University of St. Andrews

Discussion papers in Environmental Economics

http://www.st-andrews.ac.uk/gsd/research/envecon/eediscus/

Paper 2016-03

What a Difference a Stochastic Process Makes: Epidemiological-Based Real Options Models of Optimal Treatment of Disease

C.E. Dangerfield, A. E. Whalley, Nick Hanley, and C. A. Gilligan

Keywords: invasive species, pests and diseases, uncertainty, irreversibility, real options, forests.

JEL codes: G11, Q01, Q18, Q23, Q54, Q57.

What a difference a stochastic process makes: epidemiological-based real options models of optimal treatment of disease

C.E. Dangerfield[a], A.E. Whalley[b], Nick Hanley[c], and C.A. Gilligan[a]

[a] Department of Plant Sciences, University of Cambridge, Downing Street, Cambridge, CB2 3EA, UK
[b] Warwick Business School, The University of Warwick, Coventry, CV4 7AL, UK
[c] School of Geography & Geosciences, Irvine Building, University of St Andrews, North Street, St Andrews, Fife, KY16 9AL, UK

Abstract

The real options approach has been used within environmental economics to investigate the impact of uncertainty on the optimal timing of control measures to minimise the impacts of disease. Previous studies typically model the growth in infected area using geometric Brownian motion. The advantage of this is that it is simple and allows for closed form solutions. However, such a process is unbounded and so does not respect the natural upper boundary of the system which is determined by the maximum size of the host species. We show how the natural upper boundary can be incorporated endogenously into the decision problem, through the formulation of the stochastic process that describes growth in infected area. We find that ignoring the natural upper boundary of the system overestimates the value of the option to control, leading to delayed application of treatment. Indeed, when uncertainty is high or the disease is fast spreading then ignoring the upper boundary can lead to control never being deployed. Thus the results presented here have important implications for the way in which the real options approach is applied to determine optimal timing of disease control given uncertainty in future disease progression.

Keywords: invasive species; pests and diseases; uncertainty; irreversibility; real options, forests.
JEL codes: G11, Q01, Q18, Q23, Q54, Q57.

This work is funded jointly by a grant from BBSRC, Defra, ESRC, the Forestry Commission, NERC and the Scottish Government, under the Tree Health and Plant Biosecurity Initiative. We would also like to thank Dr Chris Quine, Dr Stephen Hendry and Professor John Healy for helpful comments on earlier drafts of the paper.

1. Introduction

Invasive pests and pathogens are increasing worldwide posing a threat to agriculture and forestry production (Gilligan 2008; Strange & Scott 2005). Deciding whether or not to implement expensive control strategies is complicated by uncertainties about the future spread of a pest or pathogen, and by the irreversible nature of many control strategies (felling a forest stand, for example). Therefore it may be best not to apply treatment immediately, but to wait and see how the disease progresses (Sims & Finnoff 2013).

The real options approach has been used to investigate the impact of uncertainty on the timing of control actions for disease outbreaks, as it provides a convenient way to couple uncertainty with economic analysis (Sims & Finnoff 2013). Viewing disease control as an option which can be exercised to reduce damage to the host species of a pest or pathogen, the real options framework can be used to determine the optimal timing of control (Ndeffo Mbah et al. 2010; Saphores 2000). That is, when is it optimal to use the option to control the disease? How long should we delay actions in order to learn more? This paper adds to this literature by incorporating more realistic, epidemiological models of disease spread within a real options set-up. As we show, such models change the management recommendations which emerge in terms of when it is best to intervene.

To incorporate uncertainty into the decision making process, the progress in the level of infection is described by a stochastic process. Traditionally, the real options approach assumes the change in the infected population follows a geometric Brownian motion (GBM), (Saphores 2000; Sims & Finnoff 2012). The advantage of such a simplifying assumption is that GBM is well-understood and allows for closed-from solutions to the real options problem. However, using the GBM process essentially assumes that the mean level of infection grows exponentially, and while this holds for the early part of the epidemic, it does not capture the longer term limiting behaviour effected by the finite number of susceptible hosts. In particular, such a model ignores the impact of the size of the susceptible population on the spread of infection, meaning that individual realisations of the process can violate this natural

upper boundary of the system, and so GBM is an unrealistic description of the increase in infection level for many diseases and pests. In this paper we show that sacrificing epidemiological accuracy for computational ease can significantly affect the disease management conclusions which one can draw from the real options approach.

Within the epidemiological modelling literature the evolution in the level of infection within the host population is derived based on the characteristics of infection spread (Keeling & Rohani 2008). These modelling frameworks are well established and have become an important tool in understanding the invasion and persistence of pathogens (Gilligan & van den Bosch 2008). Such an approach compartmentalises the population based on infection status (Keeling & Rohani 2008). For example, the simplest model assumes there are only two health states, either susceptible (S) or infected (I) and is termed the SI-model. Individuals move from the susceptible to infected compartment at a rate that is proportional to the current level of infection. In this paper we assume that infected hosts are immediately diseased and so throughout we use the terms infection and disease interchangeably. Since susceptibles are explicitly included within the modelling framework, the change in infection level also depends on the current level of susceptible individuals. This has the effect of accelerating the rate of infection when the proportion of susceptible hosts is high and decelerating it when the proportion is low. The evolution in the level of infection over time is described by a logistic-type term in the deterministic formulations of the models which is consistent with results from experimental studies (Large et al. 1946). The deterministic equation can be extended to include uncertainty by incorporating a drift term that represents the random fluctuations caused by variability in environmental conditions (Keeling & Rohani 2008). This is the approach taken in this paper, to ensure that the stochastic process describing the spread of infection is biologically realistic.

The principal objective of this paper is to show how epidemiological-based modelling approaches can be used within the real options framework and the impact these have relative to the standard GBM approach. By contrasting the standard approach (GBM) with two alternative, more realistic, models for infection dynamics, we show that using an inappropriate stochastic process to describe uncertainty in the level of infection leads to sub-optimal timing of control measures. The first model incorporates

logistic-type behaviour into the drift term, but has the same diffusion term as GBM and we call this the mean-reverting (MR) model. The second model incorporates a logistic-type noise term into the diffusion term as well as the drift term and is termed the logistic model. Both alternative formulations to GBM arise directly from the *SI*-type epidemiological model, according to different assumptions associated with the incorporation of the noise term into the deterministic equation. While GBM and the MR model have both been used previously within the literature, to the best of our knowledge this is the first time the logistic model has been used within the real options framework applied to the timing of disease control measures.

The main difference between these three processes is the way in which they do, or do not, incorporate the natural upper boundary of the system into the drift and diffusion coefficients of the stochastic process that describes the evolution in infection level. In particular, both the drift and diffusion coefficients of GBM are unbounded and so individual trajectories of the infection level process can become larger than the maximal number of susceptible hosts. This is also the case for the MR model, since while the drift coefficient is bounded, the diffusion coefficient is not. On the other hand, in the logistic model both the drift and diffusion terms are bounded above by the size of the host population. This means that individual trajectories respect the natural upper boundary of the system, thus ensuring realisations of the model are biologically realistic.

A key aim of this paper is thus to examine how incorporating the natural upper boundary of the infection level endogenously into the decision problem, via the stochastic process used to describe the level of infection, can alter the optimal timing of disease control within the real options framework. Since the focus here is on the comparison between the stochastic processes within the real options framework, rather than the real options model itself, we frame the problem for a simple situation of disease control application.

Our results suggest that ignoring the natural upper boundary of the system when there is a fixed size host population leads to over-valuation of the option to treat, and thus to (suboptimally) postponed deployment of control. Furthermore, in certain situations (high volatility and/or fast spreading infection) GBM and MR models result in control never being deployed, leading to losses due to disease damage.

