Scenes from a Monopoly: Renewable Resources and Quickest Detection of Regime Shifts

Neha Deopa*and Daniele Rinaldo†

The Graduate Institute of International and Development Studies, Geneva, Switzerland.

Abstract: We study the stochastic dynamics of a renewable resource harvested by a monopolist facing a downward sloping demand curve. We introduce a framework where harvesting sequentially affects the resource's potential to regenerate, resulting in an endogenous ecological regime shift. In a multi-period setting, the firm's objective is to find the profit-maximizing harvesting policy while simultaneously detecting in the quickest time possible the change in regime. Solving analytically, we show that a negative regime shift induces an aggressive extraction behaviour due to shorter detection periods, creating a sense of urgency, and higher markup in prices. Precautionary behaviour can result due to decreasing resource rent. We study the probability of extinction and show the emergence of catastrophe risk which can be both reversible and irreversible.

1 Introduction

The exploitation of renewable resources such as overfishing of the North Sea cod, deforestation in the Amazon and soil degradation due to unsustainable agricultural practices, is receiving considerable attention. The dynamic management of these resources often involve decisions about optimal extraction policies under *ecological* uncertainty, defined by Pindyck (2002) as uncertainty over the evolution of the relevant ecosystem. This raises pertinent economic questions about the behaviour and actions of a firm harvesting these resources, especially if the dynamics driving the resource growth change. One way the current literature captures this uncertainty is via stochastic bio-economic models, reflected in the variance of the fluctuations. Another is to focus on ecological regime shifts, defined as an abrupt change in the structure of the natural ecosystems supplying the resource or a change in the system dynamics such as intrinsic growth rate or the carrying capacity of the resource (Polasky et al. (2011); Arvaniti et al. (2019)).

There already exists a large literature studying the impact of stochastic fluctuations on extraction activities, mainly utilizing real options theory (Andersen and Sutinen (1984); Pindyck (1984); Reed (1988); Reed and Clarke (1990); Saphores (2003); Alvarez and Koskela (2007); Pizarro and Schwartz (2018)). An emerging literature builds on this to integrate resource management with a variety of regime shifts, such as Polasky et al. (2011), Ren and Polasky (2014), Baggio and Fackler (2016), de Zeeuw and He (2017) and Arvaniti et al. (2019)[1]. These studies, however, are limited in two respects. With the exception of Pindyck (1984), none of these works incorporate a market structure and take the price as fixed or exogenous. This is done for tractability but leads to results that may underestimate the crucial role of a market structure, which in fact often drives the firm's harvesting decisions. Second, the literature on regime shifts implicitly assumes the firm to be able to discern the change in resource dynamics and subsequently make the appropriate extraction decision[2]. We take a different approach from earlier

*neha.deopa@graduateinstitute.ch

†daniele.rinaldo@graduateinstitute.ch

[1]Refer to Li et al. (2018) for an overview.

[2]We use the words extraction and harvesting interchangeably

studies and ask: within a resource market where prices are endogenously determined, how does an ecological regime shift influence the firm's harvesting decisions? what is the profit maximizing policy of this firm who wants to *detect* this shift and where this change in regime is endogenously determined by the firm's own extraction activities?

We build a model of a monopolist firm, facing a downward sloping demand curve, who encounters two sources of uncertainty in the resources dynamics. The first takes the form of natural randomness of the environmental conditions (variance) and the other in the *timing* of the ecological regime shift. This shift, defined as a change in the resource's ability to grow, is made dependent on the firm's own extraction efforts. In our model, the firm knows with certainty that a regime shift will occur and thus the probability of the shift occurring is not under consideration. However it is the timing which matters to the firm for it's harvesting decisions. In a multi-period setting, the monopolist wants to detect this shift as soon as possible, and this detection procedure is explicitly incorporated in its profit maximizing actions. The resource dynamics are assumed to be monitored by the monopolist through sequential observations and we model the firm's detection process based on *quickest detection* methods. These build on solving the classical sequential detection problem as an optimal stopping problem with the aim to detect a change in the resource growth, if one occurs, as quickly as possible. Much of the resource literature focuses on a stationary population processes: non-stationariety generated by changes in environmental conditions, however, is more common in the physical world (Szuwalski and Hollowed (2016)). Using the sequential nature of the detection process we incorporate non-stationary dynamics in our model. This, combined with the multi-period framework, leads to a result where a change in regime alters the monopolist's time horizon for that period.

Our model has analytical solutions and shows that in the event of a negative regime shift, for low stock levels, the firm adopts a precautionary policy by extracting less. This is because the change in regime creates a physical scarcity of the resource which in turn increases the resource rent for the monopolist and this leads to reduced extraction levels. However for higher stock levels this effect is outweighed due to an interplay of two factors: (i) an altered and relatively shorter time horizon due to detecting a regime shift, which creates a sense of urgency and (ii) market preferences, where an elastic demand allows the firm to increase its markup compared to pre-regime shift period. Together this reduces the resource rent and results in the monopolist adopting an aggressive behaviour by increasing extraction. Lastly we define the risk of catastrophe where the resource maybe driven to extinction and differentiate between and irreversible and reversible catastrophe.

In the next section we lay out the different building blocks of the model and describe how it all comes together and section 3 discusses the solutions and their implications. In section 4 we define the risk and first passage time to catastrophe and section 5 provides a short note on how our model would translate to real time. Section 6 concludes.

2 The Model

We start with a model that is relatively simple to analyze, and allows for closed form solutions.

2.1 Resource Dynamics

Following the basic framework of Pindyck (1984) we start by modeling the evolution of the renewable resource stock X_t. Let X_t be the stock at time t, which behaves according to the stochastic differential equation

$$dX_t = \big(\mu - q_t\big)dt + \sigma dW_t \tag{1}$$

where $q_t \in \mathbb{R}^+$ is the resource extraction chosen by the firm, $\sigma \in \mathbb{R}^+$ is the intensity of noise in the evolution of the resource stock, $\mu \in \mathbb{R}^+$ is the constant growth rate of the resource and

$X_t \geq 0$[3]. Finally, W_t is the standard Brownian motion in the filtered probability space $(\Omega, \mathcal{F}, P)$.

In order to capture the regime shift that the dynamic system can undergo, we describe two alternative scenarios faced by the firm: one in which the resource evolves according to equation (1), and an alternative one in which the stock's ability to regenerate - the drift - changes. This is consistent with Polasky et al. (2011) in which a regime shift is defined as a change in the system dynamics such as intrinsic growth rate or the carrying capacity of the resource. The evolution for the resource stock then becomes

$$dX_t = \big(\mu + \lambda - q_t\big)dt + \sigma dW_t, \tag{2}$$

where $\lambda \in \mathbb{R}$ is the change in resource growth. If $\lambda < 0$, the growth rate of the resource is reduced, and vice versa. To provide some intuition, we give two examples where we would observe such a change in drift. In the context of fisheries the term $\big(\mu + \lambda\big)$, where $\lambda < 0$, would indicate recruitment overfishing which occurs when the parent stock (spawning biomass) is depleted to a level where it no longer has the reproductive capacity to replenish itself i.e. there are not enough adults to produce recruits (Pauly (1983)). The collapse of the Atlantic northern cod in Newfoundland in the early nineties was attributed to gross overestimation of stock sizes and the failure to recognize that recruitment overfishing was a definite possibility (Walters and Maguire (1996)). Similarly logging and timber production have a direct impact on forest recovery and tree recruitment and growth. There is evidence that the current practice of felling cycles of around 30 years may be insufficient to enable tropical forest recovery (de Avila et al. (2017)). An example of a positive regime shift with $\big(\mu + \lambda\big)$, where $\lambda > 0$, can be seen in China's forest dynamics - which has shown an increase in the forest cover by more than 46,000 square miles between 2000 and 2010 (Viña et al. (2016))[4].

Equation (2) implies that the firm's harvesting activities do not affect the resource's ability to regenerate in any way. However, this assumption does not seem grounded in empirical observation and it can be seen from the examples above that often the firm's harvesting decisions influence the resource's recruitment and growth process. Thus we rewrite (2) as:

$$dX_t = \big(\mu + \lambda(q_{ex}) - q_t\big)dt + \sigma dW_t, \tag{3}$$

where ex is the past time period that determines the magnitude of λ. We therefore study a framework in which *past extraction decisions* determine the future changes in resource growth. We want to model the scenario in which at a given change point in time θ, which is assumed deterministic but unknown, the stochastic differential equation (SDE) driving the resource stock will switch between drifts, and the growth rate of the resource will change:

$$dX_t = \begin{cases} \big(\mu - q_t\big)dt + \sigma dW_t & t < \theta \\ \big(\mu + \lambda(q_{ex}) - q_t\big)dt + \sigma dW_t & t \geq \theta. \end{cases} \tag{4}$$

The sign of $\lambda(q_{ex})$ can be both positive or negative, which represents the fact that the effect of firm extraction on the resource growth can be both positive or negative. This also implies that the firm's actions influences the magnitude of the change of regime. Note that as θ is deterministic, the question facing the monopolist is not « if » a regime shift will occur but rather when. This framework seems appropriate for today, where the focus has moved on from questions of probability of collapses and regime shifts, to that of when and how will one deal with it. At first glance it may appear odd that the occurrence of the shift itself is exogenous but the magnitude and direction is endogenous to firm's actions. A more intuitive way to understand this is even if the monopolist, say a fishery, ceased it's harvesting actions, other climactic factors such as rising sea surface temperature would still eventually alter the marine ecosystem, resulting in a

[3]This positivity constraint allows the problem to have reasonable implications and a relatively simple solution, at the expense of an increase of the hidden mathematical requirements for the solution to be sufficient and unique.

