CESifo WORKING PAPERS

7636
2019

May 2019

Should Faustmann forecast climate change?

Johan Gars, Daniel Spiro

CESifo

Impressum:

CESifo Working Papers
ISSN 2364-1428 (electronic version)
Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH
The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute
Poschingerstr. 5, 81679 Munich, Germany
Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de
Editor: Clemens Fuest
www.cesifo-group.org/wp

An electronic version of the paper may be downloaded
· from the SSRN website: www.SSRN.com
· from the RePEc website: www.RePEc.org
· from the CESifo website: www.CESifo-group.org/wp

CESifo *Working Paper No. 7636*
Category 9: Resource and Environment Economics

Should Faustmann forecast climate change?

Abstract

Climate change is predicted to substantially alter forest growth. Optimally, forest owners should take these future changes into account when making rotation decisions today. However, the fundamental uncertainty surrounding climate change makes predicting these shifts hard. Hence, this paper asks whether forecasting them is necessary for optimal rotation decisions. While climate-change uncertainty makes it theoretically impossible to calculate expected profit losses of not forecasting, we suggest a method utilizing Monte-Carlo simulations to obtain a credible upper bound on these losses. We show that an owner following a rule of thumb - ignoring future changes and only observing changes as they come - will closely approximate optimal management. If changes are observed without too much delay, profit losses and errors in harvesting are negligible. This means that the very complex analytical problem of optimal rotation with changing growth dynamics can be simplified to a sequence of stationary problems. It also implies the argument that boundedly-rational agents may behave "as if" being fully rational has traction in forestry.

JEL-Codes: C600, D810, Q230, Q540.

Keywords: climate change, decision making under uncertainty, forestry, quantitative methods.

Johan Gars
GEDB & The Beijer Institute of Ecological Economics / Royal Swedish Academy of Science / Stockholm / Sweden
johan.gars@kva.se

*Daniel Spiro**
Department of Economics, Uppsala University / Sweden & Oslo Business School / Norway
daniel.spiro.ec@gmail.com

**corresponding author*

April 24, 2019
Johan Gars' research on this article was partly funded by the The Erling-Persson Family Foundation and the Ragnar Söderberg Foundation. Daniel Spiro's research on this article was partly funded by the Research Council of Norway through the CREE project. The authors would like to thank Jette Bredahl Jacobsen, Anne-Sophie Crépin, Lars Hultkrantz, Karl-Gustaf Löfgren, Sjak Smulders and seminar participants at IIES, the Beijer Institute of Ecological Economics, the Frisch Center and the EAERE conference for helpful comments. Daniel Spiro is also affiliated with the Oslo Center for Research on Environmentally Friendly Energy (CREE).

1 Introduction

Should renewable-resource owners engage in forecasting activities? Or should they ignore future changes and base their decisions on what they observe today? These questions arise as the growth dynamics of renewable natural resources most likely will change over time. In the case of forests this may be due to technical change (e.g., genetic improvements) making trees grow both larger and faster, a spreading of disease to new environments which increases the risk of trees turning commercially worthless or due to climate change which is predicted to change the growth dynamics. These effects may be substantial. For instance, the negative effects of climate change on profitability of forestry in Europe have been predicted to be between 14 and 50 percent of net present value which corresponds to several hundred billion Euros (Hanewinkel et al., 2013). Indeed, a rather comprehensive literature (surveyed by Yousefpour et al., 2012) sets out to analyze how forest owners should adapt their decision making to future climate change. Analytically, adapting even just the Faustmann (1849) rule of optimal rotation is very hard.[1] This is possibly why the more recent literature has had a mainly computational and quantitative approach. For instance, Pukkala and Kellomäki (2012) and Schou and Meilby (2013) numerically evaluate different decision-making strategies under specific climate scenarios.[2] Other papers (e.g., Jacobsen and Thorsen, 2003) incorporate decision making with a *risk* component (i.e., a probability distribution of future climate scenarios).[3] Yet, one of the main problems with climate change (and other future changes) is precisely that assigning probabilities to different scenarios is difficult, let alone knowing which single scenario to use as a base. This is also one of the main conclusions in the survey by Yousefpour et al. (2012) – the missing piece of the puzzle is how to incorporate *fundamental uncertainty* (about the outcome and the probability distribution) into the decision making. Consequently, the recent literature has focused on how foresters can

[1] In the original setup (Faustmann, 1849; Pressler, 1860; Ohlin, 1921) the environment is stationary making it fairly easy to characterize the optimal rotation period analytically (see Newman, 2002, for a survey on extensions). With technological or environmental changes, the stationarity assumption will no longer hold. McConnell et al. (1983) and Löfgren (1985) are two early attempts at analytically addressing problems where the growth characteristics of trees change over time. They find that it is difficult to even characterize qualitatively how the decision rules should change over time, except for in some simple cases. See also Van Kooten et al. (1995) and Stollery (2005).

[2] These studies look not only at the rotation decision, but also other decisions facing forest owners. See also Susaeta et al. (2014) and Schou et al. (2015) for related studies.

[3] For management of non-forest resources under uncertainty see, for instance, Sethi et al. (2005), Tsur and Zemel (1996) and Huang and Loucks (2000).

improve their knowledge and decision making, for instance, by analyzing the arrival of information, concluding that decisions should be delayed (Brunette et al., 2014; Schou et al., 2015), that managers should update their beliefs on what climate scenario will occur as information arrives (Yousefpour et al., 2014) or that forest owners should diversify their portfolio of trees (Yousefpour and Hanewinkel, 2016). However, it remains an open question whether this effort to improve our decision making makes economic sense. That is, we do not know how much profits would be improved by having detailed knowledge about the future, to the extent that such knowledge is attainable. The uncertainty surrounding these issues implies that any forecast of the future growth of trees will be both costly to make and will most likely be inaccurate ex post.[4] Bearing this is mind our paper aims at 1) answering the question whether forecasting climate change is worthwhile for forest owners or if a less demanding rule of thumb can be used, while 2) getting around the problem of fundamental uncertainty regarding the probability distribution of future scenarios. When addressing these questions our focus is on the classic issue of rotation length.

More precisely, we compare the profits from two different decision rules. On the one hand, a hypothetical "optimal" rule that includes full information about all future changes to the trees' growth properties and the change in the risk of fire, storms and pests. On the other hand a "reactive" rule that, at every point in time, observes the prevailing growth dynamics and risks and assumes that they will be non-changing from then onwards.[5] Most forest owners have trees of many ages in their possession and hence it seems reasonable that they could at least assess the current growth properties of the trees.

Now, the fundamental problem is which climate scenario(s) we should use to evaluate losses from suboptimal decision making. Naturally, we do not ourselves know how the dynamics will change in the future. To get around the problem of uncertainty about the probability distribution the following setup is used. Instead of assessing the costs under

[4]This has also been noted by Löfgren (1985) who concludes that the assumption that forest owners can predict such changes is unlikely to hold. Indeed, Heltorp et al. (2017) and Yousefpour and Hanewinkel (2015) show that forest owners in reality find it too hard to forecast climate change and hence refrain from it.

[5]The term reactive is borrowed from Jacobsen et al. (2013). A reactive owner makes no attempt whatsoever to predict future changes – she simply adapts to observed changes. A more sophisticated owner could, for instance, be trend adaptive by inferring future changes from previous ones. A different approach in the forestry literature has been to analyze agents using Bayesian updating (see, e.g., Yousefpour et al., 2014).

particular scenarios like previous research does (e.g. Jacobsen et al., 2013; Susaeta et al., 2014; Boulanger et al., 2016; Yousefpour et al., 2017), we use scenarios of climate effects on trees that have a sufficiently broad range to cover changes which are, by far, more extreme than anything that can possibly happen in the future. Then, we calculate the profit losses from following the reactive rule for each one of a very large number of scenarios within this range (i.e. a Monte Carlo simulation). This setup enables getting an *upper bound* for what the losses may be if an owner is not forecasting the future. This upper bound is simply the losses of the reactive owner in the scenario with the largest profit losses. It is important to note that the upper bound of losses is not the same as the expected losses which by definition are smaller. The benefit of this approach is that *if* we find the upper bound to be of negligible size, which is exactly what we do, we can draw the conclusion that forecasting long-term forest dynamics is not important for optimal rotation. The benefit furthermore is that, to draw this conclusion, it is only necessary that the range of the scenario distribution is broad enough. That is, whether the underlying distribution of the scenarios is realistic *within* this range is of no concern. The analysis is performed for boreal forests (i.e. those covering the northern part of the northern hemisphere). For our simulations to credibly cover the actual outcome, which is unknown to us today, we need to ensure that our parameter variations span over at least what can possibly happen in reality. With this in mind we include very extreme scenarios where, for example, the risk of fire changes from happening every thousand years at the onset to instead happening every third year, trees growing to become four times larger and growing to 90% of their maximum size twice as fast. These scenarios are, by an order of magnitude, more extreme than those predicted in forest research. We also vary the trajectories of climate change to include, for instance, threshold effects. Since we use so extreme scenarios our conclusions are probably on the conservative side. They are also conservative since we compare the reactive rule of thumb to an optimal rule based on perfect information. A more realistic benchmark would be a decision maker using the best available estimate of future changes or using decision making under uncertainty (Sethi et al., 2005; Tsur and Zemel, 1996; Huang and Loucks, 2000). Our losses will by construction be larger than if the reactive owner would be compared to such decision makers.

Our main finding is that the reactive decision rule is very close to the optimal. Using the reactive rule yields only negligible profit losses (the upper bound is 0.2% profit losses) and implies a cutting of trees that is a very close approximation of the optimal. Even when the optimal rule is changing dramatically (e.g. cutting six times larger trees and halving the age at which to cut) the reactive rule closely follows it and yields only small relative losses. So, although the reactive owner will be constantly slightly wrong, since the actual changes are observed without delay, the owner will not be too far behind in updating.[6] In a robustness check we show that losses in the worst case scenario are negligible also when using, instead of the most recent information, a moving average of the last five years or when updating the decision rule every third year. Once the delays are more than this we can no longer, strictly speaking, conclude that losses will be small as some scenarios give losses of almost 2% when the delay is ten years (related to this Mäkinen et al., 2012; Pietilä et al., 2010, study optimal updating of forest inventories).

The implications and contributions of the paper can be summarized as follows. 1) We show how a very difficult theoretical problem – of optimal rotation with changing growth properties – can be readily collapsed to a series of stationary problems with only negligible loss of accuracy. 2) A conceptual contribution of the paper is showing that non-sophisticated (i.e., boundedly rational) forest owners will behave "as if" being fully rational. Hence, assuming full rationality when modeling rotation decisions in forestry is most likely without bias, even if in reality forest owners may or may not be fully able to predict the future. The issue of rationality is fiercely debated in economics, and ex-ante our conclusion is non obvious as previous research has shown that rationality assumptions sometimes matter for the outcome and sometimes not.[7] 3) A final contribution is methodological as we present a method for how to evaluate decision making under fundamental uncertainty that can potentially be applied to a wide range of issues both within and outside of environmental economics.