This critical difference between the GBM and MR models and the logistic model arises due to the fact that the GBM and MR models over-estimate the value that can be obtained from applying treatment since they allow the level of infection to become unrealistically large. In particular, if treatment is deployed at the 'wrong' time then a portion of the option value is never realised. The structure of the remainder of this paper is as follows. In Section 2 we review the literature on control of invasive pests and pathogens. Section 3 motivates the new stochastic differential equations (SDEs) proposed to model the spread of infection and derives the associated real options model. The results are presented in Section 4 and finally Section 5 discusses the implications of the results and concludes.

2. Literature Review

The real options approach has been used within the environmental economics literature to incorporate ecological uncertainty into the valuation of management options of renewable resources such as forests (Insley & Rollins 2005; Insley & Lei 2007) as well as agriculture (Koppl & Koppl-Tuyna 2013; Tzouramani & Mattas 2004). Here, we focus on the control of invasive pests and pathogens, since the main application of the model presented in this paper is to determine the deployment of treatment to minimise the net costs of damage due to an epidemic outbreak. Given the ecological uncertainty that characterises the future spread of infection, and the irreversibility of actions that can be taken in response to such a risk, then a real options approach is attractive. The first application of real options to pests and diseases by (Saphores 2000) considered the optimal timing of pesticide application under future uncertainty in a pest population. Varying the level of uncertainty in the pest population dynamics, Saphores (2000) showed that greater uncertainty in future pest densities increases the threshold pest density at which it is optimal to spray, since the probability that the pest population will become small is larger. The analysis gives rise to the "wait and see" approach in dealing with invasive species; the idea being that when there is great uncertainty in the future dynamics of the invasive species, there is value in waiting to learn more before investing in control.

Sims and Finnoff (2012, 2013) have since extended this initial work in a number of ways. Sims & Finnoff (2013) consider the implications of the reversibility of the control strategy on (i) how long a regulator should wait to take action and (ii) the severity of the action taken (extent of control measure). They find that if control measures are partly reversible (for example trade bans), then it is optimal always to act as soon as possible, i.e. never to adopt a wait-and-see policy. This emphasises the fact that reversibility of actions is key, and that there is a trade-off between the speed and the severity (i.e. reversibility) of actions taken. Sims & Finnoff (2012) consider the impact of a spatial boundary on timing of control by treating the maximum area that can be infected as an upper limit within the real options framework. The spatial scale, which determines the maximum size of the population, can impact the timing and stringency of control strategies, and so incorporating an upper bound is important in planning measures to minimise losses from infection and disease damage (Sims & Finnoff 2012).

The real options approach has also been used to investigate the timing of specific control measures to minimise net damage costs from disease or invasive species. Sims (2011) considers the optimal timing of salvage harvest to recoup timber values following a disease outbreak in a forest crop and finds that slower rates of forest growth delay the optimal timing of salvage harvest, while large timber and non-timber values suggest more immediate action is optimal. Multiple interacting control options, namely chemical and biological control, are considered in Marten & Moore (2011). They find that biological control is sufficient to manage the pest, so long as infestation can be detected and controlled without substantial delay. However, if the pest reaches high levels before controls can be employed a more costly combined strategy is optimal for pest management (Marten & Moore 2011).

In most studies, the economic value of the damage caused by the disease or pest is the main focus, and the aim of the control is to minimise the damage cost within the real options framework. Ndeffo *et al.* (Ndeffo Mbah et al. 2010) instead consider the value added by applying control (namely treatment for a disease or pest) in terms of the monetary gain per unit of infection treated. Furthermore they incorporate a logistic term into the drift coefficient of the stochastic process describing the spread of infection, and so the mean growth of the process is limited by a parameter that represents the carrying capacity. In particular they find the new SDE leads to a difference in the optimal time to treat when

compared with the standard GBM assumption. The difference highlights the dependence of the real options approach on the formulation of the underlying model for uncertainty in the disease dynamics. However, the disadvantage of the approach taken in Ndeffo Mbah et al. (2010) is that the stochastic process can increase above the carrying capacity (Sarkar 2009). While in certain applications (e.g. harvesting fish) it is reasonable for the population to increase above the carrying capacity, in the study of disease spread this parameter represents a physical limit such as the finite number of trees or plants within a fixed area that can be infected. Therefore it is impossible for the area infected to reach a value greater than this bound and so trajectories of the stochastic process that go beyond this point do not have any applicable meaning in this context.

Although the impact of uncertainty in disease spread on the optimal timing of control measures has been extensively studied within the economics literature, there has been an artificial separation between traditional epidemiological models and those used within the real options framework. A significant contribution of this paper is to show how the uncertainty in disease spread can be formulated directly from basic epidemiological principles for incorporation and analysis in the real options framework. This leads to two different SDEs: the first of which has previously been studied in (Ndeffo Mbah et al. 2010) and the second of which incorporates a logistic-type term into the diffusion as well as the drift coefficient. To the best of our knowledge this is the first time that such a process has been used within the real options framework to determine the optimal timing of control. In particular this new formulation provides a natural way to incorporate the natural upper bound in the level of infection directly into the model. The complexity of this new stochastic process means that the real options model no longer permits closed form solutions. We frame the problem in a similar manner to (Ndeffo Mbah et al. 2010) by assuming that treatment eradicates infection, and so damage due to disease is not permanent, (and so in the forestry or agricultural context, once trees or crops have been treated there is no loss in timber or crop value).

3. Real Options Model

3.1 Setup

Consider an agricultural or forest disease outbreak in a particular crop or tree species at the landscape scale for which there is control treatment available that would eradicate current levels of infection. In particular, we consider that the number of trees or other plant hosts remains fixed. This is often the case in agriculture or even-aged forest stands where typically crops or trees are planted and then harvested at some future time period. We assume that the control treatment can be applied at any time for a one-off fixed cost C, and this treatment (e.g. spraying of pesticide or fungicide) is completely irreversible. Furthermore there is uncertainty in the future levels of infection due to environmental and demographic noise associated with the transmission process for infection. A decision-making authority is faced with the following choice: should treatment be administered immediately or should the decision-maker, owner or regulator wait to learn more about the progression of the epidemic? Waiting allows the decision-maker to determine whether the level of infection gets worse or better over time.

Traditional net-present-value (NPV) analysis would advocate undertaking treatment providing the value of the investment (i.e. the application of treatment and the associated savings in economic losses), is greater than the cost, C. However, due to uncertainty in disease dynamics combined with the irreversibility of the decision to treat, there is value in delaying treatment so as to learn more about the progress of the disease (Dixit & Pindyck 1994). That is, there is a value associated with the option to treat.

To include uncertainty into the decision making approach, we assume that the level of infection, I, can be described by a stochastic process. Traditionally the increase in the level of infection is assumed to follow geometric Brownian motion (GBM) and so the dynamics of the level of infection is given by the following SDE (Saphores 2000; Sims & Finnoff 2012; Sims & Finnoff 2013)

$$dI = \beta I dt + \sigma I dW, \qquad (1)$$

where β is the rate of transmission of infection, σ, which we term the volatility, is a parameter that scales the amount of uncertainty, I is the current level of infection and dW is a Wiener increment. In the limit as $\sigma \to 0$ equation (1) is equivalent to assuming deterministic exponential growth in the infected area. The advantage of using the GBM is that the logarithm of the infected area follows a Brownian motion, and so analytic solutions to the real options model can be obtained. However, it does not have any real-world epidemiological-basis, other than arguably during the very early stages of an epidemic, questioning it's applicability within this wider context.