[4]Although China's forests have experienced an increase in growth, it has been at the expense of exerting negative impact on other forested areas around the world. During the period China introduced sustainable logging of natural forests, it also became one of the world's leading timber importers (Liu (2014)).

regime shift. However the added fishing pressure by the monopolist influences the magnitude of this shift.

The firm now faces two sources of uncertainty when choosing the harvesting policy to maximize its profits. The first is σ, which is the environmental variation due to the natural randomness of environmental conditions - our choice of the diffusion coefficient σ being independent of the state X_t (i.e. a drifted Brownian motion) is to include the possibility that the natural randomness of the environment may drive the resource to extinction, something that lognormal fluctuations from a geometric Brownian motion by construction cannot represent. The second is the *timing* of $\lambda(q_{ex})$, indicating which regime is the firm currently operating in.

2.2 Firm Dynamics

We consider a risk-neutral monopolist facing a linear inverse demand function of the form $p(q) = a - bq$, with quadratic harvesting costs $cq(t, X_t)^2/2$, and with fixed costs F where $c, F \geq 0$. The harvesting rate is chosen by the firm in order to maximize the expected value of the sum of discounted profits subject to the constraint (4), and the profit function takes the form

$$\Pi(q) = \left[(a - bq)q - c\frac{q^2}{2} - F\right] \tag{5}$$

We assume a profit function not directly depending on the stock level X but only on the harvested quantity: this implies a marginal cost function linear in harvesting, rather than the stock level, and fixed operating costs. This assumption can be relaxed, at the expense of a complicated form of the optimal extraction function, is omitted.

2.3 Optimal Detection

The firm's problem now involves the *detection* of the change in drift of X_t, as seen in (4). Environmental monitoring of the resource to detect for changes in the stock and its structure is quite common in real-world renewable resource management. Klemas (2013) talks about how remote sensing techniques, in near-real time, help detect changes that affect the recruitment, distribution patterns and survival of fish stocks. These techniques, combined with *in situ* measurements, constitute the most effective ways for efficient management and controlled exploitation of marine resources. Shimabukuro et al. (2019) details the monitoring of deforestation and forest degradation in the Brazilian Amazon. Another example is The Waterfowl Breeding Population and Habitat Survey in North America[5], which is a continental monitoring program providing information on spring population size and trajectory for certain duck species. This in turn is used for the annual establishment of mallard hunting regulations in the United States and Canada (Nichols and Williams (2006)).

The monopolist monitors the resource stock via sequential observations and uses quickest detection (QD) methods to detect the regime change. These methods employ a likelihood ratio for two models, one for the status quo ('no change in regime') and one for a new state ('regime change detected'). For each observation, the model likelihoods and their ratio are updated. When the updated likelihood ratio exceeds a detection threshold, an alarm is triggered. The aim is to minimize the time to detection of a change point, given the monopolist's tolerance for false alarms. When the statistical threshold is crossed and the alarm is triggered, the firm will act on it to change its extraction levels to that which will maximize its profit for the new regime[6]. This is in line with an emerging literature in ecology, addressing indicators of approaching thresholds, or impending collapse in ecosystems with a focus on QD methods (Carpenter

[5]It is conducted cooperatively by the U.S. Fish and Wildlife Service and Canadian Wildlife Service. It has been conducted every year since 1955

[6]For a short introduction to quickest detection methods refer to Polunchenko et al. (2013). For a more detailed review refer to Poor and Hadjiliadis (2008)

et al. (2014); Batt et al. (2013); Scheffer et al. (2015).

The firm searches for a "rule" (an optimal stopping time) τ adapted to the filtration $\mathcal{F}_t$, at which it detects the change point, θ, so it may reassess its harvesting decisions given the change of environment in which it operates. This gives rise to the problem of quickest change-point detection. In the period before the change point θ, the dynamics of the resource X_t are determined by the (possibly nonlinear) SDE

$$dX_t = \big(\mu - q_t\big)dt + \sigma dW_t.$$

Girsanov theory tells us that the process

$$M_t = \exp\left(-\int_0^t \frac{\mu - q_s}{\sigma} dW_s - \frac{1}{2}\int_0^t \frac{(\mu - q_s)^2}{\sigma^2} ds\right)$$

is a P-martingale. Therefore, the process

$$\tilde{W}_t = W_t + \int_0^t \frac{\mu - q_s}{\sigma}$$

is a Q-Brownian motion, where one obtains the new probability measure by $Q = \mathbb{E}_P(M_t)$. The process X_t therefore admits the representation

$$X_t = x_0 + \int_0^t d\tilde{W}_s$$

and is therefore a Brownian motion under the measure Q. The firm's detection problem now becomes

$$dX_t = \begin{cases} d\tilde{W}_t & t < \theta \\ \lambda(q_{ex}) + d\tilde{W}_t & t \geq \theta. \end{cases} \tag{6}$$

If the period ex that determines λ is outside the interval $[0, t]$, then the firm's detection problem reverts exactly to the *Brownian disorder* problem, which is the detection of the change between a martingale and a sub/supermartingale, depending on the sign of λ. This requires that the harvesting decisions, that define both sign and magnitude of the change in resource growth, be set strictly before the time of the initial condition on X (here normalized to 0, i.e. X_0). Change-point detection in the disorder problem involves the optimization of the trade off between two measures, one being the delay between the time a change occurs and it is detected i.e. $(\tau - \theta)^+$, and the other being a measure of the frequency of false alarms for events of the type $(\tau < \theta)$. This problem has been first studied by Shiryaev (1963), and the procedure of the cumulative sum process (CUSUM) has been proven to be optimal by Shiryaev (1996) and in the case of multiple drifts by Hadjiliadis and Moustakides (2006).

The firm minimizes the worst possible detection delay over all possible realizations of paths of X_t before the change and over all possible change points θ. This is given by

$$J(\tau) = \sup_{\theta} \ \operatorname{ess\,sup} \mathbb{E}_\theta[(\tau - \theta)^+ | \mathcal{F}_\theta] \tag{7}$$

and the stopping rule is obtained by minimizing (7) under a "false alarm" constraint. This stochastic control problem is given by

$$\min_{\tau} J(\tau) \quad \text{s.t.} \quad \mathbb{E}_{\theta=\infty}[\tau] = T.$$

This constraint gives the class of stopping times τ, for which the mean time $\mathbb{E}_{\theta=\infty}[\tau]$ until giving a (false) alarm is equal to T. It can be interpreted as a measure of the "quality" of the detection system, since it fixes the expected delay in the detection under a false alarm, i.e. when $\theta = \infty$

(the process never actually changes regime).

It is shown by Hadjiliadis and Moustakides (2006) that one can only focus on the constraints that bind with equality. The CUSUM procedure involves first observing the process given by the logarithm of the likelihood ratio (the Radon-Nikodym derivative) of the process X_t (note that we are under the measure Q) under the two regimes and comparing it with its minimum observed value. Define

$$u_t(\beta) = \log \frac{dQ_{\theta=0}}{dQ_{\theta=\infty}} = \lambda(q_{ex})X_t - \frac{\lambda^2}{2}t.$$

The CUSUM statistic process is then given by the difference at any instant $s \leq t$ between u_t and its minimum obtained value up to that instant, namely

$$CS_t(\lambda(q_{ex})) = u_t(\lambda(q_{ex})) - \inf_{0\leq s\leq t} u_t(\lambda(q_{ex})) \geq 0.$$

This can be interpreted simply by noticing that if the two regimes are very similar (i.e. $|\lambda|$ is very small), then the Radon-Nikodym derivative will be close to unity, implying that the CUSUM process will be most of the time close to zero, and it will be difficult to detect the presence of such a small drift. If on the other hand the two regimes are rather different, then one should be able to detect more easily when the regime changes, and the CUSUM process should reflect this change as it increases. One would therefore expect to search for a threshold in order to determine when the CUSUM process is "large enough" to reflect the change of regime: this is indeed the case. Shiryaev (1996) and Hadjiliadis and Moustakides (2006) show that the optimal CUSUM stopping rule is given by the stopping time

$$\tau(\lambda(q_{ex}), \nu) = \inf\{t \geq 0; CS_t \geq \nu), \tag{8}$$

where the threshold ν is given by the root of the equation

$$\frac{2}{\lambda(q_{ex})^2}(e^{\nu} - \nu - 1) = T.$$

It can be shown that the delay function of this procedure is given by

$$\mathbb{E}[\tau(\lambda(q_{ex}), \nu)] = \frac{2}{\lambda(q_{ex})^2}(e^{-\nu} + \nu - 1). \tag{9}$$

At the stopping time τ, therefore, the firm will detect the change in drift of λ in (6), which means that the firm will have detected a change from a Q-martingale to a Q-sub/supermartingale. Note immediately that the larger the change in drift λ, the smaller the threshold ν and the "earlier" one expects the CUSUM process to hit the threshold. If λ is very small, then ν will be very large and the firm may wait for much longer before detecting a change of regime: in such a case it may be that $\tau(\lambda(q_{ex}), \nu) \geq T$, and once T is reached the firm will assume that the regime has changed. This result implies that the firm up to time τ will program its profit maximization assuming that the non-controlled part of the drift in the SDE driving X_t is given by μ, and subsequently by $(\mu + \lambda(q_{ex}))$.

