Agenda Item 7

Committee	**Economy, Transport and Environment Scrutiny Committee**
Date	**18 March 2013**
Report By	**Director of Economy, Transport and Environment**
Title of Report	**Scrutiny review of Dutch Elm Disease in East Sussex**
Purpose of Report	**To recommend a strategy for the future management of Dutch Elm Disease in East Sussex.**

RECOMMENDATION: The Scrutiny Committee is recommended to support the proposed strategy for the future management of Dutch Elm Disease.

1. Financial Appraisal

1.1 The cost to East Sussex County Council (ESCC) of the Dutch Elm Disease (DED) sanitation programme in 2012 has been approximately £100,000, covering the rural sanitation area and highway trees. This cost varies from year to year, reflecting the natural variation in the prevalence of DED.

1.2 ESCC has been working with DEFRA's Food and Environment Research Agency (FERA) to compare the costs and effectiveness of: a) stopping the DED sanitation programme; b) returning to how the programme was delivered prior to ESCC taking it back in-house in April 2011, and; c) continuing with the prioritised approached adopted in 2012. Table 1, below, provides a summary of the estimated costs over 10 years. The conclusion is that stopping the sanitation programme is more costly over 10 years than maintaining the programme, but after approximately 20 years maintaining the sanitation programme begins to exceed the cost of stopping the programme. A detailed explanation of the costs and assumptions is set out in the draft DED Strategy, as Appendix 2.

2 Supporting Information

2.1 The objectives of the DED sanitation programme are to:

a) Ensure the long-term survival of a population of mature English elm, which makes an important contribution to the local landscape and provides a habitat to priority species (e.g. butterflies and 200 species of lichen). The Sussex elm population is considered by Natural England to be of regional importance, with Brighton & Hove housing the National Elm Collection. If ESCC were to stop its sanitation programme this would probably have a significant impact on the effectiveness and cost of Brighton and Hove City Council's (BHCC) DED sanitation programme. In addition, all public bodies have a duty to have regard to biodiversity in all of their work (the Natural Environment and Rural Communities Act, 2006).

b) Assist in managing DED on the highway, just as any other land owner is required to do under the Highways Act 1980 (section 154), and on ESCC land (eg. schools), when it poses a health and safety risk.

c) Ensure the most cost effective approach.

2.2 A Scrutiny Review of Trees and Woodland took place on 14 March 2012. It recommended a review of the current approach to managing DED, to:

a) Make an evidence-based decision as to whether to carry on the sanitation programme;

b) If the decision is to maintain the sanitation programme then develop a strategy with key partners to ensure that the approach is financially sustainable and likely to be effective in the long term.

2.3 FERA has funded modelling work by the University of Cambridge to compare the effectiveness of the three approaches outlined in 1.2 above. The key conclusions are that the prioritised approach to control leads to:

a) A significantly lower cost in the short- to-medium term than stopping the sanitation programme, because the average cost of felling highway trees and carrying out the necessary reinstatements is much higher than felling trees elsewhere (£460 compared with £60). A comparison of the costs is summarised in table 1.

b) Fewer trees becoming diseased and having to be felled, leading to a larger population of healthy trees.

Table 1. Comparison of the ten year costs of different approaches to DED.

Approach	**Total elm population over 10 years**	**Total elm population over 25 years**	**Number of elms felled over 10 years**	**Number of elms felled over 25 years**	**Cost of programme over 10 years**	**Cost of programme over 25 years**
No control	7,000	6,000	5,210	5,210	£1,228,500	£1,228,500
Historic	13,000	12,500	10,500	30,000	£638,900	£1,597,250
Prioritised	14,000	14,500	11,500	25,000	£591,100	£1,477,750

The modelling report by the University of Cambridge is included as Appendix 1.

2.4 A draft DED strategy, reflecting the conclusions from the modelling, is included as Appendix 2, as requested by Scrutiny Committee.

3. Conclusion and Reason for Recommendation

3.1 It is recommended that Scrutiny:
a) Considers the conclusions from the modelling report;
b) Endorses the draft strategy for the future management of DED, as this provides the most effective means of maintaining a significant population of English elm at least cost to ESCC; and
c) Supports the principle that ESCC will seek financial contributions from key stakeholders towards the cost of delivering the sanitation programme.

RUPERT CLUBB
Director of Economy, Transport and Environment

Contact Officer: Andy Arnold Tel. No. 01273 481606
Local Members: All

BACKGROUND DOCUMENTS
None

Modelling control of Dutch elm disease in East Sussex

Nik Cunniffe: University of Cambridge

February 2013

1 Introduction

1.1 What is being compared?

Control of Dutch elm is investigated, comparing three potential control strategies.

- **Historical**. Cut down known infected trees that are still alive, but ignore any infected trees that have died.
- **Prioritised**. Cut down any known dead infected trees if they are judged to have died recently enough to be suitable for beetles to breed in, but ignore trees that have been dead for longer than this.
- **None**. Do not cut down any trees.

1.2 How are the strategies compared?

A spatially-explicit stochastic compartmental model of Dutch elm disease in East Sussex is built and parameterised. This mathematical model is used to investigate how the different control strategies fare over ten and twenty-five year time scales, and how key outputs such as the total number of trees lost to disease or control respond to changes in key parameters.

2 Methods

2.1 Host landscape

The host landscape is taken from the GIS data sent by Anthony Becvar on 17th Jan 2013. This is a map showing the position of all known semi-mature, mature and over-mature elms within the East Sussex control zone (see Figure 1), together with metadata detailing whether trees are known to be infected and/or are in "woodland". There are approximately 16000 trees across the region, of which approximately 4000 are in woodland locations, and of which nearly 700 are known to be infected.

2.2 Host demography

The GIS data set the initial configuration of hosts in the model. However, Dutch elm disease and control both lead to removal of elm trees, and the historical rate of control ($\approx$ 1200 trees removed per year) suggests the current elm population would be totally depleted well within the timescales considered here, even if the only source of tree death were control by East Sussex Council. A simple representation of demography is therefore included in the model, mainly to ensure the pathogen can persist. In particular, every time a host is removed, either because it has been removed by control or because it was killed by the disease long enough ago that it would have become epidemiologically inert, a new host appears as the daughter of a randomly chosen woodland "mother" tree. The position of the daughter tree is chosen uniformly within a circle of radius R_d metres centred on the position of its mother (see also Table 1, in which all parameters are summarised). The mother of this "replacement" tree is chosen randomly out of the set of tree located in woodland across the entire landscape. This means that the local

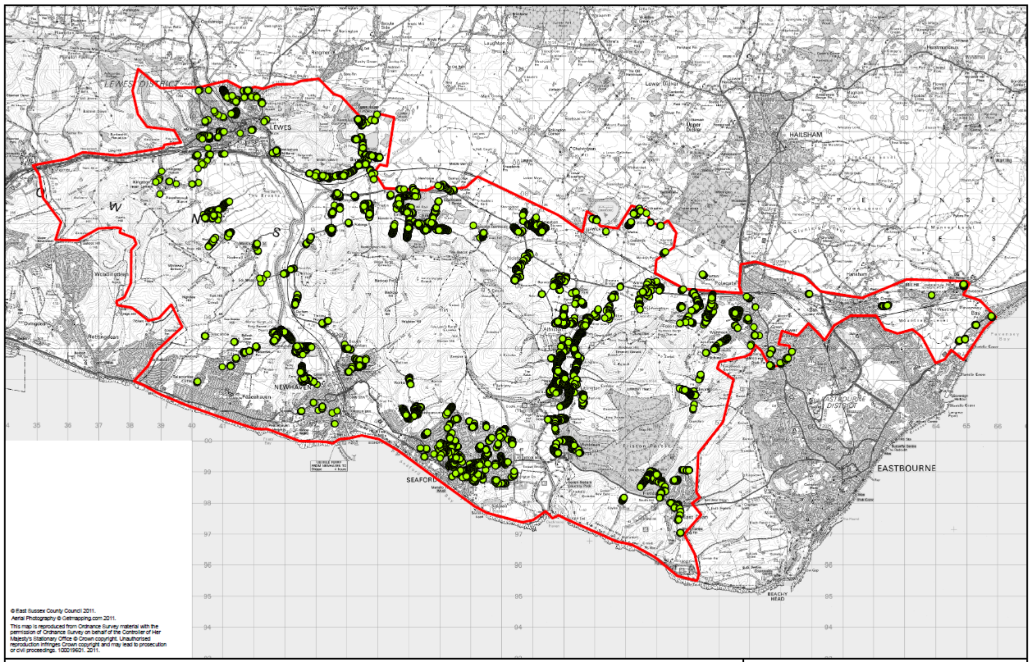

Figure 1: *The East Sussex Dutch elm control zone: note this region excludes a number of nearby areas in which Dutch elm disease is present, e.g. Brighton and Eastbourne, although note that Lewes is included in the zone.*

density of elms across East Sussex changes over time and in each run of the model. However it also avoids the immediate re-infection that would almost certainly follow putting the new tree near the dead tree it replaced. To avoid unrealistically high rates of elm replacement, only up to a maximum of N_{max} rebirths per year are permitted.

2.3 Epidemiological modelling

The modelling concentrates exclusively on semi-mature, mature and over-mature elms. At any time, any single tree can be categorised into one of the following five disjoint epidemiological classes (see also Figure 2)

1. **Susceptible (S)**. Healthy elms that have not been infected.
2. **Exposed (E)**. Very recently infected elms that are still alive, do not show symptoms, and have not yet become infectious.
3. **Live infected (LI)**. Recently infected elms that are still alive, do show symptoms, and are able to infect other trees.
4. **Dead infected (DI)**. Infected elms that have been killed by the pathogen, show extensive symptoms (since they are dead), and are able to infect other trees, primarily by acting as breeding sites for the beetles that vector the spread of the fungus.
5. **Removed (R)**. Elms that have been killed by the pathogen but have been dead for so long that they are no longer a potential beetle breeding ground, and so are epidemiologically inert.

