Scenes from a Monopoly: Quickest Detection of Ecological Regimes*

Neha Deopa[†1] and Daniele Rinaldo[‡2]

[1]The Graduate Institute of International and Development Studies, Geneva, Switzerland.
[2]Faculty of Economics, University of Cambridge, UK.

Abstract

Decisions under ecological uncertainty are a crucial part of resource management as many ecological systems undergo abrupt regime shifts, frequently triggered by the actions of the resource harvester. We study the stochastic dynamics of a renewable resource harvested by a monopolist where harvesting affects the resource's potential to regenerate, resulting in sequential endogenous regime shifts. The firm faces uncertainty in the timing of these shifts. We encapsulate in our model environmental surveillance of ecological dynamics where the firm has to find the profit-maximizing extraction policy while simultaneously detecting in the *quickest time possible* the change in regime. Our key finding is that post-detection of a negative regime shift, for low stock levels, a precautionary behaviour can result due to increasing value of in situ stock. At higher levels, this behaviour is offset by an elastic demand as declining resource rent and effective marginal costs of production result in aggressive extraction. We find that intensification of extraction is possible due to a sense of urgency caused by the prospect of resource collapse. We study the probability of resource extinction and show the emergence of catastrophe risk which can be both reversible and irreversible based on the extinction's expected hitting time.

JEL Codes: D42, Q21, Q57

Keywords: Regime shift, Monopoly, Renewable resources, Quickest change-point detection, Uncertainty, Catastrophe

---

*An earlier version of this paper has been circulated under the title "Scenes from a Monopoly: Renewable Resources and Quickest Detection of Regime Shifts".

†neha.deopa@graduateinstitute.ch

‡dr551@cam.ac.uk

Acknowledgements: We are thankful in particular to Max-Olivier Hongler for introducing us to quickest detection problems and providing invaluable insights. In addition, this paper has benefited from comments and discussions with Jean-Louis Arcand, Piergiuseppe Fortunato, Sarah Jacobson, Daniel Lee, Aakriti Mathur, Daniel Millimet, Rahul Mukherjee, Ugo Panizza, Robert Pindyck, Kritika Saxena, Cédric Tille and Lore Vandewalle. We are also grateful to conference participants at the European Winter Meeting of the Econometric Society 2020, SAEe Meeting 2020, FAERE Conference 2020, Bolivian Conference on Development Economics 2020, Italian Economic Association Conference 2020, ClimRisk2020, Applied Young Economist Webinar, OSWEET, 1st IHEID Virtual PhD Workshop, University of Geneva & IHEID PhD Day, JKU Linz Brown Bag Seminar and IHEID Brown Bag Lunch. All errors are our own.

1 Introduction

Dynamically managing renewable resources necessitates making decisions under *ecological* uncertainty, defined as uncertainty over the evolution of the relevant ecosystem (Pindyck (2002)). One way that the current literature captures this is by means of stochastic bio-economic models, reflected in the variance of the fluctuations. Another way this uncertainty can manifest is in the form of ecological regime shifts: an abrupt change in the structure of the resource ecosystem or a change in its underlying population dynamics such as the intrinsic rate of growth. Such regime shifts have been well documented both as a result of natural and anthropogenic factors.[1] Their occurrence is also reflected in the increasingly relevant and common feature of the contemporary resource market where firms are investing considerable efforts in monitoring resource stocks as the number of ecological extreme events are on the rise.

The presence of regime shifts and the act of environmental surveillance to detect them, fundamentally alters the constraints and incentives faced by firms who extract renewable resources. This paper intends to answer how the extraction policies of a monopolist, who is sequentially monitoring the resource, change in response to a new ecological regime. What is the profit-maximizing policy of this firm who wants to *detect* this shift, in a framework where this change in regime is endogenously determined by the firm's extraction activity?

There already exists a large literature studying the impact of stochastic fluctuations on firm extraction and harvesting activities utilizing real options theory.[2] 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), Costello et al. (2019) and Arvaniti et al. (2019).[3] These studies, however, are limited in two respects. First, with the exception of Pindyck (1984) and Sakamoto (2014) much of the literature does not incorporate an explicit market structure and takes the price as fixed or exogenous. This is done for tractability but leads to results that underestimate the crucial role of demand, which often drives firm harvesting decisions. For example, Myanmar Timber Enterprise (MTE), a state-run company holding monopoly over harvesting and sale of timber, has directly been responsible for the loss of more than 13,000 square miles of tree cover between 2001 and 2018 due to the lucrative and increasing demand for teak.

Second, for a firm to be able to assess its optimal policy in response to an ecological regime change, it must be able to detect when the system dynamics of the resource shift. Therefore, resource management often involves an element of monitoring and surveillance which may directly play into the firm's decisions. In order to maximize its value, it is important for the firm to quickly detect the changes in the environment and adapt its policies accordingly. However,

[1] Recent examples are the logged tropical rainforests in parts of Asia, South America and Africa which have become more fire-prone leading to a regime shift towards exotic fire-promoting grasslands (Lindenmayer et al. (2011)) or the human-induced regime change in the Baltic Sea from cod to sprat and herring as dominant species in the fish population (Österblom et al. (2007)).

[2] See, for example, Andersen and Sutinen (1984); Pindyck (1984); Reed (1988); Reed and Clarke (1990); Saphores (2003); Alvarez and Koskela (2007); Pizarro and Schwartz (2018).

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

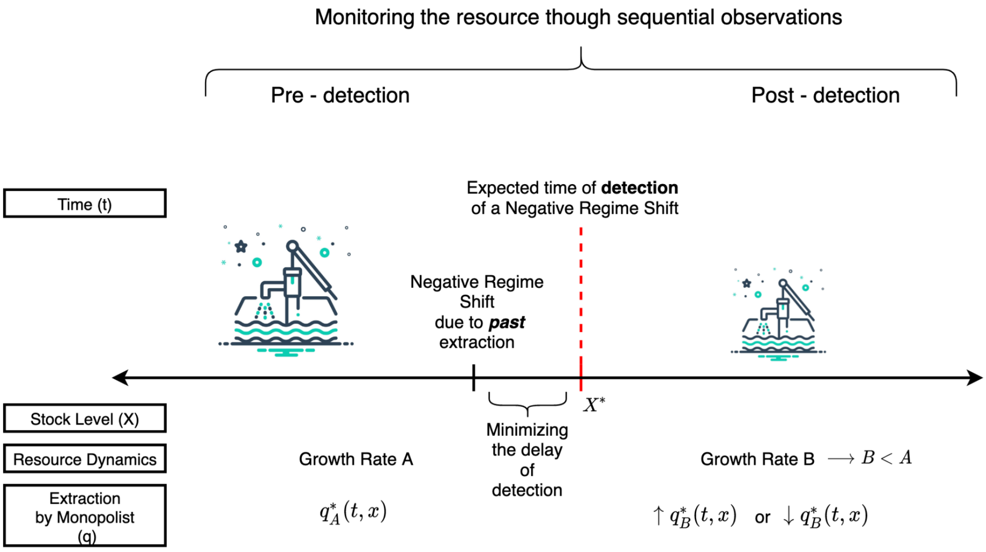

Figure 1: Model Dynamics

the current literature on regime shifts implicitly assumes the firm to be able to monitor these changes and subsequently make the appropriate extraction decision.

We build a model of a monopolist facing a linear demand curve, who encounters two sources of uncertainty in the resource dynamics. The first source is the natural randomness of the environmental conditions, here represented by Gaussian noise, and the second is 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 *past* extraction efforts. In our model, the firm knows with certainty that a regime shift will eventually occur but the timing is unknown: what matters for the firm's harvesting decision is *when* it will take place. In a multi-regime 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 by means of a *quickest detection* method. This method extends the classic change-point problems to an optimal stopping problem in a sequential framework where the stochastic process under observation is assumed to change its probabilistic characteristics at an unknown change-point in the sequence. Therefore, the aim of the firm is to maximize its present discounted value of the profits while simultaneously detecting a change in the resource growth, if one occurs, with the shortest delay possible. The sequential nature of the detection process allows us to incorporate non-stationary dynamics.

To sharpen our intuition of the model we refer to Figure 1. Suppose there exists a natural water monopoly extracting an aquifer and within the first interval it observes the resource growth rate to be A. At some unknown change-point, past extraction policies catch up to it, causing a negative regime shift. Due to continuous monitoring of the resource, the firm has an expected time of detection for this regime shift. Therefore, in this interval the firm maximizes its profits with respect to the growth rate A and over the horizon given by the detection time. Post-detection, the firm will reassess its policies as the resource dynamics have changed with a lower growth rate B: we study how the extraction policies change in this new regime. What role do the market preferences, the magnitude of the regime shift and the detection play in determining its optimal extraction policy?