2.4 Profit Maximization

If the firm's harvesting activities do not impact the growth of the resource, as seen in (2), then its actions do not influence the time and the magnitude of the change of regime. Therefore the change is entirely exogenous from the firm's point of view and the stochastic control problem is:

$$\sup_{q\in Q} \quad \mathbb{E}_0 \int_0^{\tau} \Pi(q_s)e^{-\rho s}ds + \mathbb{E}_{\tau} \int_{\tau}^{\infty} \Pi(q_s)e^{-\rho s}ds \tag{10}$$

$$dX_t = \begin{cases} \big(\mu - q_t\big)dt + \sigma dW_t & t < \tau \\ \big(\mu + \lambda - q_t\big)dt + \sigma dW_t & t \geq \tau \end{cases} \tag{11}$$

The simplest way of modeling a regime shift is to assume that the shift occurs only once, as in Polasky et al. (2011) and Ren and Polasky (2014). However, as pointed by Sakamoto (2014), regime shifts are better modeled as open-ended processes. An example being the Pacific ecosystem, where in the mid-1970s, the Pacific changed from a cool "anchovy regime" to a warm "sardine regime" and a shift back to an anchovy regime occurred in the middle to late 1990s (Chavez et al. (2003)). The above model can be straightforwardly extended to a multi-period setting, in which the firm detects multiple regime changes throughout subsequent periods and adjusts its optimal harvesting policy accordingly. The problem will then read:

$$\sup_{q\in Q} \quad \sum_{i=0}^{\infty} \mathbb{E}_{\tau_i} \int_{\tau_i}^{\tau_{i+1}} \Pi(t, q_t) e^{-\rho t} dt \tag{12}$$

$$dX_t = \begin{cases} \big(\mu + \lambda_i - q_t\big)dt + \sigma dW_t & t \in [i, i+1) \\ \big(\mu + \lambda_{i+1} - q_t\big)dt + \sigma dW_t & t \geq i+1. \end{cases} \tag{13}$$

where $i \in \mathbb{N}$ are the different periods, and the harvesting policy exists among the class of admissible controls Q. Here $\lambda_0 = 0$ and τ_i, λ_i are the subsequent detection times and relative changes in resource growth. We assume $\tau_0 = 0$ for simplicity. However to incorporate the firm's actions in a way that it affects the resource's ability to grow and regenerate, as in (3), introduces tractability issues with the model. The explicit dependence of the stopping time τ on $\lambda(q_{ex})$ makes the control variable q and the limit of integration τ simultaneous, and the model becomes intractable. In order to circumvent this issue, we formalize the structure of the firm's harvesting decisions in a sequential manner, where the firm assumes a constant $\lambda(q_{ex})$ for each period. We explain this in detail in the next section.

Up until the first detection time τ_1 the firm assumes that X_t is still driven by W_t, therefore the value of the firm is given by

$$\begin{aligned} V(0, X_0) \quad &= \quad \sup_{q\in Q} \ \mathbb{E}_0 \int_0^{\tau_1} \Pi(q) e^{-\rho t} dt \\ &\text{s.t.} \quad dX_t = (\mu - q)dt + \sigma dW_t. \\ &\qquad\ \ X_t \geq 0, \end{aligned} \tag{14}$$

Before solving the problem, let us first characterize the solution given the positivity constraint. The Hamilton-Jacobi-Bellman (HJB) equation for the firm's optimization problem reads

$$0 = V_t - \rho V + \max_{q\in Q} \Big\{ (a - bq)q - \frac{c}{2}q^2 - F - qV_x \Big\} + \mu V_x + \frac{\sigma^2}{2} V_{xx}, \quad X_t \geq 0. \tag{15}$$

where Q is the set of admissible Markov controls [7]. Once solved, this problem will yield a control in the feedback form $q(X_t)$. Because of the constraint $X_t \geq 0 \ \forall t \in [0, \tau_1]$, the value function $V(t, x)$ is not necessarily always differentiable. Using viscosity solutions, as first shown in the fundamental work by Crandall and Lions (1981), we show in the appendix that the value function V is a weak solution of the optimization problem (15), and if we obtain a form of V we can conclude it solves the firm's problem (in a weak sense).

Equation (15) implies an optimal extraction policy given by

$$q^*(t, X_t) = \left[\frac{a - V_x}{2b + c} \right]_+ . \tag{16}$$

[7] See Fleming and Soner (2006) for the full definition of control admissibility.

Note that this implies that in order for extraction to stay positive, $V_x \geq a$, meaning the resource rent cannot exceed the demand intercept parameter. This is clearly a consequence of the assumption of linear demand, which results in a quadratic criterion. It will be clear in what follows that the solution will be naturally constrained by the boundary conditions to satisfy this requirement. Substituting in (15) and grouping terms, we obtain the following partial differential equation:

$$0 = V_t - \rho V + AV_x + BV_x^2 + \frac{\sigma^2}{2} V_{xx} + C \tag{17}$$

where the constants A, B and C are given by

$$\begin{aligned} A &= \mu - \frac{a}{2b+c}, \\ B &= \frac{1}{2(2b+c)}, \\ C &= \frac{a^2}{2(2b+c)} - F. \end{aligned}$$

The natural boundary conditions of this problem are given by

$$V(t,x) = 0 \text{ for } x < 0,\ V(t,0) = 0,\ q(t,0) = 0 \tag{18}$$

without imposing a smooth pasting condition because of the viscosity argument. We also have an absorbing boundary with discontinuity for q given by

$$\lim_{q \uparrow x} q(t,x) = q^*(t,x), \lim_{q \downarrow x} q(t,x) = 0,$$

meaning that when the optimal harvesting policy reaches the actual level of the resource stock, the firm will have harvested the entire resource, and once this level is hit the harvesting shifts immediately to zero because of the natural boundary conditions (as well as common sense). Because of the separable form of $\Pi(t,q)$, we guess a solution of the HJB equation of the form

$$V(t,x) = e^{\rho(t-\tau_1)} V(x)$$

and we linearize it with the nonlinear change of variable

$$V'(x) = \frac{\sigma^2}{2B} \frac{\psi'(x)}{\psi(x)} = e^{-\rho(t-\tau_1)} V_x(t,x)$$

where $\psi(.)$ is a general twice differentiable function on $\mathbb{R}$. By this linearization, one can easily obtain the general solution

$$\psi_g(x) = c_1 e^{\alpha_1 x} + c_2 e^{\alpha_2 x}. \tag{19}$$

where $\alpha_{1,2} = \frac{-A \pm \sqrt{A^2 - 4BC}}{\sigma^2}$ and $\alpha_2 < \alpha_1$. The constants are given by the boundary conditions (18), after noticing that $V(t,0) = 0$ implies $\psi(0) = 1$. The particular solution can be computed in closed form, but its expression is lengthy and therefore omitted, and henceforth only referred to as $\psi(x)$. The optimal harvesting policy in feedback form is therefore

$$q^*(t,x) = q^m - \sigma^2 \frac{\psi'(x)}{\psi(x)} e^{-\rho(\tau_1 - t)}, \tag{20}$$

where $q^m = \frac{a}{2b+c}$ is the quantity at which the monopolist's marginal revenue equals marginal cost: it's the profit maximizing harvesting policy the monopolist would choose if there were no fluctuations in the evolution of the resource (i.e. if $\sigma = 0$). Observe that V_x here is the rent associated with a unit of the resource stock. It is the scarcity value or the market value of the marginal unit of *in situ* stock. From (16) we obtain the resource rent for the monopolist:

$$V_x = \sigma^2 \frac{\psi'(x)}{\psi(x)} e^{-\rho(\tau_1 - t)} (2b + c) \tag{21}$$

The optimal harvesting function exhibits a sigmoid-like form[8], resulting in the following limiting behavior:

$$\lim_{x\to\infty} q^*(t,x) = q^m - \nu(\Theta)e^{-\rho(\tau_1-t)}, \tag{22}$$

$$\lim_{\tau_c\to\infty} q^*(t,x) = q^m. \tag{23}$$

where ν is a general continuous and bounded function of the model parameters. This result shows that for any time $t \in [0, \tau_1]$, there is a maximum harvesting level given by a fixed amount, generated by market conditions, minus a parameter which incorporates the dynamics of the resource stock, modulated by the distance between present and the detection time representing the time horizon of the firm. If this time horizon is long enough, all resource-related parameters are ignored and the monopolist's optimal harvest is entirely driven by the market.

We therefore have an optimal harvesting policy that has two parts: one driven purely by market preferences. It is increasing on the maximum price the consumers are willing to pay i.e. representative of consumer preferences (a) and decreasing in sensitivity of price to extraction and the slope of the demand curve faced by the monopolist and its cost (b and c). The second part is variable and explicitly dependent on state X_t and current time, as well as the market and other model parameters. We can therefore write the optimal extraction policy as

$$q^*(t,x,\lambda) = q^m - q^v(t,x,\lambda), \tag{24}$$

and the resource rent as:

$$V_x(t,x,\lambda) = q^v(t,x,\lambda)(2b+c) \tag{25}$$

Note that when the market value of the marginal unit of in situ resource rises, q^* decreases

2.5 Sequential Detection and Optimal Behavior

To understand how the regime shift is endogenous to the firm's harvesting efforts and the role of the detection time, let us insert the solution to the firm's optimization problem in the sequential detection scenario. A schematic representation of which can be seen in Figure 1.