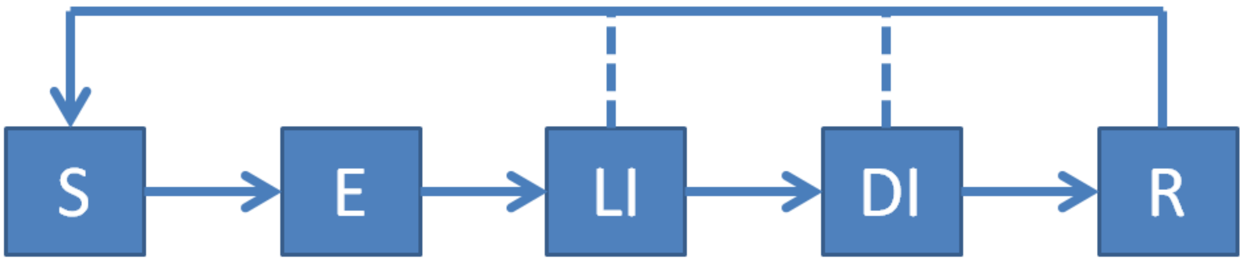

Figure 2: *The structure of the epidemiological model. Lines linking compartments show transitions that can be made by an individual host. Note the influx into the (S)usceptible compartment: this corresponds to a new replacement tree appearing when a tree is removed by entering the (R)emoved compartment (i.e. long dead and epidemiologically inert). The dotted lines indicate that tree replacement may also occur when a detected and infected tree from either of the LI (i.e. live infected) or DI (i.e. dead infected) compartments is cut down, at least when certain types of control are performed.*

In essence the epidemiological modelling fixes the rates of the following transitions between these classes.

1. $S \to E$. This corresponds to infection, and occurs at a rate which depends on the time of year and on the number and relative positions of other trees that are infected (see below).

2. $E \to LI$. This corresponds to the onset of infectivity, and takes 50 days on average.

3. $LI \to DI$. This corresponds to tree death, and takes 400 days on average.

4. $DI \to R$. This corresponds to trees becoming unsuitable for beetle breeding, and takes 365 days on average.

The rates of the $E \to LI, LI \to DI$ and $DI \to R$ transitions are taken directly from Harwood (2011). Note that sojourns are addititive, so on average a tree takes approximately 2 years, 3 months or so (i.e. $50 + 400 + 365$ days) from the time it first becomes infected to the time at which it is no longer suitable for beetles to breed in. However, because the model is stochastic, the exact time spent in any class varies from tree to tree and from run of the model to run of the model.

2.3.1 Susceptible to exposed transition (i.e. $S \to E$)

This transition controls how fast infection spreads. There are three distinct ways in which a susceptible tree can become newly infected (i.e. "exposed")

1. **LI transmission**. Live trees transmit infection to nearby trees.

2. **DI transmission**. Dead trees transmit infection widely since they act as a home for breeding beetles which go on to infect other trees when they emerge.

3. **External transmission**. Susceptible trees can be infected by beetles that fly in from outside the control zone carrying the pathogen.

In my model the rate at which susceptible tree i becomes exposed is given by

$$\lambda_i = \omega(t)\left(\beta\left(\rho \sum_{j\in\Omega_{LI}} K_{LI}(d_{ij}) + \sum_{j\in\Omega_{DI}} K_{DI}(d_{ij})\right) + \epsilon\right), \tag{1}$$

where d_{ij} is the distance between tree i and a particular infected tree j, Ω_{LI} is the set of indices of live infected (LI) trees, Ω_{DI} is the set of indices of dead infected (DI) trees and β is the global rate of infection. Note that the infection rate β is the parameter varied to match the historic rate of infection, and so to make my model "match" the spread of the pathogen in East Sussex (see "Estimating the rate of infection", below). The function $\omega(t)$ is equal to one between the start of April and the end of October and zero otherwise; this ensures that infection only occurs during this part of the year. The parameter ϵ controls the rate at which trees are infected from sources outside the control zone: based on discussions with Anthony, this was set such that an average of 200 trees are infected by this pathway per year[1]. This parameter does not vary spatially (e.g. with distance from the edge of the East Sussex control zone) in the model, since I did not have data to parameterise the fall off in rate according to the distance from the edge of the control zone that is almost certainly present in practice.

The two functions K_{LI} and K_{DI} are "dispersal kernels" associated with live and dead trees, respectively, and control how the probability of transmission drops off according to the distance between a pair of hosts. I use the Cauchy kernel to model transmission from DI trees, with

$$K_{DI}(d;\alpha_{DI}) = \frac{1}{1+\left(\frac{d}{\alpha_{DI}}\right)^2}, \tag{2}$$

and where α_{DI} is a measure of median distance of disease spread (equivalently a median distance of beetle flight), which I take to be 150 m. Although half of all dispersal is within 150m, the Cauchy kernel is a member of the broader class of so-called "thick tailed" power law kernels, and permits occasional dispersal over far longer distances (up to several kilometres). Both the form of this kernel and median dispersal come from Harwood. Transmission from live infected trees is dominated by transmission through a shared vascular system: I model this root to root pathway using the exponential kernel

$$K_{LI}(d;\alpha_{LI}) = \exp\left(-\frac{d}{\alpha_{LI}}\right). \tag{3}$$

Transmission by this route is more spatially restricted, since only nearby pairs of trees are joined by their roots. I take $\alpha_{LI} = 5$m as a typical scale for root to root transmission (based on discussions with Anthony).

The final parameter δ controls the relative infectivity of a live vs. a dead infected tree. Again I follow Harwood in setting this so that dead infected trees lead to twice the rate of infection compared to live infected trees (i.e. I take $\delta = 0.5$). However the sharp drop off in the infection kernel of LI trees means that DI trees have many more chances to infect over the entire landscape of trees, and if this is not accounted for the infectivity of LI trees is greatly

[1] Strictly speaking, ϵ is a per capita rate set such that 200 trees would be infected if the population were entirely susceptible. Assuming the disease is circulating in East Sussex, there would therefore actually be slightly fewer than 200 primary infections per year (since infection of a tree cannot occur twice in the model).

understated to the extent that LI trees would barely infect at all. I therefore normalise for this by interpreting δ as the relative rate of infection averaged across the entire landscape, and set the parameter ρ in Equation (1)

$$\rho \;=\; \delta\left(\frac{\sum_i \sum_{j,j\neq i} K_{DI}(d_{ij})}{\sum_i \sum_{j,j\neq i} K_{LI}(d_{ij})}\right). \tag{4}$$

Note the double sum itself averages over all possible interactions between all pairs of trees.

2.4 Modelling detection, control and the budget

2.4.1 Detection

The "average" tree is examined approximately yearly, and this is ensured in the model as follows.

- At the start of each year the set of all trees is randomly divided into 12 equally-sized groups[2].
- Trees in group n are examined on day $30(n-1)$ of the year, i.e.
  - trees in the first group are examined after 0 days;
  - trees in the second group are examined after 30 days;
  - ...
  - trees in the twelfth group are examined after 330 days.

  Note that a "year" in the model is actually 360 days long, for simplicity.
- Infected trees are detected with probability $p = 0.9$ on any single round of examination (note the probability of detection is independent of whether the tree is still alive or has died, based on Anthony's input).

While the mechanism used is admittedly rather simple, it does ensure that all trees are visited every year, and at roughly the correct rate.

2.4.2 Control

Whenever an infected tree is detected as decribed above, an element is added to the model's "control list". This is a list of known infected trees flagged to potentially be removed at some date in the future, kept sorted by date of potential removal. The model continuously checks the earliest element of the control list (i.e. the next control that could be performed). There are three scenarios to be considered.

- **No Control**. Although the control list is checked, it is otherwise entirely ignored, and so no infected trees are cut down.
- **Historic**. Whether or not a particular detected infected tree is added to the control list depends on its status at the time of detection (t_{detect})

[2] Note the model uses a different partitioning on each year and in each run of the model.

  - If the tree has already died by t_{detect} it is ignored and not added to the control list.
  - If the tree is still alive at t_{detect} then it is added to the control list.
    * The tree is flagged to potentially be cut down at some later date;
    * A delay t_{delay} is chosen according to a sample from the probability distribution shown in Figure 3.
    * The time of potential control, $t_{control}$, is set to be

$$t_{control} = t_{detection} + t_{delay}. \tag{5}$$

    * If the tree is still alive at $t_{control}$ it is actually removed.
    * If it has died by $t_{control}$ it is ignored.

- **Prioritised**. Here detected trees are not cut until they have entered the DI class.
  - Both living (LI) and recently dead (DI) trees are added to the control list on first detection
    * For a LI tree the historical distribution of delay times (see Figure 3) is used to set the time at which it is first considered for potential control.
    * For a DI tree, control is set to potentially occur at a random time within one month of detection.
  - In either case, at the time of potential control the tree is only cut down if it is dead (i.e. is in class DI).
  - This means that known infected LI trees are not controlled until they die.
  - In particular, after the first examination of a tree that remained LI, it is revisited monthly to check whether it has entered the DI class. Control actually occurs in the first month the tree in question is noticed to be dead.

2.4.3 Budget

The model can represent a fixed budget for cutting trees. In particular, it keeps a running count of the number of trees that have been cut down since the start of January of the current year (say n). This number is reset (i.e. $n = 0$) at the start of each year. The budget then controls the maximal number of trees that can be cut within any single year, say C. At any time a tree would be cut in either the historical or prioritised approach then the following procedure is followed.

- If $n < C$ (i.e. if budget remains)
  - Control happens.
  - The running count n is increased by one.
- If $n = C$ (i.e. if the budget is exhausted)
  - Control does not happen.
  - The running count n is unchanged.

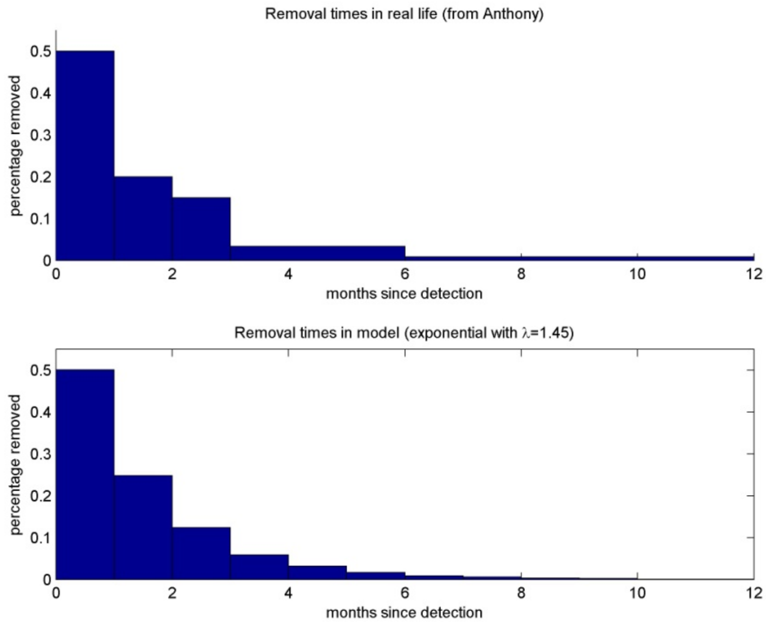

Figure 3: *The distribution used to model the delay between detection and control: the top panel shows the distribution extracted from the data given to us by Anthony; the bottom panel shows the fitted exponential distribution. The inverse scale parameter,* $\lambda = 1.45$ *is given in months.*

– The tree is added back to the control list with $t_{control}$ some time within the first month of the following year.