Our model allows for fully analytical solutions and comparative static results. First, we find that the expected time of detection of an ecological regime shift is inversely related to its magnitude. Thus providing the intuitive result that the larger is the change in the ecosystem structure, the earlier a firm is likely to detect it after its occurrence. As the monopolist maximizes its profits with respect to the resource dynamics and *over* the expected detection time, this magnitude directly determines the horizon of the firm. Further, we model the past extraction to affect the size of the regime shift in the current period. Therefore, within a multi-regime framework and non-stationary dynamics, the monopolist's horizon for each period is updated according to the size of the extraction-dependent regime shift. Second, the profit maximizing optimal extraction policy is not only a function of demand but is also explicitly dependent on the current level of the resource stock and is modulated by the distance from the detection time. We find that in the event of a detection of a negative regime shift, for low stock levels, the firm adopts a precautionary policy by reducing extraction. This is because if the amount of resource stock required by the monopolist to break even is less than the resource's rate of growth, the regime shift creates a scarcity of the resource which increases the market value of the marginal unit of in situ stock. This results in reduced extraction levels in the new regime.

At higher stock levels, however, if demand is elastic it outweighs the scarcity effect. We find that the resource rent is a concave function of the resource growth, and the value of the in situ stock falls with a negative regime shift. Therefore, after the immediate detection of the shift the firm may reduce its extraction but over the course of the new regime it continues to increase extraction eventually outpacing the levels in the previous regime, thus pursuing an aggressive approach. Moreover, due to the inter-temporal nature of the monopolist's optimization problem, when the stock levels are high, the resource rent is *negative*. This is because an incremental unit of current production reduces future production costs and thus that production of the unit brings a benefit (a negative rent). Therefore, a combination of declining resource rents and a further reduction in the full or effective marginal cost of production (marginal cost plus the resource rent) yields an incentive for the monopolist to increase extraction. Moreover, as the optimal price depends on both the demand elasticity and the effective marginal cost, increasing extraction allows the monopolist to charge higher markups.

On comparing two regime shifts of different magnitudes, *ceteris paribus*, we find an interesting result that may initially seem counter-intuitive. At higher stock levels, a shorter expected time of detection can in fact lead the monopolist to intensify its extraction efforts, as compared to a longer horizon. This behaviour is precisely due to the feature of environmental monitoring as the firm is aware of the implications of a shorter time horizon, which suggests a negative ecological regime shift of a large magnitude. Therefore, this creates a sense of urgency about the possibility of the next regime change leading to resource extinction or collapse. Lastly, we define the risk of catastrophe as the situation in which the growth rate of the resource becomes negative, thus exhibiting a net tendency for the resource to reach extinction. We further distinguish between the scenarios of irreversible and reversible catastrophe, based on whether the firm can avert resource extinction by reducing or stopping extraction. This is done by studying the first passage time to catastrophe and the distribution of the extinction's expected hitting time. We find that this hitting time follows an inverse Gaussian distribution with larger magnitudes of regime shifts resulting in thicker tails.

A novel contribution of our model is assimilating the realistic feature of environmental monitoring of the resource to detect changes in the stock and structure: a practice which is very common in real-world resource management. For example Klemas (2013) talks about how remote sensing techniques, in near-real time, help detect changes that affect 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. In ecology, using real-time remote sensing data is increasingly common, especially with indicators of approaching thresholds or impending collapse in ecosystems.[4] Thus our model is especially relevant to understand how firms operate in the modern day resource market as considerable efforts are being invested in monitoring resource stocks as the number of ecological extreme events are on the rise. The use of quickest detection methods to capture monitoring allows us to also interpret our framework in real-time detection which we discuss in further detail in section 4.

Most of the literature concludes that the risk of a regime shift often leads to a precautionary behaviour. This is evident in the results shown by Polasky et al. (2011) and De Zeeuw and Zemel (2012), where an endogenous risk of regime shift leads to an optimal management that is always precautionary. However, this is not consistent with reality, as over-exploitation is a common phenomenon observed in renewable resources. Our paper extends on the existing literature of regime shifts to explain such firm behaviour by delineating the role of demand. We highlight the mechanism underpinning the often observed over-extraction of resources even when there is a visible negative change in the resource dynamics. Building on Ren and Polasky (2014), who numerically show the possibility of aggressive extraction, we emphasise that this behavior is possible under an endogenous regime shift specifically due to the presence of an elastic demand, declining market value of the resource and increasing monopoly power. Our model differs from most by showing that explicitly considering surveillance and knowledge of the magnitude of the regime shifts can in fact motivate the firm to intensify its extraction.

The remainder of the paper is structured as follows. Section **??** highlights some motivating empirical facts, in section 2 we lay out the different building blocks of the model. Section 3 describes profit maximization within a sequential framework and additionally we define the risk and first passage time to catastrophe. In section 4 we discuss how our model could be applied

[4]See Andersen et al. (2009); Porter et al. (2012); Batt et al. (2013); Carpenter et al. (2014); Scheffer et al. (2015)

in a situation where the firm monitors the resource process in real time and section 5 discusses the model solution and its economic implications. Section 6 concludes.

2 The Model

2.1 Resource Dynamics

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 $X_t \geq 0$[5]. 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 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 it undergoes a negative regime shift, and vice versa. 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 reality. We therefore 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 happening with certainty but at time unknown, the stochastic differential equation (SDE) driving the resource stock will switch between drifts, and the growth rate of the resource will change:

[5]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.

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

Here $\lambda(q_{ex}) > 0$ or $\lambda(q_{ex}) < 0$ allowing the effect of firm extraction on the resource growth to be both positive or negative. This implies that the firm's actions influence the magnitude of the change of regime. Note that since the occurrence of θ is certain, the question faced by the monopolist is not *if* a regime shift will occur but rather *when*. This framework seems appropriate for today, since the focus has moved from questions regarding the probability of the occurrences of collapses and regime shifts, to the question of when and how such occurrences will have to be dealt with.

The firm now faces two sources of uncertainty when choosing the harvesting policy that maximizes its profits. The first is the variance of the Gaussian noise source σ^2, which is the variation inherent to the natural randomness of environmental conditions: we choose the diffusion coefficient σ to be independent of the state X_t (i.e. a drifted Brownian motion) in order to include the possibility that the exogenous environmental shocks may drive the resource to extinction, something that log-normal fluctuations in a geometric Brownian motion by construction cannot represent. The second source is the *timing* θ of the shift, at which the resource's drift changes from μ to $\mu + \lambda(q_{ex})$.

2.2 Firm Dynamics

We consider a risk-neutral monopolist facing a linear inverse demand function of the form $p(q) = a - bq$, with a cost function defined as $cq + F$, where $a > c \geq 0,\ F \geq 0$ and $b > 0$.[6] 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

$$\Pi(q) = [(a - bq)q - cq - F] \tag{5}$$

2.3 Optimal Detection

The firm's problem now involves the *detection* of the change in drift of X_t, as seen in (4). The monopolist monitors the resource stock via sequential observations and uses quickest detection (QD) method to detect the regime change. This comprises of three elements: a stochastic process under observation (the evolution of the renewable resource), an unknown change point at which the statistical properties of the process undergo a change (a regime shift), and a decision maker observing the process and wants to detect this change (the monopolist). This method builds on change-point problems and extends it to the sequential framework where as long as the behavior

[6]Cost function of this form also allows us to flexibly model a natural monopoly since the average cost $AC = c + \frac{F}{q}$ is decreasing in output. We note that the choice of a quadratic cost function of the form cq^2 leaves the results qualitatively unaffected.

of the observations is consistent with the initial state, one is content to let the process continue. However, if at some unknown time the state changes, then the observer would like to detect it as soon as possible after its occurrence. This objective must be balanced with a desire to minimize false alarms. Such problems are known as quickest detection problems. To do this 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 regime in which it operates. This can be defined as:

$$\tau = \inf\{t \geq 0; \text{Detector}_t \geq \nu\}$$

Detector_t is a test statistic based on the sequential observations via monitoring of the resource and the value ν is a critical value or threshold, which provides the decision rule. The optimal QD procedure to determine Detector_t and ν are discussed below.[7]

In the period before 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

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

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

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

set strictly before the time of the initial condition on X (here normalized to 0, i.e. X_0).

The Brownian disorder problem was first studied by Shiryaev (1963) and the QD 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). This 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 $(\tau < \theta)$.

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

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

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 unless the diffusion parameter is very small 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 after the change occurred. 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.

The effective time period in which the firm optimizes is therefore between $t = 0$ and the final time given by the minimum between T and $\tau(\lambda(q_{ex}), \nu)$, the actual time at which the regime shift occurs plus the delay of detection. In other words, the firm programs 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 $(\mu + \lambda(q_{ex}))$. The "tolerance" T is chosen by the firm; however, $\tau(\lambda(q_{ex}, \nu)$ is a random variable. Since the firm knows the average delay time of detection, as given by (9), it can assume as time horizon the sum of the expectations of both change-point and delay, which is equivalent to taking a time interval $[0, \min\{T, \tau_c = \mathbb{E}[\theta] + \mathbb{E}[\tau(\lambda(q_{ex}), \nu)]\}]$. In the baseline detection case the firm has an uniform prior on the time of the regime shift: this implies that simply $\mathbb{E}[\theta] = T/2$.