Period $[0, \tau_1]$: At the initial time $t = 0$ the resource X_t is driven by a diffusion process with natural growth rate μ and the firm begins harvesting activity at level $q^*(0, x_0)$. At a random time $\theta_1 \in [0, T]$, there is an initial exogenous change in the resource dynamics, λ_0, which we assume as $\lambda_0 < 0$ so as to start a process of subsequent adjustment. Until the detection of this change, the firm operates in an environment where the resource behaves according to the process

$$dX_t = (\mu - q_0^*(t, X_t))dt + \sigma dW_t, \quad t \in [0, \tau_1], \tag{26}$$

where $\tau_1(\lambda_0, \nu) \leq T$ is the optimal detection time. Its expectation, which is what the firm uses as a reference for its decisions, is given by

$$\mathbb{E}[\tau_1] = \frac{2}{\lambda_0^2}\left(e^{-\nu} + \nu - 1\right) \tag{27}$$

where the threshold ν solves $\frac{2}{\lambda_0^2}\left(e^{\nu} - \nu - 1\right) = T$. Within this time interval, the optimal harvesting policy for the firm is given by (20):

[8] Assume for the sake of exposition $t = \tau_c$, $\sigma = 1$ and $a = 2b + c$, one obtains

$$q^* = \frac{c_1(1-\alpha_1)e^{\alpha_1 x} + c_2(1-\alpha_2)e^{\alpha_2 x}}{c_1 e^{\alpha_1 x} + c_2 e^{\alpha_2 x}}.$$

If we have $c_1 = c_2$, $\alpha_2 < 1 < \alpha_1$, we obtain a shifted hyperbolic tangent function, directly related to the logistic function. For general parameter values, therefore, the optimal extraction policy has a modulated sigmoid form.

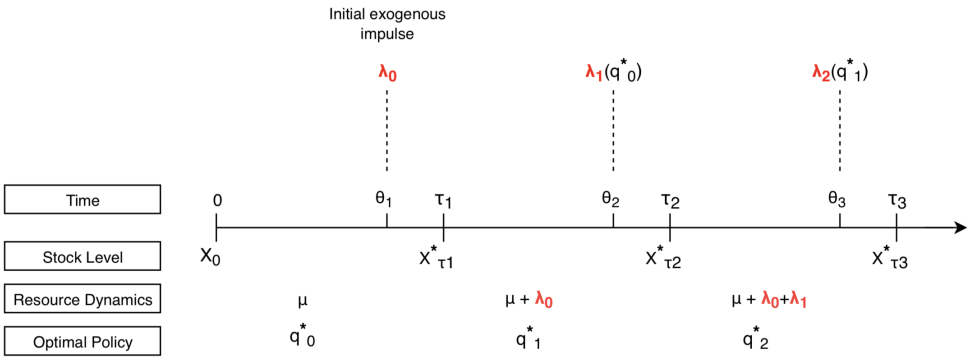

Figure 1: Sequential Detection

$$q^*(t,x) = q^m - \sigma^2 \frac{\psi'(x)}{\psi(x)} e^{-\rho(\tau_1 - t)}$$

Note that the "instantaneous" drift of the optimally controlled stock is given by

$$\mu(x,t) = \mu - q^m + \sigma^2 \frac{\psi'(x)}{\psi(x)} e^{-\rho(\tau_1 - t)}$$

and as $t \to \tau_1$ the effective discount rate reduces and the drift increases. At the end of the time period the resource stock will be given by

$$X^*_{\tau_1} = X_0 + [\mu - q^m]\,\tau_1 + \sigma^2 \int_0^{\tau_1} \frac{\psi'(X_t)}{\psi(X_t)} e^{-\rho(\tau_1 - t)} dt + \sigma \int_0^{\tau_1} dW_t. \tag{28}$$

where the second integral is to be interpreted in the Itô sense. Similar to (20), the overall dynamics of the optimally controlled resource stock also comprise of a fixed growth part, given by the natural growth μ and market preferences, and a variable growth part.

Period $[\tau_1, \tau_2]$: Once the new regime is detected at $t = \tau_1$, the firm then immediately reassesses its optimal policy to $q_1^*(t,x)$, as the dynamics of the resource stock are now

$$dX_t = (\mu - \lambda_0 - q_1^*(t, X_t))dt + \sigma dW_t, \quad t \in [\tau_1, \tau_2]$$

In the meantime, however, while the firm assumes a constant λ_0, its past decisions start to catch up. The growth of the "new" process will modify as a function of the past harvesting actions, yielding a change in drift $\lambda_1(q_0^*)$ given by

$$\lambda_1(q_0^*) = \begin{cases} \mu\ \Delta X_0 & \Delta X_0 < 0 \\ \mu\ \sqrt{\Delta X_0} & \Delta X_0 > 0 \end{cases} \tag{29}$$

$$\text{where} \quad \Delta X_0 = \frac{X^*_{\tau_1} - X_0}{X_0} \tag{30}$$

Equations (29) and (30) indicate that the magnitude of change in drift λ depends on how much the resource stock has deviated from its initial value. The observed sign will depend on whether the firm's harvesting actions have generated a net increase or decrease in the total stock of the resource. Note that the effect net change in stock levels has on the resource's capacity to regenerate is not symmetric. This is to capture the fact that an ecosystem is more vulnerable to negative shocks. When $\lambda_1(q_0^*)$ is negative due to a net decrease it has a linear impact on the

growth rate. However when $\lambda_1(q_0^*)$ is positive, the effect is concave. The sign and magnitude of the change in stock is assumed known by the firm, but *when* the regime actually changes is uncertain and is to be detected. At $\tau_2(\lambda_1(q_0^*), \nu) \leq T$ the firm will (on average) detect this change in regime of the resource dynamics. Its expectation, which is what the firm uses as a reference for its decisions, is given by

$$\mathbb{E}[\tau_2] = \frac{2}{\lambda_1(q_0^*)^2}\left(e^{-\nu} + \nu - 1\right) \tag{31}$$

where the threshold ν solves $\frac{2}{\lambda_1(q_0^*)^2}\left(e^{\nu} - \nu - 1\right) = T$.

The optimal policy for this period is easily seen to have the same form as the one for the previous period: normalizing time to $t_0 = \tau_1$, one immediately recognizes that the two problems are equivalent, with a change in drift from μ to $\mu - \lambda_0$. We therefore have

$$q_1^*(t,x) = q^m - \sigma^2 \frac{\tilde{\psi}'(x)}{\tilde{\psi}(x)} e^{-\rho(\tau_2 - t)}, \tag{32}$$

where the exponents $\tilde{\psi}(x)$ are equivalent as before but substituting $\mu - \lambda_0$ in the coefficient A. This optimal policy is maintained up to the second detection time τ_2, where the firm assesses the resource stock given by

$$X_{\tau_2}^* = X_{\tau_1}^* + \left[\mu - \lambda_0 - q^m\right]\tau_2 + \sigma^2 \int_{\tau_1}^{\tau_1+\tau_2} \frac{\tilde{\psi}'(X_t)}{\tilde{\psi}(X_t)} e^{-\rho(\tau_2 - t)} dt + \sigma \int_{\tau_1}^{\tau_1+\tau_2} dW_t.$$

Period $[\tau_2, \tau_3]$: Similar to previous periods, once the new regime is detected at $t = \tau_2$, the firm will reassess its optimal policy to $q_2^*(t,x)$, as the dynamics of the resource stock are now

$$dX_t = (\mu - \lambda_0 \pm \lambda_1(q_0^*) - q_2^*(t, X_t))dt + \sigma dW_t, \quad t \in [\tau_2, \tau_3]$$

In the meantime, the growth of the resource will modify as a function of the past harvesting actions, yielding a change in drift $\lambda_2(q_1^*)$ given by

$$\begin{aligned} \lambda_2(q_1^*) &= \begin{cases} (\mu - \lambda_0)\ \Delta X_{\tau_1}^* & \Delta X_{\tau_1}^* < 0 \\ (\mu - \lambda_0)\ \sqrt{\Delta X_{\tau_1}^*} & \Delta X_{\tau_1}^* > 0 \end{cases} \\ \text{where} \quad \Delta X_{\tau_1}^* &= \frac{X_{\tau_2}^* - X_{\tau_1}^*}{X_{\tau_1}^*} \end{aligned}$$

For subsequent periods, the problem then continues sequentially.

More generally, for the period $[\tau_i, \tau_{i+1}]$ where $i = 1, 2...n$ we can model the optimal extraction policy as:

$$q_i^*(t, x, \lambda_{i-1}) = q^m - \underbrace{\sigma^2 \frac{\psi'(x, \lambda_{i-1})}{\psi(x, \lambda_{i-1})} e^{-\rho(\tau_{i+1} - t)}}_{q^v(t,\, x,\, \lambda_{i-1})} \tag{33}$$

the resource rent as:

$$V_x(t, x, \lambda_{i-1}) = \underbrace{\sigma^2 \frac{\psi'(x, \lambda_{i-1})}{\psi(x, \lambda_{i-1})} e^{-\rho(\tau_{i+1} - t)}}_{q^v(t,\, x,\, \lambda_{i-1})} (2b + c) \tag{34}$$

the change in the growth of the resource dependent on its past harvesting actions as[9]:

[9]Note that in the first period $[0, \tau_1]$ the change in drift, λ_0 is assumed to be exogenous and not dependent on past harvest efforts.

$$\lambda_i = \begin{cases} \left(\mu + \sum_{j=0}^{i-1} \lambda_j - \lambda_{i-1}\right) \frac{X^*_{\tau_i} - X^*_{\tau_{i-1}}}{X^*_{\tau_{i-1}}} & \Delta X^*_{\tau_{i-1}} < 0 \\ \left(\mu + \sum_{j=0}^{i-1} \lambda_j - \lambda_{i-1}\right) \sqrt{\frac{X^*_{\tau_i} - X^*_{\tau_{i-1}}}{X^*_{\tau_{i-1}}}} & \Delta X^*_{\tau_{i-1}} > 0 \end{cases} \tag{35}$$

and the firm will (on average) detect the regime shift at:

$$\mathbb{E}[\tau_{i+1}] = \frac{2}{\lambda_i (q^*_{ex})^2} \left(e^{-\nu} + \nu - 1\right) \tag{36}$$

Observe that in each period the final level of resource is a random variable, and therefore so is the impact on the new growth rate, however it is continuously dependent on the optimal harvesting policy. The variation in X is conserved in the *magnitude*, the absolute value of the percentage change in the resource stock translates directly to a change in drift. The larger is the difference between initial level of stock, the larger will be the magnitude of the change in resource growth rate λ_i. This implies, on average, an *earlier* expected time of detection. Once this detection occurs, the parameter will either increase or decrease the probability of extinction of the resource by entering the SDE drift with the same sign as the difference between initial and final level of resource biomass. We discuss this in further detail in section 4.

