

Resource Management, Present Bias and Regime Shifts

Maria Arvaniti*^{1,5}, Chandra K. Krishnamurthy ^{†2,3,5}, and Anne-Sophie Crépin^{2,4}

¹Chair of Economics/Resource Economics, ETH, Zurich

²Beijer Institute for Ecological Economics, The Royal Swedish Academy of Sciences,
Stockholm

³Department of Forest Economics, Swedish Agricultural University, SLU, Umeå

⁴Stockholm Resilience Centre, Stockholm University, Stockholm.

⁵Center for Environmental and Resource Economics (CERE), Umeå, Sweden.

February 28, 2019

Work in Progress-Please do not circulate

Abstract

We investigate the decision process that present biased resource users face when they manage a resource stock that could undergo a reversible regime shift. We focus on renewable resources with some uncertainty in the growth process and evaluate various policy questions in this set-up. To that end we develop a discrete-time model of hyperbolic discounting agents. We introduce uncertainty in a novel way by focusing on the distribution of future stocks rather than on their transition equation and make no special assumptions about functional forms for utility. We show that a Stationary Markov-Nash Equilibrium (SMNE) not only exists but is unique; and further, that the optimal extraction policy is increasing in the resource stock. Finally, we do not find that the resource user should harvest less with regime shifts than without.

Keywords: Renewable resources; Regime shifts; Hyperbolic Discounting; Uncertainty; Markov Equilibrium

JEL Codes: Q20, C61, C73

*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.

[†]Corresponding author (chandra.kiran@beijer.kva.se)

1 Introduction

Many communities face renewable resource management problems that involve dynamic processes with some degree of randomness and potential critical transitions or thresholds that could trigger so-called regime shifts (abrupt changes in system structure and function, Biggs et al. (2012)) Fisheries could undergo more or less reversible collapses e.g. the Maine cod fisheries (Pershing et al. (2015)), forest regrowth could be adversely affected due to factors such as foragers (e.g. Moose, Crépin (2003)) or altered rainfall patterns drive by climate change (Scheffer et al. (2012); Van Nes et al. (2014)).

Low income communities are often more dependent on the resources they manage and if they are really poor, may also exhibit greater impatience compared to less dependent or wealthier societies. Similarly, policy makers are directly responsible to, and possibly more responsive to the concerns of, current voters, and often are explicitly unable to commit future policy makers to present policies. In any case, many situations in reality is likely to generate a present bias in dynamic decision making. This paper focuses on such situations where resource dynamics can, in addition, undergo regime shifts.

The environmental economics literature has extensively discussed issues associated with using a constant discount rate in problems with long time horizons (see Gollier and Hammitt (2014) for a review) and potential solutions, including the gamma discounting approach (e.g. Weitzman (2010); Freeman and Groom (2015)). Non-constant (e.g. hyperbolic) discounting poses the challenge of time-consistency: a policy appearing optimal from one generation's perspective may not appear so from the standpoint of a subsequent generation. Policy making must consider this aspect because current policy makers have few devices to bind future policy makers' decisions. This aspect is particularly important in many settings including climate change policy and public budgeting. Previous work related to environmental policy using non-exponential discounting has focused largely on climate change policy challenges, with studies explicitly considering some form of hyperbolic discounting largely in the context of deterministic evolution (e.g. Karp (2005); Karp and Tsur (2011)).

Dynamic systems can undergo critical transitions leading to rapid changes in system structure and function induced by crossing a threshold, which are called "regime shifts" (see the reviews of Crépin et al. (2012) and Horan et al. (2011)). To illustrate, an infectious disease can exist in a latent regime in which few individuals in a population are infected or in an outbreak regime in which the disease is rapidly spreading to larger parts of the population (Ludwig et al. (1978)); a coral reef can exist in a clear regime hosting a diverse multitude of species, in a bleached regime, or in an algae-dominated regime, which both host lower species diversity and provide different types of ecosystem services (Norström et al. (2009)). There is also evidence that essential elements of the climate system itself can exist in different regimes (Lenton et al. (2008); Steffen et al. (2018)).

The literature in environmental economics has recently focused on irreversible shifts called "catastrophes", which entail a sizeable, often fixed and perpetual cost (see e.g. Lemoine and Traeger (2014); Martin and Pindyck (2015) and references therein). There is a substantial literature on

reversible shifts, where altering some state variable like an amount of pollutant or harvest can be sufficient to exit an undesirable regime Polasky et al. (2011)). Reversibility can be more or less easy to achieve and the degree of reversibility of a particular regime is called the degree of *hysteresis* (Crépin et al. (2012)). In any case, studies embedding a regime shift framework in a non-constant discounting setting are extremely limited: to our knowledge, in fact, only one study, Karp and Tsur (2011), has done so, in the context of an irreversible regime shift pertaining to the climatic system.

Our objective is to investigate the decision process that present biased resource users face when they manage a resource stock that could undergo a reversible regime shift. We focus on renewable resources with random growth and evaluate various policy questions in this set-up. To that end we develop a discrete-time framework of hyperbolic discounting agents who manage a renewable resource stock subject to growth shocks, with the added threat of a reversible regime shift. We consider two regimes, a desirable and an undesirable one depending upon whether the stock is above or below a known threshold. The regime shift is reversible in the sense that once stock rises above the threshold, the system returns to the high growth regime. To simultaneously investigate the joint implications of regime shifts and impatient resource owners, we use a novel approach for modeling uncertainty in stock dynamics that will turn out to be technically coherent, useful in a broad variety of resource problem contexts, and a key feature of our model set up. The essence of this framework is easily grasped: while regime shifts in a deterministic setting can be represented as a change in the dynamics resulting from crossing a threshold, in our stochastic setting crossing the threshold leads to a change in the distribution function of the stock variable. More concretely, we consider a discrete-time system subject to random disturbances to stock growth, and where the distribution function for the stock exhibits a sharp change when entering (exiting) the –fixed–undesirable region of the state space. Our set-up is general enough to be used to investigate our problems of regime-shifts and time inconsistencies but could be useful for general resource extraction problems, and more common time-consistent preferences as well (including recursive utility approaches). Many problems of pollution control, resource management, and drug control to name but a few could be modelled in similar ways (see e.g. Grass et al. (2008)).

In contrast to the previous environmental literature addressing hyperbolic discounting, we focus on structural characterisation and are able to establish, that the SMNE not only exists but is unique and time-consistent. Notably, no special functional form for utility is required to establish our key results, and assumptions made regarding the stock evolution are closely related to those made in the prior literature (e.g. Polasky et al. (2011)) Moreover, we find that the optimal policy is monotonically increasing in the critical stock threshold. This is in sharp contrast with the results in Brozović and Schlenker (2011) who find a non-monotonic relationship between optimal pollutant loading and the threshold level of stock. We also find, in contrast to much of the literature, no precautionary behaviour in our model.

To our knowledge, our study is the first to establish existence and uniqueness of the SMNE, indeed an optimal policy, for the general regime shift case. Prior studies establishing similar properties for even the geometric discounting case have resorted to either special functional forms (e.g. Polasky

et al. (2011)) or special assumptions regarding uncertainty (e.g. Brozović and Schlenker (2011)) while those considering regime shifts with non-constant discounting (e.g. Karp (2005)) have faced the problem of non-uniqueness of solutions or have had to restrict policy space (Karp and Tsur (2011)). Our results pertain, in fact, to the case of either geometric or hyperbolic discounting for the case of reversible regime shifts.

The paper proceeds as follows: §2 provides an overview of the related literature, including our interpretation of hyperbolic discounting, while §3 details key aspects of our model set up, including how random growth shocks are modeled. §4 provides our main results for the formulation where regime shifts depend upon pre-extraction stock, while §5 details the more common case of post-extraction stock determining growth. Finally, §6 provides a discussion of how our results relate to resource extraction policy. Technical aspects of our model, including proofs of certain results, are provided in an Appendix.

2 Related Literature

This paper builds on several parts of the scientific literature related to management of resources with regime shifts or catastrophes; discounting and impatience; and general methods for studying and solving stochastic optimization problems. In order to aid understanding, we relate the key features of each of these aspects to those extant in the current environmental and resource economics literature.

2.1 Regime shifts and catastrophes

The most common way of modelling catastrophes in the economics literature is to consider catastrophic events as penalty functions with an associated hazard rate. In Clarke and Reed (1994)'s model, some random environmental process whose occurrence probability depends only on current pollution level can trigger irreversible catastrophic events. Instead, Tsur and Zemel (1996) and Tsur and Zemel (1998) consider events that occur as soon as the stock of pollution reaches a possible unknown threshold but do not otherwise depend on exogenous environmental conditions. Moreover, these events can be reversible: it is possible to recover from their impact (possibly with some regenerative activities). 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. It turns out that endogenous uncertainty and reversibility are both crucial for the optimal policy outcome: they always imply more conservation (references needed) In contrast, exogenous uncertainty and/or irreversible events generate a non-monotonic relationship between uncertainty and precautionary activities: an exogenous increase in the risk of catastrophe can increase or decrease the degree of precautionary activities undertaken by the resource managers behaving optimally. A different approach, perhaps more aligned with the ecological literature, is to model a regime shift as a change in the system dynamics rather than as an abrupt and irreversible catastrophic event or sudden collapse in the resource stock. That literature often uses continuous time models with convex-concave growth function generating

reinforcing dynamics (Wagener 2003; Mäler et al 2003; Brock and Starrett 2003, Crépin 2003). Even in this case the stochastic event can be modeled in a similar way, In this case, precaution is an optimal strategy as well for endogenous problems but not otherwise (Polasky et al. (2011) and see overview in Li et al. 2018 .

Earlier contributions to stochastic resource management specify the change between an old and a new level of stock as arising from the addition of an "error term" to a deterministic growth function. (e.g. Clark? and for resources with regime shifts eq. (5), Polasky et al. 2011 Crépin et al 2011, Kiseleva and Wagener). In that literature the hazard rate can be modelled in two ways: either as time distributed, with the event occurring at each point in time with a certain probability (e.g. Polasky et al. (2011)), or as state-space distributed, with the event occurring once a certain threshold related to some state variable is reached (Nævdal 2006, see also Crépin and Nævdal 2017 for an overview of both approaches). In these models, the effects of the regime shift manifests immediately once the extraction or pollutant loading occurs and the uncertain reinvestment is realised. Finally, for many ecosystems and time-scales considered, immediate regime shifts are less likely (Peterson et al. (2003); Biggs et al 2012). Only Crépin and Nævdal 2017 and Salanie and Liski 2018 explicitly consider inertia.

We introduce a regime shift into the classic framework used in discrete-time stochastic renewable resource models (reviewed in e.g. Olson and Roy (2000)): resource growth is subject to uncertainty, the regime shift threshold is fixed and known, and a regime shift can occur after the extraction decision is made. However we model uncertainty differently from the standard setting (reviewed in e.g. Olson and Roy (2000)) and follow the approach in Amir (1997) and the subsequent literature. Consequently, our set-up is somewhat different from the classic time-distributed regime shift models already discussed. The key aspect of our model set up is that the distribution function for stock exhibits a sharp change when entering (exiting) a fixed undesirable region of the state space. Thus, we shift the focus from changes (discrete or continuous) in the transition equation to similar changes in the distribution function of the stock. Clearly, adding uncertainty to a deterministic transition structure (e.g. in eq. (5)) also leads to an associated transition function. However, in many settings, dealing with the resulting system becomes difficult. Our approach has at least two substantial advantages. First it becomes much easier to analyse the system and second the regime shift depends on past stocks rather than current extraction, a realistic situation if there are lags (Peterson et al. (2003)). Many aspects of our model are standard in the stochastic resource extraction and stochastic growth model (e.g. Olson and Roy (2006), Nowak (2003), Nowak (2006)) literature, to which we direct the reader for further details. Finally, in our framework, unlike with time-distributed catastrophes, which are guaranteed to occur at some finite time, while the possibility of extinction is strictly positive in all time periods, extinction need not, and indeed will not (with most strategies), occur with probability 1.