The structure of the paper is as follows. We start by presenting the theoretical model

[6] Our main results are computed for when the growth properties improve over time. A reversed alternative, for example due to spreading of deserts, where trees' growth falls in the next 200 years has also been simulated and is presented as a robustness check. Losses are small here too.

[7] See Conlisk (1996) for a discussion and, for instance, Love (2013) and Winter et al. (2012) who evaluate welfare losses when following rules of thumb in portfolio choice and private savings decisions respectively. See also Hong et al. (2007) for an application in finance, Spiro (2014) and van Veldhuizen and Sonnemans (2011) for applications and tests in resource economics, Sims (2003) for applications in macroeconomics and Mirrlees and Stern (1972) for an early theoretical treatment. For treatments of adaptive expectations in economics see, for instance, Nerlove (1958), Marcet and Sargent (1988), Hommes (1994), Burchardt (2005), Huang et al. (2009) and Chow et al. (2011).

and describing the decision rules in Section 2. In Section 3, we present the setup of the numerical simulations of the model and in Section 4 how the decision rules differ in terms of profit loss and the sizes of trees to be cut. We present some robustness checks in Section 5. In Section 6 we analyze why profit losses are small. Finally, Section 7 concludes by discussing our results and how they may possibly change if extending the model. The appendix contains some technical considerations regarding the decision rules and a description of the numerical algorithm we use.

2 The model

We set up our model in discrete time. The decision rules will be characterized by a minimum tree size that the decision maker will choose to cut at every point in time. The change of the growth properties is exogenous and therefore not affected by the endogenous decisions made by the agent.

In addition to changes in the growth properties, we will also assume a risk of for instance a fire, pests or a storm destroying the wood. We will treat this probability as exogenous but changing over time. For brevity we will refer to it as fire risk. One could think of several ways of enriching this by, for example, letting the risk be a function of the biomass but arguably, the way it is modeled here suffices for the purpose of comparing decision rules.[8]

The intended interpretation of the changes in the growth properties and fire risk is that it is driven by changes in the surrounding environment – e.g. temperature or technology. We will, however, treat these changes as exogenous and therefore we can model the changes as depending on calendar time rather than the actual underlying driver.

2.1 The basic setup

As in earlier work, the value for the owner of the forest is due to the possibility of harvesting the trees and selling them. The price of wood is assumed to be constant over time and given by p. The cost of harvesting is a constant, c, for every harvest.[9] It will be assumed that all trees are of the same age and that they will all be harvested at the same

[8] For a more thorough treatment of the risk of fire specifically, see Stollery (2005).

[9] We assume constant p and c in order to focus on the effect of a changing growth function.

time. Since we assume that the agent maximizes discounted profits and that there is no incentive for income smoothing, the assumption of equal tree age makes no difference.[10]

In the standard formulation of the optimal-rotation problem, when the environment is not changing, the problem facing the forest owner is the same after each time the trees are cut down. This means that the decision rule is independent of calendar time and usually expressed as the time that the owner should wait between each cutting. Here our decision rule will depend on calendar time and it will be described as the minimum tree size that the owner should cut down. Correspondingly, we chose the state variables in our optimization problem as calendar time t and current biomass F_t. Generally, we can then write the evolution of biomass as

$$F_{t+1} = (1 - X_t)\left[H_t g(0,t) + (1 - H_t)\, g(F_t, t)\right], \tag{1}$$

where $g(F,t)$ is the biomass law of motion (for short, we will call it the biomass function), $H_t \in \{0,1\}$ is the decision variable regarding harvesting and $X_t \in \{0,1\}$ is a random variable that is equal to 1 if the forest is destroyed by fire. $g(F,t)$ is assumed to be continuous and increasing in F. The distribution of X_t is given by

$$P(X_t = 1) = \pi_t \text{ and } P(X_t = 0) = (1 - \pi_t). \tag{2}$$

The assumed timing within a period is such that potential harvesting takes place at the beginning of a period and the trees then grow during the period until the next harvest opportunity. If the forest is destroyed in period t, the owner starts with biomass $F_{t+1} = 0$ in period $t+1$. We can also define the τ times repeated application of the biomass function, $g_\tau(F,t)$, recursively as

$$g_0(F,t) = F \text{ and } g_\tau(F,t) = g(g_{\tau-1}(F,t), t + \tau - 1) \text{ for } \tau \geq 1. \tag{3}$$

The profit in period t is given by $H_t\,(pF_t - c)$. That is, if the trees are cut down, the profit is given by the income from selling the wood minus the harvesting cost, otherwise there is no profit in that period.

The objective of the forest owner is to maximize the expected value of the discounted

[10] An implicit assumption here is that there are no synergies between harvesting several patches of forest simultaneously.

profits, and the maximization problem can therefore be written as

$$\max_{H_t \in \{0,1\}} \mathbb{E}\left[\sum_{t=0}^{\infty} \beta^t H_t (pF_t - c)\right]$$
$$s.t. \quad (1) \quad \forall t,$$

where β is the discount factor.

Using dynamic programming, the solution of the maximization problem results in a value function $V(F,t)$ and a policy function $H(F,t)$ at each point in time. The value function satisfies the Bellman equation

$$\begin{aligned} V(F_t, t) &= \max_{H_t \in \{0,1\}} H_t (pF_t - c) + \beta \mathbb{E}\left[V(F_{t+1}, t+1)\right] \\ &\quad s.t. \quad (1). \end{aligned} \tag{4}$$

Since g is increasing and continuous in F, V will also have these properties for all t.[11] The policy function $H(F,t)$ is the profit-maximizing choice of cutting or leaving the trees. The harvesting rules will be of the form

$$H(F,t) = \begin{cases} 0 & \text{if} \quad F < \bar{F}_t \\ 1 & \text{if} \quad F \geq \bar{F}_t \end{cases} \tag{5}$$

and we will characterize the rules by $\bar{F}_t$. It is not obvious that the decision rules will always have this form. In Appendix A we describe how we verify that this is in fact the case.

We assume that the biomass function and risk of fire will change during a finite number of years and thereafter be unchanged. Hence, the optimization problem becomes stationary in some time period T. Once the problem becomes stationary, we can solve the problem from T onwards as a stationary problem and this will give us the value function $V(F,T)$ that serves as an end condition to the non-stationary part of the solution.

We will now describe the decision rules that we analyze and how to solve for them.

[11] That it is increasing follows since for any combination of future harvesting times, the first harvest will be larger the larger the initial biomass implying that you could always do better when starting from a larger initial biomass. The continuity follows since for any combination of future harvesting times the continuity of g implies that the resulting value is continuous in initial biomass. The optimal value function is then, for each F, the maximizing choice over a set of continuous functions making it continuous.

This will give us everything we need to numerically solve the problem. We consider three decision rules: optimal, reactive and non-reactive.

2.2 The optimal decision rule

The optimal decision rule yields optimal harvesting decisions in all time periods. In order to find this rule the entire optimization problem must be solved from time zero incorporating all changes to the biomass function and fire risk. We solve this problem backwards by first solving the problem facing the forest owner after period T, when the problem becomes stationary, and thereafter working backwards through the non-stationary problem.

2.2.1 The stationary problem

After period T, the problem becomes stationary and calendar time no longer plays a role (we can therefore drop time from our notation when considering the stationary problem). We will here describe how to find the decision rule and value function of a stationary problem. This is the standard optimal-rotation problem with a risk of fire and we can utilize that the problem is "reset" after each harvest. The value of having biomass zero, given that the owner waits for τ periods before cutting down the trees (provided that there has not been a fire in between) fulfills the equation

$$V_\tau(0) = \sum_{t=1}^{\tau} \beta^t (1-\pi)^{t-1} \pi V_\tau(0) + \beta^\tau (1-\pi)^\tau \left[p g_\tau(0) - c + V_\tau(0) \right]$$

where $g_\tau(F)$ is the τ times application of g. Solving this expression for $V_\tau(0)$ gives us

$$V_\tau(0) = \frac{(1-\beta(1-\pi))\, \beta^\tau \tau (1-\pi)^\tau}{(1-\beta)\,(1-\beta^\tau(1-\pi)^\tau)} \left(p g_\tau(0) - c \right).$$

We can now find the maximized value $V(0)$ that is obtained when following the optimal decision rule. It must be given by $V_\tau(0)$ maximized over τ. Let

$$V(0) = \max_\tau \frac{(1-\beta(1-\pi))\, \beta^\tau (1-\pi)^\tau}{(1-\beta)\,(1-\beta^\tau(1-\pi)^\tau)} \left(p g_\tau(0) - c \right) \tag{6}$$

and let τ^* be the maximizing choice of τ. We also want to find $\bar{F}$ in (5). From the maximization we know that

$$\bar{F} \in (g_{\tau^*-1}(0), g_{\tau^*}(0)]. \tag{7}$$

In Appendix A.3 we show that $\bar{F}$ can be found by solving the equation

$$p\bar{F} - c + V(0) = \pi\beta V(0) + (1-\pi)\beta\left(pg(\bar{F}) - c + V(0)\right). \tag{8}$$

Given $V(0)$ and $\bar{F}$ we can evaluate the value function for an arbitrary biomass F. If $F \geq \bar{F}$ (i.e., if the owner immediately cuts the trees) then $V(F) = pF - c + V(0)$. For $F < \bar{F}$, i.e. the trees are not large enough to be cut immediately, then we have that

$$V(F) = \sum_{t=1}^{\tau(F)} \beta^t \pi (1-\pi)^{t-1} V(0) + \beta^{\tau(F)}(1-\pi)^{T(F)}\left(pg_{\tau(F)} - c - V(0)\right), \tag{9}$$

where

$$\tau(F) = \min_{\tau \geq 1} g_\tau(F) \geq \bar{F}$$

is the number of time periods before the first harvest.

We now have formulas for both finding the optimal decision rule and for computing the value function for an arbitrary biomass F in a stationary setting.