Note the consequence of this is that if the budget is wildly inadequate compared to disease spread, control becomes progressively far behind itself, and more and more trees from previous years are accumulated. Note also that detection is unaffected by the budget, which in fact only affects the number of trees that can be removed by control each year.

2.5 Estimating the rate of infection

The single disease spread parameter that is impossible to estimate by looking in the literature is the underlying rate of secondary infection, (i.e. β in Equation 1, the rate at which a single infected tree would infect a susceptible tree at a distance of zero). Data on the number of removals between 2000 and 2011 indicate that approximately 1200 trees were removed per year, and this was used to fit the model. The model was repeatedly run for different values of the parameter β, simulating control under the historical approach, and searching for the value of β that lead to the correct number of tree removals due to control. The best estimate of the infection rate according to this process is $\beta \approx 3.95 \times 10^{-5}$ (see Figure 4). Unless stated otherwise, this value of β is used in all model simulations. Note that this infection rate was estimated in the absence of any budgetary constraint, since if the budget rather than the underlying disease dynamics were responsible for setting the historical rate of removal, it would be impossible to use these data to set a unique value for β (in particular, increases to larger values of β would have no effect on the apparent rate of removal if control were limited, and so

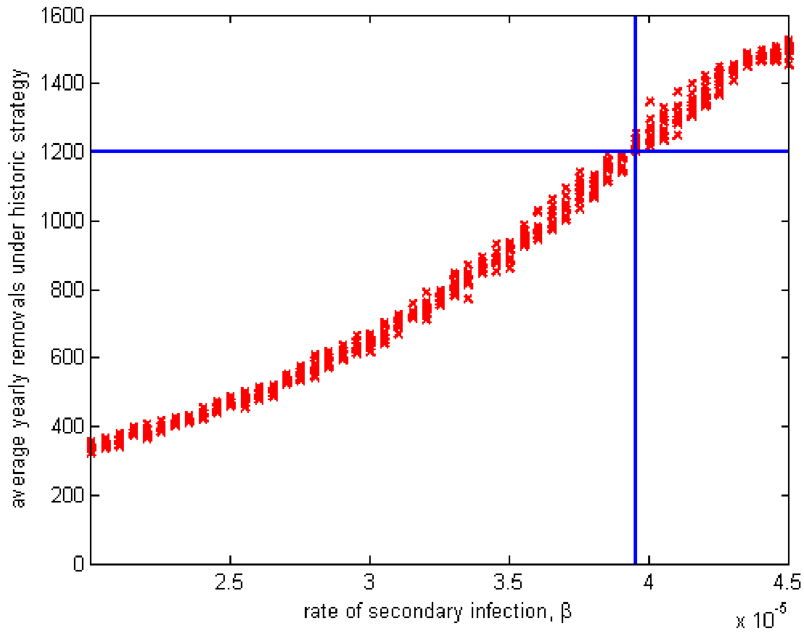

Figure 4: *The average number of trees removed by control per year over twenty-five years under the historical control strategy, showing ten replicates for each value of β. Since records indicate approximately 1200 trees were removed yearly between 2000 and 2011, the best estimate of the infection rate is $\beta \approx 3.95 \times 10^{-5}$.*

it would only be possible to put a lower bound on the parameter's value).

3 Results

The model is used to predict disease dynamics and spatial spread over ten and twenty-five year timescales, after seeding the model with the location of known infected trees.

3.1 Future disease progress

Predictions over ten and twenty-five year timescales under all three control strategies are shown in Figures 5 and 6. These predictions were obtained by averaging over 100 independent replicates of the model for each control. As well as the numbers of trees in each epidemiological compartment, a graph showing the time evolution of p(Original Alive) is presented. This is defined as the probability that any tree from the original cohort has not yet died by any particular time, allowing our attention to be restricted to those trees are currently present (i.e. not to any replacement trees included to keep the elm population size constant).

The most striking conclusion is that the prioritised approach leads to a better outcome than the historical control strategy, and a far better outcome than no control whatsoever, over both timescales considered. Focusing on p(Original Alive) at ten years, an estimate of the

Symbol	Biological meaning	Value	Source
R_d	*Regeneration distance.* A daughter tree is created within a circle of this radius around a randomly chosen woodland tree whenever a tree is removed.	10m	Anthony
N_{max}	*Maximum number of replacements per year.* No more than this many replacement trees may be created in a single year.	3000	Fitting
$1/\gamma$	*Latent period.* (Average) time taken for an exposed tree to first become infectious.	50 days	Harwood
$1/\sigma$	*Lifetime of infected tree.* (Average) time taken from a tree becoming infectious to it dying.	400 days	Harwood
$1/\mu$	*Post-mortality infectious period.* (Average) time taken for a dead tree to become unsuitable for beetle breeding.	365 days	Harwood
β	*Rate of secondary infection.* Sets the rate at which new infections are created.	3.95×10^{-5}	Fitting
ϵ	*Rate of primary infection.* Sets the rate at which new infections are imported from outside the control zone ($\approx$ 200 per year).	7×10^{-5}	Anthony
δ	*Relative infectivity.* Controls how much less infective a live tree is compared to a dead elm.	0.5	Harwood
$\omega(t)$	*Infection seasonality.* Restricts new infections to occur between April and September.	0 or 1	Anthony
$K_{DI}(d; \alpha_{DI})$	*Dispersal kernel from dead infected trees.* Sets how the probability of infection drops off with distance from dead infected trees.	$\frac{1}{1+\left(\frac{d}{\alpha_{DI}}\right)^2}$	Harwood
α_{DI}	*Scale for dispersal from dead infected trees.* Median distance of beetle flight, although occasional infection is possible over much greater ranges.	150m	Harwood
$K_{LI}(d; \alpha_{LI})$	*Dispersal kernel from live infected trees.* Sets how the probability of infection drops off with distance from infected trees that have not yet died.	$\exp\left(-\frac{d}{\alpha_{LI}}\right)$	Anthony
α_{DI}	*Scale for dispersal from live infected trees.* Typical scale of root to root transmission.	5m	Anthony
p	*Detection probability.* The probability of detecting the pathogen on a single visit to an infected tree.	0.9	Anthony
C	*Control budget.* Sets how many trees can be cut down per year.	n/a	Anthony

Table 1: Table of parameters, symbols and default values used in the simulations.

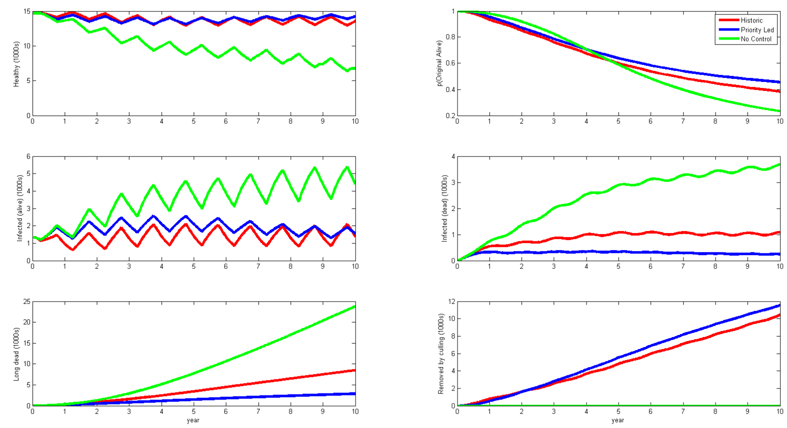

Figure 5: *Predictions from the model for the next ten years, showing the number of trees in each of the epidemiological compartments (note that "Infected (alive)" corresponds to the sum $E + LI$). The probability that a single tree from the original cohort remains alive by a given time, p(Original Alive), is shown in the panel on the top right. The "saw tooth" pattern in the graph showing the number of healthy trees is because infection only happens for six months in the year (when there is a net loss of susceptibles due to infection, and so S goes down), while control and so replacement of susceptibles occurs year round (and so with no new infection there is a net gain in S). These fluctuations in the number of healthy trees then go on to drive smaller oscillations in the numbers in all compartments.*

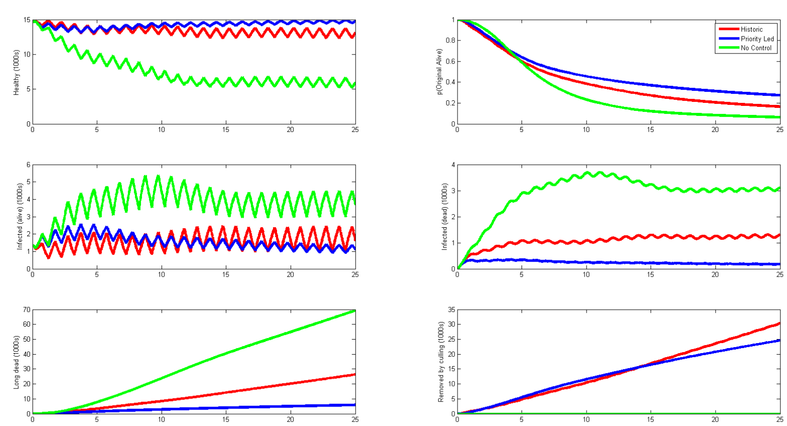

Figure 6: *As Figure 5, but showing predictions over a twenty-five year timescale. Note the relative orderings of the success of the three controls remain unchanged.*

probability of a randomly-chosen individual tree from the original cohort surviving the next ten years is 0.46 for the prioritised approach, 0.38 for the historic control, and 0.23 if control is not attempted. One way of interpreting this result is to say that a randomly chosen tree is $\frac{0.46-0.38}{0.38} \times 100\% \approx 20\%$ more likely to survive the next ten years if the prioritised approach is adopted instead of the historic control strategy. A similar calculation indicates a tree is 100% more likely to survive under the prioritised approach than with no control: i.e. its chance of survival exactly doubles. The total number of removed trees (i.e. the sum of the number of trees removed by control and the numbers of dead but still infectious and long dead trees) is approximately 5400 fewer using the prioritised approach than under the historic approach, and approximately 12800 fewer than with no control whatsoever. Note that this last calculation uses a slightly different definition of "removed" when compared to the graphs in Figures 5 and 6, since it includes infectious trees (which are dead but not epidemiologically inert), but using either definition of removed leads to the same conclusion.