3.2 Epidemiologically-Based Model of Uncertainty in Disease Spread

For a simple SI (Susceptible-Infected) epidemiological model (Keeling & Rohani 2008), the increase in the number of infected individuals is given by the product of the per capita rate at which a susceptible host contracts infection times the number of susceptible individuals. The rate at which a susceptible contracts infection is, in turn, given by the rate of transmission per infected contact, β, times the probability of contact with an infectious individual, $\mathrm{I}/\mathrm{I}_{\max}$, where $\mathrm{I}_{\max}$ is the maximum number of potential infected individuals. Assuming the total population remains constant, $\mathrm{I}_{\max}$ is equivalent to the total population size and so $\mathrm{S} = \mathrm{I}_{\max} - \mathrm{I}$. Initially, ignoring uncertainty in disease spread, the evolution in the level of infection is therefore given by the following ordinary differential equation (ODE),

$$\frac{\mathrm{dI}}{\mathrm{dt}} = \beta \mathrm{I}\left(1 - \frac{\mathrm{I}}{\mathrm{I}_{\max}}\right).$$

Uncertainty in disease spread is incorporated by assuming there is variability in the transmission parameter, β, driven by external forces. For example, fluctuations in temperature and climate have been shown to modify the infection rate (Sturrock et al. 2011). Here we assume that the fluctuations in transmission rate are stationary. There are two different approaches to incorporating this form of variation.

Firstly the 'corrected apparent infection rate' is perturbed, leading to $\beta(1 - I/I_{max}) \rightarrow \beta(1 - I/I_{max}) + \sigma\xi$ (Marcus 1991), where ξ is white noise and σ is a constant that controls the magnitude of the perturbation. The uncertain evolution of future disease spread is described by the following SDE,

$$dI = \beta I\left(1 - \frac{I}{I_{max}}\right)dt + \sigma I dW\,. \tag{2}$$

The SDE in Equation (2), which we refer to as the mean-reverting (MR) SDE, has been used within the real options framework in previous studies to describe the increase in infected area (Ndeffo Mbah et al. 2010) as well as the growth in pest populations (Marten & Moore 2011). When the level of infection reaches I_{max} , the magnitude of the diffusion term is non-zero and so there is a positive probability that the level of infection will exceed I_{max}, which is unrealistic for a fixed host population and questions the applicability of this stochastic process to the problem at hand.

Alternatively the transmission rate itself is perturbed, leading to $\beta \rightarrow \beta + \sigma\xi$ and so the evolution in the level of infection is given by the following SDE,

$$dI = \beta I\left(1 - \frac{I}{I_{max}}\right)dt + \sigma I\left(1 - \frac{I}{I_{max}}\right)dW. \tag{3}$$

We refer to the above equation as the logistic SDE. As the level of infection, I, reaches I_{max}, both the drift and diffusion term approach 0 and so the trajectories of the SDE remain within the interval $[0, I_{max}]$. Therefore the physical upper boundary of the total population size is preserved directly within the dynamics of the logistic SDE.

The logistic SDE approach provides a way of relating the uncertainty in future levels of infection to the randomness of the transmission process due to environmental factors. Therefore it provides an epidemiological-based approach to incorporating uncertainty into the decision problem. Furthermore this approach can be extended to more complex epidemiological models, for example in the case of diseases where there is an additional recovery state (termed the SIR model).

3.3 The Decision Problem

We assume that the only effect of treatment is to eradicate infection and, for simplicity, we assume treatment is applied instantaneously. Treatment could, for example, involve application of a pesticide or biocontrol agent. We only consider the gain in economic value from the timber or crop saved and not from the wider environmental damage that may be reduced. Therefore the value of applying treatment is simply:

$$V_t = pI_t, \tag{4}$$

where p is the gain in yield per unit of infected area treated, which is assumed to be constant over time and the level of infection, I_t, varies stochastically over time according to equation 1, 2 or 3. While more complicated formulations of the value from disease control have previously been used (Saphores 2000; Sims & Finnoff 2012) here we focus on a simple decision problem since the main focus of this paper is the comparison between the different stochastic processes rather than the complexity of the decision problem itself.

Viewing the application of treatment as an investment with value V_t, the decision problem can be viewed as a real option (Dixit & Pindyck 1994), which analogously to a financial option (Black & Scholes 1973) is the right but not the obligation to make an investment for a fixed price in the future. The payoff from applying treatment at time t is $V_t - C$ and so we want to maximise the expected present value,

$$F = \max \mathbb{E}\left[\left(V_{t_*} - C\right)e^{-rt_*}\right]. \tag{5}$$

Here t_* is the time in the future at which the decision is made, r is the discount rate and $\mathbb{E}$ denotes the expectation. The expectation must be taken since I_t (and therefore also V_t) is a stochastic process. This is an optimal stopping problem, and so we must find the threshold at which the value from applying treatment immediately is maximal.

Using standard methods from dynamic programming, the value of the option to apply treatment, F(V), must satisfy the following Bellman equation, (Dixit & Pindyck 1994) (see appendix for details)

$$\frac{1}{2}b(V)^2\frac{d^2F}{dV^2}+a(V)\frac{dF}{dV}-rF=0. \quad (6)$$

The functions $a(V)$ and $b(V)$ for each stochastic process used to describe the level of infection are given in Table 1.

Stochastic Process	$\mathbf{a(V)}$	$\mathbf{b(V)}$
Geometric Brownian Motion	βV	σV
Mean-Reverting SDE	$\beta V\left(1-\frac{V}{pI_{max}}\right)$	σV
Logistic SDE	$\beta V\left(1-\frac{V}{pI_{max}}\right)$	$\sigma V\left(1-\frac{V}{pI_{max}}\right)$

Table 1. Form of the functions in the Bellman equation for each infection process

$F(V)$ must also satisfy the following boundary conditions

$$F(0)=0$$

$$F(V^*)=V^*-C$$

$$\frac{d}{dV}F(V^*)=1.$$

V^* is the value at which treatment should be applied immediately. It represents the boundary between the *continuation region* and the *exercise region* (region in which treatment applied). The first condition follows from the fact that if the value of applying treatment goes to 0 it remains at 0, i.e. infection cannot be re-introduced from an outside source. The second condition is called the *value matching condition* (Dixit & Pindyck 1994), which states that when treatment is undertaken immediately the net gain is V^*-C. Finally the last condition is the *smooth pasting condition* which ensures optimality of the choice of V^*, since if F were not continuous at V^* then one could do better by investing at a different point (see Dixit & Pindyck (1994) for further discussion). We note that this is a *free-boundary problem* since the location of the boundary is unknown and must be determined as part of the solution. This threshold

in the value of treatment can easily be converted to a threshold in the level of infected area (I^*) by dividing through by p.

Due to the complexity of the logistic SDE, the problem no longer permits a closed form solution and so the problem is solved numerically in MATLAB. Since the functions $a(V)$ and $b(V)$ are non-linear for the logistic SDE, numerical solution of the ODE can be difficult. Therefore we instead find the long horizon limit of the associated finite time horizon problem using standard finite difference methods. See appendix for further details. Baseline parameter values, and ranges for those parameters that are varied are given in Table 2.

Model Parameter	***Description***	***Base case (Range)***
β	Infection transmission rate	0.05 ([0.05, 0.8])
σ	Volatility	0.5 ([0.1, 0.9])
I_{max}	Carrying capacity	100
C	Cost of treatment	20
r	Risk-free discount rate	0.1
p	Gain in yield per unit of infected area treated	1

Table 2. Parameter values used in numerical simulations.

4. Results

The solution to the free boundary problem associated with each SDE provides the value of the corresponding option to treat as a function of the treatment value (V). Figure 1 shows the value of this option as a function of the treatment value for the three different SDEs. Providing the value of the option to treat, $F(V)$, is greater than the NPV of immediate treatment, $F(V) > V - C$, there is value in retaining the option to treat, and so it is beneficial to wait. When $F(V) = V - C$, there is no additional gain in waiting and so treatment should be applied immediately. The value of treatment at which $F(V)$ first

equals the NPV, the threshold value of treatment, V^*, is the boundary between the waiting region and the immediate treatment region and is also shown for each SDE in Figure 1.

