The Optimal Control of Dynamic Pest Populations

Sean Hackett
Linacre College
University of Oxford

A thesis submitted for the degree of Doctor of Philosophy
Hilary 2018

Acknowledgements

I would like to thank Mike Bonsall for his patient and supportive supervision throughout my DPhil. I would also like to thank Nina Alphey for helping me to get started and Marc Mangel for introducing me to dynamic programming and allowing me to spend time with his group at UC Santa Cruz. I thank Neil Morrisson and Martha Koukidou for accidentally setting me on the dark path of approximate dynamic programming. My thanks to the members of MERG and my fellow students in the Department of Zoology for their support throughout and for good-naturedly abiding my gallows humour. I thank my family for encouraging me to get this far. A special thank you to Kat, Aled and Stephen, for cheerfully enduring my idiosyncrasies for over three years without complaint and for making Oxford feel like home, and to Jack for helping me to maintain a sense of perspective in the final stretch.

This thesis is dedicated to the memory of Male and Female Stick Insect. We never did get around to naming you and we can only hope that you died of old age and not because we periodically forgot to refresh your food.

This project was funded by a CASE studentship between the BBSRC and their industrial partner Oxitec®.

Publications

The following publications have arisen from this thesis and are presented in chapters 2 and 3:

- Hackett, S.C. & Bonsall, M.B. (2016) Type of fitness cost influences the rate of evolution of resistance to transgenic Bt crops. *Journal of Applied Ecology*, **53**, 1391-1401
- Hackett, S.C. & Bonsall, M.B. (2018) Management of a stage-structured insect pest: an application of approximate optimization. *Ecological Applications,* doi: 10.1002/eap.1700

Contents

Thesis Abstract

In the management of agricultural insect pests, short-term costs must be balanced against long-term benefits. Controls should be selected to account for both their immediate and downstream effects upon the demography and genetics of the pest, enabling suppression today without threatening suppression tomorrow. The iterative, algorithmic method of dynamic programming can provide optimal solutions to problems of this type, in which actions are taken sequentially and each action may influence those which follow it. However, this approach is fundamentally constrained with regards to the magnitude of the problems it may solve. As questions of insect pest management can be subject to ecological and evolutionary complexities, this may place them beyond the scope of dynamic programming. When it is the intricacies of a problem that are of interest, it may be more productive to utilise approximate dynamic programming (ADP) methods which can attempt problems of arbitrary complexity, although at the expense of no longer guaranteeing optimality. In this thesis I first challenge a dynamic programming algorithm with the management of a hypothetical insect pest feeding upon a transgenic insecticidal crop. The model explores how different realisations of fitness costs to resistance influence the algorithms suggested actions. I then apply a brute-force variant of ADP, a lookahead policy, to the management of a stage-structured, continuously reproducing pest population. This was to explore the extent to which an algorithm with a limited temporal perspective is able to balance the timetable of pest demography against the timescale over which insecticidal sprays and bisex-lethal sterile insect releases unfold. This same decision framework is then applied to a modified problem in which resistance to insecticidal toxins may evolve and releases are now male-selecting. This was used to assess the

efficacy with which simple lookahead policies utilise a control with delayed benefits (the male-selecting releases) and possible constraints on their capacity to respond to resistance evolution. Dynamic programming and ADP methods offer a versatile toolbox for accounting for the potential impacts of the evolutionary and ecological peculiarities of particular pests upon control decisions.

Chapter 1: Introduction

Agricultural insect pests remain a prominent threat to yields and to the livelihoods of growers. Agricultural ecosystems have provided ideal conditions in which a range of herbivorous insects have been able to thrive at the expense of yields, while the establishment of global trade links has facilitated the dispersal and establishment of previously endemic pests such as the diamond back moth and the Mediterranean fruit fly across the globe (Malacrida *et al.* 2007; Furlong, Wright & Dosdall 2013; Szyniszewska & Tatem 2014). Whether they threaten yield directly by feeding upon the tissues of the crop or indirectly by vectoring plant pathogens, sustainable agriculture requires dependable strategies for the suppression of insect pests. However, the diversity of insect life histories and ecologies is such that a one-size fits all approach is unattainable – given a specific combination of pest, environment and crop, some strategies will fare better than others. Additionally, it is simply not sufficient to suggest actions which will offer effective suppression as neither time nor resources are unlimited and the act of managing a pest population carries its own costs. That is, there are diminishing returns upon increasing suppression effort which can result in the costs of suppression exceeding the benefits (Oerke 2006; Mitchell & Onstad 2014; Brown & Staňková 2017). Consequently, complete eradication of a pest is rarely feasible and cost-effective management becomes not only a question of "which tools should be used and when?" but also "what level of infestation can we tolerate?" (Oerke 2006; Mitchell & Onstad 2014; Brown & Staňková 2017). Thus, the management of agricultural insect pests is a multi-faceted subject combining ideas from economics, ecology and, where pests become resistant to control, evolution.

Methods of Control for Suppression and Resistance Management

A diverse array of management tools has been developed to facilitate the suppression of insect pests in agriculture and in other fields such as public health. The most prominent tool is the application of insecticides, toxicants deliberately released into the environment to reduce the abundance of insect pests. The benefits of insecticides remain sufficient to facilitate their persistence (Oerke 2006; Cooper & Dobson 2007; Sjöberg *et al.* 2015) in spite of the complex cascade of direct and indirect effects they impose upon the agricultural environment, the environments connected to it and the non-target ecological communities which inhabit them (Guedes *et al.* 2016; Becker & Liess 2017). In contrast to insecticides, the sterile insect technique (SIT) presents an environmentally benign, species-specific means of insect population control via the release of large numbers of sterilised insects, preferably males, which has been successful in controlling pests such as Medfly and pink bollworm (Dyck, Henrichs & Robinson 2005). Traditionally, the insects to be released are sterilised via exposure to radiation, thus they are not sterile in the sense of being agametic but most or all of their progeny will bear radiation induced dominant lethal mutations (Dyck, Henrichs & Robinson 2005). Releasing sterilised insects in sufficiently large numbers impairs the reproductive potential of the target population resulting in suppression or, with appropriate area-wide planning, the eventual eradication of the pest (Enkerlin 2005). For example, area-wide SIT programmes resulted in the elimination of the screwworm *Cochilomyia hominivorax* from the United States, Mexico, Central America, and several Caribbean islands (Vargas-Teran, Hoffman & Tweddle 2005) Additionally, SIT can be utilised to create buffer zones preventing the invasion or, subsequent to eradication, reinvasion of an insect pest (Enkerlin 2005; Vargas-Teran, Hoffman & Tweddle 2005). Notably, insecticides and SIT are not mutually exclusive and can, in some instances, be complementary: sprays tend to be most effective when employed against large populations while, up to a point, the efficacy of SIT may increase as the population declines (Barclay 2005).

More recently, genetic engineering has allowed for the production and deployment of transgenic insecticidal crops – crops which have been modified to systematically express insecticidal compounds, typically derived from the Cry toxins produced by the soil bacterium *Bacillus thuringiensis*. Commonly referred to as Bt crops, transgenic crops expressing Cry toxins present a compelling alternative to insecticides as their defence is essentially automatic, requiring no input from the grower. In this respect, transgenic insecticidal crops are ecologically analogous to plants which endogenously express defensive compounds to deter herbivory (Gassmann, Onstad & Pittendrigh 2009). Furthermore, the toxins expressed by Bt crops are highly specific, having been selected for their efficacy against a particular species or taxon, and only those insects directly consuming the tissues of the crop will be exposed unlike insecticides which can come into contact with non-target organisms. This facilitates the easy integration of transgenic crops with both SIT and the release of natural enemies such as parasitoid wasps (Gurr & Kvedaras 2010; Tabashnik *et al.* 2010). Since their introduction, the popularity of Bt crops has quickly accelerated to the extent that an area encompassing approximately 98.5 million hectares was planted with Bt crops in 2016 (Tabashnik & Carrière 2017). The widespread adoption and implementation of Bt crops has generally (but not universally, see Romeu-Dalmau *et al.* 2015) yielded substantive economic and environmental benefits, in part due to a corresponding decrease in insecticide use (Hutchison *et al.* 2010; Lu *et al.* 2012; Dively *et al.* 2018).However, in managing a population, we are applying a significant selection pressure and should genetic variation exist which confers the capacity to either tolerate or otherwise overcome this pressure, the average genetic composition of the population will shift accordingly (Onstad 2014). That is, the act of controlling a population can select for resistance to the control. Thus, there is more to effective, long term control of insect pest populations than killing. Driving a pest population to low densities provides short term reductions in damage, but can ultimately accelerate selection for resistance unless measures are implemented to moderate the selection pressure (REX Consortium 2013; Brown & Staňková 2017). I adopt the terminology of Onstad (2014) in using resistance as a general term describing any heritable trait which grants tolerance to a management tactic which can include behavioural mechanisms such as that which has enabled populations of the Western Corn Rootworm (*Diabrotica virgifera*) to escape control via crop rotation (Levine *et al.* 2002).

The evolution of resistance presents a significant barrier to our capacity to continue to manage insect pests into the future. It adds an evolutionary component to questions of pest management, partially recasting them as the genetic equivalent of a harvesting problem: how can we exploit susceptibility to control without depleting it? Put another way, preventing or delaying the evolution of resistance necessitates thinking not only in terms of what we stand to gain from control today but also what we may stand to lose in future seasons should resistance become prevalent. In response to the potential for resistance evolution, the field of resistance management has emerged to address how population controls can be selected and implemented in such a way that susceptibility to control is prolonged (Onstad 2014). Broadly speaking, there are two non-mutually exclusive approaches to preserving susceptibility. Either control is implemented in such a way that part of the pest population is insulated from its effects, ensuring that some susceptible individuals survive, or the population is presented with multiple stressors at once, reducing the likelihood that the pest will develop resistance to any of the controls individually (REX Consortium 2013; Sudo *et al.* 2018). Resistance management strategies for insecticides tend to favour the latter approach by utilising rotations (the application of two or more distinct toxins in alteration), mixtures (the application of a cocktail of one of more toxins) or, at the landscape scale, mosaics (the application of distinct toxins to specific regions) (Sudo *et al.* 2018).

Being a newer technology, transgenic Bt crops were able to benefit from observations of insecticide control failures in preceding decades and were implemented in conjunction with a resistance management strategy: the high-dose/refuge strategy (Tabashnik 2004). The high-dose/refuge strategy falls into the class of resistance management strategies which maintain susceptibility by explicitly sparing some individuals from exposure to the selective agent. More specifically, the landscape to be planted with Bt crops is partitioned between the insecticidal crop and a non-toxic refuge which promotes the survival and persistence of susceptible insects, allowing them to confer their susceptibility on subsequent generations (Carrière & Tabashnik 2001; Tabashnik 2004; Carrière, Crowder & Tabashnik 2010; Sudo *et al.* 2018). The high-dose/refuge strategy is able to delay resistance further by requiring that transgenic crops express a toxin dose sufficient to kill over 95% of exposed resistant heterozygotes or 99.99% of susceptible homozygotes (Gould 1998; Tabashnik 2004). Satisfying this requirement renders the expression of a putative resistance allele functionally recessive. Given an initially rare, functionally recessive resistance allele, the high-dose/refuge strategy can delay or even halt the introgression of resistance (Tabashnik, Brévault & Carrière 2013; Carrière, Fabrick & Tabashnik 2016; Tabashnik & Carrière 2017). Where resistance to transgenic crops has emerged, it is generally linked to the violation of the high-dose assumption (Tabashnik, Brévault & Carrière 2013; Carrière, Fabrick & Tabashnik 2016; Tabashnik & Carrière 2017). For example, transgenic corn expressing the Bt toxin

Cry1F failed to satisfy the assumption of a high-dose for the fall army worm *Spodoptera frugiperda* and resistance evolved to non-trivial levels in only three years (Tabashnik, Brévault & Carrière 2013). As an additional barrier to resistance evolution, pyramided crops expressing two or more toxins are now the norm. Pyramided crops operate under the same principles as an insecticide mixture. By presenting the pest with multiple toxins it is theoretically constrained to independently evolve resistance to each toxin before it can survive on the crop (Carrière, Fabrick & Tabashnik 2016). However, even when deployed in conjunction with refugia, pyramided crops are not immune to the evolution of resistance. The presence of cross resistance, in which resistance to one toxin confers some level of tolerance to a second toxin, non-recessive inheritance of resistance, failure of the high-dose assumption for one or more of the component toxins and poor placement of pyramided crops relative to single toxin crops expressing one of the pyramid components are among the major factors which may undermine the efficacy of a pyramid (Carrière, Fabrick & Tabashnik 2016; Tabashnik & Carrière 2017).

The integration of insecticidal toxins with additional control methods which act independently of any resistance alleles can also contribute to resistance management. In particular, developments in genetic engineering have generated novel SIT variants such as gene-drive systems, male-selecting strains and sex-ratio distortion (Bourtzis *et al.* 2016; Leftwich, Bolton & Chapman 2016; Harvey-Samuel, Ant & Alphey 2017; Alphey & Bonsall 2018). Of these technologies, male-selecting strains are currently of greatest relevance to issues in resistance management as they possess the capacity to simultaneously supplement the suppression offered by insecticidal toxins while introgressing alleles for susceptibility into the population (Alphey *et al.* 2007; Alphey, Bonsall & Alphey 2009; Harvey-Samuel *et al.* 2015; Zhou *et al.* 2018). Specifically, male-selecting technologies operate by introducing a dominant lethal allele which is conditionally expressed only in female insects when specific dietary supplements, which the insects cannot obtain in the field, (such as tetracycline) are withheld (Alphey & Bonsall 2018). Such systems have the additional benefit of automatically enabling male-only releases, simplifying the process of preparing for a release. The principles underlying male-selecting releases are similar to those of conventional SIT: inundate the population with sufficient transgenic males such that females are more likely to mate with a transgenic male than a wild type male. However, unlike traditional SIT, only the female offspring arising from mating with transgenic males die, male offspring develop normally. Consequently, male-selecting releases have a reduced capacity for suppression relative to traditional SIT in which, theoretically, no viable offspring should be produced (Alphey, Bonsall & Alphey 2009). This reduction in absolute suppression potential may be compensated by the intriguing possibility of utilising male-selecting releases as a proactive resistance management tool. That is, these releases allow for the active addition of susceptibility alleles to the target population in contrast to the more passive approach represented by refugia which reduce the rate at which susceptibility is lost (Alphey *et al.* 2007; Alphey, Bonsall & Alphey 2009; Harvey-Samuel *et al.* 2015; Zhou *et al.* 2018). Furthermore, male-selecting releases are surprisingly robust to the evolution of resistance to the lethal construct as alleles conferring susceptibility to the construct are reintroduced with each release (Watkinson-Powell & Alphey 2017).

When planning for the possibility of resistance evolution, it can be useful to consider the potential for any fitness costs associated with a resistance allele to synergise with management efforts. Given a sufficiently effective control, resistance alleles of small effect will be selected against as they will be unable to confer adequate tolerance upon the pest. Instead, it is assumed that major mutations will be favoured which will be associated with negative pleiotropic effects which disadvantage resistant insects in the absence of selection (Macnair 1991; Coustau, Chevillon & ffrench-Constant 2000; Gassmann, Carrière & Tabashnik 2009; ffrench-Constant & Bass 2017). While such costs are theoretically advantageous for both suppression and resistance management, particularly with respect to the high-dose/refuge strategy in which susceptible and resistant insects may compete against one another in the refuge, it is frequently taken as a given that resistance will always decline in the absence of selection. However, this ignores the significance of fitness costs, how these manifest under field conditions, and whether the fitness cost mechanisms are relevant under field conditions (Gassmann, Onstad & Pittendrigh 2009; ffrench-Constant & Bass 2017). Furthermore, there is the implicit assumption that resistance is always acquired via a de novo mutation but this is not guaranteed as resistance may also be derived via other means which are less likely to be associated with practical fitness costs such as pre-existing adaptations or gene duplication (ffrench-Constant & Bass 2017). If a resistance allele cannot be demonstrated to carry appreciable fitness costs under field conditions then management strategies which presume the existence of fitness costs may fall short of more pessimistic alternatives which assumed fewer barriers to resistance.