These calculations can be repeated over the twenty-five year timescale, with broadly similar conclusions. However, as a consequence of the representation of demography in the current version of the model (i.e. immediate replacement of trees up to a hard limit per year, with no regard for the number of mother trees, time of year, any lag before maturity or instantaneous rate of replacement), the calculations over the ten year horizon are potentially much more reliable. I therefore prefer to emphasise results over the ten year timescale, and do so in what follows.

3.2 Maps of disease spread

GIS maps showing the spatial variation in the probability of infection are given in Figures 7 to 12. Note that the maps do not show the probability of infection for individual trees, but instead are a rasterised version of the model's results. In particular, the average probability of infection within 25m by 25m or within 250m by 250m squares is shown. This was done for two reasons (i) to make the maps easier to interpret (showing individual trees would lead to a huge number of dots on each map, each of which would have to be interrogated to understand the results and would be impossible to assimilate on a hard copy) and (ii) asking the model to correctly predict the disease status of each individual tree ten or even twenty-five years from now is simply asking a little too much and presenting the results in this fashion is misleading; amalgamating the results in space does not overstate the predictive power of the model to the same extent. Effectively, while it is fair to say that the model does a good job of predicting disease status in the future when averaged over the entire ensemble of trees and/or at the 25m by 25m scale, in the light of the paucity of data to parameterise the model, predictions on an individual tree by individual tree basis are probably best avoided. Note that the estimates of infection density are calculated using the infection status of only the original cohort of trees; i.e. replacement trees are ignored in this calculation. This is because maps showing the infection status of trees that do not yet exist are unlikely to be plausible.

3.3 Parameter scans

The model is used to scan over the values of certain epidemiological parameters, to investigate how the different control strategies behave when parameters are altered. This allows the robustness of the predictions to be tested. In each case, for each value of the parameter, ten

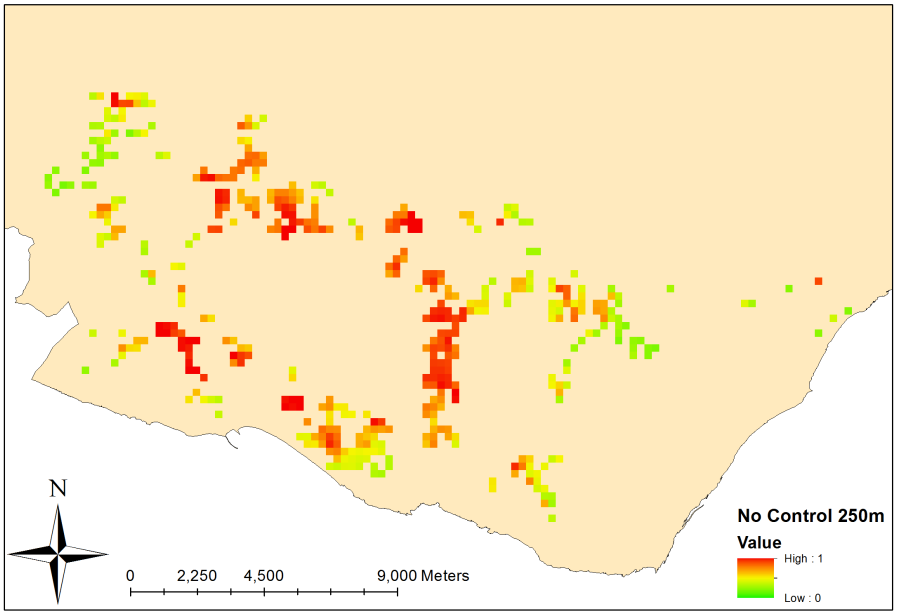

Figure 7: *Map showing probability of infection at the 250m by 250m scale after 10 years when there is no control.*

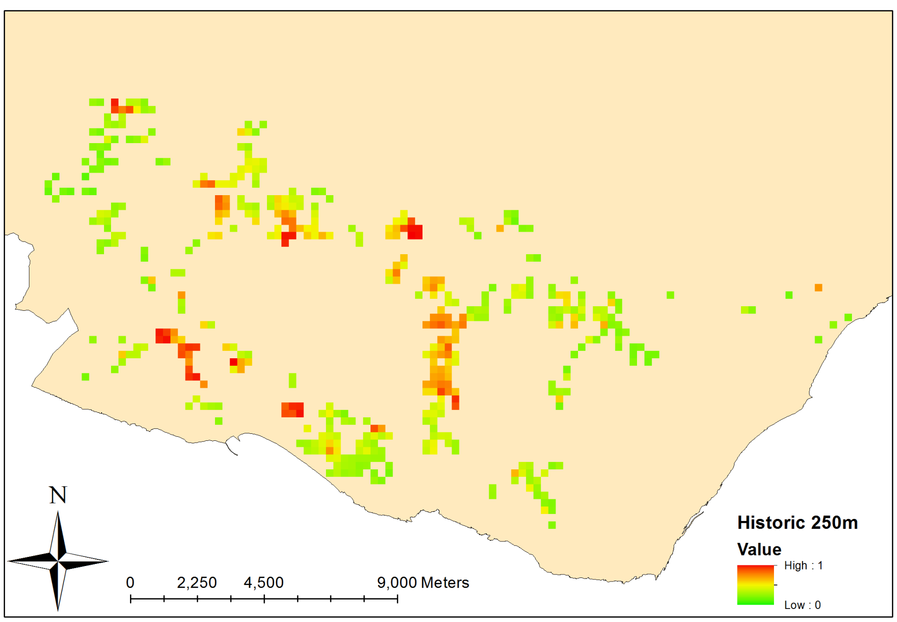

Figure 8: *Map showing probability of infection at the 250m by 250m scale after 10 years using the historical approach.*

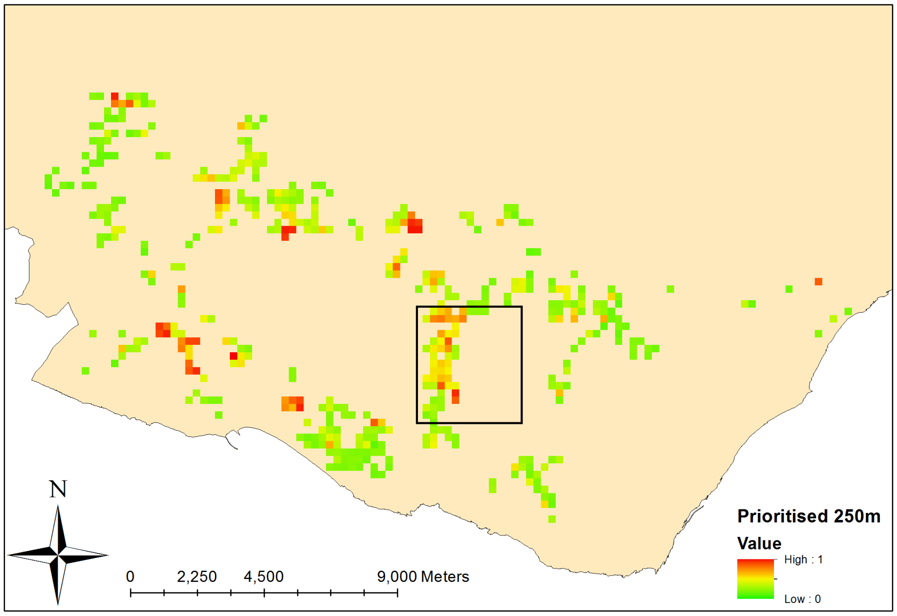

Figure 9: *Map showing probability of infection at the 250m by 250m scale after 10 years using the prioritised approach. The black box shows the smaller region that is focused upon at the 25m by 25m scale in Figures 10 to 13*

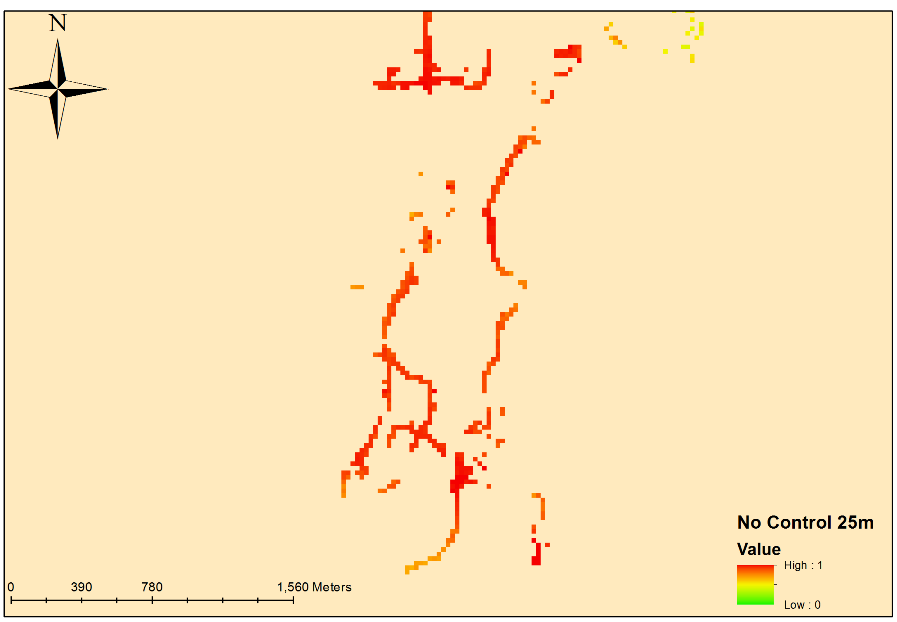

Figure 10: *Higher resolution map showing the probability of infection at the 25m by 25m scale for the region highlighted by the black box in Figure 9, after 10 years and when there is no control*