2.2.2 The non-stationary problem

Above we solved the problem in the stationary phase that begins in period T. We now move backwards into the non-stationary phase of the problem where the biomass function $g(F, t)$ and the probability of fire π_t changes over time. This also implies that the value function changes over time. The solution of the problem in the stationary phase gives us the value function in period T, $V(F, T)$, that serves as an end condition on V in the non-stationary phase. We can therefore solve the problem in period $T-1$ by using this end value. When considering the solution in any period $t < T$ we can thus assume that we have solved the problem in $t+1$ and know the value function in that period. This way we can work backward until $t = 0$. We can then formulate the Bellman equation in period t as a function of the biomass F and the continuation value of leaving either

$g(F,t)$ or $g(0,t)$ for the next period.

$$V(F,t) = \max_{H\in\{0,1\}} \left\{ \begin{array}{c} H\left[pF - c + \beta\left(\pi_t V(0,t+1) + (1-\pi_t)V(g(0,t),t+1)\right)\right] \\ +(1-H)\beta\left[\pi_t V(0,t+1) + (1-\pi_t)V(g(F,t),t+1)\right] \end{array} \right\}.$$

The solution to this equation gives us the value function and harvesting rule for any F since the right-hand side only depends on known entities when we get to period t. We want to find $\bar{F}_t$ in (5). Since the value function is continuous in F for all t, $\bar{F}_t$ must fulfill

$$\begin{gathered} p\bar{F}_t - c + \beta\left[\pi_t V(0,t+1) + (1-\pi_t)V(g(0,t),t+1)\right] \\ = \\ \beta\left[\pi_t V(0,t+1) + (1-\pi_t)V(g(\bar{F}_t,t),t+1)\right]. \end{gathered}$$

or

$$p\bar{F}_t - c = \beta(1-\pi_t)\left[V(g(\bar{F}_t,t),t+1) - V(g(0,t),t+1)\right]. \tag{10}$$

This equation gives the decision rule (as characterized by $\bar{F}_t$). We can also compute the value function of an arbitrary biomass F as

$$V(F,t) = \left\{ \begin{array}{rl} pF - c + \beta\left[\pi_t V(0,t+1) + (1-\pi_t)V(g(0,t),t+1)\right] & \text{if} \quad F \geq \bar{F}_t \\ \beta\left[\pi_t V(0,t+1) + (1-\pi_t)V(g(F,t),t+1)\right] & \text{if} \quad F < \bar{F}_t \end{array} \right. .$$

Summing up, we now have what we need to solve for the optimal decision rule and the associated value function. We first solve the stationary problem that is relevant for $t \geq T$. We then solve backwards for the decision rule and value function associated with the optimal decision rule for all $t < T$. In total this gives a sequence $\{\bar{F}_t\}_{t=0}^{T}$. We now turn to describing the "reactive" and "non-reactive" decision rules.

2.3 The reactive decision rule

The "reactive" decision rule observes the current biomass function and probability of fire. It does not take the future changes of the environment into account. In each period the decision maker solves the optimization problem that would arise if the current biomass function and risk of fire prevailed forever. Decisions are then made based on what would be optimal if the environment would not change over time. The owner is reactive in the sense that (s)he uses a rule of thumb as if noting that things have changed historically but believing that no more changes will come.

In each period t the reactive decision maker thus solves a stationary problem, as that described in Section 2.2.1, using the biomass function $g(F,t)$ and risk of fire π_t. Solving this sequence of stationary problems gives a sequence of decision rules characterized by $\{\bar{F}_t^R\}_{t=0}^T$.

The reactive decision maker, by construction, has incorrect expectations. This means that (s)he also has incorrect expected profits. When deriving the value of following the reactive decision rule we therefore want to compute the actual value of following this rule rather than the, incorrect, value that this person expects. To do this we need to calculate the value of following the reactive rule in the actual circumstances. That is, we have to calculate the value of an arbitrary biomass in period t when following the reactive decision rule, $V^R(F,t)$, as[12]

$$\begin{aligned} V^R(F,t) &= \sum_{s=1}^{\tau^R(F,t)} \beta^s \Pi_{0,s-1} \pi_s V^R(0, t+s) \\ &\quad + \beta^{\tau^R(F,t)} \Pi_{0,\tau^R(F,t)} \left[p g_{\tau^R(F,t)}(F,t) - c + V^R\left(0, \tau^R(F,t)\right)\right], \end{aligned} \tag{11}$$

where

$$\tau^R(F,t) = \min_{\tau \geq 0} g_\tau(F,t) \geq \bar{F}_{t+T}^R$$

is the number of periods before the next harvest in the absence of fire and

$$\Pi_{s_1,s_2} = \begin{cases} \prod_{s'=s_1}^{s_2-1}(1-\pi'_s) & \text{if} \quad s_1 < s_2 \\ 1 & \text{if} \quad s_1 = s_2 \end{cases} \tag{12}$$

is the probability that there is no fire in periods $s_1, \ldots, s_2 - 1$.

Since the value in period t depends on the value function in future periods we must compute the reactive value function backwards from period T where the reactive and optimal value functions coincide.

[12] We use the convention that a sum where the lower summation bound is larger than the upper is zero.

2.4 The non-reactive decision rule

For comparison we also include a "non-reactive" decision rule. This rule does not observe the current biomass function and risk of fire and instead makes decisions based on the initial biomass function and fire risk at time zero. The purpose of including such a, clearly suboptimal, decision rule is to verify that it is indeed possible to make large profit losses if ignoring changes that have already happened.

Here, in each period t the decision rule is characterized by the $\bar{F}$ that results from solving the stationary problem of Section 2.2.1 for biomass function $g(F, 0)$ and risk of fire π_0. This implies that the non-reactive decision rule will be characterized by $\bar{F}_t^{NR} = \bar{F}_0^R$ for all t.

As with the reactive decision rule we can calculate the value function of an arbitrary biomass F in period t, when following the non-reactive decision rule, $V^{NR}(F, t)$ backwards as

$$\begin{aligned} V^{NR}(F,t) \quad &= \quad \sum_{s=1}^{\tau^{NR}(F,t)} \beta^s \Pi_{0,s-1} \pi_s V^{NR}(0, t+s) \\ &\quad + \beta^{\tau^{NR}(F,t)} \Pi_{0,\tau^{NR}(F,t)} \left[p g_{\tau^{NR}(F,t)}(F,t) - c + V^{NR}\left(0, \tau^{NR}(F,t)\right)\right], \end{aligned} \tag{13}$$

where

$$\tau^{NR}(F,t) = \min_{\tau \geq 0} g_\tau(F,t) \geq \bar{F}_{t+T}^{NR}$$

is the number of periods before the next harvest in the absence of fire and where Π_{s_1,s_2} is defined in (12). The value function at T, $V^{NR}(F, T)$ is now not equal to the optimal value function and it must be computed separately.

2.5 Comparing decision rules

When evaluating the decisions rules we look at two measures. The first measure is the relative profit loss, which is

$$L^R \equiv \mathbb{E}_F \left[\frac{V(F,0) - V^R(F,0)}{V(F,0)} \right], \text{ with } F \sim U\left[0, \max\left\{\bar{F}_0^R, \bar{F}_0\right\}\right] \tag{14}$$

in the reactive case. That is, when calculating the profit loss we do it as an average over a range of initial biomasses. The reason for doing so is twofold. Firstly, most forest owners will have trees of different ages on their lands and indeed the reactive owner does not know when the changes to the growth function will start happening – the owner should thus be interested in expected losses over many initial conditions. Secondly, we do not want our results to be driven by discounting of events that happen far into the future. If we would have looked only at an initial biomass of zero, the actual mistakes would not have occurred until far into the future and the corresponding profit losses would have been discounted away. By taking the average over many initial sizes of trees (ranging all the way to a size that will be cut right away) we take into account situations where the forest owner is making actual mistakes already in the first few periods. The profit losses following such mistakes will not be discounted away.

For the non-reactive case we equivalently have

$$L^{NR} \equiv \mathbb{E}_F \left[\frac{V(F,0) - V^{NR}(F,0)}{V(F,0)} \right], \text{ with } F \sim U \left[0, \max \left\{ \bar{F}_0^R, \bar{F}_0 \right\} \right] \tag{15}$$

Our second measure is the relative errors in which size of tree to cut in each period, which are

$$D_t^R \equiv \frac{\bar{F}_t - \bar{F}_t^R}{\bar{F}_t}$$

in the reactive case and

$$D_t^{NR} \equiv \frac{\bar{F}_t - \bar{F}_t^{NR}}{\bar{F}_t}$$

in the non-reactive case. For the cutting errors we compute one value for each rule and each time period $t = 0, 1, \ldots, T$ within each scenario. By looking at the cutting errors in all periods, we capture mistakes both early and late in the transition process.

3 Numerical simulation

The problem defined in the previous section is very hard to solve analytically. It is even hard to characterize the optimal decision rule under any sort of general conditions. We therefore use a numerical approach that allows us to put an upper bound on the errors made by a reactive decision maker. As discussed earlier, it is difficult to assign probabilities to potential future climate scenarios. We address this problem by running a set of simulations (i.e., Monte Carlo) that cover a very broad range of possible future scenarios. If the errors made in all these simulations are small, which is what we find, this implies that the errors made in reality will very likely be small too. To draw such a conclusion the only requirement is that the parameters cover a wide enough range of future possibilities so that the actual realization will lie within the covered range of the parameter space. Further assumptions regarding the probability distribution within the range of scenarios covered are thus unnecessary. Note that if the resulting errors would *not* be small in all the simulations, that would still not allow us to conclude that errors would be large in reality since such a conclusion would require determining that some simulations that give large errors are in fact potentially realistic scenarios. Hence, this method allows for the following two alternative conclusions: “losses will necessarily be small” or “losses may not be small but we do not know if they will be large”.

In our main simulation we focus on a forest that changes so that it grows faster over time, where the maximum size of trees is increasing over time and where the risk of fire is increasing over time. This is meant to represent the case of climate change in boreal forests which cover the northern part of the northern hemisphere. As a robustness check we also simulate the case where growth conditions deteriorate over time. The results are no different. If anything, our main result that profit losses from following the reactive decision rule are small is strengthened.

In all simulations, we use $\beta = 0.95$ for discounting, which is a commonly used value for yearly discounting and real interest rate. Furthermore, we assume that it takes 200 years for the biomass function to become stationary.[13]

For the numerical simulation, a particular function $g(F,t)$ must be chosen. We derive our biomass function from a Bertalanffy growth function that has been converted to growth being a function of tree height rather than the age of the tree (for derivation, see Rammig et al., 2007). The resulting function for tree height h is

$$h_{t+1} = h_{\max,t}\left(1-\left(1-\left(\frac{h_t}{h_{\max,t}}\right)^{\frac{1}{D_t}}\right)e^{-A_t}\right)^{D_t}.$$

Here $h_{\max,t}$ represents the maximum height a tree can reach under the climate prevailing at time t. We interpret A_t as the initial growth of a tree and D_t as a general growth parameter, both of which prevail at time t. In the stationary phase, these parameters are constant. We calibrate the model to boreal forests (present in the northern parts

[13] One could think of the forest continuing to evolve also beyond this point. However, as $\beta^{200} \approx 0.00004$ it matters little for the discounted profits.