3 Profit Maximization

The simplest way of modeling a regime shift is to assume that the shift occurs only once, however, as pointed by Sakamoto (2014), such ecological 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

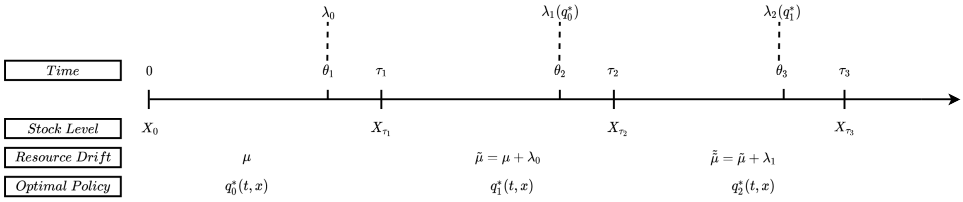

Figure 2: Sequential Detection

the middle to late 1990s (Chavez et al. (2003)). Within a multi-regime setting, in which the firm detects multiple regime changes throughout subsequent periods, the stochastic control problem of the firm will read:

$$\sup_{q \in Q} \quad \sum_{i=0}^{\infty} \mathbb{E}_{\tau_i} \int_{\tau_i}^{\tau_{i+1}} \Pi(q_t) e^{-\rho t} dt \tag{10}$$

$$dX_t = \begin{cases} \big(\mu + \lambda_i(q_{i-1}) - q_t\big)dt + \sigma dW_t, & t \in [\tau_i, \tau_{i+1}) \\ \big(\mu + \lambda_{i+1}(q_i) - q_t\big)dt + \sigma dW_t, & t \geq \tau_{i+1},\ i \in \mathbb{N}. \end{cases}$$

$$X_t \geq 0 \quad \forall t$$

$$\tau_i = \min\{T\ ,\ \mathbb{E}[\theta] + \mathbb{E}[\tau(\lambda_i), \nu)]\}$$

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 periods and relative changes in resource growth. We assume $\tau_0 = 0$ for simplicity. Here 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[8]. To analyse the firm's optimization problem in a sequential detection scenario we work through the schematic representation seen in Figure 2.

3.1 Regime $[0, \tau_1]$

At time $t = 0$ the firm believes that the resource is driven by a diffusion process with the natural growth rate μ and and begins harvesting activity at level $q^*(0, x_0)$. At a random time $\theta_1 \in [0, T]$, there is an initial exogenous change, $\lambda_0 < 0$, in the resource dynamics.[9] Until the detection of this change, the firm operates in an environment where the resource evolves according to the process

$$dX_t = (\mu - q_0^*(t, X_t))dt + \sigma dW_t, \quad t \in [0, \tau(\lambda_0, \nu)], \tag{11}$$

where $\tau(\lambda_0, \nu) \leq T$ is the detection time. The final time of the period which the firm uses as a reference for its decisions is given by

[8]The explicit dependence of the stopping time τ on λ makes the control variable q and the limit of integration τ_1 simultaneous, and the model becomes intractable. In order to circumvent this issue, we model the firm to detect a change in drift $\lambda(q_{ex})$ which is determined by extraction in the *previous* period

[9]The first change is exogenous so as to start the process of subsequent adjustment.

$$\tau_1 = \mathbb{E}[\theta] + \mathbb{E}[\tau(\lambda_0, \nu)] = \mathbb{E}[\theta] + \frac{2}{\lambda_0^2}\left(e^{-\nu} + \nu - 1\right) \tag{12}$$

where the threshold ν solves $\frac{2}{\lambda_0^2}\left(e^{\nu} - \nu - 1\right) = T$. Within this time interval $[0, \tau_1]$, the value of the firm is given by

$$\begin{aligned} V(0, X_0) &= \sup_{q \in Q} \mathbb{E}_0 \int_0^{\tau_1} \Pi(q) e^{-\rho t} dt \\ &\text{s.t. } dX_t = (\mu - q)dt + \sigma dW_t, \\ &\quad X_t \geq 0. \end{aligned} \tag{13}$$

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

$$0 = V_t - \rho V + \max_{q \in Q} \{(a - bq)q - cq - F - qV_x\} + \mu V_x + \frac{\sigma^2}{2} V_{xx}, \quad X_t \geq 0. \tag{14}$$

where Q is the set of admissible Markov controls for which $q^* \geq 0, X^*(t, q^*) \geq 0$.[10] Once solved, this problem will yield a control in the feedback form $q(t, X_t)$. 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 A that the value function V is a weak solution of the optimization problem (14), and once we obtain a solution for V we can conclude it will solve the firm's problem (in a weak sense).

Equation (14) implies an optimal extraction policy given by

$$q^*(t, X_t) = \left[\frac{a - c - V_x}{2b}\right]_+ . \tag{15}$$

Note that this implies that in order for extraction to stay positive, $V_x \leq a - c$, meaning the resource rent cannot exceed the demand intercept parameter minus the marginal cost. 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 (14) and grouping terms, we obtain the following partial differential equation:

$$0 = V_t - \rho V + AV_x + BV_x^2 + \frac{\sigma^2}{2} V_{xx} + C \tag{16}$$

where the constants A, B and C are given by

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

$$
\begin{aligned}
A &= \mu - \frac{a-c}{2b}, \\
B &= \frac{1}{4b}, \\
C &= \frac{(a-c)^2}{4b} - F.
\end{aligned}
$$

The natural boundary conditions of this problem are given by

$$V(t,x) = 0 \text{ for } x < 0,\ V(t,0) = 0,\ q(t,0) = 0 \tag{17}$$

without imposing a smooth pasting condition because of the viscosity argument.

Because of the homogeneous form of the profit function, 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

$$V'(x) = \frac{\sigma^2}{2B} \frac{\psi'(x)}{\psi(x)}$$

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

$$\psi_g(x) = c_1 e^{\alpha_1 x} + c_2 e^{\alpha_2 x}. \tag{18}$$

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 (17), 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

$$q_0^*(t,x) = q^m - \sigma^2 \frac{\psi'(t,x)}{\psi(t,x)} e^{-\rho(\tau_1 - t)}, \tag{19}$$

where $q^m = \frac{a-c}{2b}$. From (15) we also obtain the resource rent for the monopolist:

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

The "instantaneous" drift of the optimally controlled stock is given by

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

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

$$X^*_{\tau_1} = X_0 + [\mu - q^m]\,\tau_1 + \sigma^2 \int_0^{\tau_1} \frac{\psi'(X_t)}{\psi(X_t)} e^{-\rho(\tau_1 - t)} dt + \sigma \int_0^{\tau_1} dW_t. \tag{21}$$

where the second integral is to be interpreted in the Itô sense and is simply equal to the Gaussian draw $N(0, \sigma^2\tau_1)$. Similar to (19), 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.

3.2 Regime $[\tau_1, \tau_2]$

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

$$dX_t = (\mu - \lambda_0 - q^*_1(t, X_t))dt + \sigma dW_t, \quad t \in [\tau_1, \tau_2]$$

The optimal policy for this period is easily seen to have the same form as (19). Normalizing time to $t_0 = \tau_1$, one recognizes that the two problems are equivalent, with a change in drift from μ to $\mu - \lambda_0$. We therefore have

$$q^*_1(t, x) = q^m - \sigma^2 \frac{\tilde{\psi}'(t, x)}{\tilde{\psi}(t, x)} e^{-\rho(\tau_2 - t)}, \tag{22}$$

where the exponents $\tilde{\psi}(t, x)$ include the new drift in the coefficient A. In the meantime, however, while the firm assumes a constant λ_0, its past decisions start to catch up. At a random time θ_2 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

$$\lambda_1(q^*_0) = \begin{cases} \mu \, \Delta X_0 & \Delta X_0 < 0 \\ \mu \, \sqrt{\Delta X_0} & \Delta X_0 > 0 \end{cases} \tag{23}$$

$$\text{where} \quad \Delta X_0 = \frac{X^*_{\tau_1} - X_0}{X_0} \tag{24}$$

Equations (23) and (24) 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 the net change has on the resource's capacity to regenerate is not symmetric. When $\lambda_1(q_0^*)$ is negative, it has a linear impact on the growth rate. However when $\lambda_1(q_0^*)$ is positive, the effect is concave. This is to capture the fact that an ecosystem is more vulnerable to negative shocks. The sign and magnitude of this regime shift is assumed known by the firm, but *when* it occurs is uncertain and to be detected. At $\tau_2(\lambda_1(q_0^*), \nu, T)$ the monopolist will (on average) detect this change in regime of the resource dynamics. Similar to (21), the resource stock at τ_2 will be

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

3.3 Risk of Catastrophe $[\tau_2, \tau_3]$

We now illustrate the emergence of catastrophe risk. After 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[11]

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

Let us suppose that $(\mu - \lambda_0 - \lambda_1(q_0^*)) > 0$ 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 and if $\lambda_2 < 0$, the firm will realize the future emergence of catastrophe risk if

$$\mu + \sum_{j=0}^{i=2} \lambda_j < 0,$$

noting that at the next detection time τ_3 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 $(X = 0)$.

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 $t \in [\tau_i, \tau_{i+1}]$:

$$\mu + \sum_{j=0}^{i-1} \lambda_j - q^m + \sigma^2 \int_{\tau_i}^{t} \frac{\tilde{\psi}'(X_s)}{\tilde{\psi}(X_s)} e^{-\rho(\tau_{i+1}-s)} ds < 0, \tag{25}$$

which implies that $P(\lim_{t\to\infty} X_t = 0) = 1$.