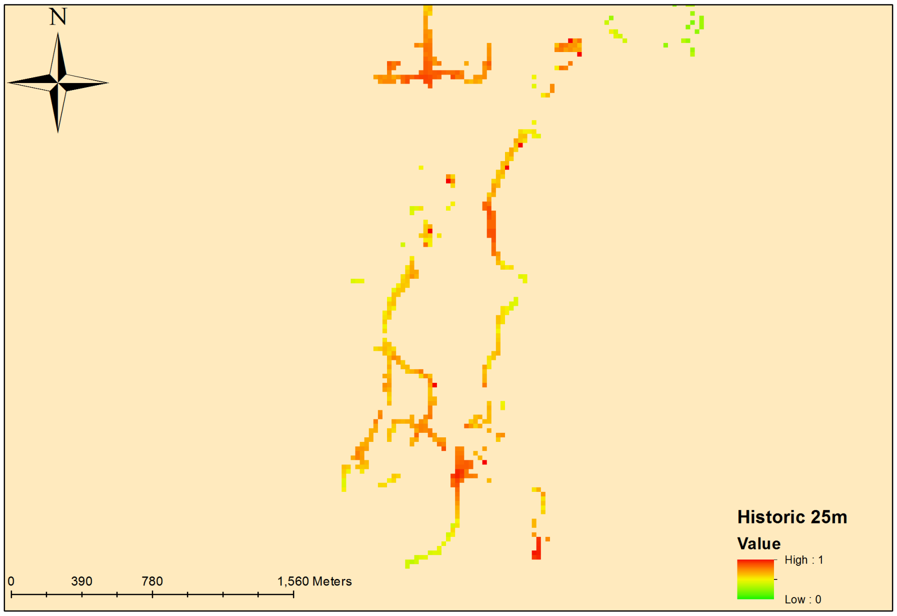

Figure 11: *Higher resolution map showing the probability of infection at the 25m by 25m scale for the region highlighted by the black box in Figure 9, after 10 years using the historic approach*

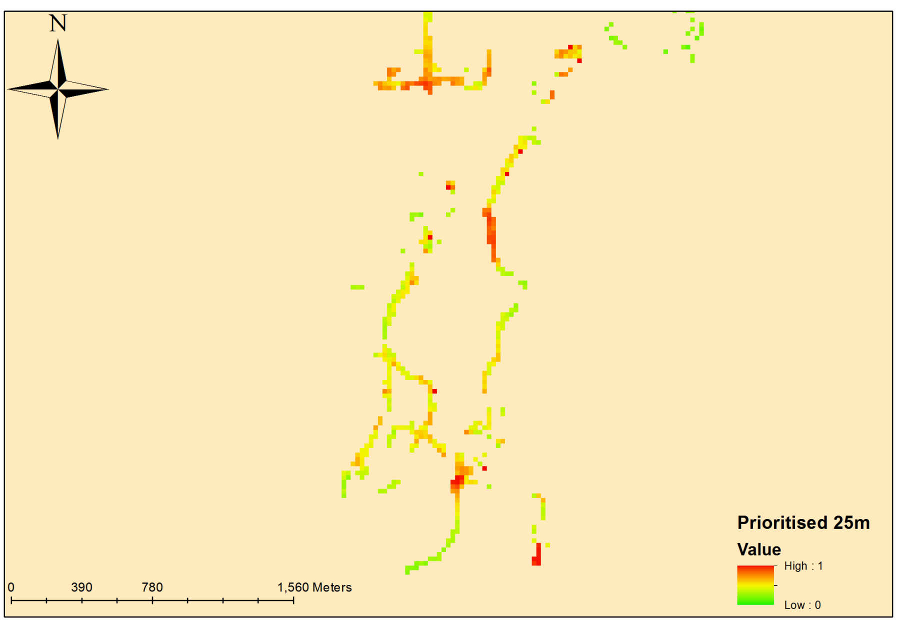

Figure 12: *Higher resolution map showing the probability of infection at the 25m by 25m scale for the region highlighted by the black box in Figure 9, after 10 years using the prioritised approach*

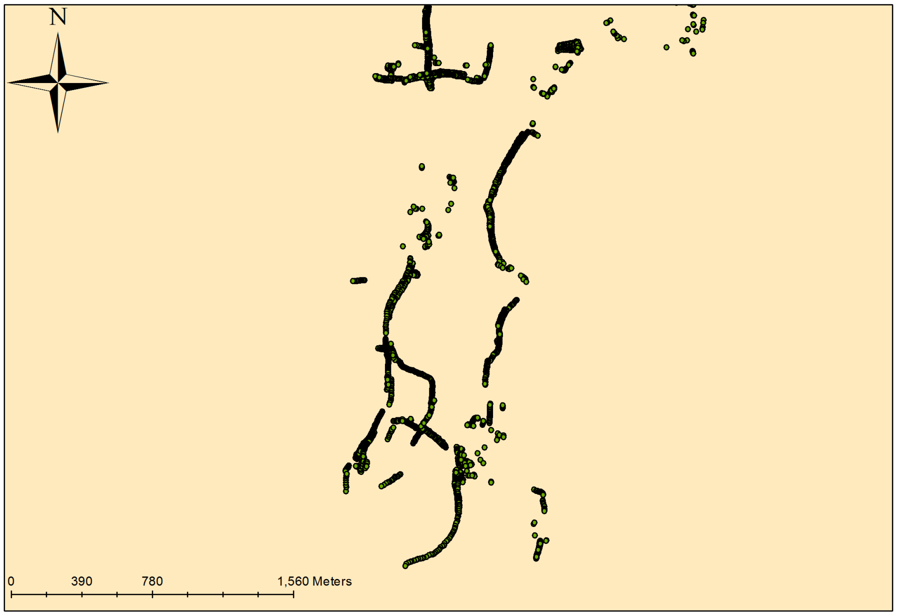

Figure 13: *Map showing the location of individual trees in the zoomed in area.*

replicates of the model for each control strategy were performed. The results presented focus on the number of tree removals and p(Original Alive), both over the ten year timescale. Note that all other parameters were fixed at the default values, and this means that the effect of one parameter changing in isolation is investigated. This is why the number of removals due to control per year diverges from 1200 in these runs; the model is not refitted each time it is run for each parameter. Instead how any alteration to one epidemiological mechanism can affect the results is considered.

3.3.1 Rate of primary infection (i.e. influx from outside)

The results in Figure 14 show the performance of the control strategies for values of ϵ between $\epsilon = 0$ (i.e. East Sussex is not subject to any influx of infected beetles) and $\epsilon = 14 \times 10^{-5}$ (i.e. there would be approximately 400 infections from outside sources per year). The prioritised approach consistently outperforms the other two control strategies across this entire range of parameters. Comparing p(Original Alive) after ten years for the prioritised approach at $\epsilon = 7 \times 10^{-5}$ and at $\epsilon = 14 \times 10^{-5}$, indicates that if the force of infection from outside the region is doubled from the original value, the probability of a randomly chosen infected tree within the control zone surviving the next ten years is reduced by about 20%. In turn this indicates that the control (or otherwise) adopted outside the East Sussex control zone can have a big impact on disease spread within it, even when intensive control is done inside the zone. Note too that in the absence of the budgetary constraint, the effects of these increases may even be understated. In particular, if it were the case that 400 new infections occured every year, then up to 1400 removals would be required, and since the budget does not in fact allow this level of intervention, more infection would presumably eventually be present in the long term, since

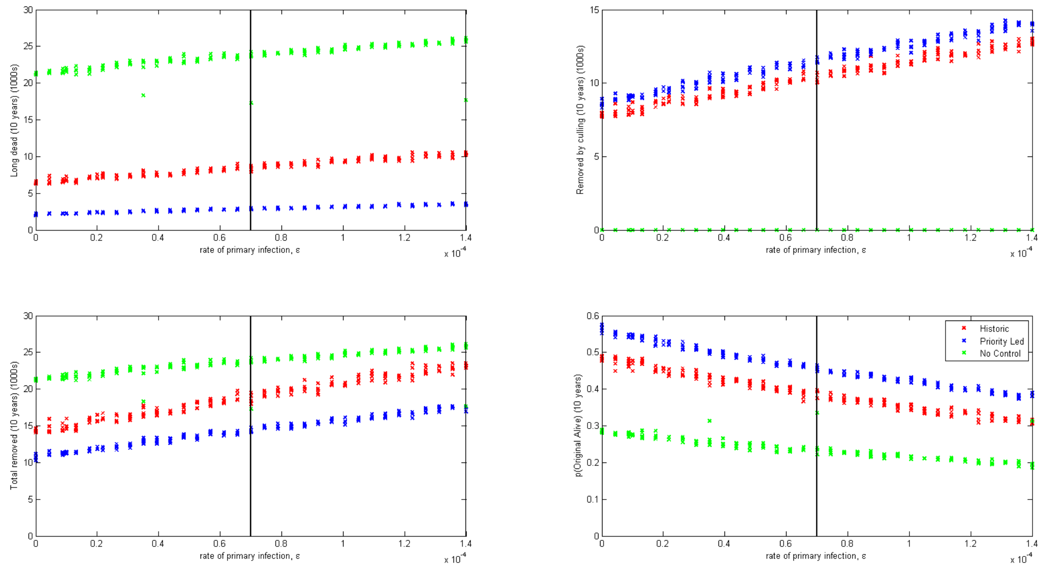

Figure 14: *Performance of the model for different values of ϵ, the rate of primary infection (note the parameter is changed to give between 0 and 400 infections from outside sources per year). The black line shows the default value of this parameter (i.e. $\epsilon = 7 \times 10^{-5}$, approx 200 infections per year) as used in all other simulations. The individual graphs show the total number of tree removals of both types after ten years (top row), the total number of removals after ten years (bottom left) and the probability that a randomly chosen tree from the original cohort is alive after ten years (bottom right).*

in practice control would not keep up with rates of disease spread.

3.3.2 Rate of secondary infection (i.e. rate of disease spread in East Sussex)

The results in Figure 15 show the performance of the control strategies for values of β between $\beta = 0$ (i.e. infected trees are not at all infectious to other trees within the zone) and $\beta = 8 \times 10^{-5}$ (i.e. the rate of infection between pairs of trees is double the best fitting value). Again the prioritised approach outperforms the other control strategies. The large variation in the number of removals and p(Original Alive) across even this relatively restricted range of values of β indicates this is an important parameter to "get right", since the rate of disease spread can have a large effect on dynamics. This is of course unsurprising, but does focus our attention on obtaining more concrete data on disease spread to parameterise the model. Note there is still some loss of trees to disease when the secondary infection rate is set to zero. This is due to (i) primary infection from outside, (ii) the loss of those trees which are initially infected at the start of the simulation.