2.2 Discounting

The choice of discounting has been a long debated issue when evaluating public projects and policies with very long time horizons and impacts that will affect future generations such as the optimal response to climate change. The discussion involves a wide-variety of approaches, from market-based approaches to those based upon ethical considerations, and also includes consideration of accounting for uncertainty and shocks (see Gollier and Hammitt (2014) for an excellent review). Although there is no consensus for the correct approach to discounting, experimental, empirical and conceptual contributions all suggest that there is little evidence for constant discounting over long horizons.

Rabin (1998) offers a psychological basis for time variant discounting focusing among others on preference reversals and a taste for immediate gratification. Phelps and Pollak (1968) provided the first formal model of such “time varying” (and inevitably time inconsistent, see Strotz (1955)) preferences. These results were supported by other studies from both psychology (Kirby and Herrnstein (1995)) and environmental economics (Cropper et al. (2014)). Apart from a taste for immediate gratification, several conceptual arguments have been used to justify the deviation from geometric discounting.¹ Evidence also suggests that heterogeneity in individual time preferences and the need of aggregation in public decisions like environmental projects can imply declining discounting.(see Gollier and Zeckhauser (2005) and Jackson and Yariv (2015)). Finally, the environmental economics literature (e.g. Sumaila and Walters (2005), Heal (2000),Karp (2005)and Karp (2007)) has focused on slightly different ways to motivate hyperbolic (or at least declining) discount rates: the Weber-Fenchel Law, the differences between inter- and intra-generational discounting or those based upon time perspective.

We interpret hyperbolic discounting as arising largely from “present-biased” policy makers. In the terminology common in the literature on hyperbolic discounting, these policy makers (be they individuals or societies) discount utility gains within their own life time (or generation) differently from those later on. In these cases, the regulator at generation/time t might procrastinate when faced with taking decisions that impose costs on the current generation, with benefits realized only later. This is particularly a problem with issues (e.g. climate change or resource extraction) that require a series of sequential, rather than single-shot, decisions; where the path of a (state) variable is to be chosen by a series of decision makers, none of whom can force the future decision maker (DM) to agree to a path chosen by the past DM. As already detailed, both individual and collective decision making can arrive at a declining discounting schedule. As a practical matter, we note that many institutions of collective decision making (e.g. parliaments) are explicitly so set up to make no binding decisions upon future DMs, amply illustrated by withdrawals from the Kyoto Protocol of

¹Rubinstein (2003) presents experimental evidence suggesting that individuals ignore small differences and focus on large differences when comparing two alternatives while the ability to distinguish differences diminishes with time. This can be relevant when it comes to “long run” environmental problems and agents who have to take into account the immediate next generation and successive generations in the far future. In fact, as Saez-Marti and Weibull (2005) have pointed out, any pure altruism expressed toward the long-term generations leads inevitably to changing utility-weighting over time.

both Australia and Japan (post a few years of operation, in the Australian case?). In any case, we interpret present-biased preferences as arising largely from the responsibility of a DM in the current generation largely to individuals in the current generation.

Few studies in the literature address environmental problems with non-constant discounting. The most closely related to our work focused on hyperbolic discounting that restricted attention to certain specific set of policies (e.g. full abatement or no abatement in Karp and Tsur (2011)) and on numerical characterisation (e.g. Karp (2005)). To our knowledge, Karp and Tsur (2011) are the only study to embed hyperbolic discounting in a model of catastrophic climate related damages modelled as irreversible and permanent welfare loss. Similar to Karp and Tsur (2011), we assume “present-biased” policy makers who discount utility gains during their own life time differently from those after their time. Hence the regulator might procrastinate in taking costly decisions, e.g. for the climate policy case, DMs might want to abate eventually, but just not yet. As is common, we assume that the regulator cannot commit to future actions and cannot dictate the decisions of future policy makers. In consequence, the regulator plays an intergenerational game with his successors, taking into account that although he cannot directly affect their decisions, he can affect the environment they inherit. Each policy maker cares for the present and future payoffs but treats bygones as bygones.

2.3 The Equilibrium: Uncertainty and equilibrium notions

Uncertainty, resource growth and regime shifts

Stochastic models differ from deterministic in many ways in economic settings, leading some times to simplifications and at other times to additional difficulties requiring the use of newer tools and approaches. Simplifications are often possible due to the ‘convexifying’ nature of uncertainty in many economic settings, aspects which have been used in game-theoretic and hyperbolic discounting-related setting (e.g. cite). In contrast for resource extraction, the addition of random terms to resource growth generates challenges involving the convexity of growth function, because unbounded shocks lead to unbounded utility. In particular, convexity requirements cannot possibly be met under uncertainty and regime shifts (except with very special assumptions). In addition, stochastic models of renewable resource extraction, surveyed in e.g., even with a single stock, are usually characterizable only under conditions which include some degree of ‘smoothness’ of the utility and transition.² Consequently, piece-wise definition of the transition functions induce difficulties and require specialized assumptions on how error terms enter the stock transition equation.

These problems largely stem from treating stochastic models as mere extensions of the deterministic one. However, such an approach is often at odds with the very different nature of models that incorporate uncertainty. To overcome these challenges, we take a different—to that currently used in models of regime shifts— approach to specifying the transition equation (detailed in). In

²Even lattice-theoretic approaches cannot easily dispense with requirements regarding smoothness of the transition, see e.g. cite.

essence, we treat stochastic models as a separate class of models, and use the tools and approaches appropriate to this setting. In the stochastic case, current action (reinvestment) leads to a probability distribution on the future state. Thus, working with such (induced) distributions is natural, and so is treating regime shifts directly as (abrupt) changes in the distribution of stock.

Equilibrium notions

The common assumption regarding the DM in a hyperbolic setting is that he cannot commit to future actions and cannot dictate decisions to future DMs: instead he is constrained to decisions within his 'political cycle'. The decision framework facing a policy maker at period t then naturally assumes a strategic aspect, in particular resembling the intergenerational game with his successors, taking into account that although he cannot directly affect their decisions, he can affect the environment they inherit. Each policy maker cares for the present and future payoffs but treats bygones as bygones. The traditional way of approaching problems of hyperbolic discounting, recursive decision theory, is challenging due to time-inconsistency: what is optimal for a decision maker at time t is not so for the decision maker at time $s \neq t$, leading to many challenges to the recursive approach, including lack of continuity in preferences.³

An alternative, and common, approach is to picture the problem as a dynamic game between selves at different points in time. In a deterministic discrete-time setting, it has been difficult to establish important properties of dynamic economic models, in particular growth-model frameworks, when the decision maker is hyperbolic, except under specific conditions (citations). Recent studies, however, have attempted to unite the approaches of stochastic games to again apply recursive methods to problems with hyperbolic decision makers, and restricting attention to the more practically useful pure strategies. Our analysis, in particular, builds upon three distinct modeling frameworks: those applying recursive approaches, via the use of methods from stochastic discounted games, to problems with hyperbolic discounting (Balbus et al. (2014), Balbus et al. (2018)—henceforth BRW18, Nowak (2003)); stochastic growth models with alternative perspectives on modelling uncertainty (Amir (1997); Nowak (2006)); and multi-player stochastic capital accumulation frameworks (Amir (1996)). In view of its origins in stochastic dynamic (exponentially) discounted games, the appropriate notion of equilibrium is that of a Stationary Markov Nash Equilibrium (SMNE). A key advantage to

³The problems arising in the hyperbolic discounting framework are discussed in many studies in economics, both theoretical and applied, see e.g. Krusell et al. (2002); Klein et al. (2008); Karp (2005).

this approach is that, if it exists at all, the SMNE is time-consistent.^{4,5}

3 Model Details

The basic set-up is a standard regime shift model in discrete-time (building on ideas from e.g. Brozović and Schlenker (2011)) embedded in a stochastic quasi-hyperbolic capital accumulation setting (e.g. Balbus et al. (2015)). It can be described as follows. Let $X_t \in \mathcal{X}$ denote a random variable for stock level at time t , with x the realized stock. The state space $\mathcal{X} \subset \mathbb{R}^+$ is not necessarily compact e.g. $[0, \infty)$. Extraction, $q(X_t)$, and “reinvestment”, $a(X_t)$, are the two key functions we define next. Consider a situation when after observing the stock $X_t = x$ at the beginning of period t , the decision maker chooses an extraction level, $q_t \in A(x) := [0, x)$, and leaves $a_t := x - q_t$ as the reinvestment.⁶ Reinvestment and starting stock levels, together with the threshold, lead to the next period stock, X_{t+1} via the transition function (also called a ‘stochastic kernel’) $Q(dX_{t+1}|x, a, \underline{X})$. In our setting, realized stock levels below \underline{X} trigger a shift to substantially lower stock regrowth compared to that above the threshold (constituting the ‘regime shift’ under question).

Instead of specifying X_{t+1} as arising from the addition of an “error term” to a deterministic growth function (as for example in Brozović and Schlenker (2011)) , we directly specify the next period stock via the following stochastic specification,

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

The transition function, Q , maps the state space to itself and defines a probability distribution over the next period stock.⁷ This function can be more easily visualized in terms of its associated distribution function, F , defined as $F(b|x, a) \equiv \Pr(X_{t+1} \leq b|x, a) := Q([0, b]|x, a)$.⁸ We return to

⁴ The notion of a SMNE encompasses three distinct concepts in the stochastic (exponentially) discounted dynamic game that the decision process for a series of hyperbolic decision makers represents: stationarity; Nash equilibrium; and Markovian strategies. By stationarity of a strategy one means that a strategy is required to depend only upon the current state (resource stock), not on past history (including stock or time t). For a strategy to be a Nash equilibrium, it is required that it dominates all potential strategies, a definition familiar from deterministic, static and repeated games. Finally, a strategy is Markovian if it depends only upon the current ‘stage’ of the game (i.e. current state and time t). By these definitions, a stationary Markovian strategy clearly does not depend upon the stage (or time variable t) of the game, only the current state variable. In consequence, a stationary Markovian (SM) strategy which also is a Nash equilibrium is time consistent, virtually by definition (since a time-inconsistent SM strategy cannot represent a Nash equilibrium which, recall, dominates all SM strategies). See Dutta and Sundaram (1998, §2) or §2.3.6 of Levy (2013) for precise definitions and detailed discussion of these notions.

⁵We note that what is termed a Markov Perfect equilibrium (MPE) in the literature (used e.g. in Karp and Tsur (2011)) is often also called the Markov Nash Equilibrium, according with the definitions provided in footnote 4. These equilibria thus need not be stationary, and consequently, for hyperbolically discounting agents, need not be time-consistent. This distinction is important to bear in mind, and is discussed in e.g. Balbus et al. (2015).

⁶Note that in models of fisheries, what we term “reinvestment” here is termed “escapement” instead. We use the terminology of stochastic growth models in the interest of generality.

⁷More formally, it is the distribution induced by the transition probability of the Markov chain that the stock series, $\{X_t, t = 0, 1, 2, \dots\}$, represents, under any stationary policy (extraction decision), $q(X)$. Since this is common in the literature on dynamic programming and stochastic optimal growth, we will use the simpler term “distribution of X_t ” for the more formal “distribution induced by the transition probability”.

⁸Since attention is only restricted to stationary policies, time subscripts on q, a will be dropped henceforth.

motivating and detailing the structure of the transition function after explaining briefly the objective function .