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

[11] λ_1 can be positive as well but for exposition we assume a negative regime shift.

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

$$\phi(x) = \Pr\left[\inf_{t\in\mathbb{R}^+} X_t \leq 0 \middle| X_0 = X^*_{\tau_i}, q^*(t, X_t) = 0\right] \tag{26}$$

and the first time to catastrophe as

$$\tau_c = \inf[\ t | X_t \leq 0, X_0 = X^*_{\tau_i}, q^*(t, X_t) = 0]. \tag{27}$$

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

$$\mathbb{E}\tau_c = \frac{X^*_{\tau_i}}{\left|\mu + \sum_{j=0}^{i} \lambda_j\right|}. \tag{28}$$

and the probability of extinction is

$$\phi(x) = \exp\left(-\frac{2\left(|\mu + \sum_{j=0}^{i} \lambda_j|\right)}{\sigma^2} x\right). \tag{29}$$

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, (28) 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 the firm stops extracting, then the first passage time to catastrophe τ_{cat}, for a resource stock starting at $X^*_{\tau_i}$, will be distributed according to the following density:

$$\begin{aligned} \mathrm{P}\{\tau_{cat} \in dt\} &= \frac{X^*_{\tau_i}}{\sqrt{2\pi\sigma^2 t^3}} \exp\left(-\frac{(X^*_{\tau_i} - (\mu + \sum_{j=0}^{i} \lambda_j)t)^2}{2\sigma^2 t}\right) dt, \\ &= IG\left(\left|\frac{X^*_{\tau_i}}{\mu + \sum_{j=0}^{i} \lambda_j}\right|, \left(\frac{X^*_{\tau_i}}{\sigma}\right)^2\right), \end{aligned} \tag{30}$$

which follows an inverse Gaussian distribution, as seen in Figure 3.

Because of the stochastic fluctuations, the firm cannot know with certainty whether the first

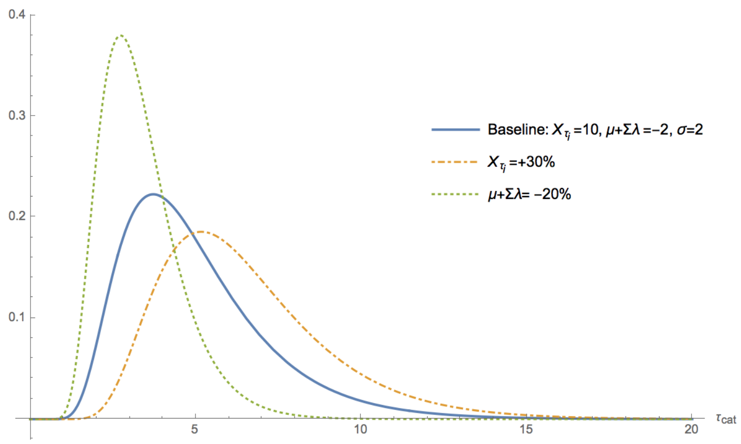

Figure 3: Distribution of the time to catastrophe and effect of a higher initial level of stock (dot-dashed) and of a larger regime shift magnitude (dashed).

passage time will happen before the next regime change, but it can have an average measurement of its probability. If $\mathbb{E}\tau_c \geq \tau_{i+1}$, so if $X_{\tau_i} > \mu\tau_{i+1}$, 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

$$1 - \psi(t) = 1 - \phi(0, t), \tag{31}$$

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

$$\frac{\partial}{\partial t}\phi(x,t) + \frac{\partial}{\partial x}\phi(x,t)\left(\mu + \sum_{j=1}^{i-1}\lambda_j - q^*(t,x)\right) + \frac{\sigma^2}{2}\frac{\partial^2}{\partial x^2}\phi(x,t) = 0 \tag{32}$$

with absorbing boundary conditions given by

$$\begin{cases} \phi(x, \tau_{i+1}) = 1 & x \leq 0 \\ \phi(x, \tau_{i+1}) = 0 & x > 0, \end{cases} \tag{33}$$
$$\phi(0, t) = 1,$$
$$\phi(t, \infty) = 0.$$

The KFE for this problem has no closed form solution, given the dependence of the extraction policy on both x and t, 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 (31) and compute its numerical first moment: if $\mathbb{E}\tau_c \geq \tau_{i+1}$ the firm continues its optimal extraction policy.

3.4 Discussion - General Solutions for Period $[\tau_i, \tau_{i+1}]$

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

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

the resource rent as:

$$V_x(t, x) = q^v(t, x) 2b \tag{35}$$

the optimally controlled stock at the time of detection as:

$$X^*_{\tau_i} = X^*_{\tau_{i-1}} + \left[\mu + \sum_{j=0}^{i-1} \lambda_j - \lambda_{i-1} - q^m \right] \tau_i + \sigma^2 \int_{\tau_{i-1}}^{\tau_{i-1} + \tau_i} \frac{\tilde{\psi}'(X_t)}{\tilde{\psi}(X_t)} e^{-\rho(\tau_i - t)} dt + \sigma \int_{\tau_{i-1}}^{\tau_{i-1} + \tau_i} dW_t. \tag{36}$$

the change in the growth of the resource dependent on its past harvesting actions as:[12]

[12]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.

$$\lambda_i = \begin{cases} \left(\mu + \sum_{j=0}^{i-1} \lambda_j - \lambda_{i-1}\right) \frac{X^*_{\tau_i} - X^*_{\tau_{i-1}}}{X^*_{\tau_{i-1}}} & \Delta X^*_{\tau_{i-1}} < 0 \\ \\ \left(\mu + \sum_{j=0}^{i-1} \lambda_j - \lambda_{i-1}\right) \sqrt{\frac{X^*_{\tau_i} - X^*_{\tau_{i-1}}}{X^*_{\tau_{i-1}}}} & \Delta X^*_{\tau_{i-1}} > 0 \\ \end{cases} \tag{37}$$

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

$$\mathbb{E}[\tau_{i+1}] = \mathbb{E}[\theta] + \mathbb{E}[\tau(\lambda_i(q^*_{i-1}), \nu, T)] \tag{38}$$

The optimal extraction (34) consists of two parts: one driven purely by market preferences as seen in $q^m = \frac{a-c}{2b}$. This 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 $\sigma = 0$). The second part not only consists of market preferences but is variable and explicitly dependent on state X_t and modulated by the distance between present and the detection time, representing the time horizon of the firm. 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. Note that when the rent of the resource rises, q^* decreases[13]. Observe that in each period the final level of resource, $X(t)$, 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 $X(t)$ is conserved in the *magnitude*, the absolute value of the percentage change in the resource stock translates directly to a change in drift. This is observed in Figure 4 where we show ten simulated time paths of an the optimally controlled stock of resource. Here the first detection time is common to all but subsequent detections are extraction-dependent.

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

[13]The optimal harvesting function exhibits a sigmoid-like form. 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. This results in the following limiting behavior:

$$\begin{aligned} \lim_{x\to\infty} q^*(t,x) &= q^m - \nu(\Theta)e^{-\rho(\tau_1 - t)}, \\ \lim_{\tau_c\to\infty} q^*(t,x) &= q^m. \end{aligned}$$

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 and the time horizon of the firm. If this horizon is long enough, all resource-related parameters are ignored and the monopolist's optimal harvest is entirely driven by the market.

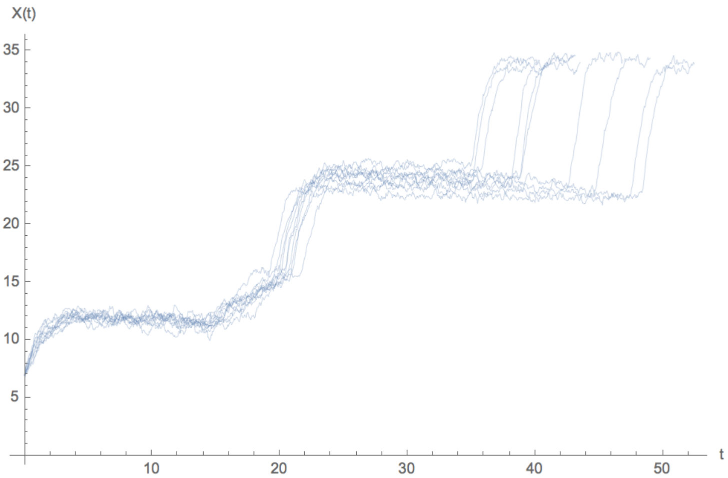

Figure 4: Simulation of 10 time paths with same parameters resulting in different extraction-dependent detection times. All simulations are done with a Shoji-Ozaki discretization method for the time-dependent drift.

as a result optimal harvest must be specified for every state that can possibly occur. Additionally, in a multi-regime setting, the detection of each regime shift alters the monopolist's time horizon. The larger is the difference between the initial and final 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 the firm detects the regime shift, the magnitude of change in the resource growth 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. This is evident in panel (a) of Figure 5, 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.