The threshold value of treatment, V^*, for each SDE corresponds to a threshold level of infection, I^* ($V^* = pI^*$), at which treatment should be applied immediately. To provide a clearer illustration of situations when the optimal time to treat violates the natural upper boundary $I^* \leq I_{max}$, rather than the threshold value of treatment, V^*, we investigate how the threshold infection level at which treatment is applied *as a proportion* of the infected area, which we term the *threshold infected proportion*, I^*/I_{max}, varies with different stochastic processes. We also explore the sensitivity of I^*/I_{max}, to the epidemiological parameters σ and β and discuss the implications for policy (note that, as in Sims & Finnoff (2013), we do not consider uncertainty over the economic parameters such as the price of agricultural commodity or timber). Finally, we consider the loss in value of the option to treat that arises from sub-optimally delaying treatment.

4.1 Impact of epidemiologically-based models on the value of treatment threshold

Using an epidemiologically-based SDE to describe uncertainty in disease spread (mean-reverting or logistic SDE) decreases the threshold value and so treatment is deployed when a lower proportion of the area is infected, compared with the standard GBM assumption (Figure 1). The effects hold for a wide range of parameter values (Figure 2). As a result, the threshold at which treatment is applied is closer to the zero NPV level (shown as the purple dotted-stared line in Figure 2) under the mean-reverting and logistic models than under GBM. As the level of infection becomes large both the drift and diffusion terms for GBM are greater than for the other two SDEs. There is thus a greater probability that the value of treatment, V, becomes large in the future which means the expected return from waiting is greater under the assumption of GBM. On the other hand, in the case of the logistic SDE as the level of infection approaches I_{max} the magnitude of both the drift and diffusion terms tend to zero and so the value of treatment cannot go above some maximal level $V_{max} = pI_{max}$, corresponding to the value of

treatment when the whole area is infected ($I = I_{max}$). Hence there is a lower value to be obtained from waiting, and so it is optimal to apply treatment at a lower threshold level of infection (i.e. at a lower value of treatment). The threshold at which to act for the mean-reverting SDE lies between GBM and the logistic SDE. This reflects the trade-off between the growth in the drift term, which becomes dampened as the level of infection approaches I_{max}, as for the logistic SDE, and the diffusion term which becomes large with increasing levels of infection, as for GBM. Therefore the expected return from waiting is greater than for the logistic SDE but less than for GBM, and so the threshold value lies between the two.

The relative magnitudes of the infection rate (β) and the discount rate r are important. If $\beta = r$, the option value under GBM never intersects the standard NPV (red dashed line) leading to the conclusion that treatment should never be applied. In fact, such a situation is equivalent to an American call in financial options and it has been shown that it is never optimal to exercise the option early, i.e. treatment should not be applied (Wilmott et al. 1995). This is typically why in previous studies the assumption is made that $\beta < r$, (Sims & Finnoff 2012) and indeed this is why we investigate the behaviour of GBM over a subset of the parameter values used for the other two processes. For the mean-reverting and logistic SDEs a finite threshold value still exists if $\beta \geq r$ (Figure 2). Therefore, using an epidemiological-based approach to modelling disease uncertainty ensures there are no restriction on the range of transmission rate (β) and discount rate (r) parameters investigated, unlike the case with GBM.

4.2 Impact of uncertainty and transmission rate on the timing of treatment and value of waiting

For all processes, the threshold value is a monotonic increasing function of volatility (Figure 2a), which is consistent with options theory (Dixit & Pindyck 1994). The treatment threshold, V^* (or I^*), represents the value of treatment (infection level) at which the benefit of immediate treatment, $V^* - C$, exactly equals the cost of immediate treatment in the form of the value of retaining the option to treat at a later date, which is forgone once treatment is applied. Increasing the volatility increases the option value of

waiting by increasing the probability of more extreme outcomes including those where the net benefit of treatment is high, whilst leaving the value of immediate treatment unchanged, and thus increases the treatment threshold.

Figure 2a shows that the rate of increase in the threshold infected proportion, (I^*/I_{max}), with increasing volatility (σ) is greater for GBM and the mean-reverting SDE than for the logistic equation. Therefore, the difference between the optimal time to treat between the three processes is greatest when uncertainty (σ) is large. Furthermore, the threshold infection level at which treatment is applied, I^*, increases above the maximum area that can be infected (I_{max}) when $\sigma > 0.4$ for our base case parameters in the case of GBM and when $\sigma > 0.6$ in the case of the mean-reverting SDE. Hence, for large volatility, the threshold at which to apply treatment obtained from GBM or the mean-reverting equation is unattainable. Use of the GBM and mean-reverting models would therefore imply it would never be optimal to apply treatment for large volatility. This is not the case for the logistic model, where the threshold I^* always remains below I_{max}, and so is attainable (Figure 2). We interpret this as follows. As volatility increases, the probability that the treatment value is large increases more for the mean-reverting or GBM-type SDE, whereas for the logistic SDE there is an upper bound to the value of treatment at $V_{max} = pI_{max}$. Therefore, as volatility increases, the threshold at which to act remains bounded since the logistic SDE cannot reach large values above $V_{max} = pI_{max}$ in the future, no matter how large the volatility becomes.

Similarly, as the transmission rate increases, the threshold infected proportion at which to act increases for both the mean-reverting and logistic SDEs as shown in Figure 2b. The threshold also increases for GBM since when $\beta \geq r$ the threshold is infinite. That is, as β increases above the discount rate, the threshold at which to act increases to a level that is unattainable in finite time. Since we assume the value of treatment is proportional to the level of infection, treatment is most valuable when the level of infection is large. As the transmission rate increases, the level of infection is growing faster and so it is beneficial to wait longer. Figure 2b also shows that, for very fast spreading diseases ($\beta > 0.2$ for our base case parameters) the threshold infection level I^* for the mean-reverting SDE can go above I_{max}, as with the case for large uncertainty. However, once again the logistic equation threshold remains

within an attainable region (i.e. $I^* < I_{max}$). For large β, the threshold infected proportion levels off close to $I/I_{max} = 1$.

4.3 Differences between the stochastic processes and implications for policy

The crucial difference between the logistic SDE compared with the GBM and mean-reverting SDE is that the logistic model allows for the natural upper boundary in the value of treatment (V_{max}) to be incorporated endogenously into the decision problem. This arises naturally from the forms of the drift and diffusion coefficients (the functions $a(V)$ and $b(V)$) in the logistic equation, which ensure that the chance that the level of infection increases above the maximum host population size is zero. On the other hand, the GBM and mean-reverting SDEs allow for the level of infection to rise above the natural boundary for the level of infection given by I_{max}. Since the value of treatment is higher the greater the level of infection, this means that these two models over-estimate the value that can be obtained from treatment when the level of infection is high. Furthermore, in both the GBM and mean-reverting models, the probability that the level of infection will increase above I_{max} increases with both volatility (σ) and transmission rate (β), thus increasing these effects for highly volatile and/or highly transmissible infections.

While the impact of an upper spatial boundary on the timing of control has previously been considered by (Sims & Finnoff 2012), the authors essentially enforce the upper boundary within the boundary conditions of the decision problem. This approach essentially truncates the stochastic process at the upper boundary, while the advantage of the method described here is that the upper boundary is incorporated directly into the process that describes the evolution in the disease dynamics. Therefore, the timing of control arises naturally from the form of the stochastic process rather than as a result of additional boundary conditions that are imposed on the system.

The implications of the different models for decision makers can be seen in policy plots (Figure 3). Under the logistic model, the waiting region is smallest, and so treatment will be applied earlier in the

course of the epidemic, compared with assumptions based upon the GBM and the mean-reverting models. Furthermore, when there is great uncertainty (σ large) or the disease is fast spreading (β large) there is no region in which treatment should be deployed under GBM or the mean-reverting equation, while under the logistic equation treatment should be applied when the proportion of infected area becomes large. Since the final size of the epidemic can be close to I_{max}, not applying treatment could lead to large losses. Therefore, when uncertainty is large or the epidemic is fast spreading, the assumption of the logistic model is more appropriate than GBM or the mean-reverting equation since it provides a realistically attainable threshold at which to apply treatment.