Sequential Decision Problems and Dynamic Programming

The effective long-term management of an insect pest requires that, of the options available to us, only that which produces the best result over the full time horizon of interest be selected. That is, if our stated management objective is to achieve cost-effective suppression over the duration of a season then the best action we can take now is the action which would leave us in the best position at the season’s end. To do this requires not only that we be aware of the ecology, demography and, if relevant, genetics of the focal pest population but also that we understand how the effects of management decisions unfold over time and how they may influence the decisions we have yet to make. Decision problems of this type, in which actions are taken recurrently, each action has the potential to influence the actions and rewards which follow it, and the components of the system evolve over time in response to decisions, the environment and properties of the system itself, are classed as dynamic programs (Powell 2011). Put another way, the future states of the system we seek to manage are potentially a function of its current state (*e.g.* the number of mated reproductive female insects), any stochastic effects (*e.g.* temperature, precipitation), what we have already done to the system (Are pesticide residues present or absent? How large is the refuge? Are any sterile male insects still present?) and what we may choose to do to the system now. Our goal when solving a dynamic program is to identify the action, or sequence of actions which guide the state trajectory of the system towards our desired outcome. However, to solve a dynamic program by the brute force application of each action to each state at each decision point produces an intractable, exponentially growing tree of state-action pairs which rapidly becomes unachievable for all but the most basic problems (Powell 2011). Instead, dynamic programs are solved using a policy which is defined as any rule for mapping permissible actions onto possible states (Powell 2014).

Any dynamic program, independent of the subject matter, can be summarised by five components (Powell 2011, 2014). In no particular order, these are the state, the transition function, the actions, the objective function and any exogenous information. The state of a dynamic program is the minimally dimensioned function of history required to model the transitions of the system from a given point in time onwards (Powell 2014). The transition function is a catch-all term for the systems of equations which define how the components of the state evolve over time in response to the state itself and the actions applied to it which, at each sequential decision point, can only be drawn from a defined set of permissible actions (Powell 2014). The objective function is a mathematical statement of intent, it defines the property of the system which we seek to maximise or minimise and the timeframe over which we seek to do this (Clark & Mangel 2000; Powell 2011, 2014; Marescot *et al.* 2013). The objective function provides the means by which the chosen policy can evaluate the performance of an action relative to other actions. Finally, exogenous information refers to any stochastic processes and variables acting upon the system (Powell 2011, 2014).

Ecological dynamic programs are typically solved using dynamic programming, an algorithmic method which sub-divides dynamic programs into a sequential series of smaller sub-problems, the solutions of which can be reassembled to produce an optimal solution to the dynamic program (Clark & Mangel 2000; Powell 2011; Marescot *et al.* 2013). Dynamic programming is able to solve problems in this way by virtue of Bellman's equation which characterises the optimal value V of a system in state S (where S may be a vector of state variables) at time t as,

$$V_t(S_t) = \max_a\{C_t(S_t, a_t) + V_{t+1}(S_{t+1})\}. \quad [1]$$

Bellman's equation identifies the action a which maximises the current $C_t(S_t, a_t)$ and future values $V_{t+1}(S_{t+1})$ of a decision we are attempting to make now. The equation assumes optimal behaviour from period $t + 1$ onwards and future values may be subject to discounting when appropriate. The Bellman equation defines an iterative, algorithmic policy which identifies the optimal mapping between states and actions by recursively identifying the action that optimises the sum of both current and future benefits for each combination of state, action and time. More specifically, the algorithm exploits our prior knowledge of the terminal conditions of the system to utilise backwards recursion. That is, the sub problems are solved in reverse, from the terminal time T to the current time t. This is made possible by the observation that it is always possible to make a statement regarding the value of the terminal period $V_T(S_T)$ whether this is a function of the terminal state or we simply assume no further benefits/costs accrue beyond the terminal period in which case $V_T(S_T) = 0$. If the terminal conditions are known for each value of S_T then Bellman's equation can find the value of $V_{T-1}(S_{T-1})$ for each value of S_{T-1}. These values can then be used to solve for $V_{T-2}(S_{T-2})$ and so on, until the backwards recursion process iterates to the current time t. By approaching the problem in this fashion, dynamic programming automatically accounts for the downstream effects of actions, guaranteeing that any suggested actions are optimal with respect to the stated time horizon (Clark & Mangel 2000; Marescot *et al.* 2013). Expressed mathematically, the optimal action at time t, a_t^* is that which satisfies

$$a_t^* = \arg\max_a \{C(S_t, a_t) + \sum_{t'=t+1}^{T} (C(S_{t'}, a_{t'}))\}. \quad [2]$$

Given the relative simplicity of dynamic programming, the guarantee that the suggested actions will be optimal and the indifference of the algorithm to stochasticity and non-linearities, dynamic programming has been widely applied in ecology, conservation and evolution (Shoemaker 1981; Mangel 1987, 2015; Plant & Mangel 1987; Houston & McNamara 1999; Clark & Mangel 2000; Shea & Possingham 2000; Marescot *et al.* 2013). Consequently, the concept of dynamic programs, the class of problem we are trying to solve, are often conflated with dynamic programming, the method we use to solve them (Marescot *et al.* 2013; Powell 2014). This is unfortunate given that dynamic programming is not a universal solution to any and all dynamic programs. The range of questions which can be solved by dynamic programming is constrained by the critical assumption that the solution, the optimal pairings of states and actions for each time period, can be represented as a lookup table. Thus, dynamic programming is subject to the curse of dimensionality, problems become intractable to it once they exceed a certain magnitude. More specifically, dynamic programming is constrained by three distinct, non-mutually exclusive curses of dimensionality: (i) the magnitude of the state space (how many state variables are there and how many values can they assume?), (ii) the number of available actions (the more actions there are, the more potential state-action pairs which must be evaluated and tabulated) and (iii) the magnitude of the outcome space (how many stochastic elements are there and how many values can they assume?) (Powell 2011).

Methods such as discretisation of continuous state variables and linear interpolation can sometimes reduce the dimensionality of a problem to the extent that dynamic programming can still be usefully applied (Shoemaker 1982; Clark & Mangel 2000) but such tools have their limitations and in simplifying a problem such that we can apply our preferred methodology we may sacrifice realism (Nicol *et al.* 2010; Powell 2011). Whether, this trade is acceptable will be contingent upon the nature of our management objectives but, given the complexity and potential intricacy of many decision problems in ecology, there may be scenarios where it is the granular details of a problem which are of greatest interest or in which reducing the scope of the problem risks distorting the meaning and relevance of the solution (Wilson *et al.* 2006; Nicol *et al.* 2010; Powell 2011). With respect to the management of agricultural insect pests, we may be interested in how insects distribute themselves across, and move within space and how management actions, such as SIT or transgenic releases are resolved across this space

(Yakob & Bonsall 2009; Miller & Sappington 2017; Watkinson-Powell & Alphey 2017). For example, larval movement within a field can potentially undermine the longevity of transgenic crops depending upon how refuge plants are distributed relative to transgenic plants (Onstad 2006; Brévault, Tabashnik & Carrière 2015; Carrière, Fabrick & Tabashnik 2016; Garcia *et al.* 2016). Similarly, the demography, and relative timings of mating and dispersal can influence the efficacy of particular spraying regimes (Sudo *et al.* 2018). Furthermore, the efficacy and acceptability of transgenic insect releases is potentially subject to a range of genetic, economic, ecological and regulatory factors which may be of interest (Baltzegar *et al.* 2018).

Dynamic programming is but one policy for the solution of dynamic programs, one which can derive optimal solutions for problems of sufficiently low dimensionality, but other policies exist which do not share its limitations. For example, the most basic policy is a myopic policy in which our decision rule is simply to identify the action at time t which produces the best outcome at time $t + 1$. By virtue of their simplicity, myopic policies are entirely agnostic to the scale of the problem and can be applied to questions of a magnitude well in excess of those which can be solved by dynamic programming (Wilson *et al.* 2006; Powell 2011). However, the obvious short-coming of myopic policies is that they completely disregard the downstream value of their decisions which, intuitively, often results in them suggesting poor quality actions and limits their capacity to evaluate any actions for which benefits are either delayed or distributed over an interval of time (Powell 2011). This focus upon only the immediate effects of a decision would, for example, impede a myopic policy attempting to evaluate a male-selecting transgenic release given that the benefits of such a release may not be fully realised until the following pest generation (Leftwich *et al.* 2014). However, myopic policies are not the only alternative to dynamic programming and in this case simply illustrate that, by relaxing the restrictive assumptions of dynamic programming, it is possible to arrive at more flexible policies. In particular, a powerful suite of policies for the solution of large-scale decision problems has been developed under the umbrella term approximate dynamic programming (ADP) (Powell 2011).

Approximate Dynamic Programming

Approximate dynamic programming encompasses four broad classes of policy which have been explicitly developed to overcome the curses of dimensionality (Powell 2011, 2014; Powell & Meisel 2016). While each policy class differs in how it approaches a given problem, the four classes are united in the approximate nature of their solutions: ADP policies are more flexible than dynamic programming but this flexibility comes at the expense of guaranteed optimality (Powell 2011, 2014; Powell & Meisel 2016). Given the complexity of many ecological management problems, the trading of assured optimality against greater flexibility in model design and the capacity to ask a broader range of questions could be advantageous (Nicol *et al.* 2010; Nicol & Chadès 2011). The four policy classes are, in no specific order, cost function approximations (CFAs), value function approximations (VFAs), policy function approximations (PFAs) and lookahead policies (Powell 2014).

Cost function approximations (CFA) represent the simplest variety of ADP policy, being an extension of a myopic policy. CFAs replace the term describing the downstream value of a decision in equation [1] with a correction term which has been designed to promote behaviours considered to be desirable (Powell 2011, 2014). Under a CFA, actions are identified as follows

$$a_t = \arg\max_a \{\bar{C}(S_t, a_t | \theta)\} \quad [3]$$

where the bar notation indicates that $\bar{C}(S_t, a_t|\theta)$ is a CFA and θ captures any tuneable parameters. A CFA is not approximating the future effects of the decision and may not necessarily even have to make reference to the state of the system. Instead, CFAs are more akin to an incentivising/disincentivising term which modifies the value of the reward/cost incurred by a particular action. As a simple example, a CFA which disincentivises insecticide application would result in a myopic policy which will either accept a higher pest burden prior to spraying or which is more likely to utilise alternative controls. CFAs perform best in scenarios where myopic policies have already been found to deliver acceptable results and, with the inclusion of tuneable parameters can be surprisingly nuanced, producing decisions which improve upon those of the original myopic policy without the need to explicitly attempt to model or approximate future information (Powell 2011, 2014; Powell & Meisel 2016).

Value function approximations (VFA) are the policy class most frequently associated with ADP, they attempt to identify appropriate actions by approximating the value of future states resulting from a decision made now (Powell 2011, 2014; Powell & Meisel 2016). The likely downstream effects of the current action, and possible future actions, are encapsulated in a function which is dependent upon the state which follows from the current decision. Thus, using a VFA, actions are selected according to

$$a_t = \arg\max_{a}\{C(S_t, a_t) + \bar{V}(S_{t+1})\} \quad [4]$$

where the bar in $\bar{V}$ indicates that this is not the true optimal value of the subsequent state as in equation [1] but is an approximation. VFAs are extremely flexible and both parametric and non-parametric methods, in addition to techniques from machine learning, can be used to design $\bar{V}(S_{t+1})$ (Nicol & Chadès 2011; Powell 2011). The design of the VFA is critically important to its performance. A poorly designed VFA, or a VFA using unrepresentative parameter values can impair policy performance and promote ineffectual or harmful actions (Powell 2011; Powell & Meisel 2016). Intuitively then, VFAs are most constructively applied to problems where the downstream values of actions can be easily, and reliably approximated (Powell & Meisel 2016). For example, if there is sufficient data to suggest that the future states of a system adhere to a particular functional form or evolve according to a specific distribution then a VFA could be developed which leverages this information.

Policy function approximations (PFA) are functions which have been designed to map states directly onto actions without solving an embedded optimisation process or attempting to forecast future trends (Powell 2011, 2014; Powell & Meisel 2016). That is, as implied by their title, PFAs are effectively an approximation of the decision making process itself, their chosen form fixing the decision rule for all time points, and are best applied to problems for which we already have an idea as to how decisions might be made and the interactions between states and actions are easily and accurately observed (Powell 2011, 2014; Powell & Meisel 2016). As a simple example, a PFA for a harvesting problem could be specified with two parameters, x and y, and the decision rule: harvest at rate x so long as the population exceeds y. The task at hand would then be to identify the values of x and y which are predicted to return the best average behaviour over the timescale of interest. This "one size fits all" approach complicates the application of PFAs to non-stationary problems(Powell 2014; Powell & Meisel 2016). To return to the harvesting example, if the recruitment rate of population being harvested varied seasonally then distinct values of x and y may need to be estimated for each season. The greater the variability in the system of interest, the more cumbersome PFAs become.

In contrast to CFAs, VFAs and PFAs, all of which utilise functional approximations and essentially attempt to use structural properties of the problem to identify exploitable patterns, lookahead policies are a comparatively blunt instrument. Lookahead policies, also known as model predictive control and the receding horizon procedure, do not utilise functional approximations. They are a brute-force, algorithmic approach which proceeds by the rote application of the same procedure to each decision point for as long as is required. However, this has the desirable property of rendering lookahead policies entirely agnostic to the structure of the problem to be solved (Nicol *et al.* 2010; Powell 2011, 2014; Powell & Meisel 2016). Lookahead policies suggest actions by constructing and explicitly solving a sequence of truncated versions of the full problem we seek to solve. The optimal action to be taken in period t is that which, in conjunction with the decisions that follow it, maximises the benefits (or minimises the costs) accruing between period t and the final time horizon T. To identify this action by systematically projecting forwards in time, all the way to the time horizon, would produce a rapidly growing tree of state-action pairs which becomes intractable for even apparently simple applications which is one of the key motivators behind the development of dynamic programming (Clark & Mangel 2000; Powell 2011; Marescot *et al.* 2013). This limitation obviously remains true for problems beyond the scope of dynamic programming but, while following the optimal path between t and T may present an unassailable problem, it may still be possible to identify satisfactory actions by solving a reduced version of the problem over some limited planning horizon, $H < T$ (Nicol *et al.* 2010; Powell 2011; Powell & Meisel 2016). If a sufficiently long planning horizon is used, of an appropriate length to capture important model behaviours, then good, but not optimal, actions may be suggested.

When using a lookahead policy, the problem we would ideally like to be able to solve (*e.g.* equation [3]) is termed the base model while the truncated versions which are solved in its place are termed lookahead models. This distinction makes it clear that we are not solving the original problem and that the actions returned by a lookahead policy are optimal only with respect to the lookahead model they were derived from: it is highly improbable that they are also optimal with respect to the base model (Nicol *et al.* 2010; Powell 2014; Powell & Meisel 2016). Thus, to identify an action to be taken at time t a lookahead policy will select the action which maximises

$$a_t = \arg\max_a \{C(S_t, a_t) + \sum_{t'=t+1}^{t'=t+H} (C(S_{\tilde{tt'}}, a_{\tilde{tt'}}))\} \quad [5]$$

where tildes are used to distinguish values only relating to the lookahead model from those in the base model. Thus, $S_{\tilde{tt'}}$ is the state experienced during period t' when solving a lookahead model to identify an action for period t. Equation [5] makes clear that lookahead policies identify desirable actions by explicitly simulating their effects up to the planning horizon H to gain a partial understanding of how the consequences of the decision may unfold. In solving a lookahead model over the planning horizon, each permissible action is applied to each state encountered at each time period which, once the best end point has been identified theoretically allows the lookahead policy to return a vector describing the best path through the lookahead model. However, in practice only the first action in this vector, the action to be taken now, is of interest and the remaining actions, those made between $t' = t + 1$ and $t' = t + H$ are discarded. Hence, once a lookahead policy identifies action a_t it implements this action and transitions to the next decision period where it experiences state S_{t+1} and constructs a new lookahead model over the interval $\{t + 1, \dots, t + 1 + H\}$ to identify the action a_{t+1}.