4 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 $\tau_{i+1} = \min\{T, \mathbb{E}[\theta] + \mathbb{E}[\tau(-\lambda, \nu)]\}$. This is therefore an *ex ante* policy: the actual detection of when the regime shift happens is only represented via a first-order stochastic criterion. The time θ at which the regime changes, however, is a random variable: the firm therefore will use the expected detection time (12) 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 thresh-

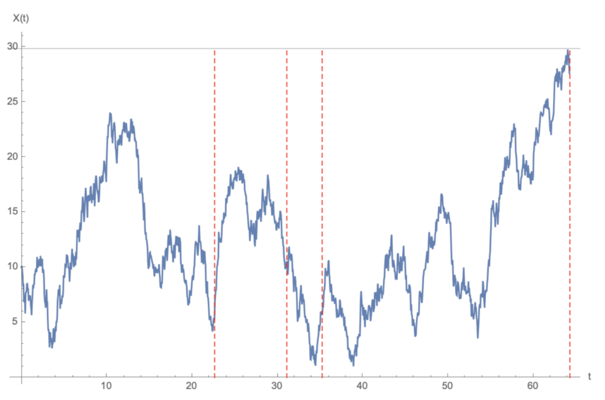

(a)

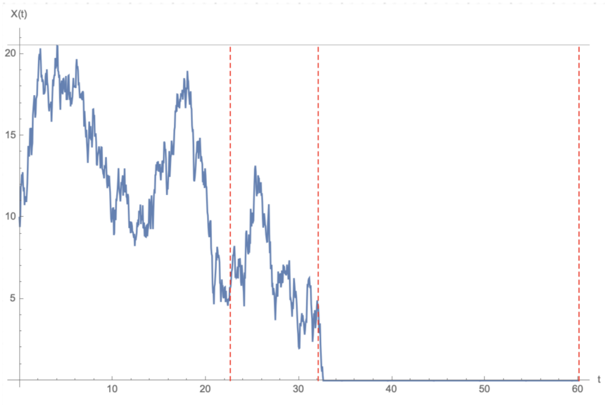

(b)

Figure 5: Simulated Time Paths of an optimally controlled biomass/stock. Red dashed lines represent detection times.

old 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, implying that the regime has not yet shifted, the firm continues the optimal extraction assuming the same underlying resource dynamics, but now in the an infinitesimal time interval as horizon. In other words, the infinitesimal optimal extraction policy if the expected detection time is exceeded is given by

$$q_i^*(t,x) = q^m - \sigma^2 \frac{\psi'(t,x)}{\psi(t,x)} e^{-\rho(T-t)}, \qquad t \in [\tau_{i+1}, T], \tag{39}$$

until either the Radon-Nikodym derivative of the measures of the two regimes (under the measure Q by which X_t is a $Q-$Brownian motion) reaches the threshold ν, or until the firm's tolerance time limit T is reached.

This notion of observation under measure changes might appear as a mathematical abstraction: we note, however, that the disorder problem (6) based on the observation of the resource stock X_t is equivalent in probability to the disorder problem based on the observation of the *residual process* given by

$$dY_t = \frac{1}{\sigma}\left(dX_t - \mu + \sum_{j=1}^{i-1}\lambda_j - q^*(t,X_t)\right)dt \tag{40}$$

under the original measure. In other words, the firm can detect the change by either observing the resource stock and changing measure appropriately, or by extracting residuals from the stock variation, the growth rate and the optimal extraction policy and then studying the original P-Brownian motion. The computational difference between the two is marginal if the extraction policy is of simple form, such as the constant part of the extraction sigmoid such that the resulting controlled resource stock effectively remains Gaussian, but the second strategy could be of substantially simpler implementation for when the extraction policy is in its nonlinear part. If real-time observations are not continuous but rather arrive at discrete times $t_i, i \in \mathbb{N}$, and assuming a constant frequency between times Δt, then the residual process on which the firm has to apply the detection procedure is the stationary process $X_{t_i} - X_{t_i-\Delta t} - (\mu + \sum_{j=1}^{i-1}\lambda_j - q^*(t, X_{t-\Delta t}))\Delta t$ (after standardization of the diffusive part).

5 Characteristics of the Solution

With (34) and (35), we can now examine how a change in regime affects the firm's extraction decisions post detection. We find:

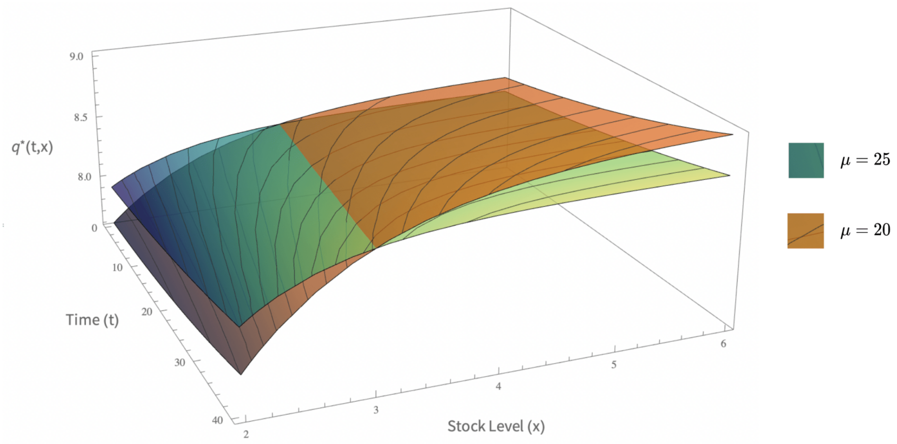

Figure 6: Optimal extraction policies for a monopolist with a demand function of the form $p(q) = 15 - 0.75q$ and cost function $3q + 30$. The intrinsic growth of the resource is $\mu = 25$, variance $\sigma = 5$. The first regime shift is detected at $\tau_1 = 40$, when the drift changes from $25 \longrightarrow 20$. The second regime shift is detected at $\tau_2 = 40$.

$$\underbrace{\frac{\partial q^*(t,x)}{\partial \mu} < 0}_{\text{Aggressive}} \quad \text{or} \quad \underbrace{\frac{\partial q^*(t,x)}{\partial \mu} > 0}_{\text{Precautionary}} \tag{41}$$

To illustrate this we choose a range of values for the model parameters which are meant to be largely illustrative. The firm incurs a variable cost with $c = 3$, fixed cost of 30 and applies a discount rate of $\rho = 0.02$. The resource has an intrinsic growth of $\mu = 25$ and $\sigma = 5$. We focus on the case of a negative regime shift as it is of more interest and relevance today. Suppose the ecological system undergoes a regime shift of magnitude $\lambda_0 = -5$ resulting in a modified drift $\tilde{\mu} = \mu + \lambda_0 = 20$. In Figure 6, the green shaded regions depict the evolution of the firm's optimal extraction policies up until it detects this regime shift at $\mathbb{E}[\tau_1] = 40$. Therefore, within the first interval $[0, \tau_1]$, the extraction levels reflect the firm's assumption that the resource is growing at its natural rate of growth as shown in (19). Once this regime is detected, the firm updates its assessment and maximizes its profits with respect to the new drift $\tilde{\mu}$. Given the current period's extraction policy q_0^*, the next regime shift is of a similar magnitude and the firm will on average detect it at the same time, meaning an $\mathbb{E}[\tau_2] = 40$.

Before we discuss how the monopolist changes its extraction after a regime shift, we characterize more precisely the nature of the firm's extraction policy. An ***aggressive*** extraction strategy is one where, for all else equal, at the level of optimally controlled stock observed at the detection of the regime shift, there exists a time in the new regime where $q_1^*(t, X_{\tau_1}^*) > q_0^*(\tau_1, X_{\tau_1}^*)$ where q_0 is the extraction policy for the parameters of the interval $[0, \tau_1]$, q_1 is for the post-regime shift interval $[\tau_1, \tau_2]$ and $t \in [\tau_1, \tau_2]$. Since q^* is monotonically increasing in time, this implies that at the subsequent expected time of detection τ_2 there exists an optimally controlled stock level at which $q_1^*(\tau_2, X_{\tau_1}^*) = q_0^*(\tau_1, X_{\tau_1}^*)$, and we call it the threshold X_{th}^*. If $X_{\tau_1}^* > X_{th}^*$ then there exists a time $\bar{t} \in [\tau_1, \tau_2]$ where the firm switches to an aggressive extraction policy, given by:

$$\bar{t} = \frac{1}{\rho} \ln\left(\frac{\psi_0{}'(X^*_{\tau_1})}{\psi_0(X^*_{\tau_1})} \frac{\psi_1(X^*_{\tau_1})}{\psi_1{}'(X^*_{\tau_1})} \right) + \tau_2 \tag{42}$$

If $X^*_{\tau_1} < X^*_{th}$ then $\bar{t} \notin [\tau_1, \tau_2]$ and the monopolist will never increase extraction in the new regime with respect to its past extraction policy, thus adopting a ***precautionary*** strategy where $q^*_1(t, X^*_{\tau_1}) < q^*_0(\tau_1, X^*_{\tau_1})$ for all $t \in [\tau_1, \tau_2]$.