We follow Brozović and Schlenker (2011); Polasky et al. (2011) in assuming that the effects of regime shifts are manifested directly, and exclusively, upon stocks, not upon utility. Hence utility U , derived only from extraction (following Karp and Tsur (2011)), is non-negative and strictly concave, possibly unbounded above. Finally, parameters $0 < \delta \leq 1$ and $0 < \beta \leq 1$ denote respectively the usual discount factor and the degree of present bias, detailed next. The objective of the decision maker can be represented by

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

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

The discounting aspect embodied in eq. (2) may be described briefly thus: 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, with larger rate of decline in the 'short-run' than in the 'long-run'. The intensity of the rate of decline increases as β decreases. Preferences exhibiting these characteristics are termed quasi-hyperbolic or present biased,⁹ which is the reason for using it in our context. Note that while discount rates that decline over time may in fact seem to be desirable from a long-run decision making perspective, this form of discounting in fact leads to strengthening the preference for immediate pay-off, relative to exponential discounting. In any case, lower values of β represent a stronger bias for the present, and any value of $\beta < 1$ captures time-inconsistent preferences.

3.1 Regime shifts

In stochastic renewable resource models, the key object of interest is the distribution (including moments) of next-period stock, X_{t+1} , since this determines aspects that directly affect decision process even in cases of no regime shifts (citations). Consequently, with regime shifts, one is interested in the effects of regime shift upon this distribution, with the anticipation that this distribution is adversely affected (e.g. probability of obtaining a higher stock is reduced) below the threshold. Paralleling the deterministic case, regime shifts in the stochastic case may be formalized as:

⁹The 'hyperbolic' aspect in the term 'quasi-hyperbolic' refers to the greater decline in the short-run compared to the long-run. This is easiest to see in the continuous time formulation (e.g. Angeletos et al. (2001)) wherein the short-run discount rate, $-\ln \beta\delta$, is larger than the long-run, $-\ln \delta$. The latter interpretation, of 'present bias', in fact follows from the observation that all future pay-offs are discontinuously discounted by the amount β , relative to the case of exponential discounting. This is easiest to see in a continuous-time formulation e.g. eq.(3) in Benhabib et al. (2010).

$$F(X_{t+1}|X_t = x, a_t = a) = \begin{cases} 0, & x \leq 0 \\ G_1(\cdot|x, a), & 0 < x \leq \underline{X}, \\ G_2(\cdot|x, a), & \text{else} \end{cases} \quad (3)$$

with G_1 and G_2 representing appropriate distribution functions. Thus, the distribution function of stock, F , is different below and above the threshold (formally, the distribution function is a piece-wise sum of two distribution functions). Before proceeding with this particular formulation, it is worth exploring the canonical transition structure in the literature, in the deterministic resource extraction (eq. (4)) and stochastic pollutant loading case (eq. (5)) respectively:

$$X_{t+1} = \begin{cases} H_1(X_t, a_t), & X_t \leq \underline{X}, \\ H_2(X_t, a_t), & \text{else} \end{cases}, \quad (4)$$

$$X_{t+1} = \begin{cases} X_t + l_t + \epsilon_t, & \text{if } X_t + l_t + \epsilon_t < X_c \\ X_t + l_t + r + \epsilon_t + u_t, & \text{else} \end{cases}. \quad (5)$$

H_1 and H_2 in eq. (4) are deterministic stock growth functions and a_t is reinvestment, l_t is the pollutant loading, X_c the exogenous threshold, r is the additional pollutant loading upon crossing the threshold and ϵ_t and u_t are "error terms" in eq. (5). In essence, eq. (4) differs from eq. (5) only in terms of a specific (linear) functional form for H and the addition of a random term affecting future stock (apart from the obvious sign change in going from extraction to loading), indicating that it is simply a special case of eq. (4). Note that eq. (5) is a version of the transition equation in Brozović and Schlenker (2011) and Peterson et al. (2003). In view of our interpretation of the region $X_t < \underline{X}$ as being undesirable, clearly $H_2 \geq H_1$; similarly, by construction, the pollutant stock X_{t+1} is larger beyond the threshold X_c . For our definition of a regime shift in eq. (3), different assumptions regarding changes post- \underline{X} lead to different relationship between G_1 and G_2 . Here we illustrate some interesting cases. Consider first the case where current stock is in the region $(0, \underline{X}]$, implying that $\mathbb{P}(X_{t+1} \geq b|x \leq \underline{X}, a) := 1 - G_1(b|x, a) \leq \mathbb{P}(X_{t+1} \geq b|x > \underline{X}, a) := 1 - G_2(b|x, a)$, $\forall b > 0$. In terms of distribution functions, then, this yields the expression $G_2(b|x, a) \leq G_1(b|x, a)$. This is equivalent to stating that the distribution of stock, X , beyond the threshold, first order stochastically dominates (FSD) that below the threshold. This requirement of FSD indeed represents the most natural extension of the deterministic notion that $H_2 \geq H_1$.¹⁰ In addition, in the deterministic setting, when the threshold value increases, the desirable region of the state space shrinks, indicating that the growth function is decreasing in the threshold. An extension of this notion to a stochastic setting is that the distribution of the next period stock is 'decreasing' in the threshold. The notion

¹⁰On the other hand, in cases of pollutant loading (which motivated the analyses of Brozović and Schlenker (2011), Peterson et al. (2003)), the contrary is true i.e. staying within the threshold is beneficial, indicating that $G_1 \leq G_2$ i.e. that the distribution G_1 FSD G_2 . Thus, the framework detailed in eq. (3) is flexible enough to accommodate both cases.

of decreasing that is appropriate here is that of stochastically decreasing i.e. that the distribution of stock, $F(y|x, a, \underline{X})$, is stochastically *decreasing* in the threshold. This notion captures, to reiterate, the very intuitive fact that increases in the threshold, for any level of starting stock (x), makes it more likely that stock regeneration is lower.

It is worth noting that for many ecosystems of interest, current extraction does not immediately determine the state of system, whose state may be only observed periodically or with substantial noise; alternatively, there may be substantial lags in the eco-system considered. In these cases, it is more plausible that rapid regime shifts are less likely. Consequently, current loading (extraction) does not directly affect threshold crossing, as is also argued in Peterson et al. (2003). This is the approach taken in our study, and encapsulated in eq. (3).¹¹ This is also in fact, as we shall later see, the more challenging case to study. The differences between this formulation and one that considers the post-extraction (or loading) stock as the basis for regime shifts (as in eq. (5)) will also be subsequently illustrated.

3.2 Stochastic Transition and Regime shifts

With state-space-based regime shift formulation, stock dynamics post-shift can differ in many different ways from pre-shift, varying from a simple (often constant) level shift (e.g. in eq. (5)) to those incorporating more general differences (involving e.g. distance from the threshold). In this formulation, in any case, crossing a threshold has additional and abrupt effects on stock evolution. For the particular formulation in eq. (3), the changes post-regime-shift are reflected in how F changes in the region of the threshold: piece-wise definition ensures a discontinuous change while a continuous formulation¹² allows for a sizeable yet continuous (in the region around the threshold) change. The shape of F , parameterized by \underline{X} (among other parameters), controls the intensity of change in the stock distribution post-regime-shift. In any case, while formulations of the form in eq. (5) also lead to distributions upon the stock (cite), the key distinction between the two formulations is this: the effects of the regime shifts upon the distribution of stock is more obscure in eq. (5), since the link between the threshold and stock distribution, the key object of interest, is only indirect.

We discuss next a few details regarding the key features of the stochastic transition, Q , introduced in eq. (1). Since we build upon the standard framework in the literature, we provide only the necessary technical preliminaries and direct the reader to the relevant literature for fuller details, in particular Balbus et al. (2014), BRW18, Nowak (2003), Amir (1997), Nowak (2006). We consider a specific

¹¹More precisely, our formulation corresponds in timing to that in Peterson et al. (2003, eq.(1)), with regime shift depending upon pre-extraction stock i.e. the regime shift is independent of current extraction (loading).

¹²One way to obtain such a formulation is to replace the indicator function for regime shifts in eq. (3) with a more smooth function. More specifically, instead of writing F as $G_1(\cdot|x, a) \mathcal{I}\{x \leq \underline{X}\} + G_2(\cdot|x, a) (1 - \mathcal{I}\{x \leq \underline{X}\})$ (simply a restatement of eq. (3)), one can also write it as $G_1(\cdot|x, a) \theta(x, \underline{X}) + G_2(\cdot|x, a) (1 - \theta(x, \underline{X}))$, with $\theta \in [0, 1]$ and $\theta(x \leq \underline{X}) > \theta(x > \underline{X})$. In other words, the measure θ puts more weight on the undesirable distribution G_1 when $x \leq \underline{X}$, ensuring that the required condition, next period stock is probabilistically lower when $x \leq \underline{X}$, is met. The degree to which it is lower is controlled by the measure θ . The continuous formulation is explored subsequently in §3.

form of the transition structure,

$$Q(\cdot|a, x) = (1 - P(\cdot|a, x)) \delta_0 + P(\cdot|a, x), \quad (6)$$

where $P(\cdot|a, x)$ is some probability measure satisfying certain properties to be detailed next and δ_0 is the Dirac measure concentrated at 0 stock. The transition function, Q , is in general a function of both the state, x , and reinvestment, a , and is fully determined by the measure P (with associated distribution function, F_P). The noteworthy point here about Q is the presence of an “atom” at 0 stock, meaning a strictly positive probability of reaching state 0 from any other stock value. In view of the fact that utility is assumed to be bounded below (by 0), in essence this strong so-called “mixing” assumption implies that (i) reaching 0 stock leads to the ‘end’ of this resource (economy), and (ii) that the probability of this event (of extinction) is strictly positive, $\forall t$.

While seemingly strong, we note that this structure bears some similarity to most models of catastrophic regime shifts, and in fact embodies weaker assumptions. In those frameworks, the catastrophe is guaranteed to occur at some finite time, leading to a permanent reduction in utility levels. In our framework, however, while the possibility of extinction is strictly positive in all time periods, extinction need not, and indeed will not with most strategies, occur with probability 1 i.e. it is certainly not the case that $\lim_{t \rightarrow \infty} P(b|a, x) = 0, \forall b > 0$. Compare, for instance, models such as blah. When combined with the fact that a judicious choice of parameters for the measure P will ensure that this probability in fact can be reduced to levels as low as required, in practice this will turn out to be a rather weak restriction on the transition.