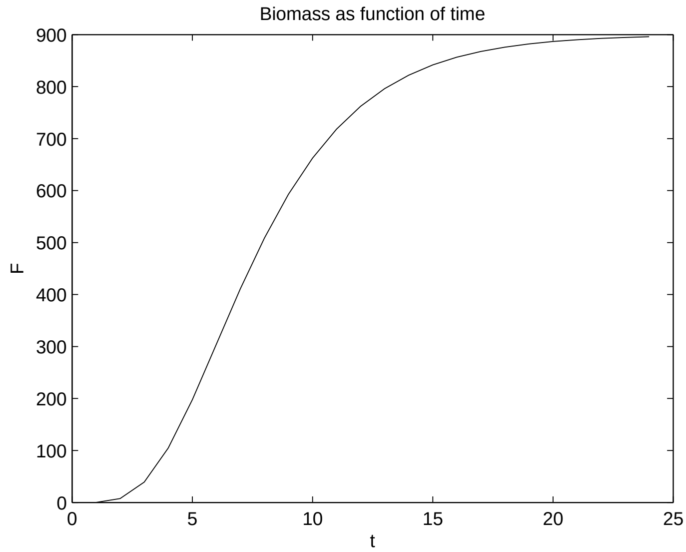

Figure 1: The typical time path of tree biomass following a Bertalanffy growth function.

of Europe, Asia and North America) and hence use the estimation of the parameters by Rammig et al. (2007) for Norway Spruce (Picea Abies) in the Swiss alps to approximately $A = 0.03$, $D = 3$ and $h_{\max} = 30$. Following Vanclay (1994, see page 70), we assume a quadratic relationship between tree height h and biomass F (implying a linear relationship between height and basal area). The biomass function is then

$$g(F_t, t) = F_{\max,t}\left(1 - \left(1 - \left(\frac{F_t}{F_{\max,t}}\right)^{\frac{1}{2D_t}}\right) e^{-A_t}\right)^{2D_t} \tag{16}$$

where $F_{\max,t} = h^2_{\max,t}$. A typical time path of the tree biomass following this biomass function with the baseline parameters (from Rammig et al., 2007) is shown in Figure 1.

Another calibration that needs to be performed is that of price p and cost c. We assume their values to be fixed over time.[14] While the relative proportions of these might be of importance for how often to harvest, their absolute values are not since they then only constitute a scaling of the absolute profits which is of no relevance here. The price is therefore normalized to one. Being interested in boreal forests we calibrate c to Swedish

[14]Solving for a general equilibrium of biomass, timber prices and harvesting costs is a very complicated endeavor, especially if it is to be done for a future under climate change, involving having estimates of the price of fossil fuels, technological change, population growth etc. Although interesting, it falls outside the scope of this paper.

forestry in 2010 (see Swedish Forest Authority, 2012, p. 283) so that $\frac{c}{pF_0} = 30\%$.[15]

We here assume that the risk of the trees being destroyed by a fire, pests or a storm, π_t, is increasing over time. This is a reasonable assumption in northern countries when the underlying driver is increased temperature. As a robustness check we also consider a case without any such risks.

In our simulations, we first set the initial parameter values (at time zero). We set $F_{\max,0} = h^2_{\max,0} = 30^2$ and randomly draw values of the remaining parameters from uniform distributions over the intervals $A_0 \in [0.025,\ 0.035]$, $D_0 \in [2.5,\ 3.5]$ and $\pi_0 \in [0.001,\ 0.3]$.[16] Second, we vary the parameter values when the biomass function stabilizes at $t = 200$. They are randomly drawn from the uniform distributions $A_{200} \in [A_0,\ 3A_0]$, $D_{200} \in [0.8D_0,\ D_0]$, $F_{\max,200} \in [F_{\max,0},\ 4F_{\max,0}]$, $\pi_{200} \in [\pi_0,\ 0.3]$. This encompasses very significant changes in the biomass function since it allows for a quadrupling of the maximum biomass (corresponding to a doubling of the maximum height) and trees growing to 90% of their maximum twice as fast as before, all happening within 200 years. To illustrate that these scenarios are indeed very extreme, we can compare them to some calibrated predictions in forest biology for boreal forests. For instance, Ge et al. (2011, table 2) have an increase in what corresponds to our $F_{\max}$ of at most 20% while we allow for a 300% increase. Another example is Kellomäki et al. (2008, table 4) who estimate an increase of the tree growing to be 100% bigger within one rotation while we have several scenarios where the tree grows to be more than 500% bigger within a rotation. Furthermore, we allow for the risk of forest fire to increase from once in a millennium to every third year. Third, we vary the trajectory for the changes during the 200 years to be either linearly increasing, concavely increasing or stepwise increasing, as illustrated in Figure 2 (all the parameters follow the same type of trajectory in each simulation). The concave and linear case represent smooth transitions with different modes of convergence. The stepwise increasing case is meant to catch climate thresholds which lead to abrupt changes. This final case is also a reduced form for abrupt technical change.

The complete Monte Carlo simulation draws 60 initial sets of parameters and 240 sets of parameter values at $T = 200$ per initial set. For each of these combinations of

[15] Andersson et al. (2013) show that this ratio has been very stable for the last 70 years. In our setting it may change in later time periods as one cuts larger or smaller trees.

[16] We do not vary the value of $F_{\max,0}$ since it is essentially a scaling of the whole problem.

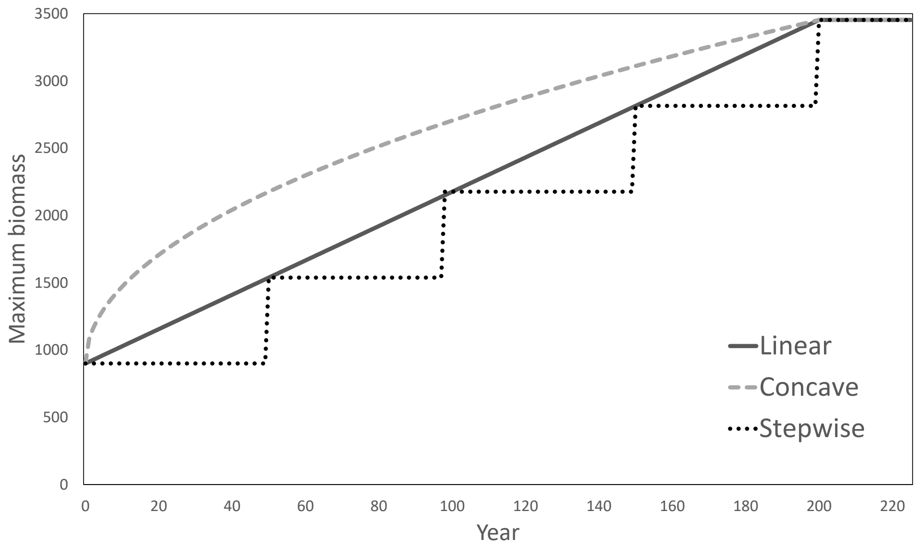

Figure 2: Example of trajectories for the maximum biomass parameter.

initial and final values it then uses all the three described trajectory shapes. In each separate simulation it calculates the optimal, reactive and non-reactive decision rules and the associated discounted profits of having a forest of any possible initial biomass, that is, for any $F_0 \in \left[0, \max\left\{\bar{F}_0^R, \bar{F}_0\right\}\right]$.[17] We then compute the expected values of L^R and L^{NR} from equations (14) and (15), that is, for a uniform distribution of initial biomass $F_0 \sim U\left[0, \max\left\{\bar{F}_0^R, \bar{F}_0\right\}\right]$. An algorithm for the numerical simulations can be found in Appendix B. The code and other materials are available upon request. In total the Monte Carlo procedure involved 43200 simulations.

4 Numerical results

The optimal and reactive decision rules of the simulation that yield the maximum expected losses (that is, our worst-case scenario) are depicted in Figure 3. The graph on the left-hand side shows how the $\bar{F}$ values that characterize the reactive and optimal decision rules change over time. As can be seen, the reactive policy is very similar to the optimal one throughout the transition period. On the right-hand side of Figure 3 we see how the minimum ages of trees that the reactive and optimal decision rules would prescribe cutting.

[17] We could instead have used initial values up to $F_{\max,0}$ but since both decision rules would prescribe immediate harvesting for $F > \max\left\{\bar{F}_0, \bar{F}_0^R\right\}$ that would most likely result in small profit losses.

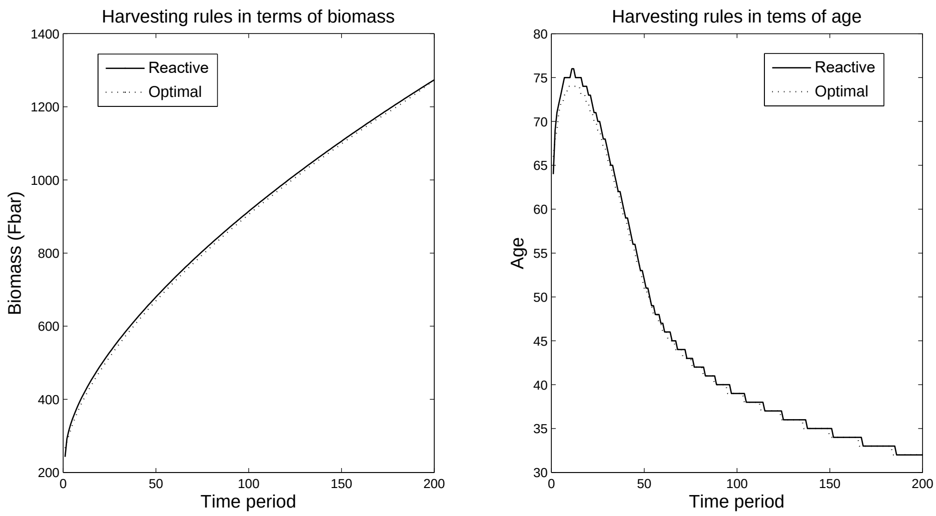

Figure 3: Decision rules of the worst case scenario. The graph on the left-hand side shows the decision rules in terms of biomass $\bar{F}$ while the right-hand side of the graphs show them in terms of the minimum age of trees that the respective decision rules prescribe harvesting. The ages are the actual ages of trees cut in each time period. For the reactive decision rule this will not be the same as the planned rotation time at planting. The initial parameters of this scenario are $A_0 = 0.0027$, $F_{\max,0} = 900$, $D_0 = 3.4975$ and $\pi_0 = 0.0017$; the parameter values after stabilization are $A_{200} = 0.0619$, $F_{\max,200} = 3453$, $D_{200} = 3.1368$ and $\pi_{200} = 0.0055$ and the trajectory is concave.

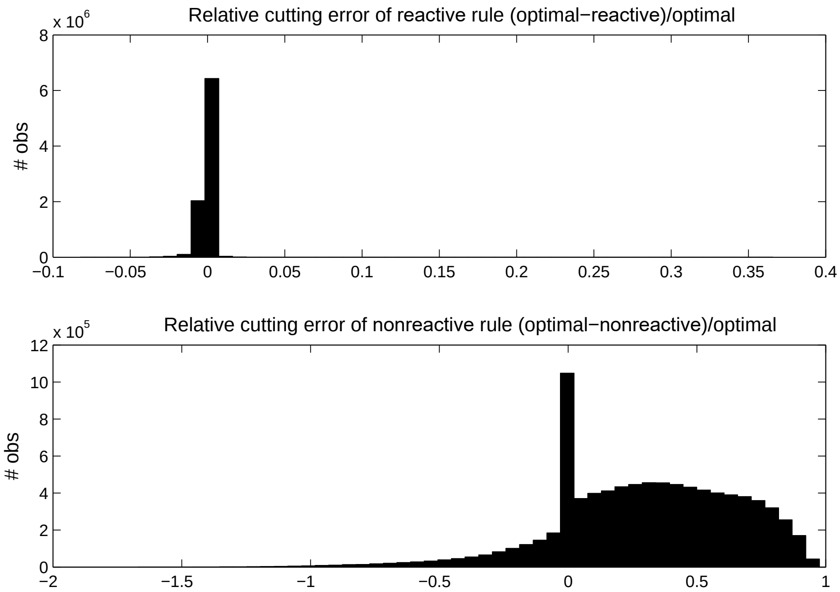

Figure 4: Histograms of the relative errors in the reactive and non-reactive decision rules.