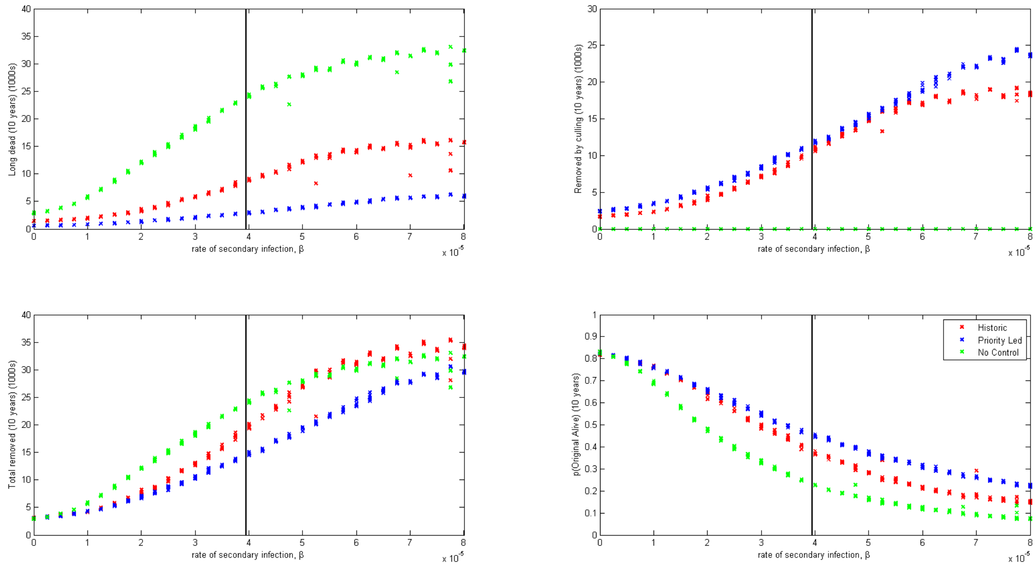

Figure 15: *Performance of the model for different values of β, the rate of secondary infection. The black line shows the default value of this parameter (i.e. $\beta = 3.95 \times 10^{-5}$) as used in all other simulations. The individual graphs are as described in Figure 14.*

3.3.3 Relative infectivity (i.e. how infectious a living tree is compared to a dead one)

The results in Figure 14 show the performance of the control strategies for values of the relative infectivity δ between 0 (i.e. live trees are not infectious) to $\delta = 1$ (i.e. live trees are just as infectious as dead trees). As before the prioritised approach performs best, with the difference between strategies increasing as δ becomes smaller (i.e. as live trees become less infectious in relative terms). This is reassuring, since the only source of information on this key parameter is Harwood, and intuitively the value $\delta = 0.5$ taken in that paper seems rather large, indicating that live trees are half as infectious as dead ones (note the value of this parameter comes from fitting the Harwood model to data rather than due to input from a biologist). For small values of δ ($\delta < 0.3$ or so) even the number of removals by culling is smaller under the prioritised approach than the historic one (for the default parameters the historic approach leads to slightly fewer removals by control intervention, but far more deaths of trees overall). Again this is expected; if live infected trees are in fact relatively unimportant epidemiologically-speaking, focusing efforts on dead trees becomes an even better idea.

3.3.4 Maximum number of trees that can be cut down per year (i.e. the budget)

Note that scans over the first three parameters were performed assuming no budgetary constraint. This allows us to focus on the underlying result, avoiding the difficulties in interpretation associated with interactions between a parameter changing and the budget being exceeded. However here the effects of a restricted budget are assessed. Figure 17 shows the results, in particular examining what happens if the number of removals per year is restricted. The most

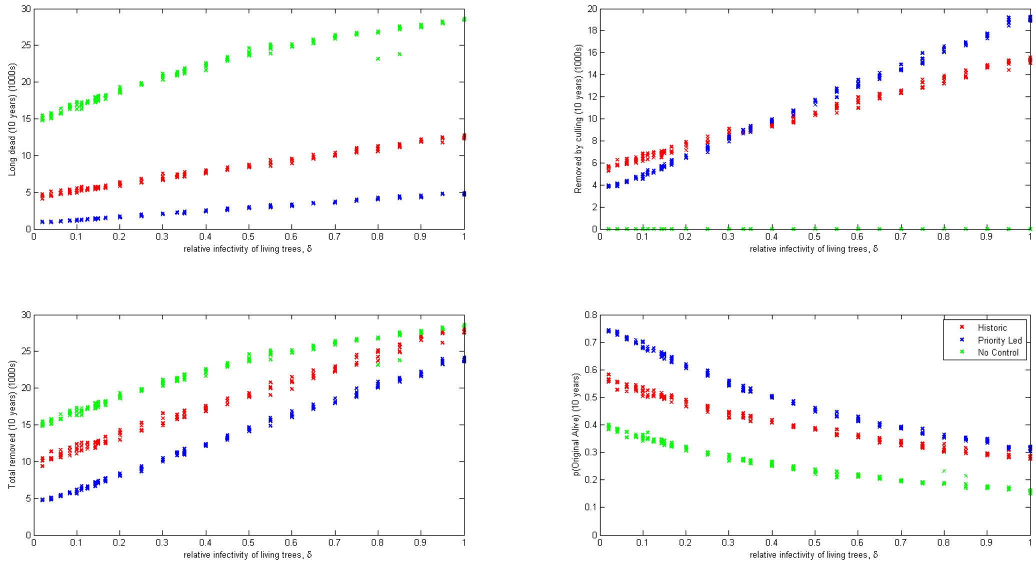

Figure 16: *Performance of the model for different values of δ, the relative infectivity of live vs. dead infected trees. The black line shows the default value of this parameter (i.e. $\delta = 0.5$, i.e. live trees are half as infectious as dead ones) as used in all other simulations. The individual graphs are as described in Figure 14.*

striking conclusion is that, in effect, "you get what you pay for": if the budget is reduced then more trees die. Note that the manner in which the responses flatten off for $C > 1200$ is because of how the model is fitted.

For the default parameters no more than 1200 trees to remove under either strategy are actually found across the landscape per year, and so even if the budget is increased, it cannot be used. We know that this does not happen in reality, and so these results do not necessarily indicate that an increased budget would be of no use. Indeed as I understand it, the budget was prematurely exhausted last year, perhaps due to an additional influx of disease from outside the control zone and/or more new susceptibles being born and leading to faster spread because there are more host trees to spread through and/or environmental conditions causing faster spread. However, since none of the effects were explicitly included in the model, nor in the model fitting, which concentrated on matching a "steady state" of disease spread, consequence(s) cannot be felt in the model's results. In reality a backlog in infection would lead to more secondary infection in the next year, and this in turn would lead to even more detections the year after, leading to an ever bigger backlog, and with insufficient budget the disease would probably increasingly get out of control. However, we did not have sufficient data to allow us to fit the model to this situation.

4 Discussion

Clearly the most important result is that the prioritised approach to control leads to fewer losses of trees overall than the historic strategy, and far fewer losses than not controlling at

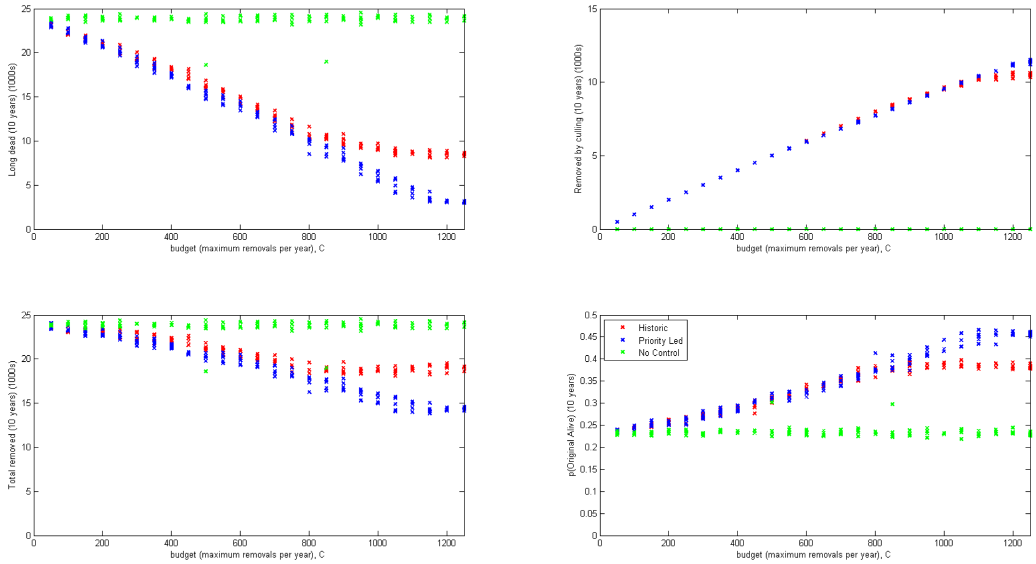

Figure 17: *Performance of the model for different values of C, the maximum number of removals allowed per year by the budget. There is no black line, since other simulations assume the budget is unlimited. The individual graphs are as described in Figure 14.*

all. Confidence in this is undoubtedly increased by the robustness of this conclusion to changes in epidemiological parameters. It is therefore fair to say – inasmuch as it is supported by the results of a relatively simple model of the type presented here – that the prioritised approach is sensible.

Nevertheless, it should be noted that there are a number of potential sources of error that could mean the conclusion is incorrect. They are summarised in the following section, in estimated order of importance. Improving these aspects of the model, together with investigating any elements of the results presented here that are particularly interesting to stakeholders in East Sussex, would form the basis of any future work.

4.1 Potential sources of inaccuracy

4.1.1 Host demography

Fundamental to any prediction of epidemiological dynamics over long time scales or when there is a large number of tree deaths is the replenishment of susceptible hosts that is necessary for pathogen persistence. However, here the treatment of this has been very simple. Certainly the predicted number of tree removals is almost certainly too large, particularly when no control is attempted. This is because our demographic assumption leads to an excessively large rate of tree replacement over East Sussex, and in turn a very fast cycling of disease. However, elm demography is complex, and data to parameterise a more realistic model (e.g. involving a juvenile class) are few. A careful consideration of demography was therefore not possible in this project. Epidemiological intuition suggest this is the most important omission.