4.4 Loss in value from delaying treatment too long

To quantify the potential economic impact from inaccurate assumptions regarding the uncertainty in future disease spread, we consider the value of the option to treat if the evolution of infection follows the logistic equation but where the decision maker applies treatment at the 'wrong' threshold. This 'wrong' threshold is derived from either the mean-reverting or GBM model[1]. Treating at the wrong threshold leads to a loss in value (Figure 4). The loss arises because treatment thresholds are higher under mean-reversion or GBM than under the logistic assumption, and hence sub-optimally high when infection actually evolves according to the logistic model. This reduces the value of treating at a sub-optimally high threshold (Figure 4, red and green lines), which fall below the standard NPV (Figure 4, dashed black line). The loss in value is greatest for GBM, where a significant proportion of the optimal value (26% for our base case parameters) is lost by exercising control too late, i.e. at too high a threshold, while treating at the lower mean-reverting threshold leads to a lower loss of 5%[2] for our base case parameters. Therefore, if the model used does not appropriately capture uncertainty in infection

[1] This is calculated by solving the differential equation associated with the logistic equation but imposing the threshold values obtained from either the mean-reverting or the GBM models (see Appendix for details).
[2] As for standard GBM, the proportion of the optimal value obtained if the treatment occurs at the wrong threshold is independent of the value of treatment V.

dynamics, then the excessive delay before treatment also implies that the full value of the option to treat is not realised. Figure 5 shows that the proportional loss from using the wrong assumptions increases with volatility for both the mean-reverting and GBM models. This is consistent with our earlier finding that the difference between the thresholds for the logistic and both the mean-reverting and GBM is greatest for large uncertainty. Also the proportional loss from using the mean-reverting thresholds is smaller than using the GBM thresholds since the mean-reverting threshold is closer to the threshold for the logistic equation (Figure 1).

5. Conclusions

The real options approach has proved useful in investigating the optimal timing of strategies to control disease and pest outbreaks. Uncertainty in disease spread is typically assumed to follow Geometric Brownian Motion due to its simplicity, rather than any epidemiological basis. However, ignoring asymptotic boundaries for infection, may result in treatment being deployed too late, and indeed not at all if the level of uncertainty is high. This has significant implications for the formulation of the real options approach used to inform disease control policy.

We have shown that the stochastic process describing uncertainty in disease spread can be derived directly from basic epidemiological principles. Comparing the mean-reverting and logistic models with the standard approach of GBM we find the following key differences:

1) When uncertainty is large or the epidemic is spreading fast, the threshold values for GBM and the mean-reverting model are unattainable, implying treatment should never be applied. This is not the case for the logistic model, where the threshold at which to treat always lies in a realistic range.

2) The threshold value at which to act is greatest for GBM and least for the logistic SDE. Therefore delaying treatment based on the results from the GBM or mean-reverting formulation may result in the late application of treatment.
3) Applying treatment at the wrong threshold leads to a loss in value, and this loss will be greatest when uncertainty is large.
4) The differences in the threshold at which to treat between the three models increases with increasing volatility and transmission rate. Therefore the implications of the different models will be most disparate when there is great uncertainty (high σ) and the infection is spreading rapidly (high β).

Describing the uncertainty in disease dynamics in terms of the logistic SDE leads to qualitatively similar implications for decision makers as GBM or the mean-reverting equation, namely that there is value to be gained from waiting to apply treatment. However, the important difference is that the logistic SDE implies that treatment should be applied earlier (i.e. at a lower level of total infected area) and that it is always optimal to apply treatment before the whole area becomes infected, even if the uncertainty is very large. This important difference between the logistic model and the GBM or mean-reverting models arises since the logistic model respects the naturally upper boundary of the system, namely the host population size, while the GBM and mean-reverting models do not. Since the level of infection under the GBM and mean-reverting model assumptions can become greater than the upper boundary of the system, this leads to the over-valuation of the option to treat within the real options frame-work. Furthermore, when uncertainty is large or the disease fast spreading, the probability of the level of infection becoming large under the GBM and mean-reverting models is greatest. On the other-hand, due to the form of the drift and diffusion coefficients within the logistic model, the level of infection remains below the natural upper boundary, even when uncertainty is large or the disease fast spreading. The result is that in these regions of parameter space, the delay in treatment under the GBM or mean-reverting models will be greatest.

If treatment is excessively delayed, by using thresholds from the GBM or mean-reverting models when uncertainty in the level of infection is more appropriately described by the logistic SDE, only a portion

of the optimal value is obtained. As uncertainty, and hence the difference between the thresholds, increases this proportion decreases, and so when uncertainty is high using predictions from the mean-reverting or GBM models leads to greater losses due to implementing treatment at the wrong threshold.

A further disadvantage of the GBM approach is that it restricts the range of parameter regimes that can be investigated since if $\beta \geq r$ the threshold value at which to apply treatment is not finite. Since the time horizon for forest management is long, typically 40-100 years, the discount rate used is usually taken to be lower than rates used for shorter-term projects[3] (Pindyck 2006). Therefore, the assumption of GBM to characterise the uncertainty in the level of infection restricts application of the real options approach to slower spreading diseases.

Since the main aim of this paper is to investigate the effect of more realistic characterisations of disease uncertainty on the conclusions of the real options model, we have used the simplest formulation of the value gained from treatment application. That is, we assume that the value of treatment is in the recovery of currently infected hosts, which would be applicable, for example, in the case of a pesticide or fungicide that kills the pest or pathogen upon application, allowing complete recovery of the previously infected hosts. However, for a number of tree diseases, for example chalara (ash dieback) and phytophthora ramorum, there is no treatment available that will completely eradicate the infection. Therefore, once a tree is infected there is an irreversible loss in value due to damage of the tree. We could extend this work to include such situations by incorporating the additional sunk benefit associated with early application of treatment due to the reduction in potential future damage as a result of disease (Dixit & Pindyck 1994, pp.412–418). The decision problem could also be extended by including economic uncertainty (market risks) which is potentially correlated with disease risk as well as considering the interaction effects between economic and ecological uncertainty. We leave these aspects for future work.

[3] What the appropriate discount rate should be is itself a frequently discussed topic, (Gollier 2011), with previous studies using discount rate values of 8% (Sims & Finnoff 2012) and 10% (Ndeffo Mbah et al. 2010). For many years, the UK Forestry Commission used a rate much lower than the standard public sector test rate of discount, partly to take into account the longer time horizons in forest investment (Hanley & Spash 1993).

One important issue with using the real options framework to assess the impact of uncertainty on the optimal timing of disease control measures is that the form of the process that best represents uncertainty in disease spread is rarely discussed. The approach described here provides a convenient method for formulating the future uncertainty since it is derived directly from epidemiological principles of pathogen transmission. Epidemiological modelling has become a well-established field that uses mathematical models to describe the evolution of a disease outbreaks based on epidemiological mechanisms of the spread of an infection. Such models have been successfully deployed to inform the structure of control measures within human (Choi et al. 2010), animal (Brooks-Pollock et al. 2014) and plant (Cunniffe et al. 2014) health. However, to date such modelling approaches have been largely ignored within the real options literature on the optimal timing of disease control.

This paper represents the first attempt to incorporate traditional epidemiological models of disease spread into a real options framework. Since this leads to greater complexity in the stochastic process describing disease spread, closed form solutions are no longer attainable and the problem must be solved numerically. However, we have shown that ignoring the epidemiological principles of disease spread can result in late application of treatment, or indeed may suggest treatment should never be applied. This result cautions against over-simplifying the description of uncertainty in disease dynamics within the real options approach to the optimal timing of disease control.