The similarity between the optimal and reactive decision rules is representative. This can be seen in Figure 4 which shows a histogram over the relative errors in the reactive decision rule (in that figure we also present the harvesting errors of following the non-reactive rule). Each simulated case (initial parameters, final parameters and trajectory) results in $T + 1$ values of $\bar{F}$ for each decision rule. Each such value of $\bar{F}$ constitutes an observation in this histogram. With the exception of a few outliers (the most extreme being 37%), the harvesting errors when following the reactive decision rule are small. While these outliers imply that we cannot discard the possibility that errors in harvesting rules can, in some cases, be large, the central question is whether they yield large profit losses.

Histograms over the relative profit losses made can be found in Figure 5. In these histograms we have one observation per simulated case for each decision rule. The profit losses are calculated using (14) and (15). That is, for each scenario we calculate the profit losses as discounted from period 0 for any given initial biomass $F_0 \in \left[0, \max\left\{\bar{F}_0^R, \bar{F}_0\right\}\right]$. We then calculate the mean profit loss for F_0 uniformly distributed in that range. It means that in each scenario we incoporate the possibility that the forest will be harvested right away. From the histogram for the reactive decision rule we can see that the expected

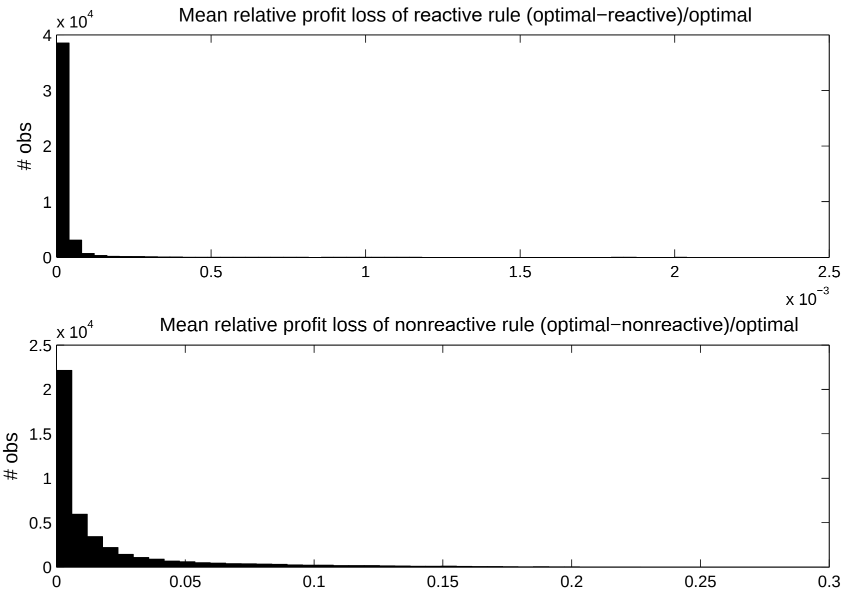

Figure 5: Histograms of the mean relative profit losses resulting from following the reactive and non-reactive rather than the optimal decision rules.

profit losses are small for all scenarios. The worst case scenario (the harvesting decisions of which are depicted in Figure 3) gives a loss of 0.2% and hence we treat this as our upper bound for losses following the reactive rule.

A forest owner following the non-reactive rule would make significant losses in many cases. We can see that many scenarios yield losses of 2-30% which, of course, makes them economically significant. This also shows that the climate scenarios we consider, indeed, can give non-negligible losses if decision making is very uninformed.

5 Robustness checks

To check the robustness of the results we have performed simulations for some extensions of the model. Firstly, our main formulation of the reactive rule implies that the owner observes the actual changes to the forest dynamics without delay. This is motivated by most forest owners having trees of various ages from which they could infer the entire biomass function of the trees. Now, while this may be reasonable, a practical problem facing an owner trying to perform such an exercise is that there is variability in the actual growth of trees from year to year depending on the weather and other temporary

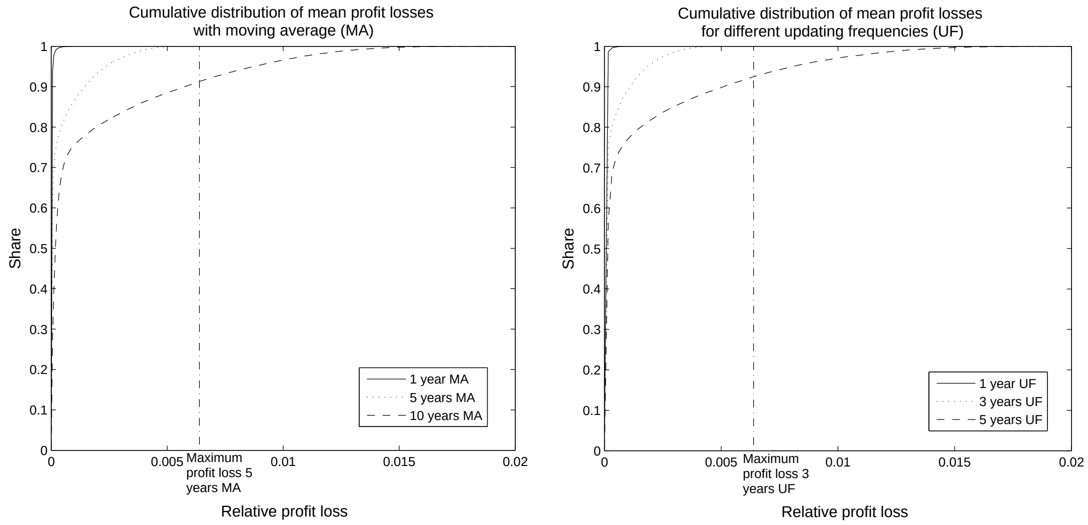

Figure 6: Cumulative distribution of mean profit losses when following lagged versions of the reactive decision rule. The vertical lines indicate the point at which the respective cumulative distributions become equal to one. That is, the largest expected profit loss in our simulations.

factors. This implies that the owner may need more than one year to infer what the changes are and hence that there will be a delay. A very extreme form of delay is obviously the non-reactive rule, which altogether does not notice the changes. We have seen that with the non-reactive rule the profit losses can be substantial so the question is what kind of delay can be allowed while keeping profit losses small. To address this issue we have computed the relative profit losses of the reactive rule when the owner observes the actual conditions with two forms of delay. In the first form, the lagged decision rules are computed as a moving average of the reactive decision rule without delay. That is, in a given period, the n years lagged decision rule is characterized by an $\bar{F}^{\text{Delay}}$ that is the average of the reactive decision rules ($\bar{F}^R$) over the n most recent periods. The results are presented in the left part of Figure 6. There we have included the profit losses following the reactive decision rule, the decision rule using the moving average of the reactive rule of the last five years and of the last ten years. Each graph represents the cumulative distribution function (CDF) for a certain number of years of delay with profit losses on the x-axis. What is of interest for us is how fast each CDF converges to one. At the point where it does, we obtain the least favorable scenario and hence the upper bound for reactive losses given a certain delay. As can be seen, as long as the delay is five years or less the maximum loss is well below one percent. However, for delays of ten years the losses are close to two percent in the least favorable scenario.

In the second form of delay the decision rule is updated not every year (like in the reactive rule) but more infrequently. So, for instance, with an updating every third year the decision rule follows the reactive rule of year one during years one to three, the reactive rule of year four during years four to six and so on. The results are presented in the right part of Figure 6. There we have included the CDF of relative losses following decision rules with updating frequency of one year (i.e., the reactive rule), three years and five years.[18] As can be seen, an updating frequency of three years yields losses of around 0.55 percent in the worst case scenario. But if the updating frequency is only every five years, losses in the worst case scenario are above 2 percent.

Our interpretation of the results following the two forms of delay is that for short enough delays (up to five year moving average or three years updating frequency) the reactive owner will still only make marginal losses while if the delay is sufficiently long we can no longer draw the conclusion that the upper bound is small. Recall, however, that this does not imply that we can draw the conclusion that profit losses will be large – we simply cannot say anything conclusive in such a case.

Secondly, we have also performed the simulations for the case where trees are growing slower over time to represent cases where the climate becomes, for instance, drier.[19] In these simulations we use the same starting values and simply use a mirror image of the end values in the basic case of trees growing faster so that trees now grow slower and to smaller sizes. The evolution of the fire risk is however not reversed – it increases over time here too. The results of such an exercise look very much the same as in our main simulation. If anything, the profit losses are now smaller.

Finally, although the case of trees growing faster and an increasing fire risk is probably the most relevant scenario for boreal forests, these two effects tend to cancel each other out. While a faster growing tree may imply cutting larger trees, the increasing fire risk implies cutting smaller trees. To see whether this is what causes the reactive rule profit

[18] This investigation is related to that of Mäkinen et al. (2012). However, while we look at updating frequency for learning about the growth function, they study the updating frequency for learning about the current biomass.

[19] For both the case where trees grow more slowly over time and the case, below, without risk of fire, the simulations draw 20 sets of initial parameters, then for each initial parameter set draws 80 sets of parameters at time T and performs the computations for each of the three possible trajectories. This gives a total of 4800 simulations.

losses to be small we have also done the simulations while shutting down the fire risk. We simply assume that there is no risk of fire in any year. Again, the profit losses of the reactive rule remain small with an upper bound of 0.35%.

While the magnitudes of the errors vary somewhat between the main simulation and the robustness checks of trees growing slower and no fire, in all cases the profit losses associated with the reactive rule are negligible and the losses associated with the non-reactive rule are about 100 times larger.

6 Why are profit losses small?

In this section we will explore and explain the intuition behind our result that reactive decision making yields negligible profit losses. Since the model with optimal decision making cannot be solved analytically, the discussion will by necessity be in the form of conjectures.

It should first be noted that the small profit losses are *not* due to discounting since we calculate the profit losses including situations where the initial biomass is large so the forest should be harvested right away (see equations (14) and (15)). In such situations the brunt of the profit losses are indeed materialized right away.

For there to be non-negligible profit losses when using a rule of thumb, two conditions have to be satisfied. First, the rule of thumb has to imply that mistakes are being made, mistakes which are by themselves non-negligible. Second, the profits have to be sufficiently sensitive to mistakes. Do reactive decision makers make mistakes? Figure 4 shows that the answer is that, for the most part, the cutting rule of what size of tree to harvest will imply virtually no error at all, but that there exist some exceptions. In particular, Figure 3 shows that, in the worst-profit-loss scenario, the reactive owner will be very close behind in updating the cutting rule. This implies that profit losses will be substantially attenuated. To see why, consider the following thought experiment. Suppose an owner follows an incorrect cutting rule in period 0 but that from period 1 and onwards this owner uses the same cutting rule as the fully forward looking one.[20] This is illustrated in Figure 7 using the parameters from the worst case scenario (those used in Figure 3). On the horizontal axis we have what size of tree the owner owns in

[20]For the sake of simplicity in illustrating our point, we consider the case of a one-time error unlike in the full simulations where the reactive owner may be wrong also later.