For the parameters given in Figure 6 we find $X^*_{th} = 3.04$ and for stock levels below this threshold a precautionary behaviour is observed. For levels greater than X^*_{th} we find that although after the immediate detection of the regime change the firm may reduce its extraction, over the course of the new regime it continues to increase its extraction eventually outpacing the levels extracted in the old regime therefore adopting an aggressive approach.

5.1 Role of the market

To understand the above result we look at the dynamics of the resource rent (35), specifically how does the slope of rent change with respect to change in stock. We find:

$$\frac{\partial V_x(x,t)}{\partial x} < 0 \quad \text{if} \quad \underbrace{\frac{F}{a-c}}_{\text{Break Even}} \; e^{-\rho(\tau - t)} < \mu \quad \longrightarrow \quad \text{Scarcity Effect} \tag{43}$$

Proof. See Appendix B

The term $\frac{F}{a-c}$ is fixed costs divided by the maximum price a consumer is willing to pay per unit minus the variable cost. This can be interpreted as the amount of resource stock required by the monopolist to break even. Therefore, as long as the drift of the resource is greater than the discounted break even amount required over the detection horizon, a scarcity of the resource stock will lead to an increase in resource rent. Therefore, when a negative regime shift occurs it reduces the growth rate of the resource and creates a physical scarcity of the stock, which in turn increases the market value of the marginal unit of *in situ* stock leading to decreased levels of extraction by the monopolist, explaining the precautionary behaviour.

At higher stocks, if demand is elastic, the scarcity effect is outweighed and the resource rent falls. Looking at Figure 7 we find that the slope of the resource rent function with respect to the drift is negative at low levels of the stock and becomes positive at high levels:

$$\frac{\partial V_x(x,t)}{\partial \mu} < 0 \quad \text{or} \quad \frac{\partial V_x(x,t)}{\partial \mu} > 0 \tag{44}$$

Moreover, as the monopolist is making intertemporal pricing and production decisions: when the stock levels are high, the resource rent is *negative*. This is because the firm is essentially

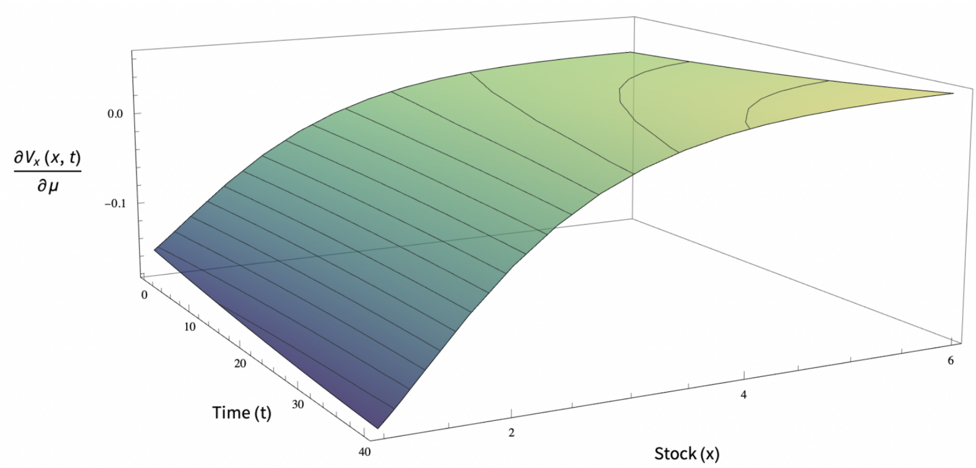

Figure 7: The slope of the resource rent function with respect to the drift.

moving down a "learning curve": that is, as they produce, learning-by-doing reduces their average and marginal costs (Pindyck (1985)). In this case, the full marginal cost (FMC) of current production $c+V_x$ is less than current marginal production cost c. The reason is that an incremental unit of current production reduces future production costs by moving the farther down the learning curve, so that production of the unit brings a benefit (a negative rent) that partly offsets its cost.

Therefore, at higher stocks, when there is a negative regime shift, the full marginal cost of production falls even further leading to increased extraction. Additionally, expressing the price set by the monopolist as in (45) we see that for levels greater than X^*_{th} increasing extraction allows the firm to charge a higher markup due to the presence of an elastic market demand E_d, as seen in Figure 8.

$$p(q^*) = \underbrace{\left(\frac{1}{1-\frac{1}{E_d}}\right)}_{\text{Markup}\uparrow} \underbrace{(c+V_x)}_{\text{FMC}\downarrow}, \qquad E_d = \frac{a}{bq^*(x,t)} - 1, \tag{45}$$

5.2 Role of the magnitude of regime shift and the detection time

The size of the regime shift λ plays an important role in the firm's extraction policy. From (9) we know that the larger is the change in drift, the "earlier" is the expected time of detection. If λ is very small then the firm may wait for much longer before detecting a change of regime. As the monopolist maximizes its profits over the expected detection time, the magnitude of λ directly influences the decision or the time horizon of the firm. Figure 9 presents two cases where the market preferences and resource dynamics observed by the firm are the same and the only difference is in λ. Panel (A) depicts a regime shift of a magnitude $\lambda_0^A = -5$ with the monopolist on average detecting this at $\mathbb{E}[\tau_1^A] = 40$. Therefore the firm maximizes its profits with respect to the resource's intrinsic growth rate $\mu = 25$ in the first interval $[0, \tau_1^A]$ over a

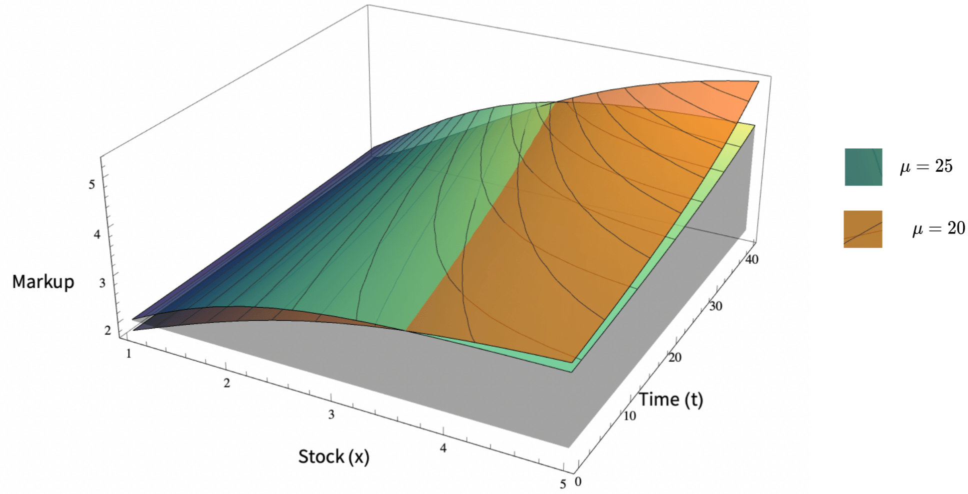

Figure 8: The price markup charged by the Monopolist..

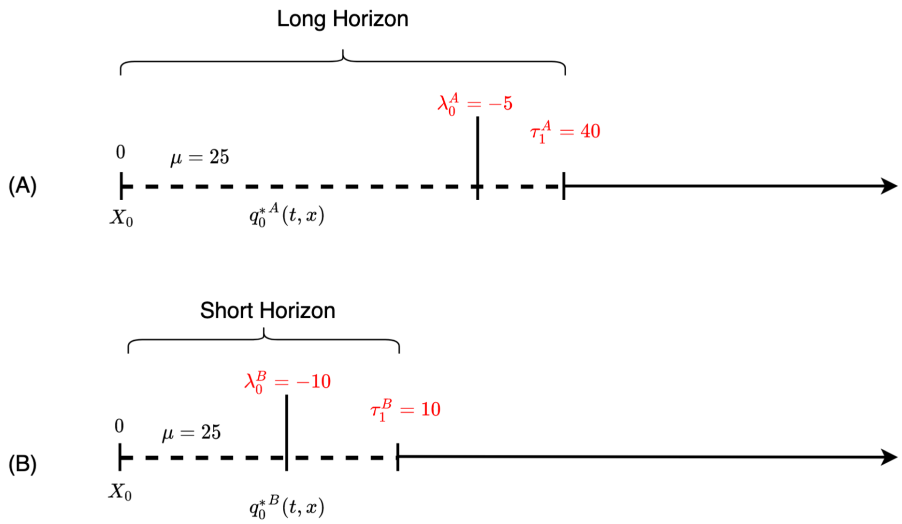

Figure 9: Regime shifts of different magnitudes

horizon of 40 months. In panel (B) the resource undergoes a large regime shift of magnitude $\lambda_0^B = -10$ resulting in a much quicker expected time of detection at $\mathbb{E}[\tau_1^B] = 10$. Accordingly the monopolist maximizes its profits over a much shorter horizon of only 10 months. Studying how the extraction $q_0^{*A}(t,x)$ differs from $q_0^{*B}(t,x)$ will emphasise the role played by the detection time.