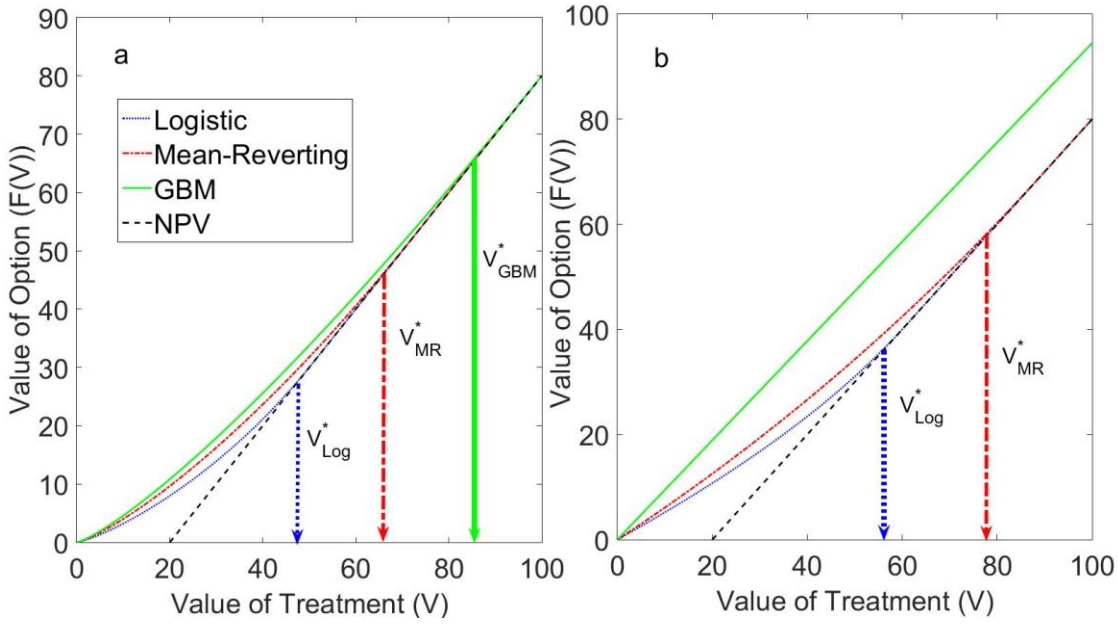

Figure 1: Value of the option to treat as a function of the value of treatment for the three different stochastic process assumptions. The standard NPV is shown as a dashed line. Plot (a) is for the case $\beta = 0.05$ and $r = 0.1$ and so $\beta < r$. Plot (b) is the case when $\beta = r = 0.1$, and it can be seen in this case that GBM never intersects the standard NPV, showing it is never optimal to apply treatment. The volatility is taken to be $\sigma = 0.5$, and the other parameter values are given in Table 2.

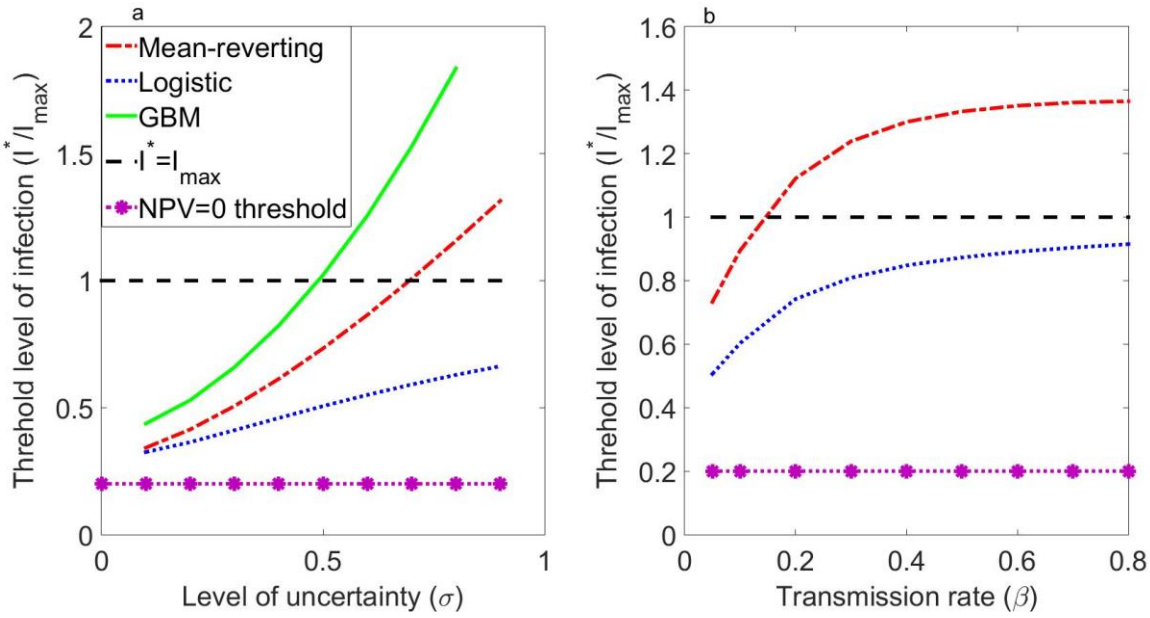

Figure 2: Threshold level of infection (I^*/I_{max}): (a) as a function of volatility (σ) for all three stochastic processes and (b) as a function of the transmission rate (β) for the mean-reverting and logistic SDEs only. The dashed black line shows when $I^* = I_{max}$, and so it is clear that the mean-reverting and GBM can go above this natural boundary while this is not the case for the logistic SDE. Other parameter values are given in Table 2.

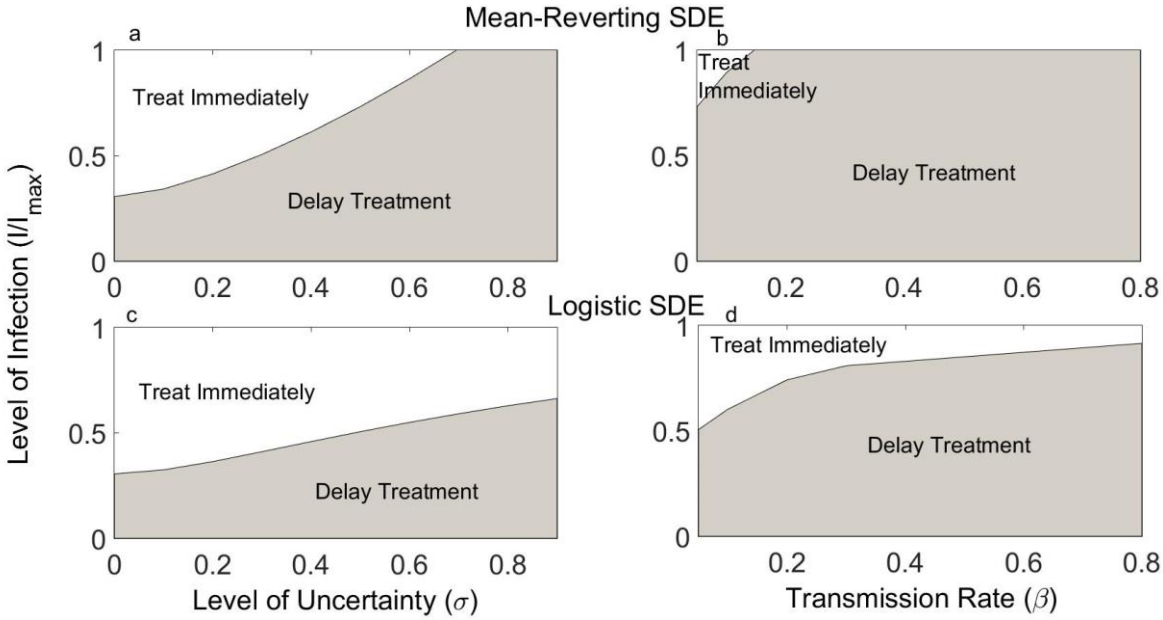

Figure 3: Policy plots showing the region in which treatment should be applied immediately and where treatment should be delayed for the mean-reverting SDE ((a) and (b)) and the logistic SDE ((c) and (d)) for different levels of uncertainty ((a) and (c)) and transmission rates ((b) and (d)). Other parameter values are given in Table 2.