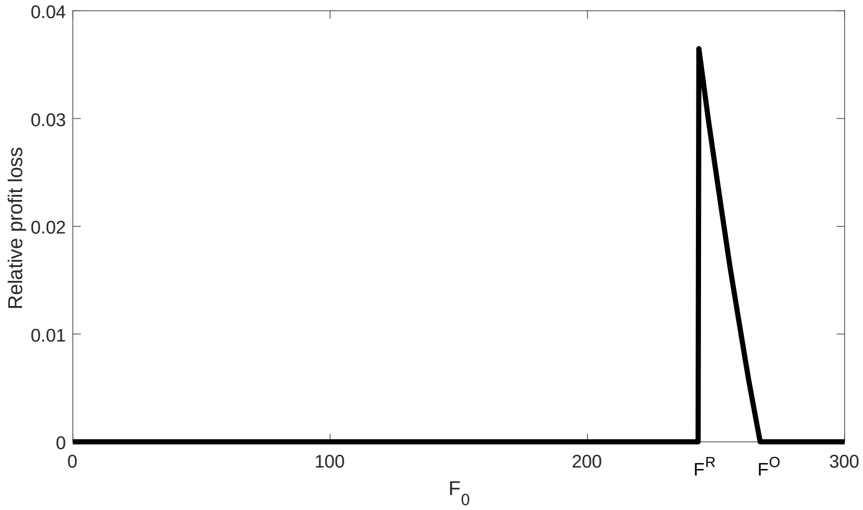

Figure 7: The relative profit loss for different initial biomasses from following the reactive decision rule in period 0 and then following the optimal rule.

period 0. The reactive and optimal cutting rules are indicated here too. As can be seen, the cutting error is about 10%. The graph shows the profit losses (in percent) using the reactive instead of the optimal rule in period 0. As can be seen, as long as the tree is smaller than both rules, the profit loss will be zero. The reason, of course, is that none of the rules prescribes cutting, hence the tree is left to grow to period 1, and from then on the owner follows the optimal rule. Likewise if the tree is very large, the profit loss will be zero since then both the reactive owner and the optimal owner will cut the tree. It is only in the range where the cutting rules do not overlap (in this case between 243 and 267) that the reactive owner will actually make a mistake. Hence, in order to make an actual mistake, the reactive owner has to have a tree precisely in this range. While this may of course happen, in particular if the error in the cutting rule is large, the *expected* profit losses will be attenuated by the fact that in many cases the error in the cutting rule will not have any actual consequences.

The next aspect to note in Figure 7 is that, even when the owner has a tree size that will lead to an actual mistake, the profit loss is still rather small. The reason for this can be understood by considering the original Faustmann-Pressler-Ohlin result. The optimal cutting rule trades off, on the one hand, cutting the tree when it is small in order to get the profit more quickly and on the other to wait for later in order to reap the benefits of the tree's growth, that is, to cut a large tree. This means that, while mistakes will obviously imply a cost, they will also have some benefits. In the scenario illustrated in Figure 7, where the reactive owner cuts too early, the cost is that the reactive owner misses out on the relatively high growth in the next period but the benefit is that profits arrive earlier. As long as the tree is not too far in size from optimal, the benefit will largely compensate for the cost and, as can be seen in Figure 7, in the region just to the left of the optimal cutting rule, the profit losses are very close to zero. In fact, this result follows from basic optimization theory: if the objective (profit) function is concave, then small deviations from the optimal choice have very small effects. Hence, large mistakes are necessary for there to be non-negligible profit losses. In summary there are two attenuating effects: 1) the erroneous cutting rule will often not lead to an actual mistake; and 2) profit losses are small as long as the actual mistake is small. Considering both of them jointly the question is how wrong the cutting rule has to be in order for profit losses to be non-negligible. This is illustrated (again for the parameters of the worst-case scenario in Figure 3) in Figure 8. On the horizontal axis we have different cutting rules with an indication of the optimal cutting rule. The graph depicts the expected relative profit losses when using an incorrect rule in period 0 and then using the optimal in period 1 and onwards. As can be seen, the profit losses are U-shaped around the optimal rule and losses are negligible for a quite broad range of rules. In particular, the reactive rule (indicated on the horizontal axis) implies a profit loss of around 0.11 percent only.[21]

7 Conclusions and discussion of the results

Our results clearly show that a forest owner who is able to gradually observe, rather than foresee, changes in the growth dynamics and in the risk of, for instance, fire will be very close to making optimal decisions of when to harvest. The worst-case reactive policy implies profit losses below 0.2% and the cutting rule does not differ much from the optimal despite considering very extreme scenarios. Even if these changes are observed with a delay of five years the profit losses of not being perfectly forward looking are only marginal. However, not noticing changes at all, like the non-reactive owner, may imply

[21]Despite using the worst-case scenario the profit losses are smaller here than what we get in the full simulation in the upper part of Figure 5. The reason is that here, for the purpose of illustration, we consider only a one-time cutting error while in the full simulation the reactive owner will make mistakes also later.

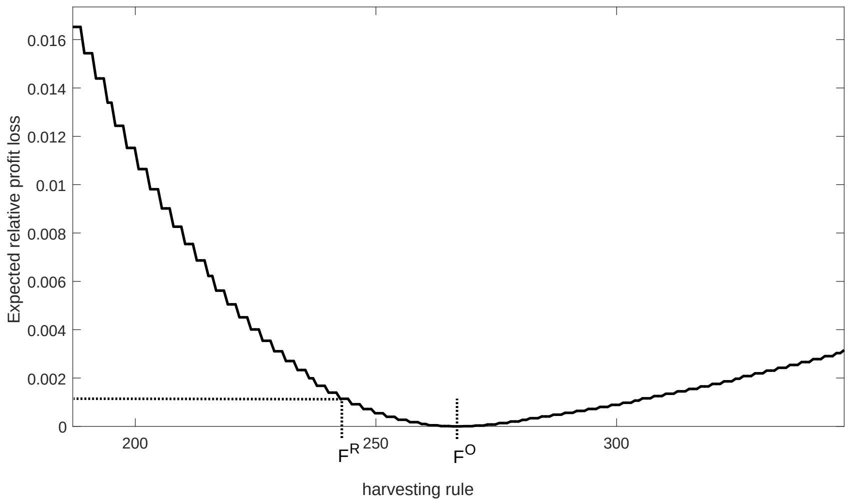

Figure 8: The expected relative profit loss from following a different harvesting rule than the optimal in the first period and then following the optimal rule. The reactive and optimal rules are indicated on the horizontal axis.

substantial profit losses and mistakes in cutting. Therefore, we conclude that observing the forest's change as it happens is important, but sufficient, for finding the optimal rotation rule. The reason is that observing the changes as they come ensures that the harvesting rule does not diverge from the optimal. For deciding how often to harvest a forest in practical forest management, it is therefore questionable whether investing in high accuracy climate and technology forecasts is worthwhile. It is hard to assess the cost of credibly forecasting climate change on a particular forest patch, but realistically it would be at least one order of magnitude larger than the maximum of 0.2% of profit losses of not forecasting we get in our simulation.[22] This, however, does not imply that foreseeing a catastrophic event next year, or short-run price changes, is not worthwhile. It is of course of great value for a forest owner to know whether, for instance, the forest will burn down next year – if such knowledge is at all possible. But this is not the same as it being valuable to know how the probability of a fire will change, which is what this paper is about.

The focus of our analysis has been on the basic rotation decision – the Faustmann rule

[22]See, for instance, Mäkinen et al. (2012) who estimate the costs of just keeping up-to-date with changes that have already happened to around 50 Euros per hectare with medium accuracy (see their Figure 2). They furthermore conclude that an owner should optimally do six inventories for every 30 year rotation. This would imply costs of roughly 3% of the revenues for a typical hectare of Scandinavian forest.

– and assessing whether it is affected by future changes. We show it is not. Naturally, our model is far more simple than the complex reality facing an actual forest owner. Our results are therefore, strictly speaking, confined to the aspects covered in the model. Whether our simplifications are of any consequence to the results crucially depend on two things. Firstly, it is not enough that an additional piece of complexity would change the harvesting rules. Since we are only interested in how wrong the reactive owner is compared to an owner with perfect foresight, what is important is whether the added component would affect one of them disproportionately more than the other and how this translates into expected profits. Secondly, as was explained in Section 6, adding a degree of complexity to the model is only of importance if it will affect losses under many of the initial conditions, rather than under only one or a few.

One possible alteration of the model would be to allow for the choice of changing the kind of tree one plants instead of having one tree species whose growth pattern changes over time. This would require a different kind of modeling where, possibly, the biomass function is constant for the currently standing tree but where planting a new tree type changes the biomass function. In this sense our way of modeling may be more representative of climate change whereas R&D may be better modeled in the alternative way. Although we see no immediate reason why this would be more of a problem for a reactive decision maker and why this would not be averaged out over the possible initial conditions, our model and simulations do not reveal whether the alternative problem makes forecasting worthwhile. As has been shown by Yousefpour et al. (2014), when making decisions about which trees to plant the losses of acting according to the wrong scenario can be severe. We can only speculate whether this difference compared to our results comes from the choice being made (when to cut compared to what species to plant) or whether it is due to that making inaccurate predictions is more problematic than not making any predictions at all. However, Jacobsen et al. (2013) show that following the wrong climate scenario when deciding on thinning and clear-felling (a decision similar to what we analyze) may imply substantial profit losses. This, together with our results, suggests that it is better to change the decision rule based on observed changes rather than wrongly predicting the change.

Another constraint that is highly relevant within forestry is that the price of timber may fall if the tree diameter becomes too large. We are fairly sure that, if anything, this should decrease the difference between the reactive and the optimal rule since it puts a cap on which size of tree to cut for both types of decision makers.

What is more uncertain is how the results would change if incorporating a decision of when to altogether stop growing trees on a certain land area. Our model only considers the direct costs and profits of growing and harvesting trees. But in reality there is also an opportunity cost since the land area may be used for other purposes. Obviously, the fact that forestry is currently taking place strongly suggests that these alternative usages should not be more profitable today. When the trees are growing faster and faster (as in our main simulations) a first guess would be that this should make forestry a relatively profitable endeavor also in the future. But in the opposite case, when trees are growing slower and slower, it might be that the alternative usages become relatively more profitable at some point. In that case losses may be incurred if one plants trees to be harvested in, say, two decades when an alternative usage is preferable already after a few years' time. A perfectly forward looking agent may be able to foresee this while a reactive agent may not. A related issue is the calculation of the present value of the forest before buying or selling it. If the effect of climate change on the growth of trees is extreme then this may affect the present value substantially. So to arrive at a correct price in such an important single decision it seems essential to at least roughly be able to foresee how climate change will affect growth. Likewise we cannot disregard that other long-lasting decisions, such as changing tree type, may affect the results. Especially in the case of stepwise climatic changes. Ultimately, our method can be used to study also such other decisions facing a forest owner.