Remark 1. *We note that an alternative interpretation of stock collapse is feasible. To see this, define Y to be the actual stock of a resource, with $X_t := \log(Y_t - b)$, where $b \geq 0$ is a parameter related to a certain base population level. In such a case, the use of X_t instead of Y_t implies that the collapse of X_t to 0 has the implication that stock Y_t collapses to some natural or base stock level, $1 + b$. Since the \log is a bijective transformation, working with X_t instead of Y_t leads to identical optimal policies and system dynamics (on the latter point, see Stachurski (2007)).*

3.3 Solution Method

Our set-up parallels the more common Bellman recursion for geometric discounting, with differences highlighted (bearing in mind the rather different notion of equilibrium). In this spirit, we denote by $w \in \mathcal{A} := \{w : \mathcal{X} \rightarrow A; w \text{ bounded}, w(x) \in A(x)\}$ a (pure strategy) SMNE for a quasi-hyperbolic agent that satisfies the functional equation Equation (7)

$$w(x) \in \arg \max_{q \in A(x)} \left[U(q) + \beta \delta \int_{\mathcal{X}} V_w(x') Q(dx'|x - q, x) \right], \quad (7)$$

where $V_w : \mathcal{X} \rightarrow \mathbb{R}^+$ is the continuation value function for the household of future selves. If such strategies exist (i.e. if $\mathcal{A} \neq \emptyset$), they are time-consistent. Similar to the case of more conventional

recursive approaches, the continuation value function satisfies the following recursion,

$$V_w(x) = U(w(x)) + \delta \int_{\mathcal{X}} V_w(x') Q(dx'|x - q(x), x) \quad (8)$$

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

$$W_w(x) = U(w(x)) + \beta \delta \int_{\mathcal{X}} V_w(x') Q(dx'|x - q(x), x), \quad (9)$$

one obtains the relation

$$V_w(x) = \frac{1}{\beta} W_w(x) - \frac{1 - \beta}{\beta} U(w(x)). \quad (10)$$

Equation (10) is the so-called generalised Bellman equation. Thus, much like the conventional Bellman equation, any fixed point, V^* , of a suitably defined operator upon V_w (see Balbus et al. (2018, §2.3) for details) corresponds to the value function of a time-consistent SMNE. In general, reflecting the well-known non-uniqueness of solutions in the hyperbolic discounting case (Karp (2005); Krusell et al. (2002); Harris and Laibson (2001)), the set of fixed points needs not be a singleton (see Balbus et al. (2015, §4.1)). If, in addition, V^* is a unique fixed point, then there is a unique, pure strategy, time-consistent SMNE, w^* . All of our results will use only generic properties of resource problems, inspired by the use of general regeneration functions in Polasky et al. (2011). Subsequently, we illustrate specific functional forms for the respective measures which satisfy many of these properties.

4 Results

4.1 Main Result

Our analysis is largely based upon lattice theory,¹³ dispensing largely with the reliance upon assumptions related to convexity of sets and existence of derivatives of the n^{th} order. Our main result is to establish the uniqueness of a time-consistent SMNE w^* , under the following set of substantive assumptions gathered under

Assumption 1. *a. $U : \mathcal{X} \rightarrow \mathbb{R}^+$ is positive, increasing and strictly concave, with $U(0) = 0$ (i.e. U is bounded below, by 0 for convenience).*

b. for $x, a \in \mathcal{X}$, the transition probability Q has the structure in eq. (6), with the measure $P(\cdot|a, x)$ satisfying the following properties,

- *for $x \in \mathcal{X} \setminus \{0\}$, $a \in [0, x]$, $P(\mathcal{X}|a, x) < 1$, $P(\mathcal{X}|0, 0) = 0$;*

¹³A very brief introduction to lattice theory in the context of natural resource economics is available from e.g. Krishnamurthy (2017); Knapp and Olson (1995) while detailed treatment with an economics focus is available from e.g. Amir (2005); Topkis (2011).

- for every bounded function $v : \mathcal{X} \rightarrow \mathbb{R}^+$, the function $(a, x) \mapsto \int_{\mathcal{X}} v(x') dF_P(dx'|a, x)$ is continuous in (a, x) and increasing and concave in a .^{14,15}

With these preliminaries concluded, we can now state our main result:

Theorem 1. *For the renewable resource extraction problem in eq. (2) and a transition function with regime shift for the resource stock in eq. (6), there is a unique, bounded value function, V^* and correspondingly, a unique time-consistent SMNE extraction policy $w^* \in \mathcal{A}$.*

Proof. This is Theorem 1, Balbus et al. (2018). □

We next discuss why Assumption 1 is satisfied in almost all models of regime shifts used in renewable resource frameworks. Denote by F_P the CDF associated with the measure in eq. (6). Note that increasing in a is virtually a defining feature of renewable resource models, while concavity (from which follows continuity), states a rather intuitive facet of resource dynamics: increasing reinvestment yields marginally decreasing benefits. Thus, these features are rather intuitive and follow from basic behaviour of (discrete-time) resource dynamics, with no special assumptions required or used. These aspects imply that the function $(a, x) \mapsto \int_{\mathcal{X}} v(x') dF_P(dx'|a, x)$ is continuous (integration preserves continuity), increasing and concave in a . To see most easily that properties of F_P carry over to the integral, choose for F_P any appropriate function separable in a .

Remark 2. *The dynamic program associated with the optimisation problem is defined, analogous to the exponential discounting case, by the operator conventionally denoted T , as*

$$TV(x) = \frac{1}{\beta}AV(x) - \frac{1-\beta}{\beta}u(BV(x)), \quad (11a)$$

with the operators A and B defined as

$$AV(x) = \max_{q \in A(x)} \left[u(q) + \beta\delta \int_{\mathcal{X}} V(x') Q(dx'|x - q, x) \right], \quad (11b)$$

$$BV(x) = \arg \max_{q \in A(x)} \left[u(q) + \beta\delta \int_{\mathcal{X}} V(x') Q(dx'|x - q, x) \right]. \quad (11c)$$

Remark 3. *Clearly, in keeping with the literature, our results pertain to the case of “no commitment” of future generations to current policies, the most plausible equilibrium notion explored in the literature (including in Karp and Tsur (2011)).*

¹⁴We follow the literature in using the notation $(a, x) \mapsto \int_{\mathcal{X}} v(x') dF_P(dx'|a, x)$ to denote that $\int_{\mathcal{X}} v(x') dF_P(dx'|a, x)$ is viewed as a function of a, x . Stated differently for this particular case, $I(a, x) := \int_{\mathcal{X}} v(x') dF_P(dx'|a, x)$ is continuous, increasing and concave in a (for every given x).

¹⁵It follows from its definition in eq.(6) that integration w.r.t. the measure (induced by) Q is identical to that w.r.t the measure P , a fact that will be repeatedly used. See Lemma 2 in Balbus et al. (2018) for a formal proof.

4.2 Structural Characterisation

We next turn to understanding two aspects of the optimal extraction policy ('optimal policy'): do increases in beginning-of-period stock and the threshold lead to unconditional increased (or reduced) reinvestment? Different papers have reached differing conclusions on these questions, as will be detailed later. In investigating these questions, properties of the measure P **defined in assumption 1b.** turn out to be vital. These will be investigated in terms of either the associated distribution function, $F_P(\cdot|a, x, \underline{X})$, or the tail probability function $\tilde{F}_P(\cdot|a, x, \underline{X}) := \mathbb{P}(X_{t+1} > b|a, x, \underline{X}) \equiv 1 - F_P(\cdot|a, x, \underline{X})$. The notation $\cdot|a, x, \underline{X}$ is intended to convey dependence of the transition upon key model parameters: reinvestment, beginning-of-period stock, and threshold, with arguments not under consideration often suppressed for convenience. In common with usage in the literature, we will refer to F_P as being e.g. stochastically increasing instead of the more cumbersome "random variable distributed as F_P " being stochastically increasing.

In the deterministic case, properties of the state evolution structure are key to obtaining structural results. Similarly, in our case, many of the structural properties are related to properties of the transition function, in particular upon the relationship between x and a or \underline{X} and a in determining next period stock. Instead of imposing very specific functional forms for transition, we consider minimal relationships that are required to yield properties of interest. In addition, we move away from assumptions related to convexity and differentiability which may not often obtain in cases of regime shifts. Viewing F as the stochastic production function for X_{t+1} (the approach advocated in e.g. Amir (1997)), the most important properties we consider relate to substitutability between the different "inputs" in this production function.

4.2.1 Monotonicity of Policies

An important question relates to the relationship between stock levels and extraction: with exponential discounting, larger stocks are often considered to lead to greater extraction, but this conclusion may no longer hold in the presence of non-convexities induced by e.g. stock-dependence in the growth function. In consequence, even with no regime shifts, it is in fact not very uncommon to be unable to establish the result that increasing stock leads to increasing extraction (see e.g. the discussion in Knapp and Olson (1995); Krishnamurthy (2017)). In the case of non-convexity induced by regime shifts, not only is there (to our knowledge) no prior result regarding this aspect but there is also reason to question whether in fact this aspect is even reasonable. This is because there are competing aspects at play: to illustrate, higher stocks can lead to enhanced incentives to extract; yet, this is counter-balanced by the degree to which stock levels are close to the threshold. Consequently, even without the question of present bias, monotonic extraction with complex regeneration function is the exception rather than the norm.

Before proceeding with the technical analysis, it is worth exploring the following interesting aspect: whether x and a are to be viewed as substitutes or complements to each other (on the margin) in terms of their contribution to future stock (recall, from Assumption 1, that the absolute returns

to reinvestment are always positive). We consider it more plausible that higher stock levels reduce the marginal benefits of reinvestment, indicating some degree of decreasing returns on marginal increases in next period stock. In terms of economic intuition, this property implies that increases in reinvestment are more beneficial (lead to larger–probabilistically–next-period stock), on the margin, at lower stock levels than at higher. Were F a deterministic function, this condition is equivalent to F satisfying decreasing differences in (or being submodular in) (x, a) .¹⁶ When F is a distribution function, this condition is formally stated as: $F_P(\cdot|a, x, \underline{X})$ satisfies stochastic decreasing differences in (a, x) . As detailed in lemma 8, this condition is equivalent to stating that $\tilde{F}_P(\cdot|a, x, \underline{X})$, viewed as a deterministic function, satisfies decreasing differences in (a, x) . For \tilde{F} , decreasing differences means that $\mathbb{P}(X_{t+1} > b|x', \cdot) - \mathbb{P}(X_{t+1} > b|x, \cdot)$ is decreasing in a , for $x' > x$. But this simply means that the effect of a upon $\mathbb{P}(X_{t+1} > b|x', \cdot)$ is smaller than upon $\mathbb{P}(X_{t+1} > b|x, \cdot)$. It is important to reiterate that decreasing differences pertain to the rate at which changes in one input affect marginal returns to the other, and are second-order properties. Note that the presence of either non-convexities or even regime shifts do not alter this logic embodied in this condition.¹⁷

Our view encompasses the structure of the transition function in much of the literature, which uses special functional forms with the growth function: depending only upon a or x but not both (e.g. Polasky et al. (2011)); or depending linearly upon a and x (Brozović and Schlenker (2011); Karp (2005)). What is surprising is that even under the form of non-convexity implied by regime shift in the distribution of X , monotonicity of extraction in stock holds. We first state the key assumption we make, relate it to the intuition regarding the transition detailed above, and then proceed to our main result for this section.

Assumption 2. *The function $(x, a) \mapsto \int_{\mathcal{X}} V(x') dF_P(dx'|a, x, \underline{X})$ is submodular (i.e. satisfies decreasing differences) in (a, x) .*

Assumption 2 implies that $\mathbb{E}[V](a, x) := \int_{\mathcal{X}} V(x') dF_P(dx'|a, x)$ satisfies decreasing differences in (a, x) , implying that for the expected value function, increases in a (reduction in extraction) are more valuable, on the margin, when x is smaller than when it is larger. Clearly, this aspect is intuitive. Additionally, we have argued above that F_P satisfies decreasing differences in (a, x) . *This aspect, in combination with the observation that the value function in resource extraction settings is unlikely to be decreasing in stock, directly establishes the required link between properties of F and*

¹⁶ A deterministic function G satisfies decreasing (increasing) differences in scalar arguments (a, x) if increases in one argument are less (more) valuable when more of the other is available i.e. that $G(\cdot, x') - G(\cdot, x)$ is decreasing (increasing) in a , for $x' > x$. When G is differentiable w.r.t. both arguments, the condition simplifies to $\frac{\partial^2 G}{\partial a \partial x} \leq (\geq) 0$. Clearly, this is also the definition of submodularity (supermodularity), to which decreasing (increasing) differences is equivalent (for the case of scalar a and x only). See Amir (2005) for an intuitive presentation of these notions.