Lookahead policies possess several features of interest to decision problems in ecology. First, while undeniably a brute-force method, their structural agnosticism renders them extremely versatile. Given the manner in which they are solved, lookahead policies are robust to non-stationary problems and may accept state spaces of arbitrary magnitude: the number and complexity of the state variables does not influence the runtime of a lookahead policy (Nicol *et al.* 2010; Powell 2011). The runtime of a lookahead policy is instead determined by the number of alternative actions, the extent of the planning horizon and, in stochastic applications, the number of alternative random outcomes (Nicol *et al.* 2010; Powell 2011). Furthermore, given that each lookahead model explicitly models the state transitions of the system over the duration of the time horizon, lookahead policies can be applied in instances where the downstream values of actions are complex functions which would be difficult to approximate reliably (Powell & Meisel 2016). Lookahead policies are also strongly predisposed to make effective use of forecasts and are robust to changes in those forecasts, as may occur in an adaptive management setting (Boettiger, Mangel & Munch 2015; Chadès *et al.* 2017). Ultimately, the blunt stupidity of lookahead policies, particularly compared to the more elegant VFAs and PFAs, may be their greatest virtue.

Thesis Outline

In what follows, this work explores how dynamic programming and lookahead policies may offer insight into the management of agricultural insect pests, particularly in scenarios where trade-offs exist between short-term and long-term benefits. Chapter 2 explores the effects of different mechanistic realisations of fitness costs to resistance upon the refuge allocation decisions of a dynamic programming algorithm. Chapter 3 then illustrates the application of lookahead policies to a pest-management problem of high dimensionality to explore the extent to which these approximate policies are able to reconcile the timeline of the pest life cycle with the timescale over which distinct control actions, insecticidal sprays and sterile male releases, unfold. Chapter 4 then extends the framework presented in chapter 3 to include male-selecting transgenic releases and the evolution of insecticide resistance to query the extent to which the limited perspective of a lookahead policy may impede its capacity to (i) utilise a management option with delayed benefits and (ii) effectively anticipate and respond to the evolution of resistance. Chapter 5 then integrates the findings of the preceding chapters and suggests how they may be built upon.

Chapter 2: Type of fitness cost influences the rate of evolution of resistance to transgenic Bt crops[1]

Sean C.Hackett[1], Michael B.Bonsall[1,2]

[1] *Mathematical Ecology Research Group, Department of Zoology, South Parks Road, Oxford, OX1 3PS, UK*

[2] *St. Peter's College, New Inn Hall Street, Oxford, OX1 2DL, UK*

Abstract

The evolution of resistance to pesticides by insect pests is a significant challenge for sustainable agriculture. For transgenic crops expressing *Bacillus thuringiensis* (Bt) crystalline (Cry) toxins, resistance evolution may be delayed by the high-dose/refuge strategy in which a non-toxic refuge is planted to promote the survival of susceptible insects. The high-dose/refuge strategy may interact with fitness costs associated with resistance alleles to further delay resistance. However, while a diverse range of fitness costs are reported in the field, they are typically represented as a fixed reduction in survival or viability which is insensitive to ecological conditions such as competition. Furthermore, the potential dynamic consequences of restricting susceptible insects to a refuge which represents only a fraction of the available space have rarely been considered. We present a generalized discrete time model which utilizes dynamic

---

- [1] Published as: Hackett, S.C. & Bonsall, M.B. (2016) Type of fitness cost influences the rate of evolution of resistance to transgenic Bt crops. *Journal of Applied Ecology*, **53**, 1391-1401

programming methods to derive the optimal management decisions for the control of a theoretical insect pest population exposed to Bt crops. We consider three genotypes (susceptible homozygotes, resistant homozygotes and heterozygotes) and implement fitness costs of resistance to Bt toxins as either a decrease in the relative competitive ability of resistant insects or as a penalty on fecundity. Model analysis is repeated and contrasted for two types of density dependence: uniform density dependence which operates equally across the landscape and heterogeneous density dependence where the intensity of competition scales inversely with patch size and is determined separately for the refuge and Bt crop. When the planting of Bt is decided optimally, fitness costs to fecundity allow for the planting of larger areas of Bt crops than equivalent fitness costs that reduce the competitive ability of resistant insects. Heterogeneous competition only influenced model predictions when the proportional area of Bt planted in each season was decided optimally and resistance was not recessive. The high-dose/refuge strategy alone is insufficient to preserve susceptibility to transgenic *Bacillus thuringiensis* (Bt) crops in the long term when constraints upon the evolution of resistance are not insurmountable. Fitness costs may enhance the delaying effect of the refuge, but the extent to which they do so depends upon how the cost is realized biologically. Fitness costs which apply independently of other variables may be more beneficial to resistance management than costs which are only visible to selection under a limited range of ecological conditions.

Introduction

The evolution of resistance to pesticides and other population control measures by insect pests poses a significant challenge to both public health and agriculture.

However, taking action to regulate or remove an insect population may introduce a selection pressure. This continuous selection pressure favours the survival of individuals expressing some level of tolerance to control, eroding our future capacity to manage the pest. As most recently commercialized insecticides are variants of previously isolated or synthesized compounds and developing and launching a new pesticide is estimated to take over a decade (REX Consortium 2013), it is unreasonably optimistic to assume that there will always be an alternative product available when control fails (Mitchell & Onstad 2014). Therefore, the sustained control of insect populations over prolonged periods requires that the strength of selection for resistance be constrained and/or reduced (REX Consortium 2013).

Recognition of the need to regulate selection for resistance has facilitated the development (and implementation) of the high-dose/refuge strategy, particularly in conjunction with transgenic Bt crops (Tabashnik 2004). This approach uses simple population genetics to delay the evolution of resistance to the *Bacillus thuringiensis* (Bt) derived Cry toxins expressed within transgenic plants by providing a non-toxic refuge to promote the survival of susceptible insects which then confer their susceptibility upon the next generation (Tabashnik 2004). The susceptible insects produced by the refuge can significantly delay the introgression of an initially rare resistance allele (Carrière & Tabashnik 2001; Tabashnik 2004; Carrière, Crowder & Tabashnik 2010). The introgression of resistance will be further delayed by the expression of a Bt toxin dose sufficient to reduce the survival of heterozygotes, rendering the resistant phenotype functionally recessive (Gould 1998; Tabashnik 2004). When resistance alleles are rare and resistance is functionally recessive, the combination of a high toxin dosage and a refuge significantly constrains the rate of resistance evolution (Carrière & Tabashnik 2001; Tabashnik 2004; Alphey *et al.* 2008). The high-dose/refuge strategy has been successful in maintaining the susceptibility of agricultural insect pests to transgenic Bt crops with failures generally being linked to either a deviation from prescribed refuge sizes or the expression of an insufficient toxin dose by the transgenic plants (Tabashnik, Brévault & Carrière 2013; García *et al.* 2015).

In spite of the strengths of the high-dose/refuge strategy, less understood threats to the efficacy of this approach are ecological effects such as intraspecific (within species) competition between insects within the refuge (Onstad, Shelton & Flexner 2014). Competition is predicted to accelerate the evolution of resistance when carrying capacities are low or density dependence is overcompensating. These conditions may be facilitated by the spatial subdivision imposed upon the landscape by the high-dose/refuge strategy (Onstad, Shelton & Flexner 2014). Superficially, we would not anticipate insects in a managed field to be able to attain numbers such that their density became self-limiting without also losing the majority of the available yield. However, confining susceptible insects to comparatively small areas may permit negative density-dependent effects at lower densities than would otherwise be predicted. Theoretical work has suggested that competition between susceptible insects may undermine the delaying effect of the refuge by reducing the production of susceptible adults (Sisterson, Antilla & Carrière 2004; Glaum, Ives & Andow 2011).

An additional consideration is the potential synergy between the high-dose/refuge strategy and any fitness costs associated with resistance alleles (Tabashnik 2004; Alphey *et al.* 2008). Assuming a high toxin dose, resistance alleles of small phenotypic effect will be selected against; only alleles conferring sufficient tolerance to overcome the expressed toxin dose will be favoured. Such tolerance is most likely to be conferred by a major mutation that may also have pleiotropic effects which negatively influence other processes pertinent to survival and reproduction(Macnair 1991; Coustau, Chevillon & ffrench-Constant 2000; Gassmann, Onstad & Pittendrigh 2009). Thus, resistant insects will likely be subject to fitness costs which may constrain the introgression of the resistance allele (Carrière & Tabashnik 2001; Tabashnik 2004; Alphey *et al.* 2008). However, the degree to which fitness costs delay resistance and synergize with a high-dose/refuge approach is contingent upon the genetic dominance of the resistance allele (Carrière & Tabashnik 2001; Tabashnik 2004; Carrière, Crowder & Tabashnik 2010). The dominance of resistance determines the extent to which the resistance of an insect with a single copy of the allele (a heterozygote) compares to that of an insect with two copies of the allele (a homozygote). If the resistance of a heterozygote is comparable to that of a homozygote for a given toxin dosage, then resistance is dominant. Dominance spans a spectrum from incomplete dominance (the heterozygote cannot tolerate toxin exposure to the same extent as a homozygote) to complete dominance (the heterozygote is functionally identical to a homozygote). If the heterozygote exhibits no resistance, then resistance is termed recessive while if it is intermediate between each homozygote, resistance is additive. Intuitively, recessive resistance is simpler to manage than dominant resistance.

A range of fitness costs have been reported, from both laboratory and field, for insects expressing resistance to chemical control agents (Gassmann *et al.* 2009). However, fitness costs explicitly associated with resistance to Bt crops remain poorly understood (Jakka, Knight & Jurat-Fuentes 2014). This shortfall in available data can be partly attributed to the success of the high-dose/refuge strategy; if instances of field-evolved resistance are rare, so studies of field-acquired resistance to Bt crops are similarly sparse. Recent efforts to categorize and quantify fitness costs associated with field-evolved resistance to Bt crops have reported variable results (Jakka, Knight & Jurat-Fuentes 2014; García *et al.* 2015; Dangal & Huang 2015; Ingber & Gassmann 2015).

Difficulties in the consistent identification of fitness costs may arise from the simplifying assumption that fitness costs are constant. However, fitness is a complex function of many factors influencing both survival and reproduction. Some of these factors will be ecological and may exaggerate or conceal pleiotropic effects associated with a resistance allele (Gassmann, Onstad & Pittendrigh 2009). That is, fitness costs may exhibit a degree of context dependence and may increase or only become visible to selection in the presence of additional environmental stressors. For example, fitness costs may be exaggerated by factors such as host plant variety (Bird & Akhurst 2007), parasites and pathogens (Raymond *et al.* 2007), temperature (Zhang *et al.* 2015) or biotic interactions (Becker & Liess 2015). When fitness costs are context-sensitive, the resistance allele frequency may not necessarily decline in the absence of selection. Unless the context is appropriate, the resistance allele may behave as if it were cost-free and resistance will evolve more rapidly (Jakka, Knight & Jurat-Fuentes 2014; García *et al.* 2015; Ingber & Gassmann 2015)

Here we develop a generalized discrete time model which utilizes dynamic programming methods to derive optimal management decisions for the control of a theoretical pest population. This approach allows us to represent managers as non-static, goal-orientated entities that adjust their decisions over time based upon the information available to them. In addition, dynamic programming captures the discrete time nature of pest management decisions. We consider a single resistance gene segregating at a diallelic locus and implement fitness costs as either a decrease in the relative competitive ability of resistant insects or as a penalty to fecundity. Costs to fecundity apply irrespective of additional factors while a decrease in competitive ability only penalizes resistant insects when the total population density is sufficiently large. We find that costs to fecundity allow for the selection of larger areas of Bt crops than costs on competitive ability. This implies that the mechanistic effect of fitness costs in the field is a significant consideration in the development of resistance management strategies.

Materials and methods

We consider the management of a hypothetical univoltine insect pest population feeding in a closed landscape with no alternative hosts. The manager seeks to suppress the pest population by selecting the area of the landscape which is to be planted with transgenic plants expressing insecticidal Bt toxins. Thus, the manager's decisions divide the landscape into two distinct connected patches of variable magnitude: the Bt crop and the refuge. The proportion of the landscape allocated to toxic plants is denoted φ and so the proportional area of the refuge is $1-\varphi$. The manager specifies the value of φ each season such that the cumulative pest burden experienced over the time horizon, T, is minimized. It is in the manager's interest to control resistance; therefore, we assume that the manager attempts to conserve the future utility of the transgenic crops by actively seeking to maintain the frequency of the resistance allele below some critical threshold. The model has two distinct components: a dynamic programming model which identifies the best decision available to the manager for a given set of conditions and a population sub-model characterizing the dynamics and genetics of the pest population.

Pest sub-model

The growth of the pest population is described by a density-dependent selection model (Roughgarden 1971) which uses the two parameter density dependence function developed by Maynard Smith and Slatkin (1973). This function was selected for its capacity to describe a range of dynamic behaviours (Bellows 1981). An insect may belong to one of three alternative genotypes: susceptible homozygotes (ss), resistant heterozygotes (sr) or resistant homozygotes (rr). The frequency of the susceptible allele, s, is denoted p and the corresponding frequency of the resistance allele, r, is denoted q. The pest sub-model proceeds as follows:

i. Pests are divided between the refuge and toxic patch (where $\varphi > 0$) in a ratio proportional to patch size.
ii. Pests within the Bt patch are exposed to toxins and mortality is quantified.
iii. Pest mortality via intraspecific competition operates either across the entire space or separately within each patch.
iv. Surviving larvae mature and mate at random to generate the next generation.

Larvae emerge into a landscape which, contingent upon the decision taken by the manager at the beginning of the season, may be a single contiguous refuge or, more probably, is subdivided into a toxic patch and a non-toxic refuge. We assume uniform oviposition so, of N_t larvae hatching during season t, φN_t will develop in the toxic patch and $(1-\varphi)N_t$ will be localized within the refuge. Larvae remain within their natal patch until maturation.

On hatching, larvae commence feeding and the φN_t larvae within the toxic patch are exposed to Bt. The proportion of larvae that survive exposure to Bt is contingent upon their genotype,g, and is denoted S_g where $S_{ss} \leq S_{sr} \leq S_{rr}$. The total population density subsequent to toxin exposure is denoted N'_t which is the sum of the refuge population and the survivors within the Bt patch. Larvae that survive toxin exposure (or were not exposed) undergo intraspecific competition. Two different implementations of competition are considered. In the first, density-dependent mortality acts uniformly, the entire landscape is treated as a single whole and differences in pest abundance between patches are ignored. However, given that there is no larval movement and thus competition will vary locally, this could underestimate the impact of competition when the population within a patch is large relative to its size. The second, heterogeneous, implementation accounts for this by specifying an inverse relationship between competitive mortality and patch size. Density dependence is then evaluated separately for each patch. To accommodate both representations we derive the number of larvae of genotype g in patch i , $N''_{t,g,i}$,surviving competition and progressing into the reproductive phase as:

$$N''_{t,g,i} = \frac{N'_{t,g,i}}{1 + (\frac{\alpha_g}{m_i} N'_{t,i})^{\beta}} \qquad [1]$$

where α_g and β are the parameters of Maynard Smith & Slatkin (1973). The value of β determines the intensity of competition and is assumed to be independent of genotype. Values of $\beta \approx 1$ are suggestive of undercompensating density dependence while $\beta > 1$ shifts the dynamics towards increasingly intense, overcompensating, competition (Bellows 1981). The value of α_g $(0 < \alpha_g)$ determines the per capita sensitivity of insects of genotype g to intraspecific competition and sets the threshold population density beyond which the net population growth rate for genotype g becomes negative. To link competition and patch size we scale the value of α_g by the factor $\frac{1}{m_i}$ where m_i is the relative area of patch i (i.e. φ for the toxic patch and $1 - \varphi$ for the refuge). Thus, the sensitivity to competition of insects within a patch increases as patch size declines. If the field has been planted with only one type of crop or competition is specified as uniform, then the subscript i is unnecessary and $m = 1$ for the sole patch.