4.1.2 Lack of data

The data available to parameterise the model was extremely limited, and basically boiled down to a table showing the number of removed trees between 2000 and 2011. There were no quantitative data concerning the spatial scale of disease spread, the times taken to transit the various epidemiological compartments, the density of beetles or influx rates of disease from outside the system. This meant we were forced to parameterise with information taken from the literature, and often these numbers were not finely resolved or were not directly applicable to the situation in East Sussex. This lack of data had a particular effect on how our model behaved when the budget was increased, as described in Section 3.3.4. Given the aims of the project, our lack of knowledge surrounding the relative infectivity of live and dead trees was also very concerning. It may in fact be useful to use the meetings planned via project partners at York and FERA to try to establish a consensus among stakeholders concerning this key parameter before going forward with the modelling.

4.1.3 Treatment of vector behaviour and density

The model does not track the density of beetle vectors, instead effectively assuming the density is proportional to the number of infected trees. This simplifies the modelling, and the same assumption was made in both the Harwood paper and the Swinton and Gilligan papers, but it presumably restricts the model's predictive power. It also means the treatment of environmental drives is necessarily in turn itself rather simple (in the model disease can spread between April and September, but between these months it is always spreading at the same rate, independent of e.g. temperature). The main reason this assumption was made was lack of concrete data on vector density. However, were more data available, it would be interesting to include it in the next iteration of the model.

4.1.4 Treatment of detection and control

Although the model faithfully represents the fundamental principles underlying both detection and the two types of control (i.e. trees are visited approximately once per year and the prioritised and historical approaches target dead and live trees, respectively), what actually happens in East Sussex is undoubtedly "richer" than the simple approach adopted in the model. This may be a fertile area of future study. Particularly interesting might be a "mixed" strategy which targets the two classes of infectious hosts to different extents (which could of course depend on the current status of the epidemic and/or the budget remaining and/or the position over the landscape). However, to do this would probably require a more careful treatment of what happens when the budget is exceeded, together with a better understanding of what causes this to happen as well as extensive discussions with staff members to better understand their behaviour, so was outside the scope of this initial investigation.

Acknowledgements

I thank Anthony Becvar for the benefits of his extensive insight into the biology of Dutch elm disease and past and present controls adopted in East Sussex, James Elderfield for assistance with programming, Ben Price for assistance with initial literature review and data processing, and James Cox and Matt Castle for assistance with GIS.

References

Harwood, T.D., Tomlinsons, I., Potter, C.A. and Knight, J.D. (2011) "Dutch elm disease revisited: past, present and future management in Great Britain". *Plant Pathology.* 60:545-555.

Swinton, J.A. and Gilligan C.A. (1996) "Dutch elm disease and the future of the elm in the U.K.: a quantitative analysis". *Philosophical Transactions of the Royal Society London, Series B*. 351:605-615.

Swinton J.A. and Gilligan C.A. (1999) "Selecting hyperparasites for biocontrol of Dutch elm disease in stochastic spatially-extended epidemics" *Proceedings of the Royal Society, Series B.* 266:437-445.

Swinton J.A. and Gilligan C.A. (1999) "A modelling approach to the epidemiology of Dutch elm disease" in The Elms: Breeding, Conservation and Disease Management (Editor C Dunn).

DRAFT ESCC DUTCH ELM DISEASE STRATEGY 2013/14 **Appendix 2**

Background:

Dutch Elm Disease (DED) is caused by a fungus which is transmitted from tree to tree by two types of elm bark beetle, or via interconnecting roots between 2 or more trees. The fungus causes elms to block their own water conduction system in an attempt to cut off the spread of infection, resulting in wilting and death of the foliage and the slow death of infected limbs. Symptoms first appear in early spring/early summer (depending on the weather) and last until the trees shed their leaves in the autumn. The beetles tend to move between trees when the temperature is between 16-20^{o}C, hence the spring-autumn period has tended to be the operational DED 'season', with trees that are suitable for breeding needing to be of a suitable size (about 15 years old) and condition; the condition being vital to the beetles' breeding success.

Currently, there are no means to eradicate DED, which means that a programme to manage the spread of DED is an open-ended commitment. The most effective means of containing the disease is through a combination of methods to minimise the beetle population. The main action is to fell infected trees, or parts of trees, and burn these to prevent further spread of the disease. Trunk girdling can also prevent the spread of the disease via the roots. Treating with chemicals (fungicidal injection) is very costly (approximately £300/tree) and not entirely effective, as many factors must be regulated to ensure potential success. Chemical treatment must also occur year on year to be successful (which can physically damage the tree).

DED is estimated to have killed approximately 25 million elms in the UK since the 1960s. ESCC established a Dutch Elm Disease (DED) control programme in 1971. The SDJC then managed the programme for the area between Brighton and Eastbourne for most of its existence, on behalf of ESCC. ESCC subsequently took over the delivery of the programme from April 2011, when the South Downs National Park Authority came into being.

The Dutch Elm Disease (Local Authorities) Order 1984 empowers, but does not require, Local Authorities to serve notice on owners of diseased trees, requiring the owner to carry out felling and appropriate disposal. Should this not happen in the time period advised, the Order permits an appropriate officer to serve notice before entering private land to enable the sanitation felling or other work considered necessary for the control of DED to be carried out, with the costs recoverable from the landowner.

In 2012 ETE Scrutiny Committee recommended that a review of the current approach to managing DED be carried out to:
1) provide an up-to-date evidence-based decision as to whether to carry on the sanitation programme;
2) if the decision is to maintain the sanitation programme then develop a strategy with key partners to ensure that the approach is:
    a) financially sustainable

b) likely to be effective in the long term.

This strategy is in response to the recommendation from Scrutiny.

What is it we're trying to achieve?

The objectives of the DED sanitation programme are to:

1) ensure the long-term survival of a significant population of mature English elm, which make an important contribution to the local landscape and provide a habitat to priority species (e.g. butterflies (red data list – white hairstreak) and 200 species of lichen (red data list – orange fruited elm-lichen and 5 others). The Sussex elm population is considered by Natural England to be of regional importance, with Brighton & Hove housing the National Elm Collection. In addition, all public bodies have a duty to have regard to biodiversity in all of their work (the Natural Environment and Rural Communities Act, 2006).

2) Assist in managing DED on the highway, just as any other land owner is required to do under the Highways Act 1980 (section 154), and on ESCC land (eg. schools), when it poses a health and safety risk.

3) Ensure the most cost effective approach.

What is the best way of doing it?

ESCC has been working with DEFRA's Food and Environment Research Agency (FERA) to compare the costs and effectiveness of:

a) stopping the DED sanitation programme;

b) returning to how the programme was delivered prior to ESCC taking it back in-house in April 2011, and;

c) continuing with the prioritised approach adopted by ESCC in 2012.

FERA has funded modelling work by the University of Cambridge to compare the costs and effectiveness of these approaches. The report is included as Appendix 1. The key conclusions are that the prioritised approach to control leads to:

Table 1, below, provides a summary of the estimated costs over 10 years and 25 years, and the effect on the population of mature elm trees. The conclusions from the table are that:

1) stopping the sanitation programme is more costly over 10 years than maintaining the programme, with the breakeven period estimated to come after approximately 20 years;

2) the prioritised approach to sanitation enables a larger population of healthy mature elm trees to survive, because fewer trees would become infected and

require felling. This, in turn, would mean that the prioritised approach is more financially sustainable than the historic approach.

Table 1. Comparison of the costs & effectiveness of the different approaches to DED.

Approach	**Total healthy elm population over 10 years (1)**	**Total healthy elm population over 25 years (1)**	**Number of elms felled in 10 years (2)**	**Number of elms felled in 25 years (2)**	**Cost of control programme over 10 years (3)**	**Cost of control programme over 25 years (3)**
No control	7,000	6,000	5,210	5,210	£1,228,050(4)	£1,228,050(4)
Historic	13,000	12,500	10,500	30,000	£638,900 (5)	£1,597,250(5)
Prioritised	14,000	14,500	11,500	25,000	£591,100 (6)	£1,477,750(6)

Notes:

1) The figures for the number of healthy elm trees are from Figure 5 of the Modelling report in Appendix 1.

2) the figures for the number of elms felled are from Figure 5 of the Modelling report in Appendix 1.

3) average costs:
    a. The numbers of trees felled under the historic or prioritised approach are taken as the proportion of the total known population of street elm (2.6%), highway elm (14.3%), ESCC estate elm (2.7%) and private elm (80.4%) felled in 2012.
    b. The cost of felling and removing urban street trees and reinstating the footway has been provided by ESCC Highways (please see below). Please note that the average figure provided by EBC for felling and reinstatement is significantly higher, at £1,750/tree, as this includes replanting.
    c. The salary costs are assumed to be:
        i. The ‘no control’ scenario: 6 month’s salary each year for 7 years, at £13.75K pa., assuming that an Officer will be required to manage sanitation felling just during the DED season.
        ii. The historic and prioritised scenarios: 1 FTE for 10 or 25 years at £27.5K p.a.

4) Cost of ‘no control’: assumes that all street elms (2048 trees at £460/tree), all rural highways elms (2631 at £60/tree) and all elms on ESCC property (531 at £60/tree) will have been felled, including salary costs of £125,000.

5) Cost of the ‘historic’ approach:
    a. Over 10 years: assumes that 195 street elms felled (£460/tree), 1070 rural highway elms felled (£60/tree), 200 elms on ESCC property (£60/tree) and 6000 rural private elms felled (£55/tree with 40% recovered from private landowners), including salary costs of £275,000.
    b. Over 25 years: assumes that 488 street elms felled (£460/tree), 2675 rural highway elms felled (£60/tree), 500 elms on ESCC property

(£60/tree) and 15000 rural private elms felled (£55/tree with 40% recovered from private landowners), including salary costs of £687,500.

6) Cost of the ‘prioritised’ approach:
   a. Over 10 years: assumes that 170 street elms felled (£460/tree), 930 rural highway elms felled (£60/tree), 175 elms on ESCC property (£60/tree) and 5200 rural private elms felled (£55/tree with 40% recovered from private landowners), includes salary costs of £275,000.
   b. Over 25 years: assumes that 425 street elms felled (£460/tree), 2325 rural highway elms felled (£60/tree), 438 elms on ESCC property (£60/tree) and 13,000 rural private elms felled (£55/tree with 40% recovered from private landowners), includes salary costs of £687,500.

It’s important to note that the limited data kept on the historic sanitation programme and the complex epidemiology of DED mean that the conclusions from the modelling work are subject to extensive caveats. In addition, it’s not possible to include an assessment of some factors, for instance the predicted effect of climate change and water stress, or the predicted effects of other tree diseases. Nevertheless, the modelling work is based on the most up-to-date knowledge of the disease and draws a reasonable set of conclusions to help inform the strategy for DED control going forward from 2013-14.