To answer this we examine the change in slope of the optimal extraction function with respect to change in the expected time of detection, seen in Figure 10. This can be written in closed form but because of the form of the boundary conditions of the problem the expression is very cumbersome and therefore we prefer to resort to numerical simulations. We note that although at very low levels of stock it is positive, at high levels the slope in fact becomes negative. Thus in the event of a large regime shift as λ_0^B, at higher stock levels, the firm increases its extraction for all x as compared to when regime of magnitude λ_0^A occurs. This implies that

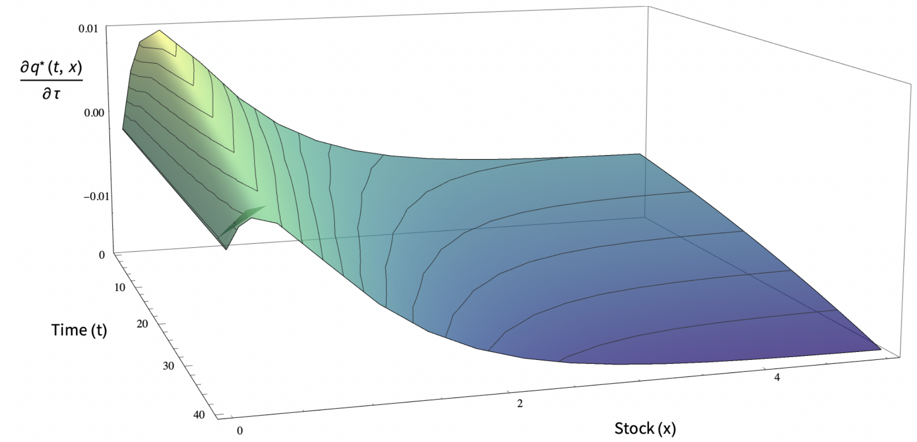

Figure 10: Change in slope of the optimal extraction function $q^*(t,x)$ with respect to change in the expected time of detection τ.

for higher stocks ${q_0^*}^B(t,x) > {q_0^*}^A(t,x)$. The escalation of extraction can be explained due to the monopolist believing that another shift in regime could happen very soon and a resource extinction or collapse may be impending thus creating a sense of urgency. The total effect of the detection delay on extraction policy is therefore the *ex ante* difference between overall policies, defined as the time integral of the optimal policy over the period between initial and detection time. This quantity is by construction a random variable, since it will be function of observed resource levels. Using Itô calculus and stochastic integration by parts we obtain

$$\begin{aligned}\left(\int_0^{\tau_2} q^*(x,t) \right. &- \left. \int_0^{\tau_1} q^*(x,t)\right) dt = \\ &= q^m(\tau_2-\tau_1) + \frac{1}{2b\rho}\left[V_x^{\tau_2}(x_0,0) - V_x^{\tau_1}(x_0,0) + \int_{\tau_1}^{\tau_2} \partial_x V_x dX_t\right], \qquad (46)\end{aligned}$$

where $V_x^{\tau_i}(x_0,0)$ is the resource rent for a firm facing an expected detection time τ_i starting at the beginning of the period with an observed level of stock x_0, and the last is an Itô integral. Equation (46) shows that the effect of a regime shift depends directly via the difference in time horizon, via the linear effect on the fixed extraction amount q^m times the difference in time the firm has available for extraction. There is an *ex ante* effect related to the difference in initial resource rent faced by the firm with different horizons. Third, there is the resource-based effect given by the integral term, which shows that the effect on the extraction policy depends on the overall variation of resource rent the rent will face in the "extra" time between the two detections $\tau_2 - \tau_1$, evaluated over all possible resource trajectories. This also tells us how the past actions, leading to the regime shift that is to be detected next, affect the firm's current extractive decisions.

5.3 Profit loss due to detection delay

The profit loss due to the delay in detection is the key element behind why a firm would want to implement such detection techniques in order to adjust production to the natural regime shifts in the quickest time possible. The reason behind this profit loss lies in the fact that between in the time interval $[\theta, \tau_c]$, the time that passes between the *actual* change in regime and when the firm detects this change and adjusts extraction, the firm is *de facto* extracting a "wrong" quantity, which would be optimal for the pre-regime shift dynamics of the resource but is suboptimal for the post-shift ones. Assume for simplicity that the regime shift occurs in the dynamics of X_t and changes μ to $\mu + \lambda$, where $\lambda \in \mathbb{R}$. This implies that there exists an extraction policy $q^\lambda := q^\lambda(x,t)$ in the time interval $[\theta, \tau_c]$ that achieves the supremum of the discounted profit function, i.e. $q^\lambda \to \sup_q \int_\theta^{\tau_c} e^{-\rho(t-\tau_c)}\Pi(q)dt$ s.t. $dX_t = (\mu + \lambda - q)dt + \sigma dW_t$. The firm, however, in this interval will continue to extract according to the policy $q := q(x,t)$ that achieves the supremum of the optimization problem constrained by the "wrong" resource dynamics $dX_t = (m - q)dt + \sigma dW_t$. The firm will therefore incur in a nonzero (discounted) profit loss due to the detection delay:

$$\int_\theta^{\tau_c} e^{-\rho t}\Big(\Pi(q^\lambda) - \Pi(q)\Big)\,dt := \int_0^{\tau} e^{-\rho(\tau - t)}L(X_t,t)dt > 0,$$

where $\Pi(q^\lambda)$ represents the "theoretical" optimal profits in $[0, \tau_c]$ obtained by switching to the new regime immediately, and $\Pi(q)$ the ones obtained by switching after the delay.

The proof that this loss exists and is positive is simple. Because of the delay, the firm chooses an extraction $q \to \sup_q \int_0^{\tau_c} e^{-\rho(t)}\Pi(q)dt$ s.t. $dX_t = (\mu - q)dt + \sigma dW_t$. Define the set of maximized profits $\bar{\Pi}^*(q,t)$ as the supremum of the maximization problem (i.e. the total volume of profits) achieved with policy q over a period t, which is a nonempty set of real numbers bounded above and below. The overall "real" supremum of the maximization problem is achieved by the following:

$$\left(\sup_q \int_0^\theta e^{-\rho t}\Pi(q)dt\right) + \left(\sup_q \int_\theta^{\tau_c} e^{-\rho t}\Pi(q^\lambda)dt\right) := \bar{\Pi}^*(q,\theta) + \bar{\Pi}^*(q^\lambda, \tau_c - \theta)$$

where $\bar{\Pi}^*(q,\theta)$ is the supremum set of the problem up to θ under the constraint with drift $\mu - q$ generated by the optimal policy q, and $\bar{\Pi}^*(q^\lambda, \tau_c - \theta)$ is the supremum set of the problem between θ and τ_c under the constraint with drift $\mu + \lambda - q^\lambda$ generated by q^λ. Due to the additive property of the supremum over bounded nonempty sets, this can be rewritten as

$$
\begin{aligned}
\bar{\Pi}^*(q^*,\tau_c) &= \sup_q \int_0^{\tau_c} e^{-\rho t}\Pi(q)dt \\
& \quad \text{s.t. } dX_t = \begin{cases} (\mu - q)dt + \sigma dW_t & t < \theta \quad (dX_t) \\ (\mu + \lambda - q)dt + \sigma dW_t & t \geq \theta \quad (dX_t^\lambda) \end{cases} \\
&= \sup_q \left(\underbrace{\int_0^{\theta} e^{-\rho t}\Pi(q)dt}_{\text{s.t. } dX_t} + \underbrace{\int_\theta^{\tau_c} e^{-\rho t}\Pi(q)dt}_{\text{s.t. } dX_t^\lambda} \right) \\
&= \sup_q \left(\underbrace{\int_0^{\theta} e^{-\rho t}\Pi(q)dt}_{\text{s.t. } dX_t} \right) + \sup_q \left(\underbrace{\int_\theta^{\tau_c} e^{-\rho t}\Pi(q)dt}_{\text{s.t. } dX_t^\lambda} \right) \\
&= \bar{\Pi}^*(q,\theta) + \Pi^*(q^\lambda, \tau_c - \theta)
\end{aligned}
$$

where q^* is the optimal policy that switches from q to q^λ exactly at θ. This implies that between θ and τ_c any admissible policy $\tilde{q} \in Q, q \neq q^\lambda$ will not achieve the supremum, and $\bar{\Pi}(\tilde{q}, \tau_c - \theta) < \bar{\Pi}^*(q^\lambda, \tau_c - \theta)$. This implies that

$$
\int_\theta^{\tau_c} e^{-\rho t} \left(\Pi(q^\lambda) - \Pi(q)\right) dt = \bar{\Pi}^*(q^\lambda, \tau) - \bar{\Pi}(q,\tau) > 0,
$$

and the detection delay induces a loss for the monopolist which is increasing in the length of the delay τ itself.