Insects carrying resistance alleles may be subject to fitness costs. Fitness costs were implemented as either an increase in the value of α_g (so that $\alpha_{ss} \leq \alpha_{sr} \leq \alpha_{rr}$) or a decrease in fecundity, λ_g. Costs to competition may be conceptualized as a decrease in the efficiency with which resistant insects acquire resources, or an increase in their resource requirement relative to susceptible insects. Costs to fecundity represent a reduction in the efficiency with which resistant insects convert acquired resources into offspring. While the mechanism differs, both costs constrain the maximum size a population of resistant insects can attain for given space and resources. Where costs increase α_g, resistant insects will experience greater levels of mortality at high population densities than susceptible insects. As density dependence is not selectively neutral when competitive fitness costs apply, allele frequencies must be recalculated prior to reproduction:

$$q'_t = \frac{N'_{t,rr} + \frac{1}{2}N'_{t,sr}}{N'_t} \quad [2]$$

$$p'_t = 1 - q'_t. \quad [3]$$

Insects which survive intraspecific competition then mature. Mating is at random with respect to both genotype and space. The number of insects in the next generation is:

$$N_{t+1} = N''_t(p'^2\lambda_{ss} + 2p'q'\lambda_{sr} + q'^2\lambda_{rr}) \quad [4]$$

where N''_t is the total number of insects in the post-competition population. The value of λ_g denotes the average fecundity of an insect of genotype g. When resistance does not influence fecundity $\lambda_{ss} = \lambda_{sr} = \lambda_{rr}$ and the intergenerational change in the number of insects simplifies to $N_{t+1} = \lambda N_t''$. The resistance allele frequency in the next generation is:

$$q_{t+1} = \frac{q'^2\lambda_{rr} + p'q'\lambda_{sr}}{p'^2\lambda_{ss} + 2p'q'\lambda_{sr} + q'^2\lambda_{rr}}. \quad [5]$$

Fitness cost magnitude is specified using the parameter c, where $0 \leq c \leq 1$,to represent the proportional decrease or increase in the value of the relevant parameter. That is, fitness costs are implemented for the resistant homozygote as:

$$\alpha_{rr} = (1 + c)\alpha_{ss} \quad [6]$$

$$\lambda_{rr} = (1 - c)\lambda_{ss} \quad [7].$$

Only a single fitness cost is considered in a given simulation (that is, penalties are applied to either fecundity or competition but never both). The dominance of resistance and its correlated fitness costs are determined by the value of the parameter h where $0 \leq h \leq 1$. A value of $h = 0$ denotes fully recessive resistance (a single copy of the r allele has no effect). As h tends to 1 resistance becomes increasingly dominant with $h = 1$ indicating fully dominant resistance (a single copy of the r allele is equivalent to two copies). The values of S_{sr}, α_{sr} and λ_{sr} are calculated as:

$$(1 - h)x_{ss} + hx_{rr} \quad [8]$$

where x represents the parameter of interest.

Values for the heterozygote then depend on the dominance of the resistance allele. This assumes that the heritability of resistance and any correlated fitness costs are identical but this need not be true (Gould 1998). To relax this assumption, unique heritability coefficients are assigned to S_{rr} and either α_{rr} orλ_{rr} depending on the fitness cost of interest. In this instance, let h_{res} denote the heritability of S_{rr} and h_{cost} represent the heritability of either α_{rr} or λ_{rr}. For reference, the use of h without a subscript refers to instances where the heritability of resistance and fitness costs are identical ($h_{res} = h_{cost}$).

Decision making

The decision process is captured using dynamic programming. We consider a manager seeking to suppress an insect pest and in doing so implicitly minimize yield loss over the considered time horizon. We acknowledge that this is a simplified interpretation of grower behaviour that excludes factors such as discounting or the likelihood that refuge plants will incur high levels of feeding damage relative to Bt plants. Such factors would bias decisions towards short-term benefits and favour the more rapid depletion of susceptibility to Bt.

From the manager's perspective, resistance management is a means to an end which prolongs the efficacy of available control mechanisms. However, in the absence of control mechanisms which can reduce pest density, N, independently of the resistance allele frequency, q, the minimization of both pest density and resistance are, to an extent, mutually exclusive outcomes. Planting mostly transgenic crops strongly suppresses population density when resistance is rare but drives selection for resistance, decreasing the crops' control efficacy in subsequent seasons. Conversely, minimal plantings of Bt crops retain a greater number of susceptible alleles but enable greater pest population densities. A simple method for capturing this trade-off in the decision model is to only consider decisions for which the r allele frequency is held below a specified critical threshold value, q_c.

The area of the landscape to be planted with transgenic crops, φ, is selected from the control set $\Phi = \{\varphi_i\}$; in our investigations we permitted fractions of the landscape ranging from 0 to 1 in increments of 0.1. This is chosen for each decision period, t, within the specified time horizon, T, such that the cumulative pest burden experienced between period t and T is minimized subject to the constraint that the resistance allele does not attain a frequency greater than q_c. This objective is captured in a dynamic programming equation which calculates the value, V, (measured as the post-reproductive pest density at the end of the period) of selecting control φ_i during time period t for a manager who began the period with a total pest population density of N_t as:

$$V_i(N_t, t) = (N_{t+1} + F(N_{t+1}, t+1)) | q_{t+1} < q_c. \quad [9]$$

The cumulative pest burden endured by a manager who selects control φ_i in period t and behaves optimally from then onwards is represented by $F(N_{t+1}, t+1)$. The current optimal decision is that which produces the minimum value of V:

$$F(N_t, t) = \min_{\varphi_i} V. \quad [10]$$

For simplicity we assume that the manager accrues no additional benefit beyond the terminal time period, T and thus is not concerned with management of the pest beyond that. This allows us to state the terminal condition:

$$F(N_T, T) = 0 \quad [11]$$

for all possible values of N (Clark & Mangel 2000). The optimal solution to this problem is derived numerically by a dynamic programming algorithm with two state variables, N and q. The state variables are presented within the algorithm as discretized vectors, $\boldsymbol{N}$ and $\boldsymbol{q}$, which have n and l divisions between their least and greatest values respectively.

The algorithm derives values for $F(N_t, t)$ from $F(N_{t+1}, t+1)$ for each combination N and q via backwards iteration. Where the state values produced by a control decision are not found within their corresponding state vector linear interpolation is used to identify the best decision (Clark & Mangel 2000)

Model exploration

(i) Fixed landscape partitioning

All simulations were carried out in R version 3.13. We begin by simulating the dynamics of the pest sub-model independently of the decision model to provide baseline estimates for the time required for resistance to evolve when the landscape partitioning is static. For all simulations we assume the per capita sensitivity of susceptible homozygotes to competition to be $\alpha_{ss} = 1 \times 10^{-6}$ with average fecundity $\lambda_{ss} = 2$. No susceptible homozygotes survive exposure to Bt toxins ($S_{ss} = 0$) and resistance is complete ($S_{rr} = 1$). Where the value of the parameter β was varied it was restricted to values between 1 and 3. The evolution of resistance was simulated over 100 generations for two fixed plantings of Bt, $\varphi = (0.95, 0.8)$. The influence of dominance and fitness costs upon resistance evolution was considered for both types of density dependence. Pest populations were initiated with $N_0 = 1000$ insects and an initial resistance allele frequency of $q_0 = 0.01$.

(ii) Dynamic partitioning

In deriving the numerical solution for the decision model the algorithm compiles a decision array containing the best permissible decision for each combination of state values within each time period. We used state vectors of length $n = 20$ and $l = 11$ which assumed values in the ranges $1 \leq n \leq 1\,000\,000$ and $0 \leq l \leq 1$ respectively.

These predictions were evaluated via forwards iteration (Clark & Mangel 2000). Beginning with an initial population density of $N_0 = 1000$ insects and an initial r allele frequency of $q_0 = 0.01$, the intergenerational change in both N and q was simulated from $t = 0$ to $t = T$. Given that the resistance allele exceeded the critical threshold of $q_c = 0.5$ within 20–30 generations for fixed patch sizes (e.g. Figures 1, S1, S2 & S3 in Supplementary figures) the time horizon for the decision model was set at $T = 60$ generations. At the beginning of each time period the element of the decision array which corresponds to the optimal decision was identified (using linear interpolation as appropriate) and implemented. Simulations were run for both uniform and heterogeneous density dependence over 20 non-zero values of each fitness cost and average Bt usage was recorded for each cost value. The proportional decrease in the fecundity of resistant homozygotes was lowered from $\lambda_{rr} = \lambda_{ss}$ (no reduction in fecundity) to $\lambda_{rr} = 0.7\lambda_{ss}$ (a 30% reduction) in increments of 0.015. The proportional increase in the sensitivity of the resistant homozygote to competition was increased from $\alpha_{rr} = \alpha_{ss}$ to $\alpha_{rr} = 2\alpha_{ss}$ (resistant homozygotes are twice as sensitive to competition as susceptible homozygotes) in increments of 0.05. Results of these simulations are reported for three dominance scenarios: recessive resistance ($h = 0$), additive resistance ($h = 0.5$) and additive resistance with a dominant fitness cost ($h_{res} = 0.5, h_{cost} = 1$). Time series of model behaviour and state dynamics were also generated for fixed fitness cost intensities and three levels of dominance: $h = 0.05$ (weakly dominant), $h = 0.5$ (additive) and $h = 1$ (dominant).

Results

Fixed landscape partitioning

For a fixed partitioning of space between the refuge and the Bt crop, the high-dose/refuge strategy delays resistance evolution and suppresses the pest population when resistance is fully recessive with larger refuges providing greater delays (Figures 1 & S1). Non-recessive resistance ($h_{res} > 0$) reduces the efficacy of refuges and accelerates resistance evolution. The type of density dependence (uniform or heterogeneous) influenced neither resistance evolution nor the growth of the pest population (Figure S1). Increasing the intensity of competition ($\beta > 1$) has a negligible influence upon population growth and resistance evolution. Neither fecundity nor competition costs strongly synergized with the high-dose/refuge strategy for the simulated refuge sizes (Figure 1, Figure S2); high levels of toxin mortality alleviate competition and mask fitness costs even when the initial population density is large (Figure S3). However, dominant fecundity costs strongly delayed resistance when the resistant phenotype was recessive (Figure 1, dashed lines). Additionally, when both resistance and fecundity costs were recessive the growth of the population was significantly slowed for larger refuge sizes (e.g. Figure 1, $\varphi = 0.8$, top row, Figure S4).

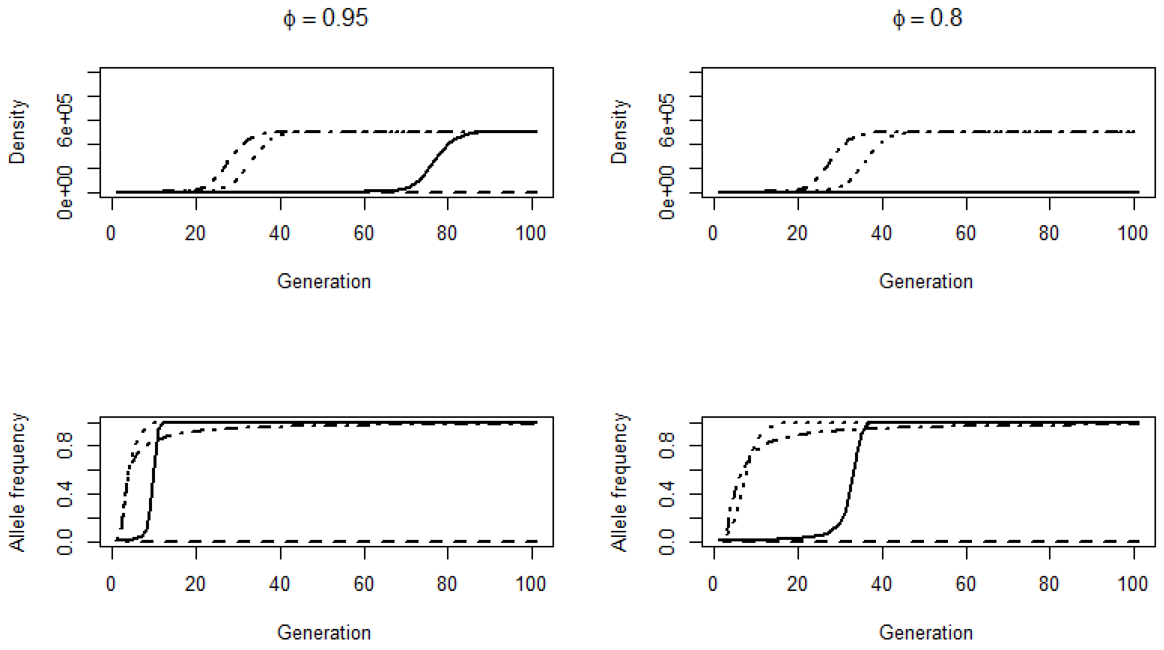

Figure 1. Population size (top row) and resistance allele frequency (bottom) over 100 generations for landscapes planted with 95% (left column) and 80% (right) Bt respectively when resistance is associated with a 25% reduction in fecundity ($\lambda_{rr} = 0.75\lambda_{ss}$). Density dependence is undercompensating and uniform ($\beta = 1$). Solid lines show recessive resistance with a recessive fitness cost ($h_{res} = 0, h_{cost} = 0$). Dashed lines depict recessive resistance with a dominant fitness cost ($h_{res} = 0, h_{cost} = 1$). Dotted lines refer to additive resistance (co-dominant) with a dominant fitness cost ($h_{res} = 0.5, h_{cost} = 1$) and dash-dot lines illustrate dominant resistance with a dominant fitness cost ($h_{res} = 1, h_{cost} = 1$). Populations were founded with $N_0 = 1000$ insects and initial resistance allele frequency $q_0 = 0.01$.

Fitness costs and average Bt usage

Fully recessive resistance may be controlled for the full extent of the considered time horizon by the continuous planting of Bt crops at a fixed level (Figure 2, row 1). This result is insensitive to the type or severity of fitness cost, the type of density dependence and the intensity of competition (not shown). The level of Bt used is sensitive to the fecundity of the susceptible homozygotes and heterozygotes (not shown). The more rapidly the susceptible homozygote (which comprises the bulk of the founding population) reproduces relative to heterozygotes and resistant homozygotes, the greater the average level of Bt which may be planted while still holding the resistance allele beneath the critical threshold of $q_c = 0.5$.

When resistance is non-recessive and carries fitness costs the mean level of Bt selected by the decision model is sensitive to the magnitude of the cost (Figure 2). Greater fitness costs promote larger optimal areas of Bt with the largest plantings being observed when fitness costs reduce fecundity (Figure 2). High levels of toxin mortality when Bt patches are large accelerate the introgression of resistance and alleviate competition. Thus, high susceptible mortality rates in large areas of Bt reduce the impact of competition costs as the insect population must be large for these costs to impede resistance evolution. In contrast, a penalty to fecundity delays the introgression of resistance even when Bt crops are prevalent by constraining the number of offspring a resistant insect may produce. However, even large penalties to fecundity are insufficient to permit a return to the average levels of Bt crops observed for recessive resistance.