¹⁷ There is seemingly an exception to this logic just underneath the threshold, where there is likely a strong additional incentive to invest. However, note carefully that the comparison—in the definition of decreasing differences—is between two scenarios with the same threshold but differing stock levels, in which case the logic outlined above continues to hold even just underneath the threshold. In any case, threshold crossing here is not affected by current reinvestment, precluding this additional incentive for reinvestment.

$\mathbb{E}[V](a, x)$.¹⁸ To summarize, Assumption 2 is not merely plausible, in respect of its implications, but follows directly from the structure of the transition function (ensuring that F_P satisfies stochastic decreasing differences) and the nature of the problem.

Our main result for this section (whose proof is provided in Appendix B) is

Theorem 2. *Under assumptions 1 and 2, the optimal time-consistent policy w^* is increasing in x .*

Remark 4. *We are unaware of any previous study in either the literature on (state-space-based) regime shifts or on stock pollution control with hyperbolic discounting to have established that the decision is monotonic in state. For instance, Brozović and Schlenker (2011), by virtue of special functional forms (e.g. linearity of transition and quadratic benefit function) find pollutant loading to be stock independent; Karp (2005) hypothesises—but is unable to establish—that carbon emissions may well be decreasing in existing stock of carbon. In the terminology of Karp (2005, §3.2.2), actions and stocks in our case are “strategic complements”.*

Remark 5. *It is important to note that theorem 2 does not imply that next period stock, X_{t+1} is increasing in current stock X_t . The chain of logic embodied in this statement is as follows: as current stock size increases, extraction increases and reinvestment decreases, with the reduction in reinvestment not sufficiently large to overwhelm the effect of the initial increase in X_t . This is a finding in e.g. Karp (2005), a (quasi-) hyperbolic case with a linear (and deterministic) transition equation (Proposition 1). In our case, as for any stochastic transition equation, the equivalent (to “ X_{t+1} increasing in X_t ”) notion is that the stochastic kernel $Q(\cdot|a, x)$ (from eq. (6)), is increasing in X . A sufficient condition for this typically involves a Lipschitz condition on extraction¹⁹ along with the transition being stock independent, which is not fulfilled in our case. It has been observed before that in resource extraction problems with exponential discounting, stock dependence can lead to many important properties (e.g. concavity of the value function; Lipschitz continuity of extraction; and next-period-stock increasing in current stock) not necessarily holding (Knapp and Olson (1995); Krishnamurthy (2017)). Our findings here suggest that even in (quasi-) hyperbolic settings, stock dependence leads to essentially the same complexity in reinvestment decisions.*

4.2.2 Monotone comparative statics

Perhaps the most interesting question in our framework is how the optimal extraction decision is affected by the threshold level of stock, \underline{X} , below which we enter the less desirable region with a

¹⁸There are two parts to this argument. The first part is simple: if the value function V is increasing in stock, then it follows (from Lemma 7) that $\mathbb{E}[V](a, x)$ satisfies decreasing differences in (a, x) if $F_P(\cdot|a, x)$ satisfies (stochastic) decreasing differences. The second part is this: intuitively, since the value function represents the value of the problem starting from stock X , and stock contributes to consumption, it is never plausible that it is everywhere (i.e. $\forall X$) decreasing in stock. However, due to many factors (including that the stock dependence of the transition kernel being rather moderate), it is likely that regions of the state space over which the (expected) value function is strictly decreasing is small; in these cases, it can be shown that if the value function over this region is “small enough” then arguments from Lemma 7 continue to hold.

¹⁹Conventionally, the condition is: $\frac{\partial w}{\partial X} \in (0, 1)$, which follows if w , the extraction policy, is Lipschitz continuous.

stochastically dominated regime for the level of the resource stock. We investigate here the effect of the presence of a non-zero threshold upon optimal decision (extraction or reinvestment). One may anticipate, for instance, that increases in the threshold, leading to enlarged adverse region of the state space, will lead to reduced extraction (or increased investment). In essence, motivated by the finding of a monotone (in stock) extraction policy, interest centers upon understanding whether a similar effect of the threshold upon extraction can be established.

Before proceeding with the technical analysis, it is worth articulating the intuition regarding two key properties of a transition function encapsulating a regime shift. The first is that increases in the threshold reduce the probability of reaching a larger next-period stock i.e. $\tilde{F}_P(\cdot|a, x, \underline{X}) := \mathbb{P}(X_{t+1} > b|\underline{X}, a)$ is decreasing in \underline{X} . This property is equivalent to stating that the transition distribution, F_P , is stochastically decreasing in the threshold, \underline{X} .²⁰ The second one pertains to the substitutability between (a, \underline{X}) on the margin. Stated simply, this is a question of whether the marginal returns to reinvestment are increasing or decreasing with the threshold (recalling again that absolute returns are always increasing, by Assumption 1). Following previous literature, we also assume that the marginal reinvestment rate exhibits a degree of 'decreasing returns' i.e. that as the threshold increases, marginal reinvestment levels required to reach a fixed next-period stock are decreasing (implying that $\mathbb{P}(X_{t+1} > b|\underline{X}', \cdot) - \mathbb{P}(X_{t+1} > b|\underline{X}, \cdot)$ is decreasing in a , for $\underline{X}' > \underline{X}$).

Note that the most commonly used explicit functional forms for transition (e.g. that in Polasky et al. (2011); Brozović and Schlenker (2011)) satisfy these two properties. Consider the encompassing functional form in eq. (4): a key feature is separation between \underline{X} and (a, x) and the assumption that the growth function is either linear (Brozović and Schlenker (2011)) or that it depends upon one of x, a but not both (e.g. Polasky et al. (2011) see eq.(2)). The first assumption directly implies in the deterministic case that X_{t+1} is decreasing in \underline{X} (whose extension to the stochastic case in the first property detailed above). The first and the second (together, if H is not linear, with the assumption that $H_{1,a} < H_{1,a}$ indicating that marginal benefits of reinvestment are higher in the desirable region above the threshold, e.g. p.231, Polasky et al. (2011)) together imply that the marginal returns to reinvestment are decreasing in threshold levels.

These assumptions pertain to the structure of the transition function and, we reiterate, preclude neither monotonic nor non-monotonic extraction (loading) in general.²¹

We next state the key assumptions under which the main result for this section will be proved. After a brief discussion of this assumption, we state the key result whose proof will be presented in the Appendix. We begin first with some notation. For a partially ordered set $\underline{\mathcal{X}}$, with $\underline{X} \in \underline{\mathcal{X}} \subset \mathcal{R}^+$, we define the unique SMNE as $w_{\underline{X}}^*$, explicitly focusing on the dependence between extraction and

²⁰A random variable indexed by parameter α , with distribution function $F(\cdot|\alpha)$, is stochastically decreasing in the parameter α if $1 - F(\cdot|\alpha)$ is decreasing in α (in the sense of a function). This is formalized in lemma 8. To understand the intuition, consider an equivalent definition: α shifts F in the sense of stochastic dominance, meaning that increases in α shift probability mass to the left, putting more weight on lower outcomes.

²¹To illustrate, with a very similar structure of transition—in terms of the relationship between (a, \underline{X}) , Brozović and Schlenker (2011) and Peterson et al. (2003) report opposing findings: the former finds (under certain conditions) "precautionary" behaviour while the latter finds no such behaviour.

the threshold (suppressing the argument x in extraction for clarity). Note that due to dependence of the distribution of the measure P , F_P , upon the threshold \underline{X} (viewed as a parameter), key model aspects, including the optimal policy and the operators associated with the generalized Bellman equation, are all explicitly parameterized by the threshold.

Our subsequent analysis stands upon the following additional assumption:

Assumption 3. *We assume:*

- a. u does not depend on \underline{X} and satisfies Assumption 1.
- b. for any $x, a \in \mathcal{X}$ and $\underline{x} \in \underline{\mathcal{X}}$, the transition probability Q has the structure in eq. (6), with the measure $P(\cdot|a, x, \underline{X})$ satisfying Assumption 1.
- c. for every bounded function $V : \mathcal{X} \rightarrow \mathbb{R}^+$, the function $(a, \underline{X}) \mapsto \int_{\mathcal{X}} V(x') dF_P(dx'|a, x, \underline{X})$ is submodular in (a, \underline{X}) and $\underline{X} \mapsto \int_{\mathcal{X}} V(x') dF_P(dx'|a, x, \underline{X})$ is decreasing on $\underline{\mathcal{X}}$.

Assumption 3c implies that $\mathbb{E}[V](a, \underline{X}) := \int_{\mathcal{X}} V(x') dF_P(dx'|a, x, \underline{X})$ satisfies decreasing differences in (a, \underline{X}) , and is decreasing in \underline{X} . Clearly, the latter is intuitive: increases in the threshold lead to reduction of the expected value function by increasing the size of the undesirable region of state space. The former implies that for the expected value function, increases in a (reduction in extraction) are more valuable when \underline{X} is smaller than when it is larger. Put another way, the marginal benefit (in terms of the expected value function) to reinvestment is larger at lower thresholds. As regards the reasonability of Assumption 3, we note that the remarks immediately succeeding Assumption 2 in section 4.2.1 are directly applicable: properties of $\mathbb{E}[V](a, \underline{X})$ follow from those related to $F_P(\cdot|a, \underline{X})$ and the nature of our problem.

Next, we modify the operators defined in remark 2 to explicitly allow them to depend upon the parameter \underline{X} . These are denoted $T_{\underline{X}}$, $A_{\underline{X}}V(x)$ and $B_{\underline{X}}V(x)$ (see Appendix B for details). With these definitions and Assumption 3 in place, we state an important result (proved in Appendix B) which is key to the proof of our main result.

Lemma 3. *Let $\phi : \mathcal{X} \times \underline{\mathcal{X}} \mapsto \mathbb{R}^+$ be a bounded function for each $\underline{X} \in \underline{\mathcal{X}}$ and $\phi(x, \cdot)$ is decreasing for each $x \in \mathcal{X}$. Then $\underline{X} \mapsto T_{\underline{X}}(\phi(\cdot, \underline{X}))(x)$ is a decreasing function of \underline{X} .*

With Assumption 3 and Lemma 3 in place, we are now ready to prove our main result on comparative statics whose proof is relegated to Appendix B.

Theorem 4. *Let Assumption 3 be satisfied. Then the mapping $\underline{X} \mapsto w_{\underline{X}}^*$ is increasing on $\underline{\mathcal{X}}$.*

In words, this result unambiguously indicates that the presence of a threshold leads to increased extraction, for any threshold level. In particular, then, the presence of a threshold leads to increased extraction than without, for any level of threshold.

Remark 6. Notice that all the results hold with weak submodularity of $(a, \underline{X}) \mapsto \int_{\underline{X}} V(x') dF_P(dx'|a, x, \underline{X})$ in (a, \underline{X}) . In particular, by Theorem 2.8.1 in Topkis (2011) $B_{\underline{X}}V(x) = \arg \max_{q \in A(x)} [G(q, V, \underline{X})]$ is constant in \underline{X} but Lemma 3 still holds.

Remark 7. It will be useful to at this point to note that our finding of a monotonic relationship between the threshold and extraction is somewhat 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 (see Fig. 1). 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 never leads to “precautionary behaviour” i.e. a reduction in extraction with a hope of avoiding the regime shift in question. This is what the simulation study by Peterson et al. (2003) reports, in an investigation of emission loading.

5 Regime shifts based upon post-extraction stock

As mentioned earlier, it has been common in the literature to consider ecosystems with very rapid dynamics (see for example Polasky et al. (2011)). In these cases, current extraction (or pollutant loading) immediately determines the state of the system and threshold crossing. Therefore, regime shifts are conditioned on post-extraction stock or reinvestment a_t and the conditions on the distributions functions described in eq.(3) are modified as follows

$$F(X_{t+1}|X_t = x, a_t = a) = \begin{cases} 0, & a \leq 0 \\ G_1(\cdot|x, a), & a \leq \underline{X} \\ G_2(\cdot|x, a), & \text{else} \end{cases} \quad (12)$$

This formulation generates more insights. It turns out that Theorems 1, 5, and 6 are still valid, and in addition we can prove the existence of an Euler equation, which is helpful in characterizing a few more structural properties of the problem.