3 Characteristics of the Solution

Due to the sequential nature of the detection process and the stochastic dynamics of the resource, there is no steady state in our model. The system is non stationary and is randomly changing and as a result optimal harvest must be specified for every state that can possibly occur. Additionally, in a multi-period setting, the detection of each regime shift alters the monopolist's time horizon. This is evident in Figure 2(a), which shows a possible time path of the stock biomass, for the first four periods, being harvested under the profit maximizing policies of the firm. As the monopolist detects each regime shift, represented by the red dashed lines, it's horizon for the period changes and it pursues the appropriate optimal policy. Panel (b) shows an example of the firm extracting the resource to extinction, with a collapse occurring in the third period. The varying time horizon of each period plays into the firm's extraction decisions.

With (33) and (34), we can now examine how a change in regime and its detection affects the firm decisions. To do this we choose a range of values for the model parameters, which are meant to be largely illustrative. In section 4 we simulate our model with values taken from real examples. Figure 3 shows the optimal extraction policy of the firm in the first period $[0, \tau_1]$ as shown in (20). The resource has a natural growth $\mu = 6000$ *tonnes/year* and $\sigma = 3250$ *tonnes/year*. The firm incurs a variable cost with $c = 1.25$, a fixed cost of 500 *$/tonne/year* and applies a discount rate of $\rho = 0.02$. We illustrate the case of a negative regime shift as it is of more interest and relevance today. Suppose the ecological system undergoes a change, resulting in a modified drift with $\lambda = -2500$ *tonnes/year*. In Figure 3 panel (a), we observe that at lower stock levels the firm adopts a precautionary behaviour for all detection periods. This is due to the negative regime shift reducing the growth rate of the resource and creating a physical scarcity. This, consecutively, increases the resource rent and reduces the extraction levels by the monopolist. Panel (c) depicts the resource rent for the monopolist at various detection periods for two different stock levels. Note that for a low stock, a decline in resource growth i.e. a negative regime shift always leads to an increase in the value of the marginal unit of *in situ* stock.

On the contrary for higher stock levels the monopolist adopts an aggressive approach and increases it's extraction. To understand this behaviour, remember the rule of thumb for pricing

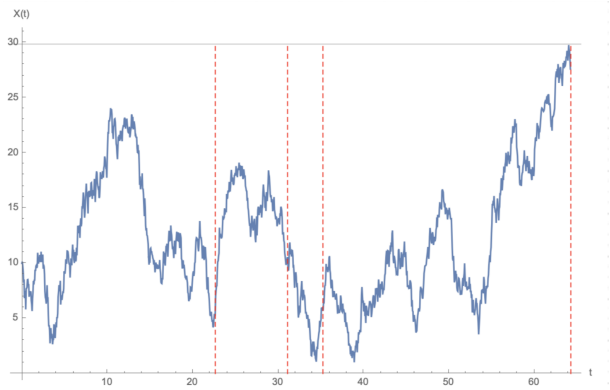

(a)

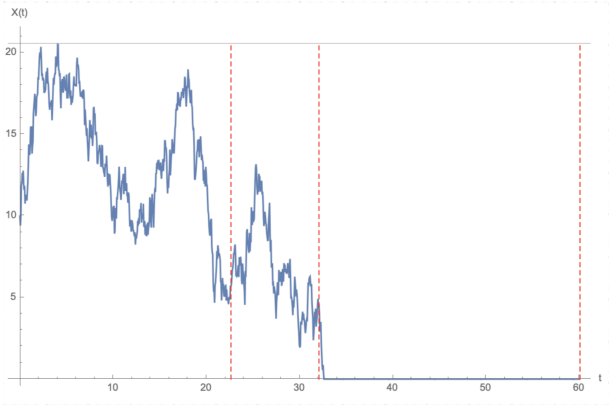

(b)

Figure 2: Possible Time Paths of the Optimally Controlled Stock Biomass. Red dashed lines represent detection times. For a demand function of the form $p(q) = 3 - 0.1q$, variable cost function $c(q) = 0.5q^2/2$ and fixed cost $F = 0.25$. The natural growth of the resource $\mu = 5$, variance $\sigma = 3$ and $T = 50$. $\lambda_0 = -1.5$ and $X_0 = 10$. The values are in thousand tonnes and years.

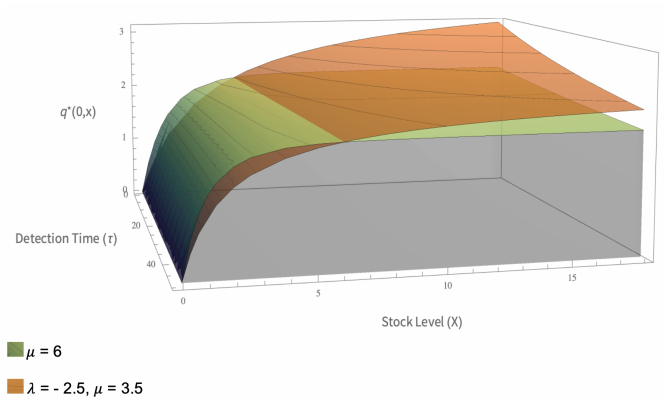

(a)

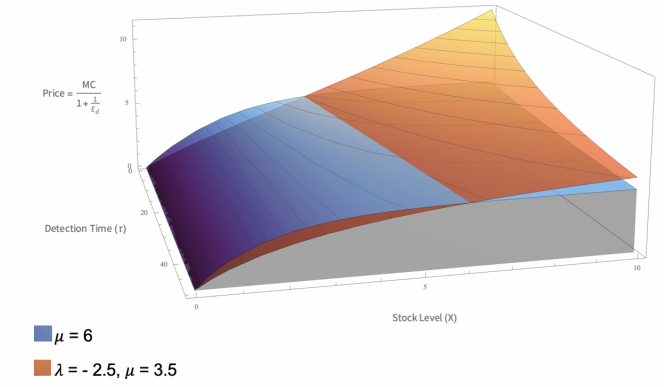

(b)

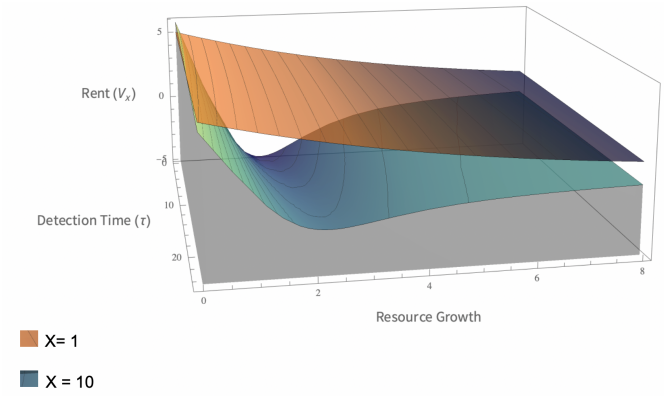

(c)

Figure 3: Panel (a): Optimal extraction policies for the monopolist. For a demand function of the form $p(q) = 5 - 0.75q$, variable cost function $c(q) = 1.25q^2/2$ and fixed cost $F = 0.5$. The natural growth of the resource $\mu = 6$, variance $\sigma = 3.25$. The values are in thousand tonnes and years. Panel (b): Price set by the monopolist. Panel (c): Resource rent for the monopolist at low ($X = 1$) and high ($X = 10$) stock level.

for a monopolist:

$$P = \frac{MC}{1 + \frac{1}{E_d}} \tag{37}$$

The firm in our model is a pure monopolist thus its demand curve is the market demand curve and $E_d = 1 - \frac{a}{bq^*(t,X_t)}$. For a given slope, unchanged consumer preferences and large resource stocks, by extracting higher levels, compared to pre-regime shift , the monopolist can in fact charge a higher price. This is shown in panel (b). Therefore at higher stock levels, the scarcity effect is outweighed by the increase in the markup over the marginal cost which reduces the resource rent and results in higher extraction by the monopolist. There is another factor driving the increased extraction efforts by the monopolist - the changing horizon of the firm. In panel (a) observe the increase in extraction is not uniform across all detection periods. Remember from section 2.3, the model implies, larger is the regime shift, the quicker a firm is likely to detect and shorter is the time horizon for that period. If the monopolist is faced with a short time horizon and high levels of resource stock, it increases its extraction even more than it would for the same stock levels and a longer detection period. This is because the monopolist believes that another shift in regime could happen very soon and resource extinction or collapse may be impending. This behaviour is also reflected in the price setting decision by the monopolist, as seen in panel (b).