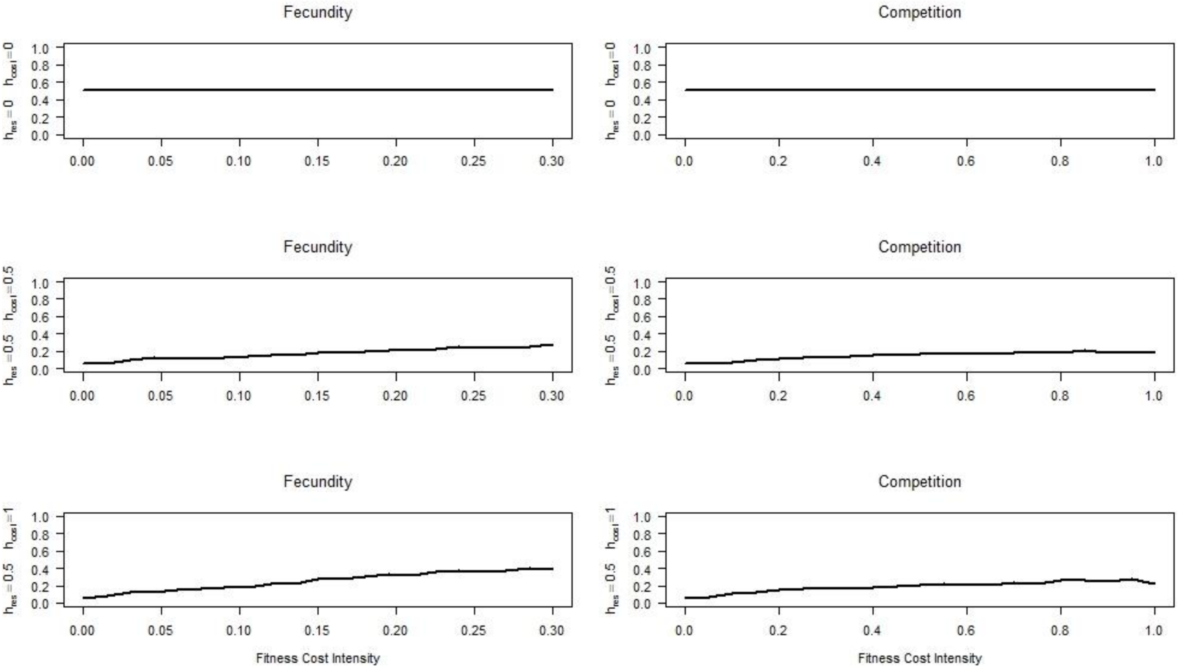

Figure 2. The mean proportion of the landscape allocated to Bt crop over $T = 60$ generations plotted against twenty non-zero levels of fitness costs to fecundity (λ_g, left column) or competitive ability (α_g, right) for a hypothetical pest population. Density dependence is undercompensating ($\beta = 1$) and uniform. Plots show (top to bottom): fully recessive resistance ($h_{res} = 0, h_{cost} = 0$), additive resistance ($h_{res} = 0.5, h_{cost} = 0.5$) and additive resistance with a dominant fitness cost ($h_{res} = 0.5, h_{cost} = 1$). The values on the x-axis describe either the proportional decrease in the value of λ_{rr} relative to λ_{ss} (left column) or the proportional increase in the value of α_{rr} relative to α_{ss} (right column). $N_0 = 1000$ and $q_0 = 0.01$.

Density dependence, fitness costs and time series

When resistance is costless, with an initial population of $N_0 = 1000$ insects and an initial resistance allele frequency of $q_0 = 0.01$, fully recessive resistance ($h = 0$) is always treated with a fixed level of Bt as predicted by Figure 2. Weakly dominant, cost-free resistance ($h = 0.05$) may still be controlled but requires that Bt usage vary with time resulting in smaller average Bt areas (Figures 3 & S5). Further increases in the heritability of resistance see Bt usage decline to zero within 20 generations. Thus, average Bt usage against non-recessive cost-free resistance is low (Figure 2).

The existence of fitness costs offers the decision model additional flexibility. Time series for the resistance allele frequency, genotype abundance and Bt usage over $T = 60$ generations for a pest population with uniform density dependence and a fitness cost that reduces the competitive ability of resistant insects, α_{rr}, by 25% relative to susceptible insects ($\alpha_{rr} = 1.25\alpha_{ss}$) are shown in Figure 3. Inheritance of resistance and fitness cost are assumed to be identical and three levels of dominance are considered; weakly dominant ($h = 0.05$ $i.e.$ near recessive), additive ($h = 0.5$) and dominant ($h = 1$). Decision trajectories for fully recessive resistance are unresponsive to fitness costs (not shown). As with cost-free resistance, the pattern of Bt usage is principally determined by the dominance of resistance. The response to weakly dominant resistance is equivalent to that for cost-free resistance (Figure S5). When resistance is additive or dominant, Bt usage initially declines, as for cost-free resistance. However, the competitive penalty renders the continued planting of Bt optimal either at a fixed low level (Figure 3, middle row, $h = 0.5$) or at higher levels in acute periods interspersed with multiple consecutive generations in which Bt plants are absent (Figure 3, bottom row, $h = 1$).

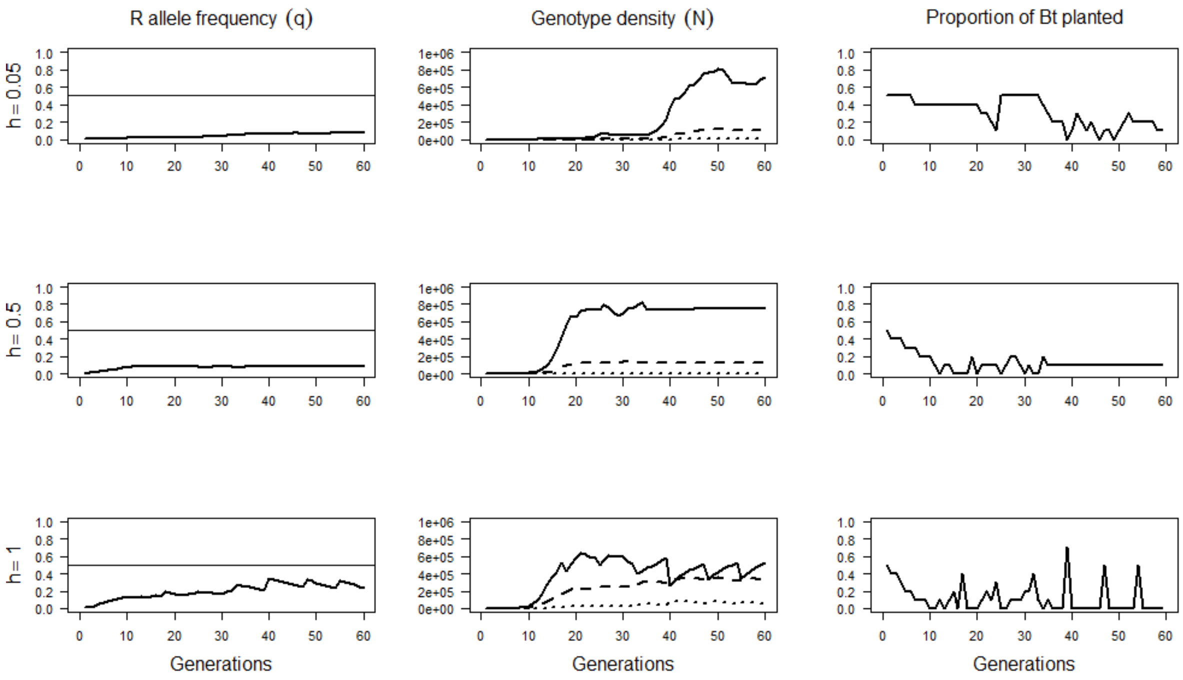

Figure 3. Resistance allele frequency (left-most column), genotype abundance (central column) and proportional allocation of space to Bt crops (rightmost column) over $T = 60$ generations when resistance is associated with a 25% reduction in competitive ability ($\alpha_{rr} = 1.25\alpha_{ss}$). Density dependence is undercompensating ($\beta = 1$) and uniform. The dominance of resistance was increased from $h = 0.05$ (top row), to $h = 0.5$ (centre) and $h = 1$ (bottom). The solid horizontal lines (left) are the critical allele frequency, $q_c = 0.5$. On plots of genotype specific density, solid lines represent susceptible homozygotes (ss), dashed lines represent heterozygotes (sr) and dotted lines represent resistant homozygotes (rr). $N_0 = 1000$ and $q_0 = 0.01$.

Fitness costs on fecundity have a more pronounced effect upon the decision trajectories than competition costs of equivalent magnitude (Figure 4). The impact of fecundity costs is sufficiently great that the strategy adopted to manage weakly dominant resistance ($h = 0.05$) is almost equivalent to that observed for fully recessive resistance (*cf.* Figure 4, top row and Figure 2, top row). Thus, fecundity costs influence the pest population even when expressed weakly. For additive($h = 0.5$) and dominant ($h = 1$) resistance, moderate to large areas of Bt are planted more frequently than when fitness costs impaired competitive ability. However, Bt-free recovery periods are still necessary to maintain the resistance allele frequency beneath the critical threshold (q_c) and are most frequent for dominant resistance (Figure 4, bottom row).

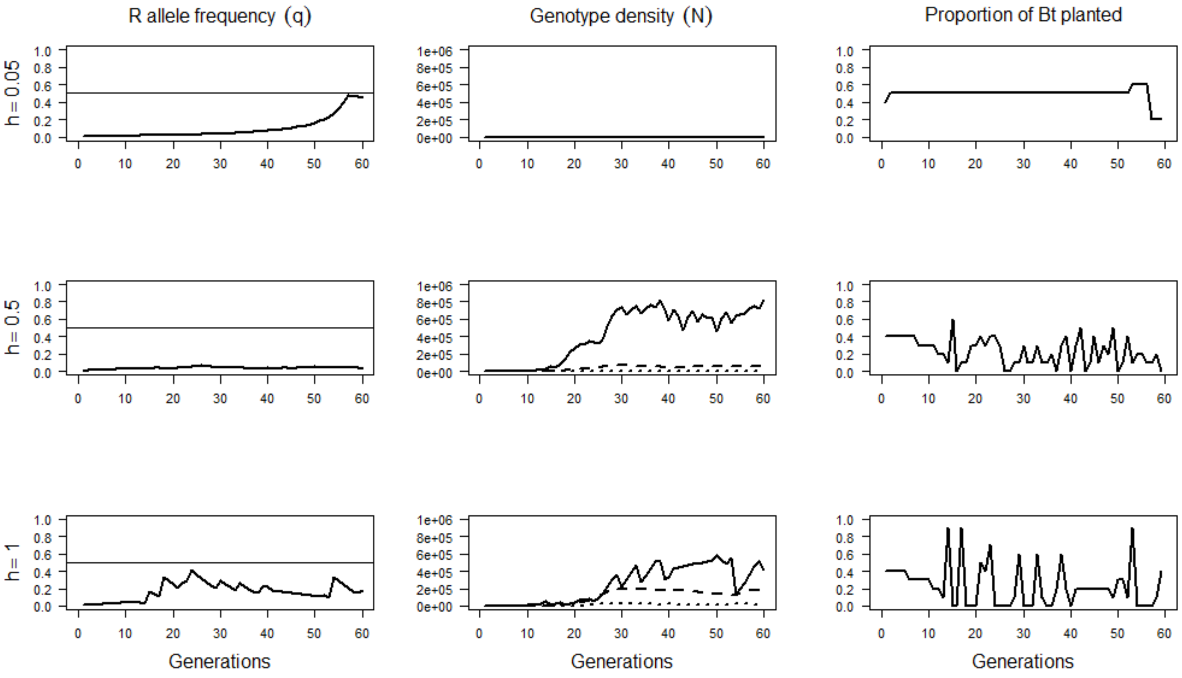

Figure 4. Resistance allele frequency (left-most column), genotype abundance (central column) and proportional allocation of space to Bt crops (right-most column) over $T = 60$ generations when resistance is associated with a 25% reduction in fecundity ($\lambda_{rr} = 0.75\lambda_{ss}$). Density dependence is undercompensating ($\beta = 1$) and uniform. The dominance of resistance was increased from $h = 0.05$ (top row), to $h = 0.5$ (centre) and $h = 1$ (bottom). The solid horizontal lines (left) are the critical allele frequency, $q_c = 0.5$. On plots of genotype-specific density, solid lines represent susceptible homozygotes

(ss), dashed lines represent heterozygotes (sr) and dotted lines represent resistant homozygotes (rr). $N_0 = 1000$ and $q_0 = 0.01$.

While heterogeneous density dependence does not influence average Bt usage (not shown), it does influence the decision trajectories (illustrated in Figure 5, with parameters as in Figure 3). Heterogeneous density-dependent mortality frequently replaces periods of static Bt levels with oscillations between larger and smaller areas of Bt crops. Similar effects are also observed for cost-free resistance (Figure S6) and fecundity costs (Figure S7). In spite of these fluctuations, the general trends observed for heterogeneous density dependence remain similar to those for uniform density dependence.

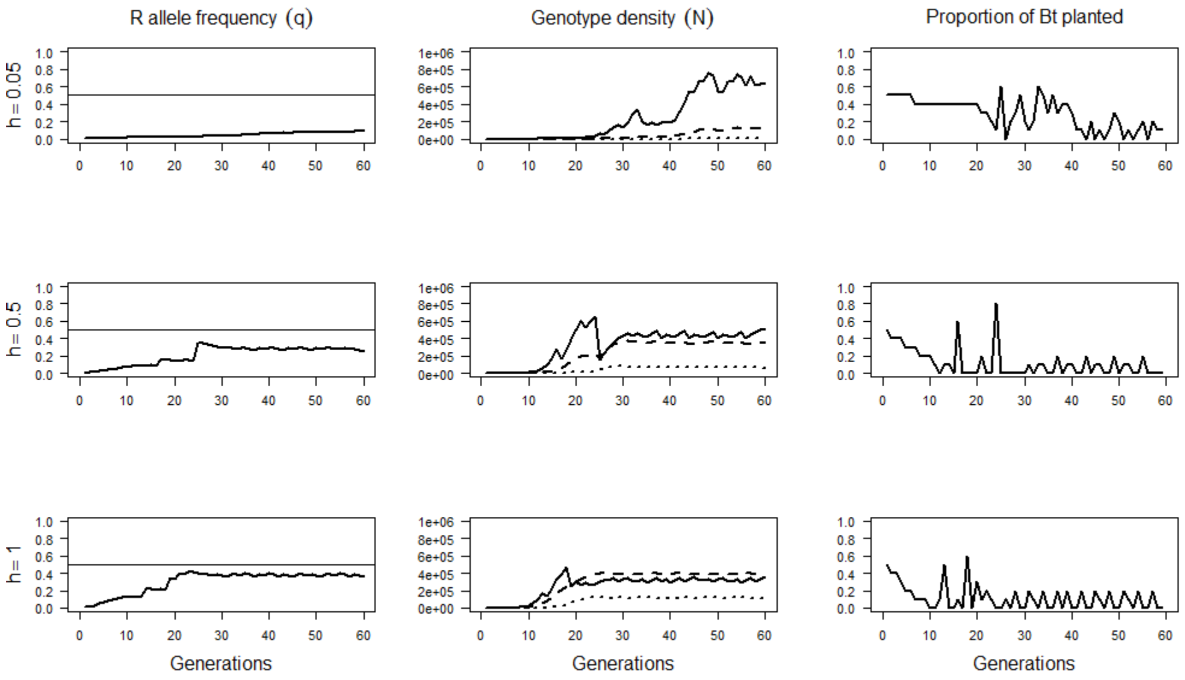

Figure 5. Resistance allele frequency (left-most column), genotype abundance (central column) and proportional allocation of space to Bt crops (right-most column) over $T = 60$ generations when resistance is associated with a 25% reduction in competitive

ability ($\alpha_{rr} = 1.25\alpha_{ss}$). Density dependence is undercompensating ($\beta = 1$) and heterogeneous. The dominance of resistance was increased from $h = 0.05$ (top row), to $h = 0.5$ (centre) and $h = 1$ (bottom). The solid horizontal lines (left) are the critical allele frequency, $q_c = 0.5$. On plots of genotype-specific density, solid lines represent susceptible homozygotes (ss), dashed lines represent heterozygotes (sr) and dotted lines represent resistant homozygotes (rr). $N_0 = 1000$ and $q_0 = 0.01$.

Thus, the most significant constraint upon model behaviour is the dominance of the resistant phenotype, with greater dominance decreasing Bt usage. For a given level of dominance of resistance, the type and dominance of any fitness costs may relax these constraints, allowing for larger areas of Bt crop than would otherwise be optimal. Nested within this, the type of density dependence may alter the transient dynamics of the decision set. These findings are qualitatively summarized in Table 1.

Table 1. Summary of the general effect of the genetic dominance of resistance, fitness costs, the type of fitness cost and heterogeneous density dependence on the planting of Bt crops by the decision model. Downward pointing arrows indicate a reduction in the planting of Bt. Upward pointing arrows indicate an increase in the planting of Bt. Arrow thickness is indicative of the strength of the effect with thicker arrows representing stronger effects. Factors may interact e.g. dominant fecundity costs provide greater increases in Bt usage than dominant competition costs. Heterogeneity in competitive mortality between the refuge and Bt crop does not change average model behaviour but promotes oscillations in Bt usage which is indicated by the wave

Factor	Influence on planting of Bt	Explanation
Dominance of resistance	↓	High heterozygote survival when resistance is dominant requires large refuges.
Fitness costs on fecundity (λ)	↑	Fecundity costs delay resistance independently of other factors. They have a positive effect on Bt planting.
Fitness costs on competitive ability (α)	↑	Competition costs delay resistance when the population is large. They increase Bt planting less effectively than fecundity costs.
Dominance of fitness cost	↑	Increasing the dominance of either fitness cost promotes planting of Bt. The extent of the increase is determined by the type and dominance of the cost.
Heterogeneous density dependence	∿	Heterogeneous density dependence promotes oscillating levels of Bt crops to correct for additional competitive mortality of susceptible insects.