5.1 Main Results

When regime shifts depend upon post-extraction stock crossing a threshold, we can make the additional assumption that the current state(stock) has no direct effect on the transition probability. This implies that the transition function has the following structure

$$Q(\cdot|a, \underline{X}) = (1 - P(\cdot|a, \underline{X})) \delta_0 + P(\cdot|a, \underline{X}), \quad (13)$$

Under Assumptions 1 and 3, Theorems 1 and 3 remain valid. In particular, the results related to monotonicity of the optimal policies require no further assumptions. In fact, we are able to establish more properties for the optimal policies with fewer assumptions than before. The key feature is of

course that the transition function has no direct dependence on the current stock x . Theorem 5 is an extended version of Theorem 2.

Theorem 5. *Under Assumption 1, the optimal time-consistent policy w^* is increasing in x and Lipschitz with modulus 1.*

Proof. See the Appendix for a proof. □

In other words under the standard concavity assumptions for utility and the rather general properties of the transition probability Q defined in Assumption 1, the optimal extraction policy is increasing in the earlier stock and is continuous, everywhere differentiable, and the absolute value of its derivative is bounded above by 1. Hence despite the threshold, the optimal policy is rather smooth.

5.2 Euler Equation

It has been common in the literature to assume that a smooth Markovian equilibrium exists so that one can apply a generalised Euler equation approach in order to compute it. The problem is that it is not always true that smooth equilibria exist. Below we provide conditions under which our SMNE is unique, interior and differentiable. Then we can state our version of the Generalised Euler equation and characterise a few more structural properties of the problem. For every bounded and increasing function $V : \mathcal{X} \rightarrow \mathbb{R}^+$, let $F_V(a) = \beta\delta \int_{\mathcal{X}} V(x') dF_P(dx' | a(x), \underline{X})$.

Assumption 4. *We assume :*

- a. u is twice continuously differentiable
- b. the function F_V is twice continuously differentiable on $\mathcal{X} \setminus \{0\}$
- c. $\lim_{q \rightarrow 0} u'(q) = \infty$ and $\lim_{a \rightarrow 0} F_V'(a) = \infty$
- d. $|u''| > 0$ and $|F_V''| > 0$ for any $x \in \mathcal{X}$

With Assumption 4 in place we are ready to state our theorem for the existence of a Generalised Euler Equation.

Theorem 6. *Under Assumptions 1 and 4, the optimal time-consistent policy w^* and value function V^* are differentiable on $(0, \infty)$.*

Proof. See Balbus et al. (2018) for a proof. □

With Theorem 6 at hand, we can now write our version for the generalised Euler equation characterising the SMNE investment. Assume that $a^*(x)$ is a differentiable SMNE reinvestment such that $a^* = x - w^*$. Then we have

$$u'(x - a^*(x)) = \beta\delta \frac{d}{da(x)} \int_{\mathcal{X}} V^*(x') dF_P(dx' | a^*(x), \underline{X}). \quad (14)$$

From the definition of the value function we get

$$V'^*(x) = u'(x - a^*(x))(1 - a'^*(x)) + \delta a'^*(x) \frac{d}{da(x)} \int_{\mathcal{X}} V^*(x') dF_P(dx' | a^*(x), \underline{X}). \quad (15)$$

Integrating (15) wrt to x using Lemma 4.2 from Amir (1997) we get after some simplification the following expression

$$u'(x - a^*(x)) = -\beta\delta \int_{\mathcal{X}} u'(x' - a^*(x')) \left(1 + \left(\frac{1}{\beta} - 1\right) a'^*(x')\right) \frac{\partial F_P(dx' | a(x), \underline{X})}{\partial a(x)} \quad (16)$$

The above expression is the generalized Euler equation in a model with hyperbolic discounting. The difference with its counterpart when discounting is exponential is that the constant exponential discount factor δ of the latter model is replaced with $\beta\delta(1 - a'^*(x')) + \delta a'^*(x')$. Following the discussion of Harris and Laibson (2001), this effective discount factor is endogenous and stochastic as it depends on the marginal propensity to reinvest. Since $\beta < 1$, the effective discount factor is positively related to the future marginal propensity to reinvest. The intuition behind this result is the following. 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 lesser the regulator values the future.

A closer inspection of (16) reveals additional insights regarding the effect of β , the short term discount factor, on the time consistent equilibrium level of extraction and reinvestment. One of the fundamental questions in this framework is whether a present biased regulator would extract more than if he had a constant exponential discount factor. We summarize our findings in the following proposition:

Proposition 1. *In the renewable resource extraction problem with regime shifts conditioned on post-extraction stock as described by the transition function in (12), the time consistent equilibrium level of extraction w^* is decreasing in β , the short term discount factor.*

Proof. See the Appendix for a proof. □

In other words, a present biased regulator always extracts more: and the lower β is, the more short-sighted he is and the more he extracts. With Theorem 2 and Proposition 1, we have established the respective effect of introducing regime shifts as represented by the critical threshold \underline{X} respectively a present bias, described by the short term discount factor β in a standard renewable resource extraction problem. It turns out that both aspects, a more impatient regulator (lower β) and a more likely regime shift (higher \underline{X}), lead to an increase in the time consistent equilibrium level of extraction. An interesting question arises regarding the interaction between these two features of our model. We summarise our findings in the following proposition.

Proposition 2. *In the renewable resource extraction problem with regime shifts conditioned on post-extraction stock as described by the transition function in (12),*

- a. *the effect of present bias on the equilibrium extraction is larger the less prominent a regime shift is as represented by a lower \underline{X}*
- b. *the effect of regime shifts on equilibrium extraction is larger the less present biased a regulator is as represented by a higher β*

Proof. See the Appendix for a proof. □

In other words, when combining the two adverse features of our model, namely present bias and regime shifts, their effect on the optimal policy weakens: a lower β leads to a smaller increase on optimal extraction when regime shifts are more prominent. Similarly, regime shifts lead to a smaller increase on optimal extraction when faced with a very impatient regulator. On the other hand, present bias leads to a larger increase on extraction when regime shifts do not pose a big threat. In the same logic, regime shifts increase extraction more, the less present biased a regulator is.

6 Discussion

Although our way of incorporating uncertainty and regime shifts in a resource extraction model is fundamentally different from how the relevant literature has treated these issues, it bears mentioning the relation of our results with the main findings of the literature. In Polasky et al. (2011), they consider a renewable resource model with regime shifts modelled as changes in the system dynamics with or without stock collapse. The time of the regime shift is stochastic and could potentially depend on the resource stock. They find that only in the case of endogenous regimes shifts with changed system dynamics, does the optimal policy become precautionary in the sense that the optimal harvest decreases as a result of a potential regime shift. However, no precautionary action is taken for exogenous regime shifts while in the case of a stock collapse, it is in fact optimal to increase harvesting. Instead, Brozović and Schlenker (2011) model regime shifts with equations of motion that differ by an additive term once a possibly unknown, reversible threshold is crossed. They find a non-monotonic relationship between precautionary activity and uncertainty of two types: exogenous uncertainty embedded in the system dynamics over which the policy maker has no control and endogenous uncertainty about where the threshold is located. An increase in uncertainty of any type (an increase in the variance of the associated error term) may first increase precautionary behaviour in terms of reduced pollutant loading but it will always eventually decrease it. Moreover, they find that optimal loading changes non-linearly as the critical threshold changes. For extremely low or extremely high thresholds, the probability of moving in or out of any region is very low and practically unaffected by precaution activities. Instead, for intermediate levels of threshold, there is a possibility of moving between states in either direction and additional precautionary activity has a potential economic benefit: expected pollutant stock will drop as we get closer to the threshold.

Our results are quite different. In our framework, regime shifts are modelled as changes in the distribution function of the future stock level while there is always a positive probability of a resource collapse. To start with, we find a monotonic relationship between the optimal resource extraction and the stock level: extraction is always increasing in the stock while in the case of post-extraction stock regime shifts, we are able to establish an increasing relationship between reinvestment and the stock level. Quite surprisingly, this is an issue in which most of literature has remained silent or has been unable to establish a result (see e.g. the discussion in Knapp and Olson (1995) and Krishnamurthy (2012)). Regarding uncertainty and the critical threshold, we show that optimal extraction is in fact increasing in the threshold: the larger the undesirable region becomes, the more you would like to extract. This is in sharp contrast with the results in Brozović and Schlenker (2011) who find a non-monotonic relationship between optimal pollutant loading and the threshold level of stock. Although not directly translatable, our results indicate that in contrast to most of the literature, there is no precautionary behaviour in our model in the sense of reduced extraction to avoid a regime shift.

The reason for such a difference could come from two sources. We have introduced impatience in a model with regime shifts. It is intuitively reasonable that this could influence the propensity for precaution. Indeed if resource users care less about the future they will be more interested in harvesting now and not bother too much about the risk of a regime shift in a future that may be far. In fact, in the case of regime shifts that depend upon post extraction stock, the generalised Euler equation verifies that when $\beta < 1$, the equilibrium extraction level is higher than with $\beta = 1$. On the other hand, we have introduced specific assumptions regarding uncertainty and regime shifts, in particular when it comes to the relationship between reinvestment and the threshold: reinvestment is less beneficial, the higher is the threshold. This assumption is crucial in the determination of the effect of a change in the size of the desirable region on the equilibrium level of extraction. Again, the Euler equation indicates that the direction of change only depends on the sign of $\frac{\partial F_P(dx'|a(x), X)}{\partial a(x)\partial X}$, while β only plays a role on the magnitude of the effect. Moreover, the Generalized Euler equation gives us additional insights regarding the interaction between present bias and regime shifts. We find that the effect of these adverse features on optimal extraction gets weaker when combined. A more present biased regulator will always choose a higher extraction level but this increase will be smaller, the more prominent regime shifts are. Similarly, a regulator faced with a more threatening regime shift will increase extraction by less, the more present biased he is.

7 Conclusions

Although there is no consensus for the correct approach to discounting, experimental, empirical and conceptual contributions all suggest that there is little evidence for constant discounting over long horizons. However, non-constant (e.g. hyperbolic) discounting poses the challenge of time-consistency: a policy appearing optimal from one generation's perspective may not appear so from the standpoint of a subsequent generation. This is particularly relevant for issues related to climate change and the optimal management of renewable resources where costs are incurred by the current

generation while benefits are typically realised in the future. Policy making must consider this aspect since current policy makers have typically few devices to bind future policy makers' decisions.

In this paper, we investigate the decision process that present biased resource users face when they manage a renewable resource stock subject to growth shocks with the added threat of a reversible regime shift. In contrast to the previous environmental literature addressing stock collapse and regime shifts, we focus on the distribution rather than on a transition equation for the future stock while a regime shift constitutes a sharp change in that distribution when stock falls below a critical threshold. We are able to establish that a SMNE not only exists but it is unique and time consistent. Notably, no special functional form for utility is required to establish our key results, and assumptions made regarding the stock evolution are closely related to those made in the prior literature. Moreover, we find that the optimal policy is monotonically increasing in the critical stock threshold. This is in sharp contrast with the results in Brozović and Schlenker (2011) who find a non-monotonic relationship between optimal pollutant loading and the threshold level of stock. We also find, in contrast to much of the literature no precautionary behaviour in our model: the possibility of a regime shift does not lead to reduced extraction as a way to prevent a potential shift.