In panel (c) the behaviour of the resource rent for a higher stock level is more varied compared to that of a lower stock. A decline in resource growth rate reduces the rent which increases the firm's extraction. This effect is amplified for shorter detection periods. However if the regime shift is exceptionally large, the rent starts to increase and the firm may decide to lower its extraction. For this reason, at high stock levels, the extraction decisions of the firm is a combination of market dynamics and the altered time horizon from the detection of the regime. The interplay between these two factors is evident in Figure 4. The resource undergoes a regime shift similar to Figure 3 and additionally in panel (a) and (b) the market preferences change as the slope of the monopolist's demand curve becomes steeper indicating that the demand for the resource is now less elastic. The extraction levels for longer detection periods and all stock levels continue to be lower than the pre-regime shift levels. However, for short detection periods and higher stock levels, the firm increases its extraction, and continues to charge a higher price. This highlights the effect of urgency or impending doom that drives the monopolist's decisions. In panel (c) and (d) the regime shift is accompanied with a change in consumer preferences as shown by a decline in the maximum price the consumers would be willing to pay for the resource. The monopolist exhibits an aggressive extraction behaviour even for longer detection periods and higher stock levels as it is able to charge a higher price.

4 Risk of Catastrophe

We define the **risk of catastrophe** as the situation in which the instantaneous drift of the resource stock X_t is negative in period i at any time $t \in [\tau_i, \tau_{i+1}]$:

$$\mu + \sum_{j=0}^{i-1} \lambda_j - q^m + \sigma^2 \int_{\tau_i}^{t} \frac{\tilde{\psi}'(X_s)}{\tilde{\psi}(X_s)} e^{-\rho(\tau_{i+1}-s)} ds < 0, \tag{38}$$

which implies that $P(\lim_{t\to\infty} X_t = 0) = 1$. To illustrate the emergence of catastrophe risk, let us suppose the firm finds itself at $t = \tau_i$, with an optimally controlled stock level $X^*_{\tau_i}$. Expliciting again the dependence of the sign λ_i on the difference between initial and final levels of stock for each period, for the subsequent period the resource stock will follow the controlled process given by

$$dX_t = \left(\mu + \sum_{j=0}^{i-1} \lambda_j \cdot \text{sign}(X^*_{j-1} - X^*_{j-2}) - q^*_i(X,t) \right) dt + \sigma dW_t$$

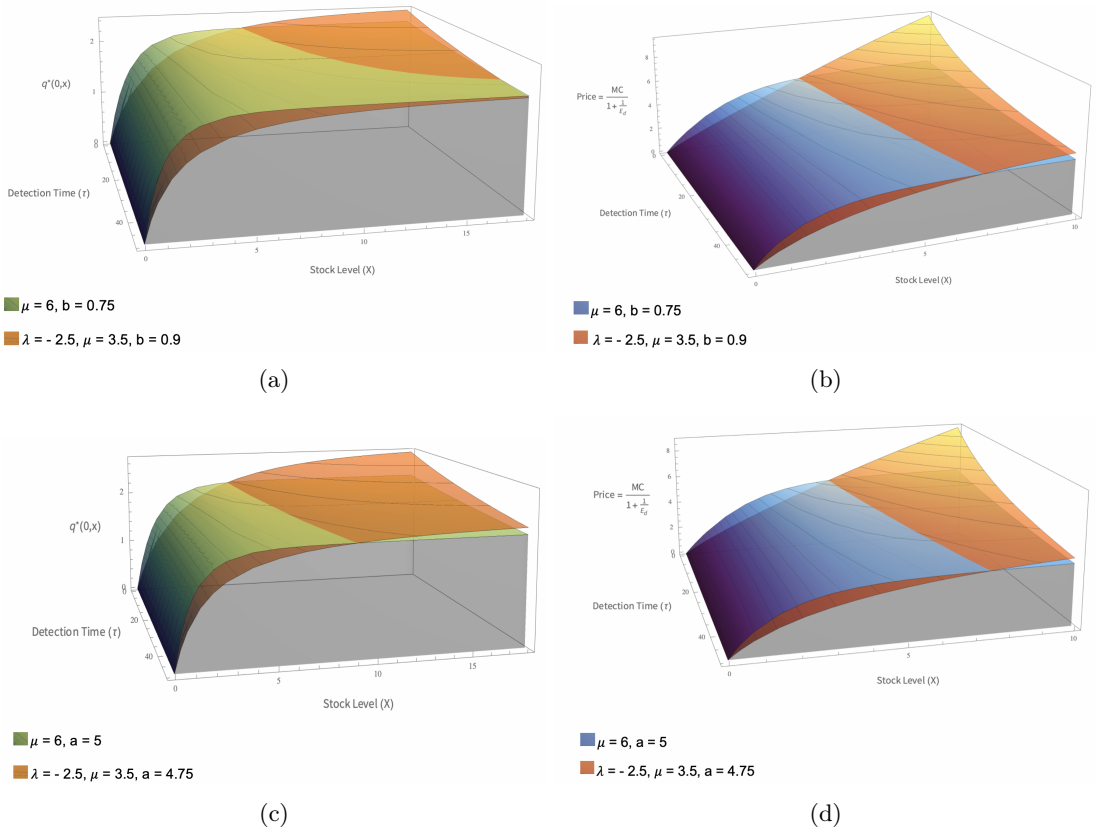

Figure 4: Panel (a): Optimal extraction policies for the monopolist facing a regime shift in the resource dynamics and a change in the slope of the market demand. Panel (c): Optimal extraction policies for the monopolist facing a regime shift in the resource dynamics and a change in the maximum price the consumers would pay for the resource. Panel (b) and Panel (d): Price set by the monopolist. Firm parameters are the same as Figure 3

until the next detection time τ_{i+1}, and for all $t \in [\tau_i, \tau_{i+1}]$ the profit-maximizing extraction policy will be $q_i^*(x,t)$. Let us suppose that $\mu + \sum_{j=0}^{i-1} \lambda_j \cdot \text{sign}(X_{j-1}^* - X_{j-2}^*) > 0$ so that the firm does not find itself under risk. The firm at this point begins the detection process for the next change of regime, which will be given by (35) and if $\lambda_i < 0$, the firm will realize the future emergence of catastrophe risk if

$$\mu + \sum_{j=0}^{i} \lambda_j < 0,$$

noting that the sum goes to period i, meaning that at the next detection time τ_{i+1} the new regime will be one in which the drift of the resource stock process will be negative, meaning that the resource will have a net tendency to be driven towards an extinction state ($X = 0$).

First passage time to catastrophe: At this moment the firm may have to reassess its extraction policies, due to the fact that the resource growth rate has been affected by its past extraction decisions to a point where extinction is likely. In fact, the probability of the resource being zero in infinite time is unity, which means that the resource eventually *will* be depleted. The firm, however, can now exploit the non-stationary nature of the time intervals in which it operates: a first immediate analysis should be what happens if it stops extracting. Normalizing time to $\tau_i = 0$, we define the probability of extinction as

$$\phi(x) = \Pr\left[\inf_{t \in \mathbb{R}^+} X_t \leq 0 \middle| X_0 = X_{\tau_i}^*, q^*(t, X_t) = 0\right] \tag{39}$$

and the first time to catastrophe as

$$\tau_c = \inf[\ t | X_t \leq 0, X_0 = X_{\tau_i}^*, q^*(t, X_t) = 0]. \tag{40}$$

Then X_t follows simply a drifted Brownian motion and the problem is equivalent of finding where a standard Brownian motion crosses the line $x - \mu - \sum_{j=0}^{i} \lambda_j$ (remember that $\mu + \sum_{j=0}^{i} \lambda_j$ is negative). It's a classic stochastic analysis problem, and it allows the firm to realize that if it stops extracting the expected time to catastrophe is

$$\mathbb{E}\tau_c = \frac{X_{\tau_i}^*}{\left|\mu + \sum_{j=0}^{i} \lambda_j\right|}. \tag{41}$$

and the probability of extinction is

$$\phi(x) = \exp\left(-\frac{2\left(|\mu + \sum_{j=0}^{i} \lambda_j|\right)}{\sigma^2} x\right). \tag{42}$$

If $\mathbb{E}\tau_c \leq \tau_{i+1}$, on average the resource will be depleted within the detection period even if the firm stops extracting altogether: we are therefore in a situation of **irreversible catastrophe**, where even the most precautionary of extraction behavior cannot avoid on average the resource from being depleted. In other words, since extraction always reduces the drift, (41) gives the upper bound on all first times to catastrophe. Since deviation from the optimal policy is costly, it is likely that the firm will continue its extraction policy until extinction.