Discussion

We used a decision model to explore the implications of different types of both fitness costs and density dependence on the management of resistance in a theoretical insect pest population feeding on an insecticidal Bt crop. We found that the heritability of resistance and any associated fitness costs had the greatest effect upon model decisions (e.g. Figures 3 & 4), while the type of density dependence only became relevant under particular conditions (non-recessive resistance with non-recessive costs, Figure 5). These predictions indicate that differences between susceptible and resistant insects may not necessarily translate into management benefits. Thus, long-term resistance management using the high-dose/refuge strategy will benefit from additional control measures which are less sensitive to the heritability and pleiotropic effects of resistance genes.

The dominance of the resistant phenotype was the most significant determinant of model behaviour, significantly constraining the usage of Bt crops. This is congruent with the long established sensitivity of the high-dose/refuge strategy to the dominance of resistance genes (Carrière & Tabashnik 2001; Tabashnik 2004; Carrière, Crowder & Tabashnik 2010). The phenotype of heterozygotes is crucial to the spread of rare alleles; high heterozygote survival rates require large refuges to maintain susceptibility. Notably, planting some Bt remained optimal even when cost-free resistance was additive ($h = 0.5$) or dominant ($h = 1$) but this could only be justified for a few generations before the optimal decision was to plant the entire landscape as refuge (Figure S5), indicating the failure of Bt crops. In contrast, fully recessive resistance ($h = 0$) was managed by planting Bt patches of fixed magnitude (Figure 2, top row). These fixed values of Bt reflected the trade-off between long- and short-term killing required for sustainable resistance management and were sensitive to the fecundity of the susceptible homozygote, the principal driver of population growth when the resistance allele is rare. In effect, the model behaves as if susceptibility is a renewable resource (Mitchell & Onstad 2014) and allocates space between the refuge and Bt crop in a manner that regulates the number of susceptible insects while maintaining the resistance allele beneath the critical threshold. The more rapidly susceptible insects are replaced, the more that can be exposed to Bt in a given generation without prompting rapid selection for resistance. Results for dominant resistance may also be understood in this context. When resistance is dominant the rate at which susceptibility is depleted is accelerated, narrowing the range of optimal Bt values, which indicates that refuges alone are insufficient to control dominant resistance.

Neither type of fitness cost significantly delays the evolution of resistance under constant selection (Figures 1 & S2) but both types of cost relaxed constraints upon the decision model. However, for a given level of dominance of resistance, costs to fecundity enabled greater average Bt usage than competition costs (e.g. Figures 2, 3 & 4). Costs to fecundity most closely resemble the "few genes of large effect" classically associated with resistance evolution (Macnair 1991; Coustau, Chevillon & ffrench-Constant 2000; Gassmann, Onstad & Pittendrigh 2009). Fecundity costs apply irrespective of other factors such as competition; for any given set of conditions, resistant insects always produce fewer progeny than susceptible insects. Therefore the resistance allele is slower to spread, more susceptible insects may be killed within each generation, and the fitness cost synergizes with the high-dose/refuge strategy (Carrière & Tabashnik 2001; Tabashnik 2004; Alphey *et al.* 2008).

In contrast, the expression of costs to competition is contextual. Competition costs only impair the spread of the resistance allele when the population exceeds a threshold density beyond which resistant insects fail to compete and decline. Such a cost is more reflective of the weaker fitness costs reported in field-evolved resistance to Bt crops (Jakka, Knight & Jurat-Fuentes 2014; Dangal & Huang 2015; Ingber & Gassmann 2015). Competition costs were insufficient to delay the evolution of resistance under constant selection (Figure S2) and provided only marginal benefits when the model was free to control the Bt exposure of the pest (Figures 3 & 5). When resistant insects were less competitive than susceptible insects, the decision model consistently allowed the total population to grow towards this threshold density before adopting a decision set that regulated the total population around this value (Figures 3 & 5). This maintains the population of susceptible insects within a range for which the further growth of the resistant population is blocked, allowing for a degree of killing without fixing the resistance allele. However, to generate sufficient numbers of susceptible insects that the reduced competitive ability of resistant insects becomes pertinent, extremely large refuges must be planted for much of the time horizon. While this requirement for large refuges runs contrary to the desires of growers, who favour smaller refuges, these decision sets highlight that fitness differences between resistant and susceptible insects need not translate into management benefits. For fixed refuges competition costs did not delay resistance due to the high rate of susceptible mortality; competition only becomes a limiting factor once the population is comprised of primarily resistant insects and thus the cost provides no hindrance. Similarly, fecundity costs only become limiting under fixed selection when refuges are large enough to enable the survival of a sufficient number of susceptible insects that they are able to outpace the growth of the resistant genotype.

Heterogeneous density dependence neither accelerated nor impeded the evolution of resistance when refuges were fixed (Figure S1). High rates of toxin mortality mitigate any effect of competition upon the dynamics and genetics of the population. Under the dynamic programming model, heterogeneous density dependence influenced transient features of the decision set but average model behaviour was unaffected (Figures 5, S6 & S7). Thus spatial variation in the strength of density dependence as implemented here was not a significant factor in the evolution or management of resistance. However, if the model were to be extended to include additional behavioural and ecological idiosyncrasies such as larval dispersal or oviposition biases, then local variation in the intensity of competition could become significant. For example, female fall armyworm *Spodoptera frugiperda* preferentially oviposit on undamaged transgenic Bt crops in response to larval feeding on refuge plants which is predicted to accelerate the evolution of resistance (Téllez-Rodríguez *et al.* 2014). Therefore, it remains significant that additional thought be given to the role of intraspecific (Okuyama & Hsu 2013) and interspecific (Becker & Liess 2015) interactions in shaping the evolution of resistance in the field.

Although the decision sets proposed by this model were generated for a simple agricultural system, a key implication of our results is that additional insect control which operates independently of Bt resistance genes, such as sterile insect releases (Harvey-Samuel *et al.* 2015; Tabashnik *et al.* 2010) may be beneficial in prolonging susceptibility, suppressing pests and protecting yields (Alphey, Bonsall, & Alphey 2009), potentially making smaller refuges optimal. Furthermore, we have assumed transgenic crops express a single toxin but pyramided crops expressing multiple toxins are now available (Tabashnik, Brévault, & Carrière 2013) and may present an additional barrier to resistance by necessitating the occurrence of two distinct mutations. While the risk of cross resistance by major agricultural pests is non-trivial, the possibility remains that, as supported by our results, the suppression of pest populations to preserve susceptibility to Bt crops may require larger refuges (Carrière, Fabrick, & Tabashnik 2016).

The pleiotropic effects of resistance alleles may have a crucial role to play in the evolution of resistance to insecticidal toxins such as those expressed by Bt crops, contributing to both the rate of spread and management of resistance. If the fitness costs associated with resistance alleles do not consistently affect the survival or reproduction of resistant insects, then the effects of these genes can be masked by selection. Resistance may then evolve more rapidly than originally projected and prescribed refuges will not provide adequate defence, necessitating either a revision of refuge sizes or the incorporation of additional tactics to maintain effective suppression. Long-term resistance management will benefit from an improved understanding of the interplay between the realization of fitness costs and the ecology of pests in the field.

Data accessibility

The R code (version 3.13) used in this study has been deposited in the Dryad data repository: http://dx.doi.org/10.5061/dryad.4s405 (Hackett and Bonsall 2016)

Supplementary figures

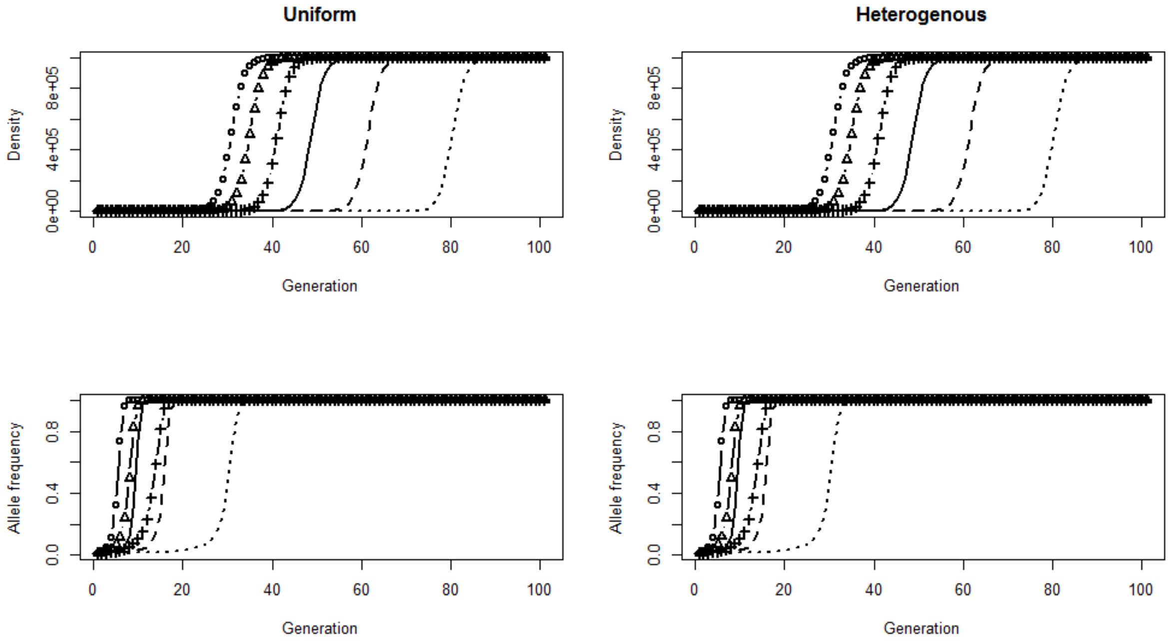

Figure S1: Time series of the change in pest density (top row) and resistance allele frequency (bottom row) over 100 generations of selection. In the left column, density dependent mortality acts uniformly across both the toxic and refuge patches while in the right column density dependent mortality is patch specific and scales inversely with patch size. Values are plotted for two levels of dominance, fully recessive resistance, $h = 0$ (lines without points), and partially recessive resistance, $h = 0.05$ (lines with points). Results are plotted for three Bt cropping regimes. Time series for a 5% refuge ($\varphi = 0.95$) are depicted as solid lines ($h = 0$) and lines with circular points ($h = 0.05$). For a 10% refuge ($\varphi = 0.9$) results are plotted as dashed lines ($h = 0$) and lines with triangular points ($h = 0.05$). For a 20% refuge results are plotted with dotted lines ($h = 0$) and crosses ($h = 0.05$). In all instances the population is founded by $N_0 = 1000$ insects with an initial resistance allele frequency of $q_0 = 0.01$.

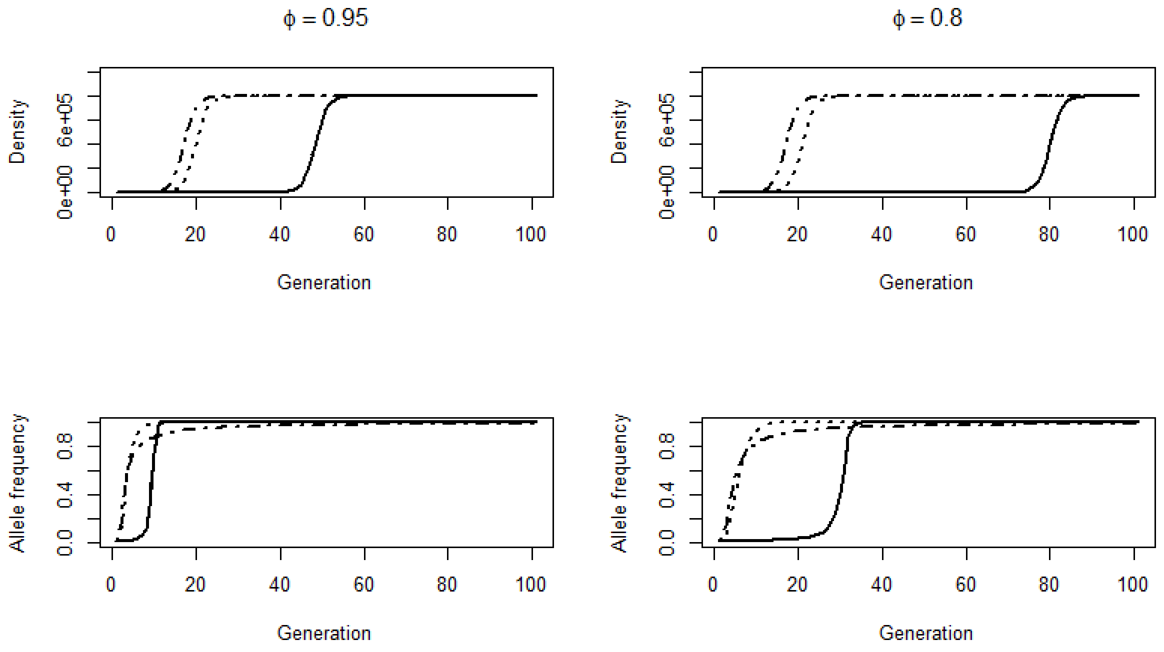

Figure S2: Population size (top row) and resistance allele frequency (bottom row) over 100 generations for landscapes planted with 95% (left column) and 80% (right) Bt respectively when resistance is associated with a 25% reduction in competitive ability ($\alpha_{rr} = 1.25\alpha_{ss}$). Density dependence is uniform and undercompensating ($\beta = 1$). Solid black lines refer to recessive resistance with a recessive fitness cost ($h_{res} = 0, h_{cost} = 0$). Dashed lines depict recessive resistance with a dominant fitness cost ($h_{res} = 0, h_{cost} = 1$). Dotted lines refer to additive resistance with a dominant fitness cost ($h_{res} = 0.5, h_{cost} = 1$) and lines with both dots and dashes illustrate dominant resistance with a dominant fitness cost ($h_{res} = 1, h_{cost} = 1$). Populations were founded with $N_0 = 1000$ insects and an initial resistance allele frequency of $q_0 = 0.01$.