The instantaneous loss function L_t is a function of both stock X_t and time t. Omitting arguments for clarity, its behavior in the interval $t \in [0,\tau]$ can be obtained by defined by the stochastic differential

$$
\begin{aligned}
dL(X_t,t) &= \left[V_x^\lambda(\mathcal{A}^\lambda q^\lambda) - V_x(\mathcal{A}^\lambda q) - B\sigma^2\left(V_{xx}^{\lambda\,2} - V_{xx}^2\right)\right]dt + \sigma\left(V_x^\lambda q_x^\lambda - V_x^\lambda q_x\right)dW_t \quad (47)\\
&= A(X_t,t)dt + B(X_t,t)dW_t,
\end{aligned}
$$

obtained by standard Itô calculus methods and where the subscripts in the drift and diffusion coefficients indicate partial derivatives. The term $\mathcal{A}^\lambda$ is the infinitesimal generator of the controlled resource stock given by

$$
\mathcal{A}^\lambda\phi := \mathcal{A}^\lambda\phi(x,t) = (\mu + \lambda - \phi(x,t))\phi_x(x,t) + \frac{\sigma^2}{2}\phi_{xx}(x,t) + \phi_t(x,t),
$$

and all the extraction policy terms q have the λ exponent in order to represent whether the policy is evaluated at the post-shift drift $\mu + \lambda$ or not. Lastly, the terms V_x^λ, V_x indicate the resource rents as given by (35), evaluated at the respective extraction policies q^λ, q. This equation has an intuitive interpretation: the deterministic part of the instantaneous loss evolves according to two differential terms. The first is the difference between the instantaneous expected change in extraction $\mathcal{A}^\lambda q = \mathbb{E}dq$ of the "theoretical" extraction policy q^λ with the suboptimal policy generated by the detection delay τ, expressed in units of resource rent V_x. The inefficiency is generated by the fact that the policy q is applied to a resource stock that evolves according to a process that is *not* what the firm considers in its optimization. i.e. $\mathcal{A}^\lambda q$. Intuitively, this represents how the firm applies the "wrong" extraction policy in the interval $[\theta, \tau_c]$: the chosen policy q is optimal for a resource stock that grows at a rate μ, and is applied to a resource stock that however grows at the post-shift rate $\mu + \lambda$. This is the first of the sources of the firm's profit losses, and it represents the loss of profits derivating from the fact that the firm cannot observe the regime shift during the detection delay, and thus effectively optimizes with respect to the "wrong" dynamics. The second source of loss stems from the difference between the squared sensitivities of the resource rent to changes in stock stemming from each extraction policy, and is a measure of the change in value of the non-extracted resource derivating from the two different policies.

6 Concluding Remarks

We introduce a model of a monopolist firm that operates in a resource market where the prices are endogenously determined and in which ecological uncertainty takes the form of both Gaussian noise and regime shifts. These shifts are allowed to be dependent on the the monopolist's extraction efforts: unlike the previous literature, we explicitly model the firm's detection process of the regime change and incorporate it in the profit-maximizing policies. Our closed form solutions help us pin down the economic mechanisms that drive the extraction behaviour of the firm. In the event of a negative regime shift, for low resource stock levels, an increase in the resource rent results in the firm adopting a precautionary policy by reducing extraction. 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.

To conclude, some caveats are in order. Our model framework is intentionally simple and stylized in order to be able to obtain analytical solutions, allowing us to characterize the importance of the market structure whilst still allowing for a rich solution behavior. Furthermore, 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 obtaining an extraction policy only in numerical form. Lastly, a potential criticism may be the assumption of a monopoly: a pure monopoly is rare, and a game theoretic approach of several powerful players interacting could be more appropriate to the renewable resource market. However, our primary aim is to see how a firm, whose prices are not exogenously determined, adjusts its extraction levels in the presence a regime shift that it can attempt to detect in real time. The case of a monopoly can then be used as a first step towards richer competition structures.

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.

Andersen, T., J. Carstensen, E. Hernandez-Garcia, and C. M. Duarte (2009). Ecological thresholds and regime shifts: approaches to identification. *Trends in Ecology & Evolution 24*(1), 49–57.

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.

Costello, C., B. Nkuiya, and N. Quérou (2019). Spatial renewable resource extraction under possible regime shift. *American Journal of Agricultural Economics 101*(2), 507–527.

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 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.

De Zeeuw, A. and A. Zemel (2012). Regime shifts and uncertainty in pollution control. *Journal of Economic Dynamics and Control 36*(7), 939–950.

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.

Lindenmayer, D. B., R. J. Hobbs, G. E. Likens, C. J. Krebs, and S. C. Banks (2011). Newly discovered landscape traps produce regime shifts in wet forests. *Proceedings of the National Academy of Sciences 108*(38), 15887–15891.

Österblom, H., S. Hansson, U. Larsson, O. Hjerne, F. Wulff, R. Elmgren, and C. Folke (2007). Human-induced trophic cascades and ecological regime shifts in the baltic sea. *Ecosystems 10*(6), 877–889.

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

Pindyck, R. S. (1985). The measurement of monopoly power in dynamic markets. *The Journal of Law and Economics 28*(1), 193–222.

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.

Porter, J. H., P. C. Hanson, and C.-C. Lin (2012). Staying afloat in the sensor data deluge. *Trends in ecology & evolution 27*(2), 121–129.

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.

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.

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 (14), 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 (14) if it satisfies the equation

$$-\frac{\partial}{\partial t} V + A_t[V(t, .)](x) = 0, \tag{48}$$

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

$$-\frac{\partial}{\partial t} W(t, x) + \mathcal{H}(t, x, DW(t, x), D^2W(t, x)) = 0. \tag{49}$$

A function $V(t, x) \in \mathcal{C}([0, \tau_c] \times \mathbb{R})$ is a viscosity subsolution of (49) 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 (49) 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 (49) 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 (14). Let us now show that V is a viscosity solution of our problem (14).

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 Itô'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 (14). 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 (14). We can conclude that V is a viscosity solution of (14). 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 (14) 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 (14). Following the previous definitions, we must have

$$
\begin{aligned}
v_t(t,0) &\leq -\mathcal{H}(t,0,Dv,D^2v) && (50)\\
&\leq \sup_{q\in Q}\left\{\Pi(t,x,q) + v_x(0)(\mu-q) + v_{xx}(0)\frac{\sigma^2}{2}\right\} && (51)
\end{aligned}
$$

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 (51) 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

$$
\begin{aligned}
V_x(t,0) &\quad \geq \quad \Pi(t,0,q) && (52)\\
V(t,0) &\quad = \quad 0 && (53)\\
V(t,x) &\quad \text{solves} \quad V_t - \mathcal{H}(t,x,DV(t,x),D^2V(t,x)) = 0 \quad x \in [(0,\tau_c]\times\mathbb{R}]
\end{aligned}
$$

is a solution to the problem (14). 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.

B Resource Rent and Stock

$$
\frac{\partial V_x(x,t)}{\partial x} = \frac{M\times N\times O}{P^2\sigma^2} \tag{54}
$$

Where

$$
M = 8e^{\rho(t-\mathrm{tc})+\frac{2\sqrt{F+\mu(-a+c+b\mu)}x}{\sqrt{b}s^2}} \tag{55}
$$

$$
\begin{aligned}
N = &-a^2\left(e^{\rho t}-e^{\rho\tau}\right)^2 - c^2\left(e^{\rho t}-e^{\rho\tau}\right)^2 + 2ac\left(e^{2\rho t}+e^{2\rho\tau}\right) + \\
&4be^{2\rho t}F + 4bce^{\rho(t+\tau)}\mu - 4ae^{\rho(t+\tau)}(c+b\mu)
\end{aligned} \tag{56}
$$

$$
O = F + \mu(-a+c+b\mu) \tag{57}
$$

$$
\begin{aligned}
P = {} & ae^{\rho t} - ce^{\rho t} - ae^{\rho\tau} + ce^{\rho\tau} - 2be^{\rho t}\mu + 2\sqrt{b}e^{\rho t}\sqrt{F + \mu(-a + c + b\mu)} + \\
& 2\sqrt{b}e^{\rho t + \frac{2\sqrt{F+\mu(-a+c+b\mu)}x}{\sqrt{b}s^2}}\sqrt{F + \mu(-a + c + b\mu)} + \\
& e^{\frac{2\sqrt{F+\mu(-a+c+b\mu)}x}{\sqrt{b}s^2}}\left((a-c)e^{\rho\tau} + e^{\rho t}(-a + c + 2b\mu)\right)
\end{aligned} \tag{58}
$$

We can easily see that the denominator is positive and $M > 0$.

For $O > 0$:

$$
(F + c\mu) > (a - b\mu)\mu \tag{59}
$$

This is a reasonable assumption which implies that if the firm extracts as much as the drift, the cost it will incur will be greater than its revenue.

Simplifying N gives us:

$$
N = -\left(e^{\rho t} - e^{\rho\tau}\right)^2 (a-c)^2 + 4be^{\rho t}\left(Fe^{\rho t} - \mu(a-c)e^{\rho\tau}\right) \tag{60}
$$

For $N < 0$ we find the necessary condition to be:

$$
\frac{F}{a-c}e^{-\rho(\tau-t)} < \mu + q^m e^{\rho(\tau-t)}\left(e^{\rho t} - e^{\rho\tau}\right)^2 \tag{61}
$$

and the sufficient condition is:

$$
\frac{F}{a-c}e^{-\rho(\tau-t)} < \mu \tag{62}
$$

If (62) holds then $\frac{\partial V_x(x,t)}{\partial x} < 0$