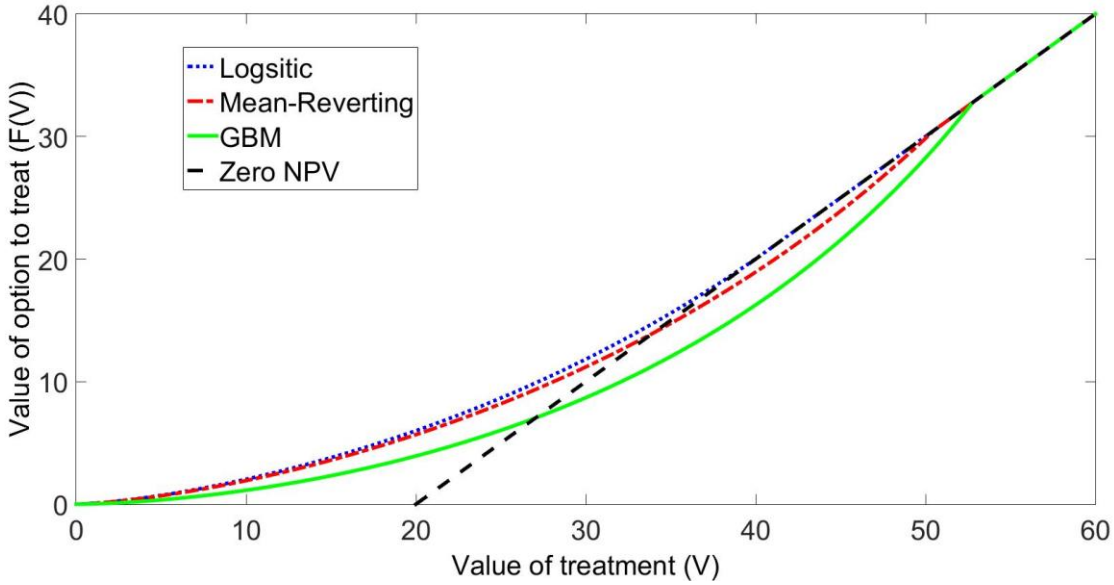

Figure 4: Value of the option to treat ($F(V)$) as a function of the treatment value (V) for the logistic model when treatment is applied at the wrong threshold, namely the threshold obtained from the mean-reverting model (red dot-dashed line) and the GBM model (green solid line). Also shown are the values of the option when treatment is applied at the optimal threshold (blue dotted line) and the standard NPV (black dashed line), i.e. the value of the option when NPV is zero. The volatility is taken to be $\sigma = 0.3$ and the transmission rate is taken to be $\beta = 0.05$.

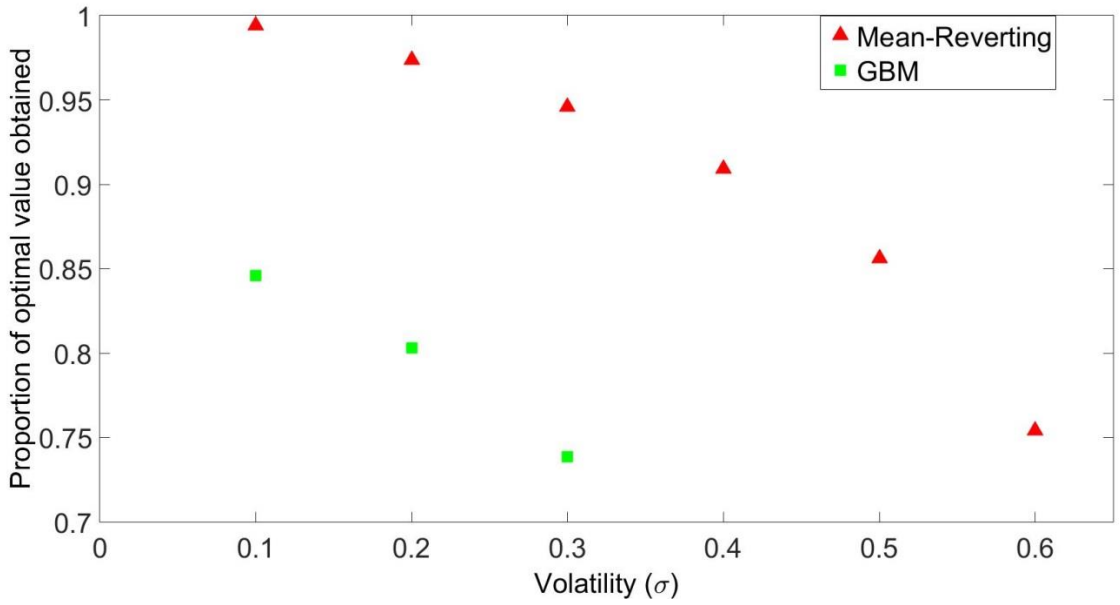

Figure 5: The proportion of the optimal value that is obtained under the logistic model when the treatment threshold from the mean-reverting (red triangles) and the GBM (green squares) models are enforced. The proportion of value obtained is shown as a function of the volatility (σ).

Appendix

Formulation of the Bellman Equation

The evolution in the value of treatment is given by the SDE,

$$\mathrm{dV} = \mathrm{a(V)dt} + \mathrm{b(V)dW},$$

where the functions $\mathrm{a(V)}$ and $\mathrm{b(V)}$ differ depending of the stochastic process assumed (see Table 1 in the main text).

At any point in time t, the decision authority either applies treatment (i.e. invests in treatment) and obtains the net gain $\mathrm{V(t) - C}$, or waits and obtains the expected change in the value of the option. Therefore, we can formulate the decision of whether or not to apply treatment as an optimal stopping problem, described by the following

$$\mathrm{F(V, t)} = \max[(\mathrm{V - C}), \mathrm{e}^{-\mathrm{r\,dt}}\mathbb{E}[\mathrm{F(V + dV, t + dt)|V}]].$$

The first term on the right is the value the decision authority would obtain if treatment were applied immediately while the second term is the expected value in retaining the option to treat in the future, discounted at rate r. Therefore, in the region where it is optimal to wait and not treat immediately, referred to as the *continuation region*, the second term on the right hand side of the equation will be larger. It follows that in the continuation region

$$rF(V,t)dt = \mathbb{E}[dF].$$

Since $V(t)$ is a stochastic process, then using Ito's Lemma we can evaluate dF as follows,

$$dF = \frac{\partial F}{\partial t} + \frac{1}{2}\frac{\partial^2 F}{\partial V^2}(dV)^2 + \frac{\partial F}{\partial V}dV.$$

Substituting in the equation for dV, and noting that $\mathbb{E}[dW] = 0$, we obtain the following

$$\mathbb{E}[dF] = \frac{\partial F}{\partial t}dt + \frac{1}{2}b(V)^2\frac{\partial^2 F}{\partial V^2}dt + a(V)\frac{\partial F}{\partial V}dt,$$

where $a(V)$ and $b(V)$ are the drift and diffusion coefficients, respectively, of $V(t)$ (given in Table 1 in the main text).