If $\mathbb{E}\tau_c \geq \tau_{i+1}$, catastrophe is on average avoidable within the first detection period if the firm stops extraction, therefore the firm can study whether its optimal extraction policy allows to avoid it as well. In other words, the firm wants to check whether

$$\begin{aligned} \mathbb{E}\tau_c &\leq \tau_{i+1}, \\ \tau_c &= \inf[\ t | X_t \leq 0, t \in [0, \tau_{i+1}], X_0 = X_{\tau_i}^*]. \end{aligned}$$

Define $\psi(t) = \psi(t; X_{\tau_i}, 0)$ the density function of the first time to catastrophe: then we have that

$$1 - \psi(t) = 1 - \phi(0, t), \tag{43}$$

where $\phi(x,t)$ is the probability that the optimally controlled resource stock X_t^* hits the absorbing barrier at 0, and can be written as

$$\phi(x,t) = \Pr\left[\inf_{s\in[t,\tau_{i+1}]} X_s^* \leq 0 \middle| X_t = x\right],$$

for $0 \leq t \leq \tau_{i+1}$. The firm therefore has to solve the Kolmogorov forward equation given by

$$\frac{\partial}{\partial t}\phi(x,t) + \frac{\partial}{\partial x}\phi(x,t)\left(\mu + \sum_{j=1}^{i-1}\lambda_j - q^*(t,x)\right) + \frac{\sigma^2}{2}\frac{\partial^2}{\partial x^2}\phi(x,t) = 0 \tag{44}$$

with absorbing boundary conditions given by

$$\begin{aligned} &\begin{cases} \phi(x,\tau_{i+1}) = 1 & x \leq 0 \\ \phi(x,\tau_{i+1}) = 0 & x > 0, \end{cases} \\ &\phi(0,t) = 1, \\ &\phi(t,\infty) = 0. \end{aligned} \tag{45}$$

The KFE for this problem has no closed form solution, given the highly nonlinear form of the extraction policy, and needs to be solved numerically with standard methods. Once the solution is obtained, the firm can recover the density of the first time to catastrophe τ_c from (43) and compute its expectation: if $\mathbb{E}\tau_c \geq \tau_{i+1}$ the firm continues its optimal extraction policy.

5 Real-time detection and optimal extraction

The optimal extraction policy in each time interval $[\tau_i, \tau_{i+1}]$ is obtained by assuming as time horizon the expectation of the optimal stopping time $\mathbb{E}[\tau(-\lambda, \nu)]$. This is therefore an *ex ante* policy: the actual detection of when the regime shift happens is not included except for the first moment of the procedure. The time θ at which the regime changes, however, is a random variable: the firm therefore will use the expected detection time (31) to evaluate the boundary conditions, but simultaneously observe continuously the optimally controlled level of stock X_t, change to the measure Q and compute the Radon-Nikodym derivative of the two measures (before and after the regime change) and check whether its value exceeds the threshold value ν. If the threshold is reached *before* the expected detection time τ_i+1, then the firm simply switches to the subsequent period with the modified drift, since the regime shift has been detected. If the expected detection time τ_{i+1} is reached and the threshold has not yet been reached (i.e. the regime has not yet shifted), the firm continues the optimal extraction and does not switch to the next period, but using an infinitesimal time interval as horizon. In other words, the infinitesimal optimal extraction policy if the expected detection time is exceeded is given by

$$q_i^*(t, x, \lambda_{i-1})dt = q^m dt - \sigma^2 \frac{\psi'(x, \lambda_{i-1})}{\psi(x, \lambda_{i-1})} e^{\rho dt} \tag{46}$$

(note that $e^{\rho dt} \sim 1$, and for numerical simulations dt is the size of the time mash), until either the Radon-Nikodym derivative of the measures of the two regimes reaches the threshold ν, or until the firm's tolerance time limit T is reached.

6 Concluding Remarks

Much of the literature on natural resources assumes a fixed and exogenous price, we introduce a model where a monopolist firm operates in a resource market where the prices are endogenously determined. Ecological uncertainty is introduced in the form of regime shifts that several

ecosystems undergo, for example lakes may shift from a clear to a turbid state, affecting the fish population and water quality. These shifts are made to be dependent on the the monopolist's extraction efforts. Unlike previous work, we explicitly model the firm's detection process of the regime change and incorporate it in it's profit maximizing policies. Our closed form solutions help us pin down the mechanisms driving the extraction behaviour of the firm. In the event of a negative regime shift, for low stock levels, an increase in the the resource rent results in the firm adopting a precautionary policy by extracting less. For higher stock levels, a regime shift leads to an increase in extraction due to an altered and relatively shorter time horizon and demand elasticity - which reduces the resource rent and results in the monopolist adopting an aggressive behaviour.

In concluding, some caveats are clearly in order. Our model is intentionally simple and stylized. An initial criticism maybe the assumption of a monopoly. Although a pure monopoly maybe rare, and a game theoretic approach of several powerful players interacting, maybe more appropriate to the renewable resource market - our primary aim with the model is to see how a firm, whose prices are not exogenous, decides its extraction levels in the event of detecting a regime shift. Starting with a monopoly is the first step and a novel benchmark before extending research in that direction. Lastly, we have made two simplifying assumptions in the form a constant growth rate and cost function that is not directly dependent on stock. Both these assumptions can be relaxed at the expense of greater mathematical and numerical demand.

References

Alvarez, L. H. and E. Koskela (2007). Optimal harvesting under resource stock and price uncertainty. *Journal of Economic Dynamics and Control 31*(7), 2461–2485.

Andersen, P. and J. G. Sutinen (1984). Stochastic bioeconomics: A review of basic methods and results. *Marine Resource Economics 1*(2), 117–136.

Arvaniti, M., A.-S. Crépin, and C. K. Krishnamurthy (2019). Time-consistent resource management with regime shifts. *Economics Working Paper Series 19*.

Baggio, M. and P. L. Fackler (2016). Optimal management with reversible regime shifts. *Journal of Economic Behavior & Organization 132*, 124–136.

Batt, R. D., S. R. Carpenter, J. J. Cole, M. L. Pace, and R. A. Johnson (2013). Changes in ecosystem resilience detected in automated measures of ecosystem metabolism during a whole-lake manipulation. *Proceedings of the National Academy of Sciences 110*(43), 17398–17403.

Carpenter, S. R., W. A. Brock, J. J. Cole, and M. L. Pace (2014). A new approach for rapid detection of nearby thresholds in ecosystem time series. *Oikos 123*(3), 290–297.

Chavez, F. P., J. Ryan, S. E. Lluch-Cota, and M. Ñiquen (2003). From anchovies to sardines and back: multidecadal change in the pacific ocean. *science 299*(5604), 217–221.

Crandall, M., H. Ishii, and P. Lions (1992). User's guide to viscosity solutions of second order partial differential equations. *Bulletin of the American Mathematical Society 27*(1), 1–67.

Crandall, M. G. and P.-L. Lions (1981). Condition d'unicité pour les solutions généralisées des équations de Hamilton-Jacobi du premier ordre. *C. R. Acad. Sci. Paris Sér. I Math. 292*(3), 183–186.

de Avila, A. L., G. Schwartz, A. R. Ruschel, J. do Carmo Lopes, J. N. M. Silva, J. O. P. de Carvalho, C. F. Dormann, L. Mazzei, M. H. M. Soares, and J. Bauhus (2017). Recruitment, growth and recovery of commercial tree species over 30 years following logging and thinning in a tropical rain forest. *Forest ecology and management 385*, 225–235.

de Zeeuw, A. and X. He (2017). Managing a renewable resource facing the risk of a regime shift in the ecological system. *Resource and Energy Economics 48*, 42–54.

Fleming, W. and H. Soner (2006). *Controlled Markov Processes and Viscosity Solutions*. Stochastic Modelling and Applied Probability. Springer New York.

Hadjiliadis, O. and V. Moustakides (2006). Optimal and asymptotically optimal cusum rules for change point detection in the brownian motion model with multiple alternatives. *Theory of Probability & Its Applications 50*(1), 75–85.

Klemas, V. (2013). Fisheries applications of remote sensing: an overview. *Fisheries Research 148*, 124–136.

Li, C.-Z., A.-S. Crépin, C. Folke, et al. (2018). The economics of resilience. *International Review of Environmental and Resource Economics 11*(4), 309–353.

Liu, J. (2014). Forest sustainability in c hina and implications for a telecoupled world. *Asia & the Pacific Policy Studies 1*(1), 230–250.

Nichols, J. D. and B. K. Williams (2006). Monitoring for conservation. *Trends in ecology & evolution 21*(12), 668–673.

Pauly, D. (1983). *Some simple methods for the assessment of tropical fish stocks*. Number 234. Food & Agriculture Org.

Pindyck, R. S. (1984). Uncertainty in the theory of renewable resource markets. *The Review of Economic Studies 51*(2), 289–303.

Pindyck, R. S. (2002). Optimal timing problems in environmental economics. *Journal of Economic Dynamics and Control 26*(9-10), 1677–1697.

Pizarro, J. and E. S. Schwartz (2018). The valuation of fisheries rights: A real options approach. Technical report, National Bureau of Economic Research.

Polasky, S., A. De Zeeuw, and F. Wagener (2011). Optimal management with potential regime shifts. *Journal of Environmental Economics and management 62*(2), 229–240.

Polunchenko, A. S., G. Sokolov, and W. Du (2013). Quickest change-point detection: A bird's eye view. *arXiv preprint arXiv:1310.3285*.

Poor, H. V. and O. Hadjiliadis (2008). *Quickest detection*. Cambridge University Press.

Reed, W. J. (1988). Optimal harvesting of a fishery subject to random catastrophic collapse. *Mathematical Medicine and Biology: A Journal of the IMA 5*(3), 215–235.

Reed, W. J. and H. R. Clarke (1990). Harvest decisions and asset valuation for biological resources exhibiting size-dependent stochastic growth. *International Economic Review*, 147–169.

Ren, B. and S. Polasky (2014). The optimal management of renewable resources under the risk of potential regime shift. *Journal of Economic Dynamics and Control 40*, 195–212.

Sakamoto, H. (2014). Dynamic resource management under the risk of regime shifts. *Journal of Environmental Economics and Management 68*(1), 1–19.

Saphores, J.-D. (2003). Harvesting a renewable resource under uncertainty. *Journal of Economic Dynamics and Control 28*(3), 509–529.