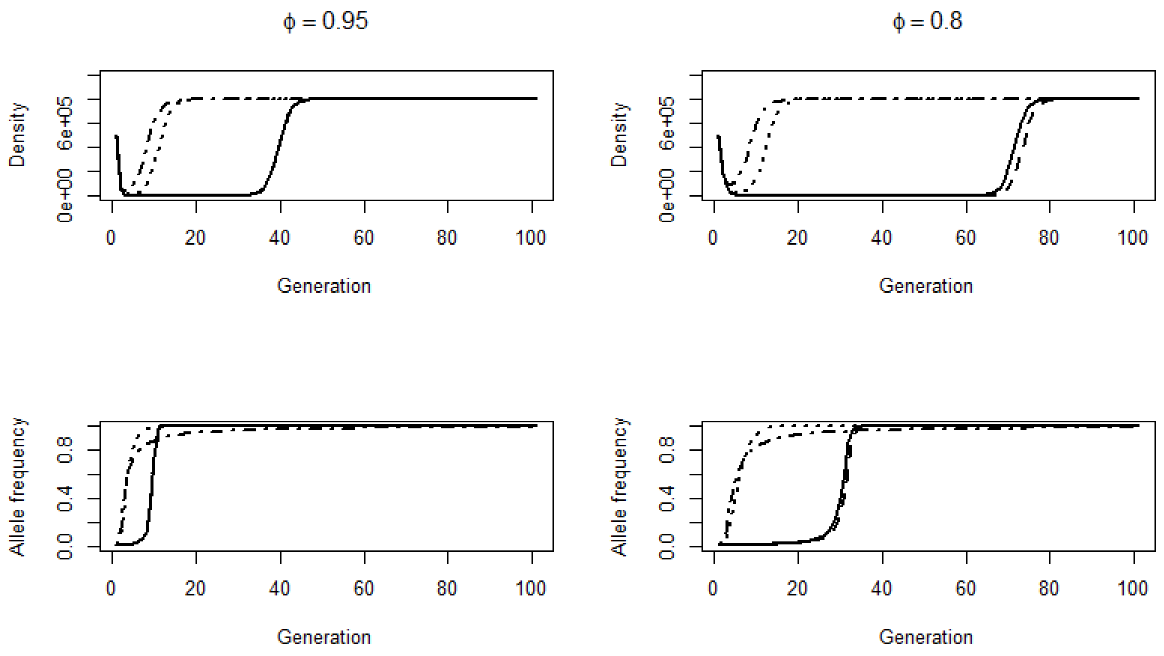

Figure S3: Population size (top row) and resistance allele frequency (bottom row) over 100 generations for landscapes planted with 95% (left column) and 80% (right column) Bt respectively when resistance is associated with a 25% reduction in competitive ability ($\alpha_{rr} = 1.25\alpha_{ss}$). Density dependence is uniform and undercompensating ($\beta = 1$). Solid black lines refer to recessive resistance with a recessive fitness cost ($h_{res} = 0, h_{cost} = 0$). Dashed lines depict recessive resistance with a dominant fitness cost ($h_{res} = 0, h_{cost} = 1$). Dotted lines refer to additive resistance with a dominant fitness cost ($h_{res} = 0.5, h_{cost} = 1$) and lines with both dots and dashes illustrate dominant resistance with a dominant fitness cost ($h_{res} = 1, h_{cost} = 1$). Populations were founded with $N_0 = 500{,}000$ insects and an initial resistance allele frequency of $q_0 = 0.01$.

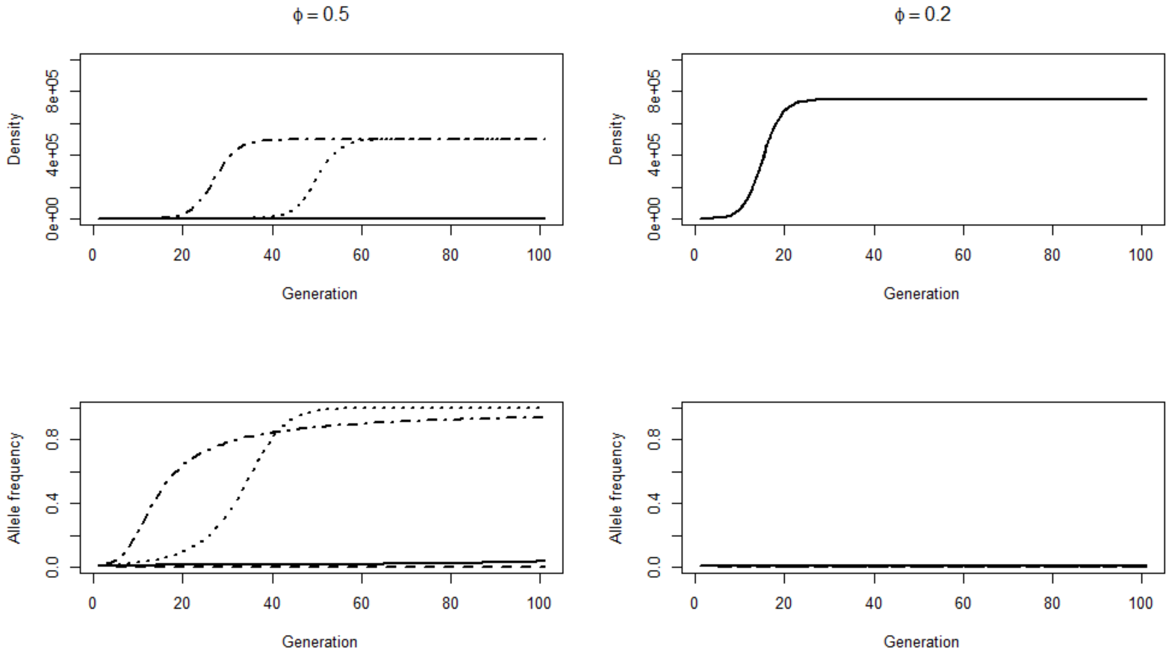

Figure S4: Population size (top row) and resistance allele frequency (bottom) over 100 generations for landscapes planted with 50% (left column) and 20% (right) Bt respectively when resistance is associated with a 25% reduction in fecundity ($\lambda_{rr} = 0.75\lambda_{ss}$). Density dependence is uniform and undercompensating ($\beta = 1$). Solid lines show recessive resistance with a recessive fitness cost ($h_{res} = 0, h_{cost} = 0$). Dashed lines depict recessive resistance with a dominant fitness cost ($h_{res} = 0, h_{cost} = 1$). Dotted lines refer to additive resistance (co-dominant) with a dominant fitness cost ($h_{res} = 0.5, h_{cost} = 1$) and dash-dot lines illustrate dominant resistance with a dominant fitness cost ($h_{res} = 1, h_{cost} = 1$). Populations were founded with $N_0 = 1000$ insects and initial resistance allele frequency $q_0 = 0.01$

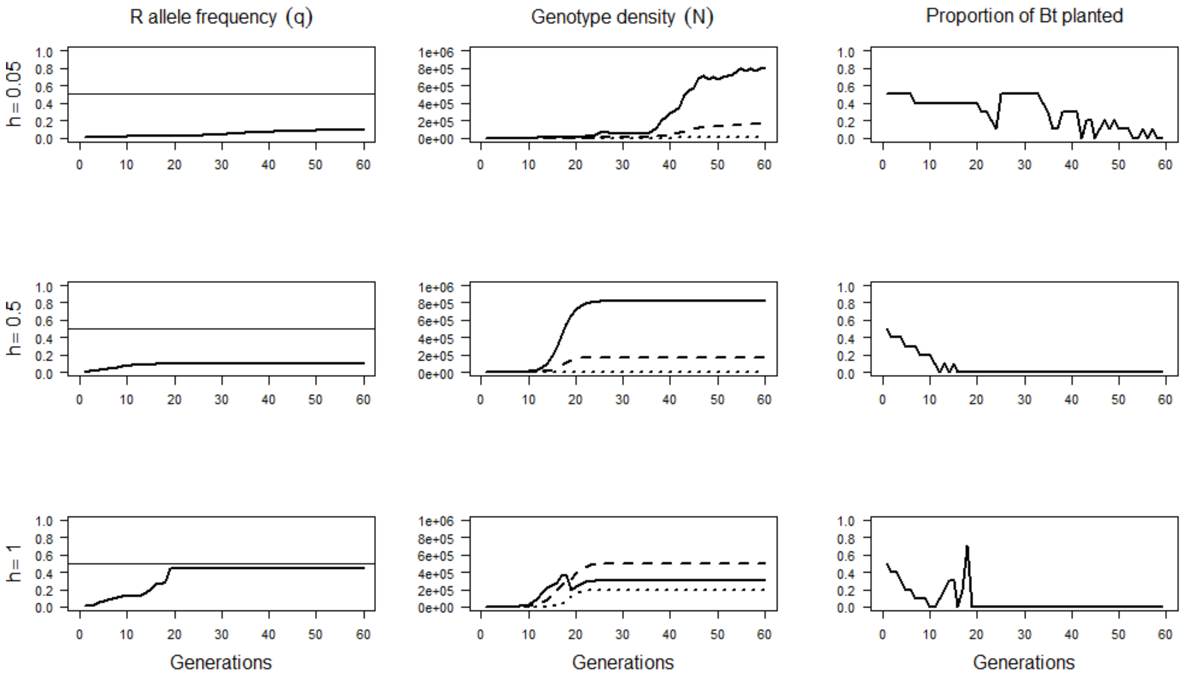

Figure S5: Resistance allele frequency (leftmost column), genotype abundance (central column) and proportional allocation of space to Bt crops over $T = 60$ generations when resistance is cost free. Density dependence is uniform and undercompensating with $\beta = 1$. The dominance of resistance was increased from $h = 0.05$ (top row), to $h = 0.5$ (central row) and $h = 1$ (bottom row). The solid horizontal line on plots of resistance allele frequency indicates the critical allele frequency, $q_c = 0.5$ beyond which resistance management is deemed to have failed. For plots of genotype specific density, solid lines represent susceptible homozygotes (ss), dashed lines represent heterozygotes (sr) and dotted lines represent resistant homozygotes (rr). Populations were founded by $N_0 = 1000$ insects with a resistance allele frequency of $q_0 = 0.01$.

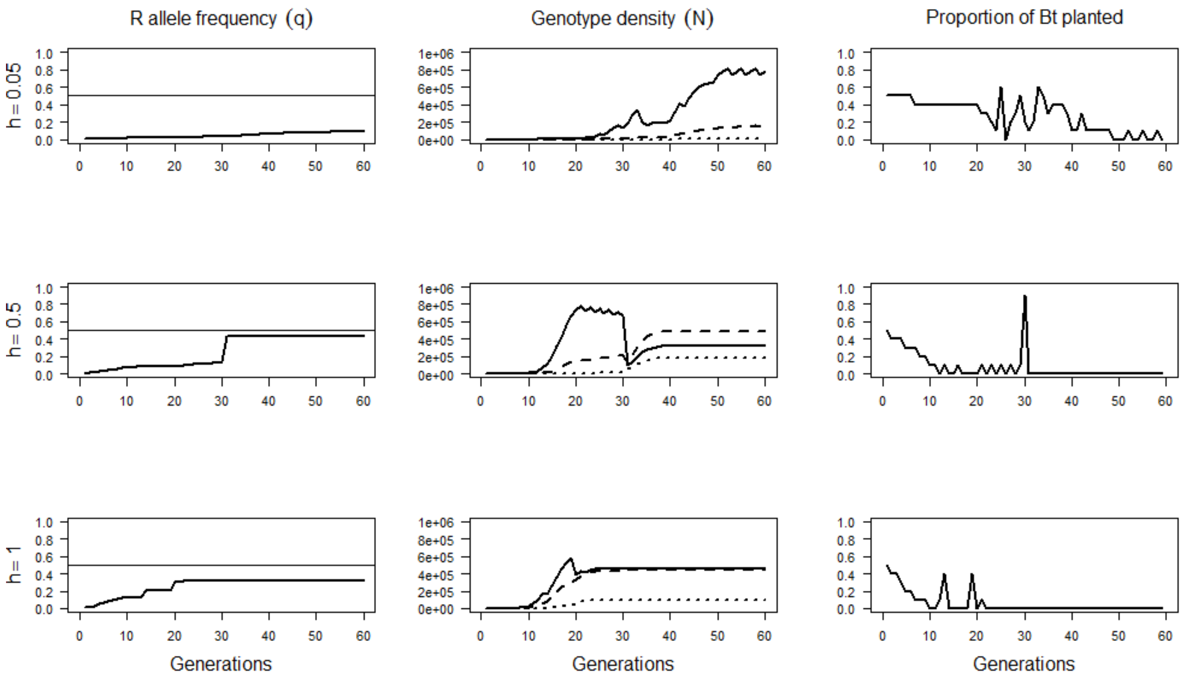

Figure S6: Resistance allele frequency (leftmost column), genotype abundance (central), and the proportional allocation of space to Bt crops over $T = 60$ generations when resistance is cost free. Density dependence is heterogeneous and undercompensating ($\beta = 1$). The dominance of resistance was increased from $h = 0.05$ (top row), to $h = 0.5$ (central row) and $h = 1$ (bottom row). The solid horizontal line on plots of resistance allele frequency indicates the critical allele frequency, $q_c = 0.5$ beyond which resistance management is deemed to have failed. For plots of genotype specific density, solid lines represent susceptible homozygotes (ss), dashed lines represent heterozygotes (sr) and dotted lines represent resistant homozygotes (rr). Populations were founded by $N_0 = 1000$ insects with a resistance allele frequency of $q_0 = 0.01$.

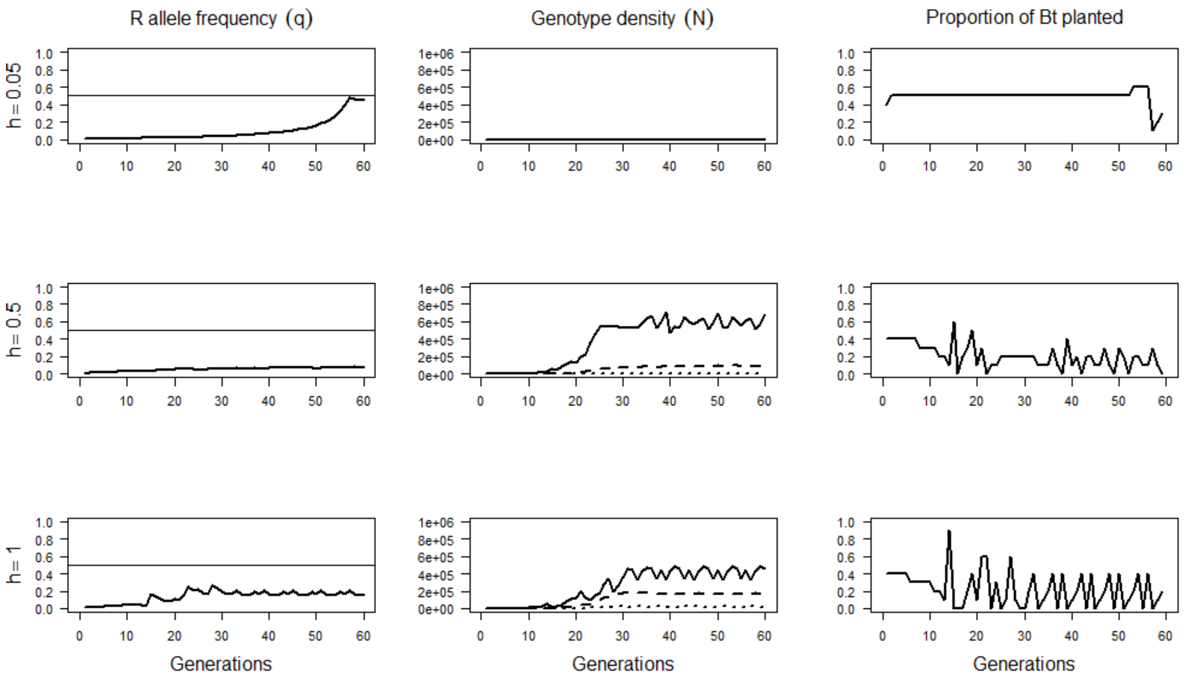

Figure S7: Resistance allele frequency (leftmost column), genotype abundance (central) and the proportional allocation of space to Bt crops over $T = 60$ generations when resistance is associated with a 25% reduction in fecundity ($\lambda_{rr} = 0.75\lambda_{ss}$). Density dependence is heterogeneous and undercompensating. The dominance of resistance was increased from $h = 0.05$ (top row), to $h = 0.5$ (central row) and $h = 1$ (bottom row). The solid horizontal line on plots of resistance allele frequency indicates the critical allele frequency, $q_c = 0.5$ beyond which resistance management is deemed to have failed. For plots of genotype specific density, solid lines represent susceptible homozygotes (ss), dashed lines represent heterozygotes (sr) and dotted lines represent resistant homozygotes (rr). Populations were founded by $N_0 = 1000$ insects with a resistance allele frequency of $q_0 = 0.01$

Chapter 3: Management of a stage-structured insect pest: an application of approximate optimization[2]

Sean C. Hackett[1] and Michael B.Bonsall[1,2]

[1] *Mathematical Ecology Research Group, Department of Zoology University of Oxford, South Parks Road, Oxford, OX1 3PS*

[2] *St. Peter's College, New Inn Hall Street, Oxford, OX1 2DL, UK*

Abstract

Ecological decision problems frequently require the optimization of a sequence of actions over time where actions may have both immediate and downstream effects. Dynamic programming can solve such problems only if the problem dimensionality is sufficiently low. Approximate dynamic programming (ADP) provides a suite of methods applicable to problems of arbitrary complexity at the expense of guaranteed optimality. The most easily generalised method is the lookahead policy: a brute-force algorithm which identifies reasonable actions by constructing and solving a series of temporally truncated approximations of the full problem over a defined planning horizon. We develop and apply this approach to a pest management problem inspired by the Mediterranean fruit fly, *Ceratitis capitata*. The model aims to minimise the cumulative costs of management actions and medfly-induced losses over a single 16-week season.