Using the fact that $rFdt = \mathbb{E}[dF]$, we arrive at the following partial differential equation (PDE)

$$\frac{\partial F}{\partial t} + \frac{1}{2}b(V)^2\frac{\partial^2 F}{\partial V^2} + a(V)\frac{\partial F}{\partial V} - rF = 0.$$

Whilst, in principle, the value of the option to treat can be a function of time, t, as well as the current value of treatment, *V*, since there is no terminal date beyond which treatment cannot be applied the payoff is independent of time and the problem is time invariant, so $\frac{\partial F}{\partial t} = 0$ and *F* is the solution of the ODE (6):

$$\frac{1}{2}b(V)^2\frac{d^2F}{dV^2} + a(V)\frac{dF}{dV} - rF = 0,$$

subject to the following boundary conditions

$$F(0,t) = 0$$

$$F(V^*,t) = V^* - C$$

$$\frac{dF}{dV}(V^*,t) = 1.$$

Solution Method

The problem described in Section 3.3 assumes that the decision maker can apply treatment at any point in the future, i.e. the time horizon for the problem is infinite. We solve equation (6) (see main text) and associated boundary conditions by finding the long horizon limit of the associated finite horizon problem

$$\frac{\partial \mathrm{F}}{\partial \mathrm{t}} + \frac{1}{2}\mathrm{b}(\mathrm{V})^2 \frac{\partial^2 \mathrm{F}}{\partial \mathrm{V}^2} + \mathrm{a}(\mathrm{V}) \frac{\partial \mathrm{F}}{\partial \mathrm{V}} - \mathrm{rF} = 0, \qquad (A\ 1)$$

such that $\mathrm{F}(0,\mathrm{t}) = 0$, $\mathrm{F}(\mathrm{V}^*(\mathrm{t}),\mathrm{t}) = \mathrm{V}^*(\mathrm{t}) - \mathrm{C}$, $\frac{\partial \mathrm{F}}{\partial \mathrm{V}}(\mathrm{V}^*(\mathrm{t}),\mathrm{t}) = 1$, $\mathrm{F}(\mathrm{V},\mathrm{T}) = \max(0, \mathrm{V} - \mathrm{C})$, which assumes treatment can only be applied until time *T*, so both the differential equation and the threshold value, $\mathrm{V}^*(\mathrm{t})$ depend explicitly on time. This free boundary problem can be solved straightforwardly using the Crank-Nicolson finite-difference method with successive-over-relaxation (SOR) (Wilmott et al. 1995).

In order to solve the problem numerically, an upper boundary condition must be stipulated in the case where it is not worthwhile treating immediately. We assume that as V becomes large $\partial^2 \mathrm{F}/\partial \mathrm{V}^2 = 0$ which is used to obtain the finite difference scheme at the upper boundary (Insley 2002).

We take *T* sufficiently large that $\frac{\partial \mathrm{F}}{\partial \mathrm{t}} \to 0$. We find that taking $\mathrm{T} = 100$ is more than a sufficient length of time to ensure steady state is reached for all three stochastic processes.

The values of model parameters and parameters within the numerical method used in simulations are given in Table 2 in the main text.

Exercising the option at the wrong threshold

In order to determine the value of the option to treat if the wrong threshold is implemented in the logistic model we solve the PDE given by (A1) with $\mathrm{a}(\mathrm{V})$ and $\mathrm{b}(\mathrm{V})$ taken from the logistic SDE, again using

the Crank-Nicolson finite-difference scheme. However, we no longer have a free boundary problem, since we know the upper boundary at the threshold value (where the threshold value is taken from that obtained for the mean-reverting or GBM model). Let V^W be the threshold value obtained from the mean-reverting or GBM model. The boundary conditions for (A1) are thus

$$F(0) = 0 \text{ and } F(V^W) = V^W - C.$$

References

Black, F. & Scholes, M., 1973. The Pricing of Options and Corporate Liabilities. *Journal of Political Economy*, 81(3), pp.637–654.

Brooks-Pollock, E., Roberts, G.O. & Keeling, M.J., 2014. A dynamic model of bovine tuberculosis spread and control in Great Britain. *Nature*, 511(7508), pp.228–231.

Choi, Y.H. et al., 2010. Transmission dynamic modelling of the impact of human papillomavirus vaccination in the United Kingdom. *Vaccine*, 28(24), pp.4091–4102.

Cunniffe, N.J. et al., 2014. Cost-Effective Control of Plant Disease When Epidemiological Knowledge Is Incomplete: Modelling Bahia Bark Scaling of Citrus. *PLoS Computational Biology*, 10(8).

Dixit, A. & Pindyck, R., 1994. *Investment Under Uncertainty*, Princeton University Press.

Gilligan, C. A, 2008. Sustainable agriculture and plant diseases: an epidemiological perspective. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, 363(1492), pp.741–759.

Gilligan, C. A & Van den Bosch, F., 2008. Epidemiological models for invasion and persistence of pathogens. *Annual review of phytopathology*, 46, pp.385–418.

Gollier, C., 2011. Pricing the future : The economics of discounting and sustainable development. *Princeton University Press*, (230589), pp.1–210.

Hanley, N. & Spash, C.L., 1993. *Cost-Benefit Analysis and the Environment*, Cheltenham: Edward

Elgar.

Insley, M., 2002. A Real Options Approach to the Valuation of a Forestry Investment 1. *Journal of Environmental Economics and Management*, 44, pp.471–492.

Insley, M. & Lei, M., 2007. Hedges and Trees: Incorporating Fire Risk into Optimal Decisions in Forestry Using a no-Arbitrage Approach. *Journal of Agriculture and Resource Economics*, 32(3), pp.492–514.

Insley, M. & Rollins, K., 2005. On solving the multirotational timber harvesting problem with stochastic prices: A linear complementarity formulation. *Amer. J. Agr. Econ.*, 83(3), pp.735–755.

Keeling, M.J. & Rohani, P., 2008. *Modeling Infectious Diseases In Humans and Animals*, Princeton University Press.

Koppl, S. & Koppl-Tuyna, M., 2013. Real Options for Agricultural Investments. *Quantiative Methods in Economics*, XIV(1), pp.253–264.

Large, E.C., Beer, W.J. & Patterson, J.B.E., 1946. Field trials of copper fungicides for the control of potato blight; spray retention. *The Annals of applied biology*, 33(1), pp.54–63.

Marcus, R., 1991. Deterministic and stochastic logistic models for describing increase of plant diseases. *Crop Protection*, 10(2), pp.155–159.

Marten, A.L. & Moore, C.C., 2011. An options based bioeconomic model for biological and chemical control of invasive species. *Ecological Economics*, 70(11), pp.2050–2061.

Ndeffo Mbah, M.L. et al., 2010. Economically optimal timing for crop disease control under uncertainty: an options approach. *Journal of the Royal Society, Interface / the Royal Society*, 7(51), pp.1421–8.

Pindyck, R.S., 2006. Uncertianty in Environmental Economics. *NBER Working Paper Series*.

Saphores, J., 2000. Pest Population : A Real Options Approach. *American Journal of Agricultural Economics*, 82, pp.541–555.

Sarkar, S., 2009. Optimal fishery harvesting rules under uncertainty. *Resource and Energy Economics*,

31(4), pp.272–286.

Sims, C., 2011. Optimal timing of salvage harvest in response to a stochastic infectation. *Natural Resource Modeling*, 24(3), pp.383–408.

Sims, C. & Finnoff, D., 2012. The role of spatial scale in the timing of uncertain environmental policy. *Journal of Economic Dynamics and Control*, 36(3), pp.369–382.

Sims, C. & Finnoff, D., 2013. When is a "wait and see" approach to invasive species justified? *Resource and Energy Economics*, 35(3), pp.235–255.

Strange, R.N. & Scott, P.R., 2005. Plant disease: a threat to global food security. *Annual review of phytopathology*, 43, pp.83–116.

Sturrock, R.N. et al., 2011. Climate change and forest diseases. *Plant Pathology*, 60(1), pp.133–149.

Tzouramani, I. & Mattas, K., 2004. Employing real options methodology in agricultural investments: the case of greenhouse construction. *Applied Economics Letters*, 11(6), pp.355–359.

Wilmott, P., Howison, S. & Dewynne, J., 1995. *The Mathematics of Financial Derivatives*, Cambridge University Press.