To our knowledge, our study is the first to establish existence and uniqueness of the SMNE, indeed an optimal policy, for the general regime shift case. Our set-up is general enough to be used to investigate our problems of regime-shifts and time inconsistencies but could be useful for general resource extraction problems, and more common time-consistent preferences as well (including recursive utility approaches). Many problems of pollution control, resource management, and drug control to name but a few could be modelled in similar ways (see e.g. Grass et al. (2008)).

References

- Amir, R. (1996). Continuous stochastic games of capital accumulation with convex transitions. Games and Economic Behavior, 15(2):111–131.
- Amir, R. (1997). A new look at optimal growth under uncertainty. Journal of Economic Dynamics and Control, 22(1):67–86.
- Amir, R. (2005). Supermodularity and complementarity in economics: An elementary survey. Southern Economic Journal, pages 636–660.
- Angeletos, G.-M., Laibson, D., Repetto, A., Tobacman, J., and Weinberg, S. (2001). The hyperbolic consumption model: Calibration, simulation, and empirical evaluation. Journal of Economic Perspectives, 15(3):47–68.
- Balbus, Ł., Reffett, K., and Woźny, Ł. (2014). A constructive study of markov equilibria in stochastic games with strategic complementarities. Journal of Economic Theory, 150:815–840.
- Balbus, Ł., Reffett, K., and Woźny, Ł. (2015). Time consistent markov policies in dynamic economies with quasi-hyperbolic consumers. International Journal of Game Theory, 44(1):83–112.

- Balbus, Ł., Reffett, K., and Woźny, Ł. (2018). On uniqueness of time-consistent markov policies for quasi-hyperbolic consumers under uncertainty. Journal of Economic Theory, 176:293–310.
- Benhabib, J., Bisin, A., and Schotter, A. (2010). Present-bias, quasi-hyperbolic discounting, and fixed costs. Games and Economic Behavior, 69(2):205–223.
- Biggs, R., Blenckner, T., Folke, C., Gordon, L., Norström, A., Nyström, M., and Peterson, G. (2012). Regime shifts. Sourcebook in theoretical ecology. University of California Press, Berkeley.
- Brozović, N. and Schlenker, W. (2011). Optimal management of an ecosystem with an unknown threshold. Ecological economics, 70(4):627–640.
- Clarke, H. R. and 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.
- Crépin, A.-S. (2003). Multiple species boreal forests—what faustmann missed. Environmental and Resource Economics, 26(4):625–646.
- Crépin, A.-S., Biggs, R., Polasky, S., Troell, M., and de Zeeuw, A. (2012). Regime shifts and management. Ecological Economics, 84:15–22.
- Cropper, M. L., Freeman, M. C., Groom, B., and Pizer, W. A. (2014). Declining discount rates. American Economic Review, 104(5):538–43.
- Dutta, P. K. and Sundaram, R. K. (1998). The equilibrium existence problem in general markovian games. Organizations with incomplete information, 5:159–207.
- Freeman, M. C. and Groom, B. (2015). Positively gamma discounting: combining the opinions of experts on the social discount rate. The Economic Journal, 125(585):1015–1024.
- Gollier, C. and Hammitt, J. K. (2014). The long-run discount rate controversy. Annu. Rev. Resour. Econ., 6(1):273–295.
- Gollier, C. and Zeckhauser, R. (2005). Aggregation of heterogeneous time preferences. Journal of political Economy, 113(4):878–896.
- Grass, D., Caulkins, J. P., Feichtinger, G., Tragler, G., and Behrens, D. A. (2008). Optimal control of nonlinear processes. Berlino: Springer.
- Harris, C. and Laibson, D. (2001). Dynamic choices of hyperbolic consumers. Econometrica, 69(4):935–957.
- Heal, G. (2000). Valuing the future: economic theory and sustainability. Columbia University Press.

- Horan, R. D., Fenichel, E. P., Drury, K. L., and Lodge, D. M. (2011). Managing ecological thresholds in coupled environmental–human systems. Proceedings of the National Academy of Sciences, 108(18):7333–7338.
- Jackson, M. O. and Yariv, L. (2015). Collective dynamic choice: the necessity of time inconsistency. American Economic Journal: Microeconomics, 7(4):150–78.
- Karp, L. (2005). Global warming and hyperbolic discounting. Journal of Public Economics, 89(2):261–282.
- Karp, L. (2007). Non-constant discounting in continuous time. Journal of Economic Theory, 132(1):557–568.
- Karp, L. and Tsur, Y. (2011). Time perspective and climate change policy. Journal of Environmental Economics and Management, 62(1):1–14.
- Kirby, K. N. and Herrnstein, R. J. (1995). Preference reversals due to myopic discounting of delayed reward. Psychological science, 6(2):83–89.
- Klein, P., Krusell, P., and Rios-Rull, J.-V. (2008). Time-consistent public policy. The Review of Economic Studies, 75(3):789–808.
- Knapp, K. and Olson, L. (1995). The economics of conjunctive groundwater management with stochastic surface supplies. Journal of Environmental Economics and Management, 28(3):340–356.
- Krishnamurthy, C. (2012). Optimal management of groundwater under uncertainty: A unified approach. Working Paper 2012:19 (Revised June 2015), Center for Environmental and Resource Economics, Umeå.
- Krishnamurthy, C. K. B. (2017). Optimal management of groundwater under uncertainty: a unified approach. Environmental and Resource Economics, 67(2):351–377.
- Krusell, P., Kuruşçu, B., and Smith, A. A. (2002). Equilibrium welfare and government policy with quasi-geometric discounting. Journal of Economic Theory, 105(1):42–72.
- Lemoine, D. and Traeger, C. (2014). Watch your step: optimal policy in a tipping climate. American Economic Journal: Economic Policy, 6(1):137–66.
- Lenton, T. M., Held, H., Kriegler, E., Hall, J. W., Lucht, W., Rahmstorf, S., and Schellnhuber, H. J. (2008). Tipping elements in the earth's climate system. Proceedings of the national Academy of Sciences, 105(6):1786–1793.
- Levy, Y. (2013). Discounted stochastic games with no stationary nash equilibrium: two examples. Econometrica, 81(5):1973–2007.

- Ludwig, D., Jones, D. D., and Holling, C. S. (1978). Qualitative analysis of insect outbreak systems: the spruce budworm and forest. The Journal of Animal Ecology, pages 315–332.
- Martin, I. W. and Pindyck, R. S. (2015). Averting catastrophes: the strange economics of scylla and charybdis. American Economic Review, 105(10):2947–85.
- Norström, A. V., Nyström, M., Lokrantz, J., and Folke, C. (2009). Alternative states on coral reefs: beyond coral–macroalgal phase shifts. Marine ecology progress series, 376:295–306.
- Nowak, A. (2003). On a new class of nonzero-sum discounted stochastic games having stationary nash equilibrium points. International Journal of Game Theory, 32(1):121–132.
- Nowak, A. S. (2006). On perfect equilibria in stochastic models of growth with intergenerational altruism. Economic Theory, 28(1):73–83.
- Olson, L. J. and Roy, S. (2000). Dynamic efficiency of conservation of renewable resources under uncertainty. Journal of Economic Theory, 95(2):186–214.
- Olson, L. J. and Roy, S. (2006). Theory of stochastic optimal economic growth. Handbook on Optimal Growth 1, pages 297–335.
- 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., and Brock, W. A. (2003). Uncertainty and the management of multistate ecosystems: an apparently rational route to collapse. Ecology, 84(6):1403–1411.
- Phelps, E. S. and Pollak, R. A. (1968). On second-best national saving and game-equilibrium growth. The Review of Economic Studies, 35(2):185–199.
- Polasky, S., De Zeeuw, A., and Wagener, F. (2011). Optimal management with potential regime shifts. Journal of Environmental Economics and management, 62(2):229–240.
- Rabin, M. (1998). Psychology and economics. Journal of economic literature, 36(1):11–46.
- Rubinstein, A. (2003). “economics and psychology”? the case of hyperbolic discounting. International Economic Review, 44(4):1207–1216.
- Saez-Marti, M. and Weibull, J. W. (2005). Discounting and altruism to future decision-makers. Journal of Economic theory, 122(2):254–266.
- Scheffer, M., Hirota, M., Holmgren, M., Van Nes, E. H., and Chapin, F. S. (2012). Thresholds for boreal biome transitions. Proceedings of the National Academy of Sciences, 109(52):21384–21389.

- Stachurski, J. (2007). Log-linearization of stochastic economic models. Journal of Difference Equations and Applications, 13(2-3):217–222.
- 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, page 201810141.
- Strotz, R. H. (1955). Myopia and inconsistency in dynamic utility maximization. The Review of Economic Studies, 23(3):165–180.
- Sumaila, U. R. and Walters, C. (2005). Intergenerational discounting: a new intuitive approach. Ecological Economics, 52(2):135–142.
- Topkis, D. M. (2011). Supermodularity and complementarity. Princeton university press.
- Tsur, Y. and 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. and Zemel, A. (1998). Pollution control in an uncertain environment. Journal of Economic Dynamics and Control, 22(6):967–975.
- Van Nes, E. H., Hirota, M., Holmgren, M., and Scheffer, M. (2014). Tipping points in tropical tree cover: linking theory to data. Global change biology, 20(3):1016–1021.
- Weitzman, M. L. (2010). Risk-adjusted gamma discounting. Journal of Environmental Economics and Management, 60(1):1–13.

Appendix A Technical Appendix

We introduce the following two lemmas, which will form the cornerstones of our analysis. The first lemma in essence provides a set of sufficient conditions to verify many claims which relate parametric properties of a distribution function, $F(\cdot|\theta_1, \theta_2)$, to expected value of a function, v , w.r.t to it (with θ_1 being a scalar parameter and θ_2 being a vector, two-dimensional in our case). All integrals defined hence forth will be assumed to satisfy the usual conditions (i.e. they exist, are well-defined and finite). Lemma 8, which is well known in the literature (Lemma 3.9.1, Topkis (2011)), provides verifiable conditions upon the distribution function $F(\cdot|\theta_1, \theta_2)$, allowing properties defined in lemma 7 to be easily verified. Applicability to our case can be directly seen by identifying $F(\cdot|\theta_1, \theta_2)$ with $F_P(\cdot|a, x, \underline{X})$ and the expectation of interest with $\mathbb{E}(V) := \int_{\mathcal{X}} V(x') dF_P(dx'|a, x, \underline{X})$.

Lemma 7. For $v(s)$ increasing, if $F(s|\theta_1, \theta_2)$ is

1. stochastically increasing (decreasing) in θ_1 ;
2. stochastically concave in θ_1 ;
3. stochastically supermodular (submodular) in θ_2 ;

then it is the case that $I(\theta_1, \theta_2) := \int_{\mathcal{X}} v(s) dF(s|\theta_1, \theta_2)$ is

1. increasing (decreasing) in θ_1 (for fixed θ_2);
2. concave in θ_1 (for fixed θ_2);
3. satisfies increasing (decreasing) differences in θ_2 (for fixed θ_1);

Proof. see lemma 3.9.1, Topkis (2011). □

Denote by $\{F_\theta(s); \theta \in \Theta\}$, with Θ the parameter space for a vector of parameters θ , a parameterised collection of distribution functions. For concreteness, consider $\theta = (a, x)$.

Lemma 8. The collection of distribution functions, $\{F_\theta(s); \theta \in \Theta\}$ is stochastically increasing (decreasing), concave (convex) or supermodular (submodular) in θ iff $1 - F_\theta$ (viewed purely as a function of θ) is increasing (decreasing), concave (convex) or supermodular (submodular) in θ .

Remark 8. Note that lemma 7 is most useful in applications where distribution function properties are known, but the function v cannot be explicitly written down. In particular, the only property required of the unknown function v is its non-decreasing nature.

