C., et al., 2018. The economics of resilience. *International Review of Environmental and Resource Economics* 11 (4), 309–353.
- Lindahl, T., Crépin, A.-S., Schill, C., 2016. Potential disasters can turn the tragedy into success. *Environ. Resour. Econ.* 65 (3), 657–676.
- Lindahl, T., Jarunggrattanapong, R., 2022. Avoiding catastrophic collapse in small-scale fisheries through inefficient cooperation: evidence from a framed field experiment. *Environment and Development Economics* 1–19.
- Ludwig, D., Jones, D.D., Holling, C.S., 1978. Qualitative analysis of insect outbreak systems: the spruce budworm and forest. *J Anim Ecol* 315–332.
- Mäler, K.-G., Xepapadeas, A., De Zeeuw, A., 2003. The economics of shallow lakes. *Environ. Resour. Econ.* 26 (4), 603–624.
- Mitra, T., Roy, S., 2007. On the possibility of extinction in a class of Markov processes in economics. *Journal of Mathematical Economics* 43 (7–8), 842–854.

- Norström, A.V., Nyström, M., Lokrantz, J., Folke, C., 2009. Alternative states on coral reefs: beyond coral–macroalgal phase shifts. *Mar. Ecol. Prog. Ser.* 376, 295–306.
- Pershing, A.J., Alexander, M.A., Hernandez, C.M., Kerr, L.A., Le Bris, A., Mills, K.E., Nye, J.A., Record, N.R., Scannell, H.A., Scott, J.D., et al., 2015. Slow adaptation in the face of rapid warming leads to collapse of the gulf of maine cod fishery. *Science* 350 (6262), 809–812.
- Peterson, G., Carpenter, S., Brock, W.A., 2003. Uncertainty and the management of multistate ecosystems: an apparently rational route to collapse. *Ecology* 84 (6), 1403–1411.
- Polasky, S., De Zeeuw, A., Wagener, F., 2011. Optimal management with potential regime shifts. *J Environ Econ Manage* 62 (2), 229–240.
- Ren, B., Polasky, S., 2014. The optimal management of renewable resources under the risk of potential regime shift. *Journal of Economic Dynamics and Control* 40, 195–212.
- Schill, C., Lindahl, T., Crépin, A.-S., 2015. Collective action and the risk of ecosystem regime shifts: insights from a laboratory experiment. *Ecol. Soc.* 20 (1).
- Schill, C., Rocha, J.C., 2019. Uncertainty can help protect local commons in the face of climate change. Available at SSRN 3468677 doi:10.2139/ssrn.3468677.
- Stachurski, J., 2009. *Economic Dynamics: Theory and Computation*. MIT Press.
- Steffen, W., Rockström, J., Richardson, K., Lenton, T.M., Folke, C., Liverman, D., Summerhayes, C.P., Barnosky, A.D., Cornell, S.E., Crucifix, M., et al., 2018. Trajectories of the earth system in the anthropocene. *Proceedings of the National Academy of Sciences* 201810141.
- Topkis, D.M., 2011. *Supermodularity and Complementarity*. Princeton university press.
- Tsur, Y., Zemel, A., 1996. Accounting for global warming risks: resource management under event uncertainty. *Journal of Economic Dynamics and Control* 20 (6), 1289–1305.
- Tsur, Y., Zemel, A., 1998. Pollution control in an uncertain environment. *Journal of Economic Dynamics and Control* 22 (6), 967–975.
- Wagener, F.O., 2003. Skiba points and heteroclinic bifurcations, with applications to the shallow lake system. *Journal of economic dynamics and control* 27 (9), 1533–1561.
- de Zeeuw, A., He, X., 2017. Managing a renewable resource facing the risk of a regime shift in the ecological system. *Resource and Energy Economics* 48, 42–54.