The DED control strategy

To achieve the objectives set out above, the key deliverable of the sanitation programme will be to reduce the beetle population by adopting a prioritised approach to sanitation felling within the existing control zone covered by the Dutch Elm Disease (Local Authorities) Order 1984. This will entail:

1) Removing trees containing grubs (brood trees) or in condition to contain grubs (host trees).

2) Using host trees spotted late in the season as trap trees, which will be felled during colder weather.

The following sections set out the proposed control strategy

1. Finance
- Maintain the current ESCC budget level at approximately £100,000 p.a.
- Continue to seek 50% of the costs of the sanitation programme from private landowners (including public bodies).
- Deal appropriately with private individuals who are unable or unwilling to contribute 50% of the costs, on a case-by-case basis.
- Keep the 50% contribution rate under review.
- Seek external funding from key stakeholders to deal with clearing any DED to commit to clearance of DED work back-log.

2. Control zone
- Maintain the current control zone boundary, shown in figure 1.

- Keep a check on areas on the periphery of the control zone where elms are present, to avoid ‘flare-ups’ that could affect our programme. This may entail felling trees that pose a significant threat to important areas on the edge of the control zone. This would be exercised as in the control zone.

3. Spotting

- Maintain a volunteer database across the control zone.
- Hold registration/training sessions early in the season.
- Use volunteers to help mapping of elms.
- Create more effective way to collect and collate elm information from volunteers/public if possible or practical.
- Implement the communications actions below to improve spotting & reporting.

4. Contractors

- We will use a number of local and/or regional contractors to ensure a wide spread of continuous work can occur across the control zone
- Contractors completed a tendering process in April 2011. We continue to use this same contractor base with addition of local SME’s should they meet the high standards of health and safety and insurance ESCC require. Their work is monitored continuously.
- Where possible, keep contractors working in areas they know or where land owners know them.
- Use larger, less local contractors to cover extra workload and more specific jobs, e.g. requiring specialist work, machinery or legislation.

5. Felling process

- Prioritise felling of trees currently being used as a breeding habitat within the shortest period of time possible.
- Prioritise felling of trees able to be used as a breeding or over-wintering habitat during that season within the shortest period of time possible (unless trees can be used as ‘trap trees’ (i.e. to ‘trap’ breeding beetles).
- Prioritise felling trees depending on the risk they pose, e.g. to EBC and BHCC, to healthy, mature specimens, to street elms.
- Assess any site risks before work is commissioned and monitor hazards, e.g. before and after photos, photos of buildings or other risks, check contractors risk assessments.
- Figure 2 sets out the steps in the control process.

6. Surveying and mapping

- Continue to survey and map trees between the B2124 and A27 to decide if boundary readjustment is necessary (see point 2 of control zone paragraph).
- Continue to readjust mapping using handheld GIS device to maintain population data.

Figure 2. Steps in the DED control process.

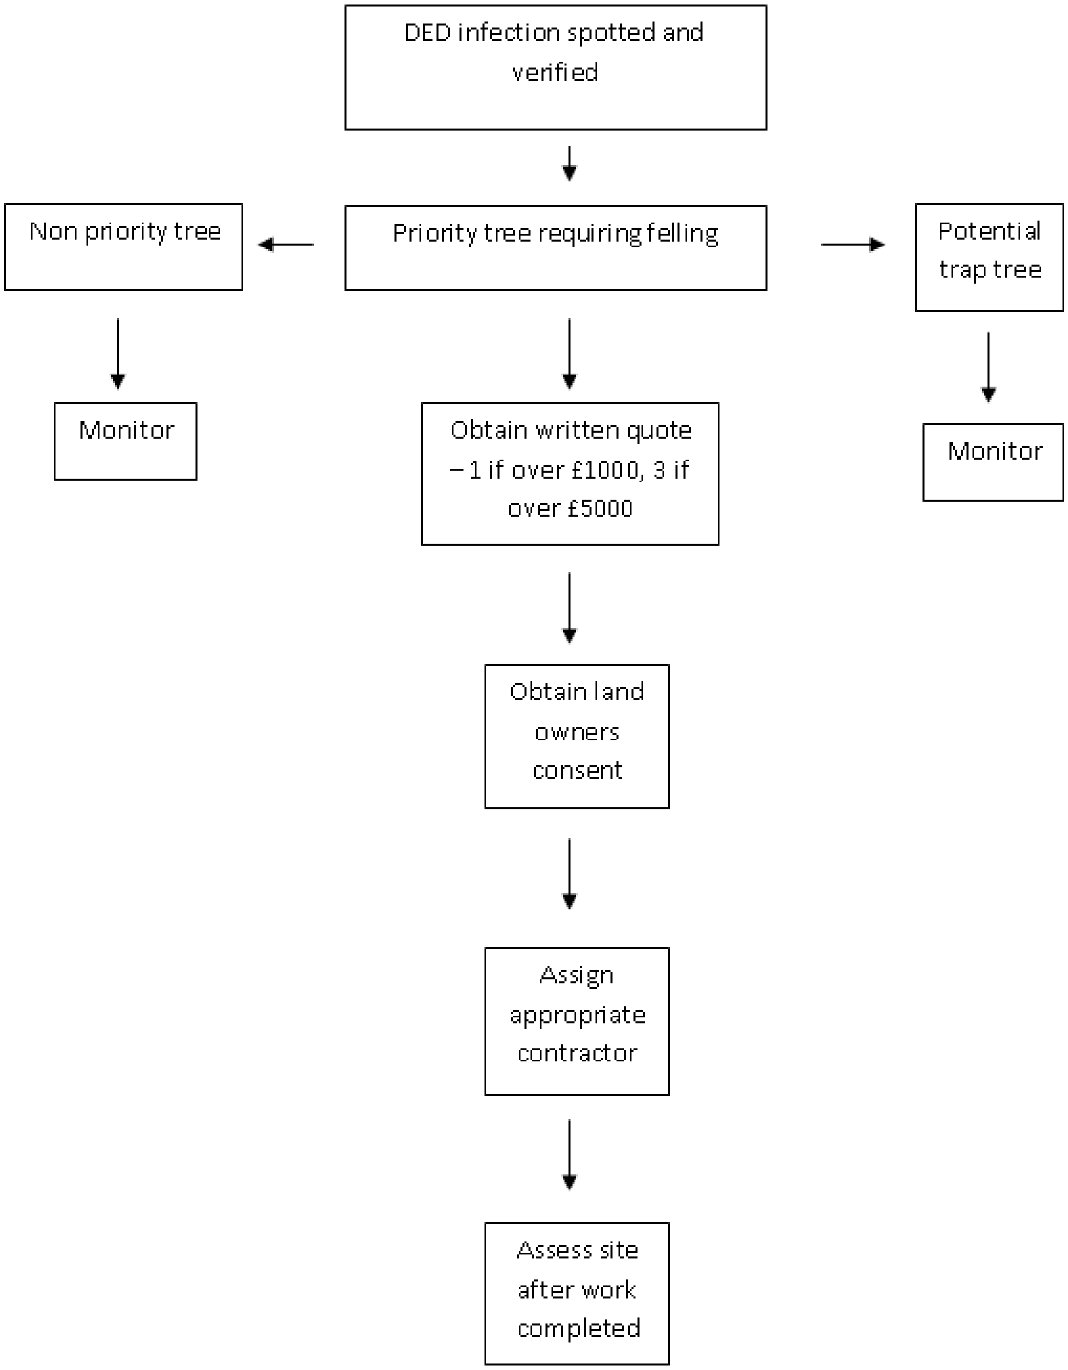

7. Communications

- Use the media to actively promote the control programme and its success over the past 30 years.
- Keep local and national journalists fully informed by issuing standard press releases and offering suitable news and feature stories linked to our key messages.
- engender a sense of personal and community responsibility and encourage residents and staff to get involved with protecting the elm population of East Sussex by:
    - Promoting our website
    - Promoting the contact centre
    - Using the intranet and departmental newsletters
    - Using the East Sussex Elms Facebook site
    - Providing members with written or verbal briefings
- exercise control over the tone of media coverage and combat myths about the control of Dutch elm disease.
- Key spokespeople:
    - East Sussex County Council Lead Cabinet Member for Economy, Transport and Environment, for messages relating to policy.
    - Dutch Elm Disease Officer, for technical messages.

8. Partnership working

- Continue with regular meetings of the Elm Partnership (6 monthly).
- Work with Conservation Foundation to seek additional funding for sanitation and replacement planting.
- Work closely with the South Downs National Park Authority, for instance regarding the use of volunteers and staff, and potential sources of funding.
- Maintain close links with Plumpton College (e.g. felling, monitoring, volunteers).

9. Monitoring & reporting

- maintain the existing database with information on each infection site (e.g. land owner name, address, number of trees, quotes, etc).
- Continually assess and monitor contractor's work throughout season.
- Provide an annual report on progress with implementing this strategy, covering:
    - Numbers of trees felled and at what cost
    - Changes in the total elm population
    - Review of the assumptions made above, to incorporate relevant new evidence.

10. Risks

Risk	Potential impact	Measure to address the risk
Beetle population increases as elm population increases	Increased beetle population can increase the amount of infection that can occur annually due to more potential breeding ground in mature elms.	The prioritised approach aims to reduce the potential breeding ground directly by felling elms in the specific condition required by the beetle for breeding.
Length of season increases due to climate change	Warmer and earlier springs could see beetle emergence before budgets are decided allowing the disease to begin spreading earlier. A longer DED 'season' can increase the number of breeding cycles the beetles successfully complete, increasing beetle population during the season and number of infections.	Having enough budget to be able to clear all brood tree back log will reduce the emerging population and resulting infection. By prioritised targeting, monitoring and felling of brood and host trees, breeding cycles can be anticipated and halted.
Disease spread by storing and transporting diseased wood	Unknown beetle breeding sites cause 'flare ups' in infection that are unpredicted. This could be in high priority areas of significant landscape value elms. Large stores of brood wood can create vast amounts of new infection, increasing annual costs. Transporting brood wood from areas of high infection to areas of low infection will increase costs. If infection is spread to areas where money has already been spent on clearing infection there is an added cost.	Careful monitoring by the DED Officer will help reduce the potential occurrence of infected log piles. Education of the public through the communication plan will also help reduce this means of infection. Checks on contractor's facilities and methods should occur through the season to ensure they are not part of the problem.

Figure 1. DED Control Zone.