Scheffer, M., S. R. Carpenter, V. Dakos, and E. H. van Nes (2015). Generic indicators of ecological resilience: inferring the chance of a critical transition. *Annual Review of Ecology, Evolution, and Systematics 46*, 145–167.

Shimabukuro, Y. E., E. Arai, V. Duarte, A. Jorge, E. G. d. Santos, K. A. C. Gasparini, and A. C. Dutra (2019). Monitoring deforestation and forest degradation using multi-temporal fraction images derived from landsat sensor data in the brazilian amazon. *International journal of remote sensing 40*(14), 5475–5496.

Shiryaev, A. (1963). On optimum methods in quickest detection problems. *Theory of Probability & Its Applications 8*(1), 22–46.

Shiryaev, A. (1996). Minimax optimality of the method of cumulative sums (cusum) in the case of continuous time. *Russian Math. Surveys 51*(4), 750–751.

Szuwalski, C. S. and A. B. Hollowed (2016). Climate change and non-stationary population processes in fisheries management. *ICES Journal of Marine Science 73*(5), 1297–1305.

Viña, A., W. J. McConnell, H. Yang, Z. Xu, and J. Liu (2016). Effects of conservation policy on china's forest recovery. *Science advances 2*(3), e1500965.

Walters, C. and J.-J. Maguire (1996). Lessons for stock assessment from the northern cod collapse. *Reviews in fish biology and fisheries 6*(2), 125–137.

A Viscosity solutions

In all that follows we will use as a reference Fleming and Soner (2006) as well as follow its notations. What we want to achieve is to show that the value function V is a weak solution of the optimization problem (15), and if we obtain a form of V we can conclude it solves the firm's problem (in a weak sense).

We write the HJB equation in form of its infinitesimal generator. Define the set $\mathcal{D} \in C([0, \tau_c] \times \mathbb{R})$. Then $V(t, x) \in \mathcal{D}$ is a classical solution of the optimization problem (15) if it satisfies the equation

$$-\frac{\partial}{\partial t} V + A_t[V(t, .)](x) = 0, \tag{47}$$

where A is the generator of the HJB equation. If X_t were modeled as a geometric Brownian motion, the state constraint would not need to apply, since the multiplicative nature of the noise would naturally allow the resource stock to be positive, and because of the well-behaving nature of the functional forms of the problem we expect a smooth solution for all $X_t > 0$. But imposing $X_t \geq 0$ does not imply that that the value function has to be differentiable at $X = 0$. Now, define a continuous function $\mathcal{H}$ (the Hamiltonian) such that

$$A_t[\phi](x) = \mathcal{H}(t, x, D\phi(x), D^2\phi(x))$$

and consider the equation

$$-\frac{\partial}{\partial t} W(t, x) + \mathcal{H}(t, x, DW(t, x), D^2W(t, x)) = 0. \tag{48}$$

A function $V(t, x) \in \mathcal{C}([0, \tau_c] \times \mathbb{R})$ is a viscosity subsolution of (48) if for all $v \in C^\infty(\mathcal{D})$

$$-\frac{\partial}{\partial t} v(\bar{t}, \bar{x}) + \mathcal{H}(\bar{t}, \bar{x}, Dv(\bar{t}, \bar{x}), D^2v(\bar{t}, \bar{x})) \leq 0$$

for every point $(\bar{t}, \bar{x})$ which is a local maximum of $V - v$. Similarly, $V(t, x)$ is a viscosity supersolution of (48) if if for all $v \in C^\infty(\mathcal{D})$

$$-\frac{\partial}{\partial t} v(\bar{t}, \bar{x}) + \mathcal{H}(\bar{t}, \bar{x}, Dv(\bar{t}, \bar{x}), D^2v(\bar{t}, \bar{x})) \geq 0.$$

for every point $(\bar{t}, \bar{x}) \in \mathcal{D}$ which is a local minimum of $V - v$. The function $V(t, x)$ is a viscosity solution of the equation (48) if it is both a viscosity subsolution and a viscosity supersolution. This implies that the function $V(t, x)$ is a weak solution of the optimization problem (**??**). Let us now show that V is a viscosity solution of our problem (15).

Let $v \in C^2([0, \tau_c] \times \mathbb{R})$, let $V - v$ be maximized at the point $(\bar{t}, \bar{x}) \in ([0, \tau_c] \times \mathbb{R})$ and let us fix an optimal control (extraction rate) $q \in Q$. Let $X(.) = X(.; t, q)$ be the controlled stochastic process that drives the resource stock. For every time $\tau > \bar{t}$ for which $X_\tau > 0$, we have, using Ito's lemma and Bellman's principle of optimality,

$$\begin{aligned} 0 &\leq \frac{\mathbb{E}_{\bar{t}}\big[V(\bar{t}, \bar{x}) - v(\bar{t}, \bar{x}) - V(\tau, x(\tau)) + v(\tau, x(\tau))\big]}{\tau - \bar{t}} \\ 0 &\leq \frac{1}{\tau - \bar{t}} \mathbb{E}_{\bar{t}} \left[\int_{\bar{t}}^{\tau} \Pi(t, x, q) dt - v(\bar{t}, \bar{x}) + v(\tau, x(\tau)) \right]. \end{aligned}$$

This implies

$$0 \leq v_t(\bar{t}, \bar{x}) + \Pi(\bar{t}, x, q) + v_x(\mu + q) + \frac{\sigma^2}{2} v_{xx}$$

for all $q \in Q$: we can then write

$$
\begin{aligned}
0 &\leq v_t(\bar{t},\bar{x}) + \sup_{q\in Q}\left[\Pi(\bar{t},x,q) + v_x(\mu+q) + \frac{\sigma^2}{2}v_{xx}\right] \\
0 &\leq v_t - \mathcal{H}(\bar{t},\bar{x},Dv(\bar{t},\bar{x}),D^2v(\bar{t},\bar{x})).
\end{aligned}
$$

This proves that V is a viscosity subsolution of the problem (15). Proceeding similarly proves that V is a viscosity supersolution of the problem: if $V - v$ attains a minimum at $(\bar{t},\bar{x})$ then for any $\epsilon > 0$ and $\tau > \bar{t}$ we can find a control $q \in Q$ such that

$$
0 \geq -\epsilon(\tau - \bar{t}) + \mathbb{E}\left[\int_{\bar{t}}^{\tau} \Pi(t,x,q)dt - v(\bar{t},\bar{x}) + v(\tau,x(\tau))\right]
$$

which implies

$$
\epsilon \geq \frac{1}{\tau-\bar{t}}\mathbb{E}_{\bar{t}}\left[\int_{\bar{t}}^{\tau} \Pi(t,x,q)dt - v(\bar{t},\bar{x}) + v(\tau,x(\tau))\right].
$$

Proceeding equivalently as before, one shows that V is a viscosity supersolution of (15). We can conclude that V is a viscosity solution of (15). Note that for every time $\tau_e \in [0,\tau_c]$ for which $X_\tau > 0$, since for optimality we have $\Pi_q(.,q^*) - V_x = 0$ and Π is continuous and twice differentiable in q, it can be easily shown that the inequalities of the definition of sub- and supersolution are satisfied with equality, which means that $V(t,x)$ is also a classical solution of (**??**) for each $t = \tau_e$. We now need to deal with the positivity constraint. Given the "feasible" set $\mathcal{D}' = ([0,\tau_c] \times O \subset \mathbb{R}^+)$, we cannot impose that the value function $V(t,x)$ is differentiable (or continuous, for that matter) at 0 at the left boundary of $\partial\mathcal{D}'$. Following Fleming and Soner (2006), we need to impose a boundary inequality, which does not require neither V nor the boundary $\partial\mathcal{D}'$ to be differentiable at 0. This implies that the value function $V(t,0)$ must be a viscosity subsolution of (15). Following the previous definitions, we must have

$$
\begin{aligned}
v_t(t,0) &\leq -\mathcal{H}(t,0,Dv,D^2v) && (49)\\
&\leq \sup_{q\in Q}\left\{\Pi(t,x,q) + v_x(0)(\mu-q) + v_{xx}(0)\frac{\sigma^2}{2}\right\} && (50)
\end{aligned}
$$

for all continuous functions for which $V - v$ is locally maximized around $x = 0$. Given a natural boundary condition given by the fact that when the resource is zero, the extraction must be zero and consequently the objective Π must be zero. Since $V - v$ has to be maximized around 0, we have

$$
\mathcal{H}(t,0,a,a_x) \geq \mathcal{H}(t,0,v_x(t,0),v_{xx}(t,0)) \quad \forall a \geq v_x(t,0).
$$

The proof is simple, one just needs to write $\mathcal{H}(t,0,\alpha,\alpha_x)) = \sup_{q\in Q}\Pi(t,q) + \alpha(\mu+q) + \alpha_x\frac{\sigma^2}{2}$ and use $\alpha \geq v_x(t,0)$ to show the inequality holds. Given this result, condition (50) is easily seen to be satisfied by $V(t,0) = 0$, which we choose because of its immediate intuitive economic interpretation. We therefore can say that the constrained viscosity solution given by

$$
\begin{aligned}
V_x(t,0) &\geq \Pi(t,0,q) && (51)\\
V(t,0) &= 0 && (52)\\
V(t,x) &\text{ solves } V_t - \mathcal{H}(t,x,DV(t,x),D^2V(t,x)) = 0 \quad x \in [(0,\tau_c]\times\mathbb{R}]
\end{aligned}
$$

is a solution to the problem (15). Uniqueness of the solution is proven by means of the comparison principle, and since the proof follows closely the one provided by Crandall et al. (1992), is omitted.