These potential caveats aside, it is worth noting that the conclusions we draw regarding the Faustmann rule under climate change are probably on the conservative side. First of all our decision maker with perfect foresight is indeed a utopian agent. In practice no one is able to foresee climate change and its consequences and economic agents are left to either believing in a specific scenario or weighting different scenarios and letting expected profit maximization determine the harvesting rule. Secondly, the scenarios we have included are very extreme. Since this paper is about dealing with uncertainty rather than quantifiable risk, and since we do not ourselves know what probability distribution to use, we have chosen to include these extreme scenarios as to be sure that the future possibilities all fall within our simulations.

A Monotonicity of the harvesting rule

We will here derive sufficient conditions for the harvesting rule to have the form (5). We will start by verifying that our biomass function is increasing and concave. We will then draw some implications from this and show how it can be used to verify the form of the harvesting rule.

A.1 Concavity of the biomass function

Regarding the biomass function g we will assume that it is continuous, increasing and concave for all $F \geq 0$ in all time periods. This will hold for our biomass function as can be seen by differentiating it. Starting from (16), we rewrite the biomass function as[23]

$$g(F) = F_{\max}\left(1 - e^{-A} + e^{-A}\left(\frac{F}{F_{\max}}\right)^{\frac{1}{2D}}\right)^{2D}.$$

Differentiating it once we get

$$\begin{aligned}
g'(F) &= F_{\max}2D\left(1 - e^{-A} + e^{-A}\left(\frac{F}{F_{\max}}\right)^{\frac{1}{2D}}\right)^{2D-1} e^{-A}\frac{1}{2D}\frac{1}{F}\left(\frac{F}{F_{\max}}\right)^{\frac{1}{2D}} \\
&= e^{-A}\left(\frac{g(F)}{F_{\max}}\right)^{\frac{2D-1}{2D}}\left(\frac{F}{F_{\max}}\right)^{\frac{1-2D}{2D}} = e^{-A}\left(\frac{g(F)}{F}\right)^{\frac{2D-1}{2D}} > 0.
\end{aligned}$$

Differentiating again we get

$$\begin{aligned}
g''(F) &= \frac{2D-1}{2D}g'(F)\left(\frac{g'(F)}{g(F)} - \frac{1}{F}\right) \\
&= \frac{2D-1}{2D}\frac{g'(F)}{F}\left(\frac{F}{g(F)}e^{-A}\left(\frac{g(F)}{F}\right)^{\frac{2D-1}{2D}} - 1\right) \\
&= \frac{2D-1}{2D}\frac{g'(F)}{F}\left(e^{-A}\left(\frac{g(F)}{F}\right)^{-\frac{1}{2D}} - 1\right).
\end{aligned}$$

[23]For practical purposes it does not matter much what we assume about the biomass function for $F > F_{\max}$ but for completeness we need to assume something. Assuming that the biomass function is given by the same expression from there as well implies that the biomass can become smaller over time $g(F) < F$. While it is not obvious what the right assumption is, this assumptions implies that the biomass function will always be increasing and strictly concave.

Assuming $2D > 1$ this is negative since $g'(F) > 0$ and

$$
\begin{aligned}
e^{-A}\left(\frac{g(F)}{F}\right)^{-\frac{1}{2D}} &= \frac{e^{-A}}{\left(\frac{F_{\max}}{F}\right)^{\frac{1}{2D}}\left(1-e^{-A}+e^{-A}\left(\frac{F}{F_{\max}}\right)^{\frac{1}{2D}}\right)} \\
&= \frac{e^{-A}}{\left(1-e^{-A}\right)\left(\frac{F_{\max}}{F}\right)^{\frac{1}{2D}}+e^{-A}} < 1
\end{aligned}
$$

This shows that g is increasing and concave.

A.2 The value of waiting a prescribed number of periods

In order to verify the monotonicity of the harvesting rule for both stationary and non-stationary problems, we will utilize the shape, as a function of F, of the value of waiting a prescribed number of time periods before harvest. We assume now that we are in period t and that the value of having biomass zero is known for all future time periods including period t (the value $V(0,t)$ can easily be computed given the decision rule and the value of zero biomass in all future time periods). We can then define the value of waiting for τ time periods before harvesting as

$$V_0(F,t) = pF - c + V(0,t)$$

and, for $\tau \geq 1$

$$
\begin{aligned}
V_\tau(F,t) &= \sum_{s=1}^{\tau} \beta^s \Pi_{0,s-1} \pi_s V(0,t+s) \\
&\quad + \beta^\tau \Pi_{0,\tau}\left[p g_\tau(F,t) - c + V(0,\tau)\right]
\end{aligned} \tag{17}
$$

with Π_{s_1,s_2} defined in (12). We also define

$$w_\tau(F,t) = V_\tau(F,t) - V_0(F,t)$$

the net value of waiting for τ periods before harvesting rather than harvesting directly. From the concavity of g it follows that $w_\tau(F,t)$ is concave in F if $\tau \geq 1$. The concavity implies that it can change sign at most twice, once from negative to positive and once, for a larger F, from positive to negative. In particular we have that if $\tau \geq 1$, $F_2 > F_1 \geq 0$

then

$$w_\tau(F_1,t) > 0 \;\&\; w_\tau(F_2,t) < 0 \Rightarrow w_\tau(F,t) < 0 \;\forall\; F > F_2 \tag{18}$$

and

$$w_\tau(F_1,t) \geq 0 \;\&\; w_\tau(F_2,t) \geq 0 \Rightarrow w_\tau(F,t) \geq 0 \;\forall\; F \in [F_1,F_2]. \tag{19}$$

Suppose now that we have found a biomass $\tilde{F}$ and waiting time $\tilde{\tau}$ such that

$$w_{\tilde{\tau}}(\tilde{F},t) = 0 \text{ and } w'_{\tilde{\tau}}(\tilde{F},t) < 0.$$

Assume, furthermore, that we can find sequences of biomass $\{F_i\}_{i=0}^{I}$ and waiting times $\{\tau_i\}_{i=1}^{I}$ such that

$$F_0 = \tilde{F},\; F_{i+1} < F_i,\; \&\; F_I = 0;\; \tau_1 = \tilde{\tau} \;\&\; \tau_{i+1} > \tau_i \tag{20}$$

and that these sequences fulfill

$$w_{\tau_1}(F_1,t) > 0 \text{ and for } i > 1 \; w_{\tau_i}(F_i,t) > 0 \;\&\; w_{\tau_i}(F_{i-1},t) > 0. \tag{21}$$

Then (19) implies that for each $F < \tilde{F}$ there is a τ such that $w_\tau(F,t) > 0$ and consequently that the harvesting rule will not prescribe harvest for any $F < \tilde{F}$.

A.3 The stationary problem

We will start by deriving sufficient conditions for the harvesting rule to have the form (5) in a stationary problem. We do not need to keep track of the calender time and the biomass, harvesting and value functions will be functions of only biomass.

We start by noting that in a stationary problem we will always want to harvest immediately if $F > F_{\max}$ since otherwise the biomass would decrease. We therefore need only analyze the case where we start from a biomass $F < F_{\max}$ which will also imply that the biomass will remain below $F_{\max}$ and we can assume that $g(F) > F$ always holds.

Assume now that we have found an $\tilde{F}$ such that $w_1(\tilde{F}) = 0$ and $w_1'(\tilde{F}) < 0$. This is what our numerical algorithm will give us as a candidate for the $\bar{F}$ of the decision rule (5). The concavity of w_1 implies that we then will have $w(F) < 0$ for all $F > \tilde{F}$. This, in turn, tells us that the decision rule will always prescribe harvest for $F > \tilde{F}$ (note that this includes $g(F)$ for any $F \in (\tilde{F}, F_{max})$). What remains in order to verify that the decision rule has the form (5) with $\bar{F} = \tilde{F}$ is that the decision rule does not prescribe harvest for any $F < \tilde{F}$. We can do this by finding sequences $\{F_i\}_{i=0}^{I}$ and $\{\tau_i\}_{i=1}^{I}$ that fulfill conditions (20) and (21) with $\tilde{\tau} = 1$.

A.4 The non-stationary problem

For the non-stationary problem, the biomass, harvest and value functions all depend on calendar time t. We assume that we are in period t and we know that the harvesting rules in all future periods are monotone and we also know their $\bar{F}$. We will here derive sufficient conditions for the harvesting rule to be monotone in period t as well. Since, in our simulations, the stationary situation follows the non-stationary, we can use the conditions derived above to verify that the harvesting rule is monotone in the first period of the stationary problem and we can work backwards from that.

Assume that we have found an $\tilde{F}$ such that

$$p\tilde{F} - c = \beta(1 - \pi_t)\left[V(g(\tilde{F}, t), t + 1) - V(g(0, t), t + 1)\right]$$

and

$$p > \beta(1 - \pi_t)V'(\tilde{F}_t, t)$$

where prime denotes derivative with respect to F. This is what our numerical algorithm will give us. Let

$$\tilde{\tau} = \min_{\tau \geq 1} g_\tau(\tilde{F}, t) \geq \bar{F}_{t+\tau}$$

be the number of time periods before harvest if the trees are not cut down in the current period.

We now want to derive sufficient conditions for $\tilde{F}$ to be the $\bar{F}$ of the harvesting rule (5). An alternative way to characterize $\tilde{F}$ is to say that

$$w_{\tilde{\tau}}(\tilde{F}, t) = 0 \text{ and } w'_{\tilde{\tau}}(\tilde{F}, t) < 0.$$

We now want to derive conditions under which for all $F > \tilde{F}$, $w_\tau(F, t) \leq 0$ for all $\tau \geq 1$ and under which for each $F < \tilde{F}$ there is a $\tau \geq 1$ such that $w_\tau(F, t) > 0$.

The monotonicity of the harvesting rule in all future time periods (and in particular in period $t+\tilde{\tau}$) implies that $w_\tau(F,t) < 0$ for all $\tau > \tilde{\tau}$ and $F > \tilde{F}$. A sufficient condition for $w_\tau(F,t) < 0$ for all $\tau \in [1, \tilde{\tau}-1]$ (assuming that there are such τ-values) is that for each such τ, there is an $F < \tilde{F}$ such that $w_\tau(F,t) > 0$. This follows from (18) since $w_\tau(\tilde{F},t) < 0$ for such τ. Numerically, we can verify this by simply testing values until we have found such an F for each τ.

Finally, we want to find conditions that guarantee that for each $F < \tilde{F}$ there is a $\tau \geq 1$ such that $w_\tau(F,t) > 0$. As for the stationary problem, we can do this by finding sequences $\{F_i\}_{i=0}^I$ and $\{\tau_i\}_{i=1}^I$ that fulfill conditions (20) and (21).

B Numerical algorithm

We start by drawing initial values of A, $F_{\max}$, D and π. Based on these we calibrate the harvesting cost c. We do this by iteratively solving (6) and updating c until it is 30% of the revenues in a stationary problem based on the initial parameters.