Appendix B Proofs

Theorem 2. Under assumptions 1 and 2, the optimal time-consistent policy w^* is increasing in x .

Proof. Let $w^* = BV^*$ and $V^* = TV^*$. Consider the function

$$G(q, x, V^*) = u(q) + \beta\delta \int_{\mathcal{X}} V^*(x') dF_P(dx'|x - q, x). \quad (\text{B.1})$$

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. It is easy to see that submodularity of $(x, a) \mapsto \int_{\mathcal{X}} v(x') dF_P(dx'|a, x)$ in (a, x) implies supermodularity in (q, x) (since $a = x - q$), which suffices for $G(q, x, V^*)$ to be supermodular in (q, x) . Then, by a standard result in parametric optimisation (Theorem 2.8.1, Topkis (2011)), $w^*(x, V^*) := \arg \max_{q \in A(x)} G(q, x, V^*)$ is increasing in x on \mathcal{X} . \square

We first define the relevant operators used. $T_{\underline{X}}$ is defined on the space of bounded functions V as:

$$T_{\underline{X}}V(x) = \frac{1}{\beta}A_{\underline{X}}V(x) - \frac{1-\beta}{\beta}u(B_{\underline{X}}V(x))$$

where the pair of operators $A_{\underline{X}}$ and $B_{\underline{X}}$ defined on the same space are given by :

$$A_{\underline{X}}V(x) = \max_{q \in A(x)} \left[u(q) + \beta\delta \int_{\mathcal{X}} V(x') Q(dx'|x - q, x, \underline{X}) \right],$$

$$B_{\underline{X}}V(x) = \arg \max_{q \in A(x)} \left[u(q) + \beta\delta \int_{\mathcal{X}} V(x') Q(dx'|x - q, x, \underline{X}) \right],$$

Lemma 3. Let $\phi : \mathcal{X} \times \mathcal{X} \mapsto \mathbb{R}^+$ be a bounded function for each $\underline{X} \in \mathcal{X}$ and $\phi(x, \cdot)$ is decreasing for each $x \in \mathcal{X}$. Then $\underline{X} \mapsto T_{\underline{X}}(\phi(\cdot, \underline{X}))(x)$ is a decreasing function of \underline{X} .

Proof. It is easy to see that for every bounded function V , Assumption 3 guarantees that a mapping $A_{\underline{X}}(V)$ is a decreasing function. Next we show that $B_{\underline{X}}V(x)$ is increasing in \underline{X} . For each $x \in \mathcal{X}$, let us define

$$G(q, V, \underline{X}) := u(q) + \beta\delta \int_{\mathcal{X}} V(x') P(dx'|x - q, x, \underline{X}) \quad (\text{B.2})$$

Then, for $q_1 < q_2 \leq x$, we have

$$\begin{aligned} G(q_2, V, \underline{X}) - G(q_1, V, \underline{X}) &:= u(q_2) - u(q_1) \\ &+ \beta\delta \int_{\mathcal{X}} V(x') P(dx'|x - q_2, x, \underline{X}) - \beta\delta \int_{\mathcal{X}} V(x') P(dx'|x - q_1, x, \underline{X}) \end{aligned}$$

From Assumption 3, we know that the function $(a, \underline{X}) \mapsto \int_{\mathcal{X}} V(x') P(dx' | a, x, \underline{X})$ has decreasing differences with (a, \underline{X}) so that the RHS of the above expression is increasing in \underline{X} (note that $a = x - q$) which implies that $G(q, V, \underline{X})$ has increasing differences in (q, \underline{X}) . (from the definition of increasing differences). Then by Theorem 2.8.1 in Topkis (2011), the $B_{\underline{X}}V(x) = \arg \max_{q \in A(x)} [G(q, V, \underline{X})]$ is increasing in \underline{X} .²² By Lemma 2 in Balbus(2016) $V \mapsto B_{\underline{X}}(V)$ is decreasing in V . As a result, $B_{\underline{X}}(\phi(\cdot, \underline{X}))$ is increasing in \underline{X} and hence $\underline{X} \mapsto T_{\underline{X}}(\phi(\cdot, \underline{X}))$ decreasing in \underline{X} for any $x \in \mathcal{X}$. \square

Theorem 3. Let Assumption 3 be satisfied. Then the mapping $\underline{X} \mapsto w_{\underline{X}}^*$ is increasing on $\underline{\mathcal{X}}$.

Proof. Observe that by Theorem 1, $V_{\underline{x}}^*(x) = \lim_{n \rightarrow \infty} T_{\underline{X}}^n(\mathbf{0})(x)$ (where $\mathbf{0}$ is a zero function). Applying Lemma 3, we see $T_{\underline{X}}(\mathbf{0})(x) \in \mathcal{V}$ and this expression is decreasing in \underline{X} . Applying Lemma 3 again, we can show that all $T_{\underline{X}}^n(\mathbf{0})(x)$ are decreasing in \underline{X} , hence $V_{\underline{x}}^*(x)$ is decreasing in \underline{X} . Finally, observe that $w_{\underline{x}}^*(x) = B_{\underline{X}}(V_{\underline{x}}^*)(x)$ and using similar arguments as in the proof of Lemma 3, we find that $w_{\underline{x}}^*(x)$ is decreasing in \underline{X} . \square

Theorem 9. Under Assumption 1, the optimal time-consistent policy w^* is increasing in x and Lipschitz with modulus 1.

Proof. See the Appendix for a proof. Let $w^* = BV^*$ and $V^* = TV^*$. Consider the function

$$G(q, x, V^*) = u(q) + \beta\delta \int_{\mathcal{X}} V^*(x') dF_P(dx' | x - q). \quad (\text{B.3})$$

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. Moreover, by concavity of $a \mapsto \int_{\mathcal{X}} v(x') dF_P(dx' | a)$ in a , we conclude that $G(q, x, V^*)$ is supermodular in (q, x) . Then, by a standard result in parametric optimisation (Topkis(1978)–citation!), $w^*(x, V^*) := \arg \max_{q \in A(x)} G(q, x, V^*)$ is increasing in x on \mathcal{X} .

Next, we rewrite the problem as function of investment a

$$H(a, x, V^*) = u(x - a) + \beta\delta \int_{\mathcal{X}} V^*(x') dF_P(dx' | a). \quad (\text{B.4})$$

where H is now supermodular in the choice variable a on a lattice $[0, x]$, the feasible action set, which is also increasing in Veinott's strong set order. Then, by concavity of u , it is evident that $H(a, x, V^*)$ is supermodular in (a, x) , therefore by Topkis(1978) theorem, the optimal solution a^* is increasing in x on \mathcal{X} .

Observe that since $a^* = x - w^*$ and both a^* and w^* are increasing in x , it follows that a^* and w^* are Lipschitz with modulus 1. \square

Euler Equation

²²SDF

Assume that $a^*(x)$ is a differentiable SMNE reinvestment such that $a^* = x - w^*$. We can then write the generalised Euler equation characterizing the SMNE investment $a^*(x)$:

$$u'(x - a^*(x)) = \beta\delta \frac{d}{da(x)} \int_{\mathcal{X}} V^*(x') dF_P(dx'|a^*(x), \underline{X}). \quad (\text{B.5})$$

$$V^*(x) = u(x - a^*(x)) + \delta \int_{\mathcal{X}} V^*(x') dF_P(dx'|a^*(x), \underline{X}).$$

so that

$$V'^*(x) = u'(x - a^*(x))(1 - a'^*(x)) + \delta a'^*(x) \frac{d}{da(x)} \int_{\mathcal{X}} V^*(x') dF_P(dx'|a^*(x), \underline{X}). \quad (\text{B.6})$$

Next we multiply (B.6) with $\frac{\partial F_P(dx|a^*(s), \underline{X})}{\partial a(s)}$ and integrate wrt to x .

$$\begin{aligned} \int_{\mathcal{X}} V'^*(x) \frac{\partial F_P(dx|a^*(s), \underline{X})}{\partial a(s)} dx &= \int_{\mathcal{X}} u'(x - a^*(x))(1 - a'^*(x)) \frac{\partial F_P(dx|a^*(s), \underline{X})}{\partial a(s)} dx \\ &+ \delta \int_{\mathcal{X}} a'^*(x) \left(\frac{d}{da(x)} \int_{\mathcal{X}} V^*(x') dF_P(dx'|a^*(x), \underline{X}) \right) \frac{\partial F_P(dx|a^*(s), \underline{X})}{\partial a(s)} dx \end{aligned} \quad (\text{B.7})$$

Using integration by parts and Lemma 4.2 from Amir (1997) (See also proof of theorem 3.2), we can show that

$$\frac{d}{da(x)} \int_{\mathcal{X}} V^*(x') dF_P(dx'|a^*(x), \underline{X}) = - \int_{\mathcal{X}} V'^*(x') \frac{\partial F_P(dx'|a^*(x), \underline{X})}{\partial a(x)} \quad (\text{B.8})$$

Then using (B.5) and (B.8), we can write (B.7) as

$$u'(s - a^*(s)) = - \int_{\mathcal{X}} u'(x - a^*(x))(1 - a'^*(x)) \frac{\partial F_P(dx|a^*(s), \underline{X})}{\partial a(s)} dx - \delta \int_{\mathcal{X}} u'(x - a^*(x)) a'^*(x) \frac{\partial F_P(dx|a^*(s), \underline{X})}{\partial a(s)} dx \quad (\text{B.9})$$

Simplifying the above expression we get

$$u'(x - a^*(x)) = -\beta\delta \int_{\mathcal{X}} u'(x' - a^*(x')) \left(1 + \left(\frac{1}{\beta} - 1\right) a'^*(x')\right) \frac{\partial F_P(dx'|a(x), \underline{X})}{\partial a(x)} \quad (\text{B.10})$$

Proposition 3. *In the renewable resource extraction problem with regime shifts conditioned on post-extraction stock as described by the transition function in (12), the time consistent equilibrium level of extraction w^* is decreasing in β , the short term discount factor.*

Proof. We differentiate the right hand side of (B.10) with respect to β to get

$$-\delta \int_{\mathcal{X}} u'(x' - a^*(x')) \left(1 + \left(\frac{1}{\beta} - 1\right) a'^*(x')\right) \frac{\partial F_P(dx'|a(x), \underline{X})}{\partial a(x)} + \frac{\delta}{\beta} \int_{\mathcal{X}} u'(x' - a^*(x')) a'^*(x') \frac{\partial F_P(dx'|a(x), \underline{X})}{\partial a(x)} =$$

(B.11)

$$\delta \int_{\mathcal{X}} u'(x' - a^*(x')) (a'^*(x') - 1) \frac{\partial F_P(dx'|a(x), \underline{X})}{\partial a(x)} > 0$$

by Assumption 1 and Theorem 5. Then, concavity of the utility function implies that a decrease in β from $\beta = 1$ to $\beta < 1$ would result in an increase in the equilibrium level of extraction. \square

Proposition 4. *In the renewable resource extraction problem with regime shifts conditioned on post-extraction stock as described by the transition function in (12),*

- a. the effect of present bias on the equilibrium extraction is larger the less prominent a regime shift is as represented by a lower \underline{X}*
- b. the effect of regime shifts on equilibrium extraction is larger the less present biased a regulator is as represented by a higher β*

Proof. See the Appendix for a proof. \square

From Assumption 3 follows that the differentiation of (B.11) with respect to \underline{X} gives us

$$\delta \int_{\mathcal{X}} u'(x' - a^*(x')) (a'^*(x') - 1) \frac{\partial F_P(dx'|a(x), \underline{X})}{\partial a(x) \partial \underline{X}} < 0 \quad (\text{B.12})$$