We then draw parameter values at time T and generate the parameter trajectories from 0 to T. Based on the parameters for each $t \in [0,T]$ we solve a stationary problem based on these parameters. We do this by first finding $V(0)$ from the solution to (6) and then by finding the $\bar{F}$ in the interval from (7) that fulfills the condition (8) using a search algorithm. We now have the reactive and non-reactive decision rules. Using (9) we can also get $V(F,T)$, which is the value function of both the reactive and optimal decision rules, for an arbitrary F. We also check the sufficient condition for the decision rule to be monotone in each of the stationary problems using the method described in Section A.3.

We then find the optimal decision rule by moving backwards from $t = T$ to $t = 0$ and in each step finding the $\bar{F}_t$ that fulfills (10). When computing the value function we, rather than storing the value function, use the formula

$$\begin{aligned} V(F,t) \quad = \quad & \sum_{s=1}^{\tau(F,t)} \beta^s \Pi_{0,s-1} \pi_s V(0, t+s) \\ & + \beta^{\tau(F,t)} \Pi_{0,\tau(F,t)} \left[p g_{\tau(F,t)}(F,t) - c + V\left(0, \tau(F,t)\right)\right], \end{aligned} \tag{22}$$

where

$$\tau(F,t) = \min_{\tau \geq 0} g_\tau(F,t) \geq \bar{F}_{t+T}$$

is the number of periods before the next harvest and where Π is defined in (12). We also store $V(0,t)$ for each t. When we get to $t=0$ we have the optimal decision rule and we check the monotonicity assumption as described in A.4.

The next step is to compute the value function associated with each decision rule at $t=0$ and for biomasses $F \in [0, \max\{\bar{F}_0^R, \bar{F}_0\}]$. To do this we need the value of biomass 0, $V(0,t)$, for all decision rules. For the optimal decision rule we already have $V(0,t)$ for all t. For the reactive decision rule we have $V^R(0,T)$ (which, at $t=T$, is equal to the optimal value function) and for the non-reactive decision rule we compute $V^{NR}(0,T)$ using (13). We then move backwards from $t=T-1$ to $t=0$ computing the reactive and non-reactive value functions $V^R(0,t)$ and $V^{NR}(0,t)$ using (11) and (13) respectively. We can now compute the value functions at $t=0$ for different biomasses for the optimal, reactive and non-reactive decision rules using (22), (11) and (13) respectively.

The last thing we do is to compute the value of following lagged reactive decision rules the same way as we computed the reactive decision rules but using a shifted version of the reactive decision rule.

References

Andersson, L., Hultkrantz, L., Mantalos, P., et al. (2013). Stumpage prices in sweden 1909-2011: Testing for non-stationarity. Technical report.

Boulanger, Y., Gray, D. R., Cooke, B. J., and De Grandpré, L. (2016). Model-specification uncertainty in future forest pest outbreak. *Global change biology*, 22(4):1595–1607.

Brunette, M., Costa, S., and Lecocq, F. (2014). Economics of species change subject to risk of climate change and increasing information: a (quasi-) option value analysis. *Annals of forest science*, 71(2):279–290.

Burchardt, T. (2005). Are one man's rags another man's riches? identifying adaptive expectations using panel data. *Social Indicators Research*, 74(1):57–102.

Chow, G. C. et al. (2011). *Usefulness of adaptive and rational expectations in economics.* Center for Economic Policy Studies, Princeton University.

Conlisk, J. (1996). Why bounded rationality? *Journal of economic literature*, 34(2):669–700.

Faustmann, M. (1849). Berechnung des werthes, welchen waldboden, sowie noch nicht haubare holzbestande fur die waldwirthschaft besitzen [calculation of the value which forest land and immature stands possess for forestry]. *Allgemeine Fotst-und Jagd-Zeitung*, 25:441–455.

Hanewinkel, M., Cullmann, D. A., Schelhaas, M.-J., Nabuurs, G.-J., and Zimmermann, N. E. (2013). Climate change may cause severe loss in the economic value of european forest land. *Nature Climate Change*, 3(3):203.

Heltorp, K. M. A., Kangas, A., and Hoen, H. F. (2017). Do forest decision-makers in southeastern norway adapt forest management to climate change? *Scandinavian Journal of Forest Research*, pages 1–13.

Hommes, C. H. (1994). Dynamics of the cobweb model with adaptive expectations and nonlinear supply and demand. *Journal of Economic Behavior & Organization*, 24(3):315–335.

Hong, H., Stein, J. C., and Yu, J. (2007). Simple forecasts and paradigm shifts. *The Journal of Finance*, 62(3):1207–1242.

Huang, G.-H. and Loucks, D. P. (2000). An inexact two-stage stochastic programming model for water resources management under uncertainty. *Civil Engineering Systems*, 17(2):95–118.

Huang, K. X., Liu, Z., and Zha, T. (2009). Learning, adaptive expectations and technology shocks. *The Economic Journal*, 119(536):377–405.

Jacobsen, J. B. and Thorsen, B. J. (2003). A danish example of optimal thinning strategies in mixed-species forest under changing growth conditions caused by climate change. *Forest Ecology and Management*, 180(1-3):375–388.

Jacobsen, J. B., Yousefpour, R., and Thorsen, B. J. (2013). Chapter viii: Climate change and practical forest management-when to worry and when not? *Adapting to climate change in European forests–results of the MOTIVE project*, page 83.

Kellomäki, S., Peltola, H., Nuutinen, T., Korhonen, K. T., and Strandman, H. (2008). Sensitivity of managed boreal forests in finland to climate change, with implications for adaptive management. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 363(1501):2339–2349.

Löfgren, K. G. (1985). Effect on the socially optimal rotation period in forestry of biotechnological improvements of the growth function. *Forest Ecology and Management*, 10(3):233–249.

Love, D. A. (2013). Optimal rules of thumb for consumption and portfolio choice. *The Economic Journal*, 123(571):932–961.

Mäkinen, A., Kangas, A., and Nurmi, M. (2012). Using cost-plus-loss analysis to define optimal forest inventory interval and forest inventory accuracy. *Silva Fennica*, 46(2):211–226.

Marcet, A. and Sargent, T. J. (1988). The fate of systems with" adaptive" expectations. *The American Economic Review*, 78(2):168–172.

McConnell, K. E., Daberkow, J. N., and Hardie, I. W. (1983). Planning timber production with evolving prices and costs. *Land Economics*, 59(3):292–299.

Mirrlees, J. A. and Stern, N. H. (1972). Fairly good plans. *Journal of Economic Theory*, 4(2):268–288.

Nerlove, M. (1958). Adaptive expectations and cobweb phenomena. *The Quarterly Journal of Economics*, 72(2):227–240.

Newman, D. H. (2002). Forestry's golden rule and the development of the optimal forest rotation literature. *Journal of Forest Economics*, 8(1):5–27.

Ohlin, B. (1921). Till frågan om skogarnas omloppstid. *Ekonomisk Tidskrift*, pages 89–113.

Pietilä, I., Kangas, A., Mäkinen, A., Mehtätalo, L., et al. (2010). Influence of growth prediction errors on the expected losses from forest decisions. *Silva Fenn*, 44(5):829–843.

Pressler, M. (1860). Aus der holzzuwachlehre (zweiter artikel). *Forst-und Jagd-Zeitung*, 36:173–191.

Pukkala, T. and Kellomäki, S. (2012). Anticipatory vs adaptive optimization of stand management when tree growth and timber prices are stochastic. *Forestry: An International Journal of Forest Research*, 85(4):463–472.

Rammig, A., Bebi, P., Bugmann, H., and Fahse, L. (2007). Adapting a growth equation to model tree regeneration in mountain forests. *European Journal of Forest Research*, 126(1):49–57.

Schou, E. and Meilby, H. (2013). Transformation of even-aged european beech (fagus sylvatica l.) to uneven-aged management under changing growth conditions caused by climate change. *European journal of forest research*, 132(5-6):777–789.

Schou, E., Thorsen, B. J., and Jacobsen, J. B. (2015). Regeneration decisions in forestry under climate change related uncertainties and risks: Effects of three different aspects of uncertainty. *Forest Policy and Economics*, 50:11–19.

Sethi, G., Costello, C., Fisher, A., Hanemann, M., and Karp, L. (2005). Fishery management under multiple uncertainty. *Journal of environmental economics and management*, 50(2):300–318.

Sims, C. A. (2003). Implications of rational inattention. *Journal of monetary Economics*, 50(3):665–690.

Spiro, D. (2014). Resource prices and planning horizons. *Journal of Economic Dynamics and Control*, 48:159–175.

Stollery, K. R. (2005). Climate change and optimal rotation in a flammable forest. *Natural Resource Modeling*, 18(1):91–112.

Susaeta, A., Carter, D. R., and Adams, D. C. (2014). Sustainability of forest management under changing climatic conditions in the southern united states: Adaptation strategies, economic rents and carbon sequestration. *Journal of environmental management*, 139:80–87.

Swedish Forest Authority (2012). Swedish statistical yearbook of forestry, 2012.

Tsur, Y. and Zemel, A. (1996). Accounting for global warming risks: resource management under event uncertainty. *Journal of Economic Dynamics and Control*, 20(6-7):1289–1305.

Van Kooten, G. C., Binkley, C. S., and Delcourt, G. (1995). Effect of carbon taxes and subsidies on optimal forest rotation age and supply of carbon services. *American Journal of Agricultural Economics*, 77(2):365–374.

van Veldhuizen, R. and Sonnemans, J. (2011). Nonrenewable resources, strategic behavior and the hotelling rule: An experiment. Discussion Paper TI 2011-04/1, Tinbergen Institute.

Vanclay, J. K. (1994). Modelling forest growth and yield: applications to mixed tropical forests. *School of Environmental Science and Management Papers*, page 537.

Winter, J. K., Schlafmann, K., and Rodepeter, R. (2012). Rules of thumb in life-cycle saving decisions. *The Economic Journal*, 122(560):479–501.

Yousefpour, R. and Hanewinkel, M. (2015). Forestry professionals' perceptions of climate change, impacts and adaptation strategies for forests in south-west germany. *Climatic Change*, 130(2):273–286.

Yousefpour, R. and Hanewinkel, M. (2016). Climate change and decision-making under uncertainty. *Current Forestry Reports*, 2(2):143–149.

Yousefpour, R., Jacobsen, J. B., Meilby, H., and Thorsen, B. J. (2014). Knowledge update in adaptive management of forest resources under climate change: a bayesian simulation approach. *Annals of forest science*, 71(2):301–312.

Yousefpour, R., Jacobsen, J. B., Thorsen, B. J., Meilby, H., Hanewinkel, M., and Oehler, K. (2012). A review of decision-making approaches to handle uncertainty and risk in adaptive forest management under climate change. *Annals of forest science*, 69(1):1–15.

Yousefpour, R., Temperli, C., Jacobsen, J. B., Thorsen, B. J., Meilby, H., Lexer, M., Lindner, M., Bugmann, H., Borges, J., Palma, J., et al. (2017). A framework for modeling adaptive forest management and decision making under climate change. *Ecology and Society*, 22(4).