---

[2] Published as: Hackett, S.C. & Bonsall, M.B. (2018) Management of a stage-structured insect pest: an application of approximate optimization. *Ecological Applications,* doi: 10.1002/eap.1700

The medfly population is stage-structured and grows continuously while management decisions are made at discrete, weekly intervals. For each week, the model chooses between inaction, insecticide application or one of six sterile insect release ratios. Lookahead policy performance is evaluated over a range of planning horizons, two levels of crop susceptibility to medfly and three levels of pesticide persistence. In all cases, the actions proposed by the lookahead policy are contrasted to those of a myopic policy which minimises costs over only the current week. We find that lookahead policies always out-performed a myopic policy and decision quality is sensitive to the temporal distribution of costs relative to the planning horizon: it is beneficial to extend the planning horizon when it excludes pertinent costs. However, longer planning horizons may reduce decision quality when major costs are resolved imminently. ADP methods such as the lookahead policy based approach developed here render questions intractable to dynamic programming amenable to inference but should be applied carefully as their flexibility comes at the expense of guaranteed optimality. However, given the complexity of many ecological management problems, the capacity to propose a strategy that is “good enough” using a more representative problem formulation may be preferable to an optimal strategy derived from a simplified model.

Introduction

Ecological management problems span a range of disciplines from conservation to public health and are frequently presented as a sequential decision problem in which decisions must be made recurrently and actions taken now can influence the actions we may take in the future. Problems of this type can be represented as a dynamic program, a sequential decision problem (which may be stochastic) in which the system evolves over time and that satisfies key technical assumptions (Powell 2014). The system is summarised by a set of state variables that describe features -which may be qualitative or quantitative- required to model changes in the system induced by the passage of time and management actions and to evaluate decisions. For example, a simple single-species conservation problem could use the abundance of the focal species and remaining funds as state variables. In solving a dynamic program, the goal is to identify an optimal strategy for each state. That is, to find for each state the action that, when applied to that state produces the best outcome with respect to the specified time horizon, which may be finite or infinite (Powell 2011, 2014; Marescot *et al.* 2013). Thus, the chosen action must be optimal, not only with respect to the current state, but also the sequence of actions and states that will follow from this state as a consequence of the action. Classically, dynamic programs in ecology have been approached using dynamic programming, a method which partitions a complex, multi-period, decision problem into a sequential series of simpler sub-problems (Clark & Mangel 2000; Powell 2011; Marescot *et al.* 2013). The combined solutions for each sub-problem describe the optimal strategy for the full problem (Clark & Mangel 2000; Powell 2011; Marescot *et al.* 2013).

Traditionally, dynamic programming problems in ecology have been solved using Bellman's equation (also known as the dynamic programming equation) which characterises the optimal value V of a system in state S at time t by solving an expression of the form

$$V_t(S_t) = \max_a \{C_t(S_t, a_t) + V_{t+1}(S_{t+1})\} \quad [1]$$

to identify the action a which maximises the current $C_t(S_t, a_t)$ and future value $V_{t+1}(S_{t+1})$ of the decision (which assumes optimal behaviour from $t+1$ onwards and may be discounted when appropriate). The Bellman equation defines an iterative algorithm which identifies the best mapping between states and actions by recursively identifying the action that optimises the sum of both current and future benefits for each combination of state, action and time. As equation [1] defines the optimal value of a state recursively in terms of both immediate and long term rewards, the downstream effects of a decision made now are accommodated automatically (Clark & Mangel 2000; Powell 2011; Marescot *et al.* 2013). Thus, the selected action will be optimal with respect to the time horizon of interest. Equation [1] is typically solved within a backwards recursion algorithm (for finite horizon problems) or as part of a value or policy iteration algorithm (for infinite horizon problems) all of which readily accommodate the stochasticity and non-linearity inherent to many ecological questions (Marescot *et al.* 2013).

However, irrespective of the chosen horizon, to solve a dynamic program using Bellman's equation requires that the problem be of sufficiently low dimensionality that the solution for each state, at each time (for finite horizon problems) can be compiled in a lookup table (Clark & Mangel 2000; Powell 2011; Marescot *et al.* 2013). Three curses of dimensionality that limit the potential applications of dynamic programming arise from this condition: should any one of (i) the state space, (ii) the action space or (iii) (for stochastic models) the outcome space grow sufficiently large the problem becomes intractable to solution by traditional dynamic programming algorithms (Powell 2011). In some instances it is possible to reduce the dimensionality of the problem, e.g. by discretising a continuous state variable and then using linear interpolation (Clark & Mangel 2000), such that traditional algorithms may still be used (Shoemaker 1982). However, in reducing the dimensionality of the problem, we sacrifice realism which may be of use when making predictions in a management context and constrain the range of questions we can explore and exclude problems where dimensionality reduction is not possible.

In circumventing the curses of dimensionality, it is useful to distinguish between a dynamic program and the method used to solve it. Bellman's equation and dynamic programming are frequently conflated which constrains how dynamic programs are designed and solved (Marescot *et al.* 2013; Powell 2014). A dynamic program is solved using a policy, which is any rule for mapping feasible actions onto states (Powell 2014). Thus, Bellman's equation is not itself a dynamic program, but is both a means of describing an optimal policy (actions which satisfy equation [1] will always be optimal) and a valid algorithm for the solution of a dynamic program of sufficiently low dimensionality (Powell 2014). However, alternative policies for decision-making exist which are more permissive of high dimensional problems and capable of accommodating complications such as continuous state variables or continuous random outcomes. Such policies are collected together under the umbrella term 'approximate dynamic programming' (ADP). At the expense of guaranteed optimality, ADP methods can accommodate problems of arbitrary complexity (Powell 2011). In the context of complex ecological management problems, the trade of sub-optimal decisions against the ability to consider a broader range of questions, and to incorporate greater realism, may be a fair one (e.g. Nicol *et al.* 2010; Nicol & Chadès 2011).

Approximate dynamic programming policies can be broadly subdivided into four policy classes based upon the type of approximation used and the feature of the problem approximated (Powell 2014). The simplest policy type is the cost function approximation (CFA) which expands upon a myopic policy, a policy which optimises only with respect to the current benefits of an action ($C_t(S_t, a_t)$ in equation [1]) by introducing a correction term (which may contain tuneable parameters). The correction term is designed to promote decisions deemed “good” from the perspective of the current objective and to increase the robustness of the myopic policy to uncertainty (Powell 2014; Powell & Meisel 2016). A CFA is neither approximating nor explicitly modelling the downstream effects of an action: it remains a myopic policy, concerned only with the single period effect of an action. However, myopic policies can prove surprisingly effective for some problems (Wilson *et al.* 2006) and a well-designed CFA may be of most utility where a myopic policy already produces satisfactory results (Powell & Meisel 2016).

The policy type most commonly associated with ADP is the value function approximation (VFA). A VFA employs either parametric or non-parametric methods to define a function approximating the future value of downstream states resulting from an action we take now (the term $V_{t+1}(S_{t+1})$ in equation [1]) (Powell 2011, 2014; Powell & Meisel 2016). The dynamic program is then solved and actions suggested as if the VFA holds true. VFAs are extremely versatile and can be tailored to a diverse range of problems however, they are most useful when the value of the future is easily approximated and we can be confident of the fidelity of the approximation as a poorly specified VFA can significantly underperform (Powell 2014; Powell & Meisel 2016). Policy function approximations (PFAs) also utilise functional approximations but, unlike VFAs, they do not attempt to approximate an action’s downstream value. Instead, PFAs define a function representing the policy itself and then tune this function until it produces satisfactory outcomes (Powell 2014; Powell & Meisel 2016). Thus, while both PFAs and VFAs utilise functional approximations, they differ significantly in how they approach the optimization problem. VFAs explicitly solve the optimization problem as if the constructed VFA provides a true representation of future costs and benefits. That is, a VFA identifies an optimal strategy with respect to the chosen functional approximation. PFAs do not solve the original optimization problem, they are not searching for an optimal strategy as the decision rule for each period is fixed by the architecture of the PFA. Instead, a PFA is searching for the parameter values which produce the best long-run behaviour. As such, PFAs are most usefully applied to problems where the relationship between state and action is clearly defined and easily observable (for example, determining how much water to add to a reservoir) or where a good decision rule can be easily conceived (for example, top up the reservoir by X litres whenever it falls below level Y) (Powell 2011, 2014; Powell & Meisel 2016). However, it follows from their comparative rigidity that PFAs struggle to accommodate non-stationary processes and parameters and are best suited to stationary applications (Powell 2014; Powell & Meisel 2016).

The fourth policy type is the lookahead policy, also referred to as the rolling horizon procedure and model predictive control. Lookahead policies identify actions by solving a series of truncated approximations of the full problem (Nicol *et al.* 2010; Powell 2011, 2014; Powell & Meisel 2016). For example, the optimal action that could be taken during the current period t for a system in state S_t is that which maximises the benefits accrued from the current period t to the terminal period T where, for the purposes of this example, the values of t and T are arbitrary (with the caveat that $t < T$). Where the curses of dimensionality render this full problem intractable it may still be possible to identify good actions by optimising between the current period t and some truncated planning horizon $H < T$ (Powell 2014; Powell & Meisel 2016). Thus, lookahead policies do not depend upon functional approximations and instead essentially approximate the problem itself. Consequently, while the actions suggested by a lookahead policy are unlikely to be optimal with respect to the full horizon of interest, lookahead policies are extremely versatile. The runtime of a lookahead policy is independent of the magnitude of the state space (Nicol *et al.* 2010; Powell 2011), they readily accommodate non-stationarity (Powell 2014; Powell & Meisel 2016) and they can be applied to problems where the value of a state is a complex function or is subject to sufficient uncertainty that it is difficult to reliably approximate (Powell 2014; Powell & Meisel 2016). Thus we follow the example of Nicol *et al.* (2010) and Nicol & Chadès (2011) in suggesting that lookahead policies are an excellent candidate for application to ecological dynamic programs that are otherwise intractable to exact solution methods.

In what follows we more fully explain the operation of a lookahead policy and then construct and apply one to the intraseasonal management of a stage structured pest population inspired by the Mediterranean fruit fly (or medfly), *Ceratitis capitata.* The medfly is a widely distributed highly invasive fruit pest of substantive economic significance (Diamantidis *et al.* 2009, 2011; Szyniszewska & Tatem 2014). Medfly can be long-lived (Carey 2011; Diamantidis *et al.* 2011), reproduce continuously and exhibit overlapping generations (Barclay 2005, 2016), properties which are not conducive to traditional dynamic programming methods. At present, medfly management is highly reliant upon insecticidal sprays (Vontas *et al.* 2011) but sterile insect releases – in which pest suppression results from sustained mass-releases of sterilised male insects - have also provided effective areawide suppression in some regions (Enkerlin 2005). Thus, we consider a lookahead policy seeking to integrate these tools to suppress the simulated medfly population. We find that lookahead policies are able to produce reasonable action sequences which effectively integrate both bi-sex lethal insect releases and insecticidal sprays for the management of medfly. Furthermore, they offer significant potential for future extension to accommodate more complex problems such as those presented by the recent diversification of sterile insect methods, facilitated by developments in genetic engineering (Alphey *et al.* 2010; Bourtzis *et al.* 2016; Leftwich, Bolton & Chapman 2016; Harvey-Samuel, Ant & Alphey 2017) .

Methods

Overview

We present a sequential decision model for the combined application of self-limiting insect releases (male-only releases, the progeny of which do not survive to the reproductive stage) and insecticidal sprays to the intraseasonal management of a polyvoltinous pest inspired by the medfly. As medfly populations typically reproduce continuously and exhibit overlapping generations (Barclay 2005, 2016), we represent the medfly population with a set of stage-structured delay differential equations. This representation creates a large state space with multiple continuously valued state-variables, rendering the problem intractable to traditional dynamic programming methods. Thus, we approach this problem using a lookahead policy, a variant of approximate dynamic programming that identifies good (if sub-optimal) actions by solving the optimization problem over a truncated time horizon (Nicol *et al.* 2010; Powell 2011, 2014).

We assume a manager seeking to minimise the cumulative costs of medfly induced losses and medfly management over the course of the season. A season is sub-divided into W discrete weeks. At the start of each week w, a decision is made, based upon the current state, and applied to the medfly population instantaneously. Note that we adopt the convention of Powell (2016) in defining the state of the system to be the minimally dimensioned function of history required to model the transitions of the system from a given point in time onwards. Using this definition, the state can include historical values from prior time points (Powell 2014), which enables us to represent the medfly population with delay equations. While actions are taken at discrete, one-week intervals, their effects and the dynamics of the medfly population are resolved in continuous time. Each week spans t_w time increments (hence, a day spans $\frac{t_w}{7}$ time increments), giving a total season duration of $t_w W$ time increments. That is, the state after t_w time increments informs the decision for the subsequent week.

For any given week w the manager may elect to do nothing or to choose from one of two alternative actions: insecticidal foliar sprays and self-limiting releases with an early-acting bisex lethal effect. The decision to spray or to perform a self-limiting release instantaneously updates the relevant control variable and incurs an immediate point cost. Spraying sets φ the instantaneous insecticide-induced mortality rate to its maximum value $0 < \varphi_{max} < 1$ (note that this implementation prevents spraying eradicating the population, there will always be survivors) from which it subsequently decays with time at instantaneous rate α. Spraying is assumed to be uniform across the entire landscape; there are no refugia. For each spraying action the manager incurs a constant point cost c^{spray}. If the decision is made to release self-limiting insects, then the manager may choose from a range of permissible release ratios. Thus, the costs of self-limiting releases scale with both the magnitude of the medfly population and the release ratio. If the basic cost of a single self-limiting insect is c^{ste} then the cost of releasing m^{ste} self-limiting males is $c^{ste}m^{ste}$. Irrespective of the release ratio, released insects are assumed to be distributed uniformly across the landscape. In addition to management costs, feeding by larval medfly continuously devalues the crop from one week to the next at instantaneous rate γ which broadly reflects the susceptibility of the crop to damage. Thus, large values of γ portray a more vulnerable crop while smaller values indicate a greater tolerance to medfly.

The medfly model

We consider a closed medfly population with a 1:1 sex ratio. We assume that male and female medfly exhibit equivalent parameter values and so we explicitly track only female numbers. The medfly population is subdivided into the following sequential stages: eggs e, larvae l, pupae p and pre-ovipositional females f^{pre}which then enter one of two reproductive stages, f^{wild} or f^{ste}based upon whether they are mated to a wild-type male (superscript $wild$) or a self-limiting male (superscript ste). Neither sprays nor self-limiting releases can directly affect the abundance of the juvenile stages. Thus, in the absence of competition and other exogenous sources of mortality, eggs, larvae and pupae are subject only to recruitment, maturation and natural mortality (which reflects viability as opposed to competition or predation). For any given life stage i, μ_i denotes the instantaneous natural mortality rate and τ_i defines the duration of the stage in days. The proportion of individuals surviving stage i to enter stage $i+1$ is $\sigma_i = \exp(-\mu_i \tau_i)$. If wild-mated females oviposit at an age-invariant rate of of λ eggs per female per day then the instantaneous changes in egg, larval and pupal abundance are found as

$$\frac{de}{dt} = \lambda f^{wild}(t) - \mu_e e(t) - \lambda f^{wild}(t-\tau_e)\sigma_e \quad [2]$$

$$\frac{dl}{dt} = \lambda f^{wild}(t-\tau_e)\sigma_e - \mu_l l(t) - \lambda f^{wild}(t-\tau_e-\tau_l)\sigma_e\sigma_l \quad [3]$$

$$\frac{dp}{dt} = \lambda f^{wild}(t-\tau_e-\tau_l)\sigma_e\sigma_l - \mu_p p(t) - \lambda f^{wild}\left(t-\tau_e-\tau_l-\tau_p\right)\sigma_e\sigma_l\sigma_p \quad [4].$$

Changes in adult abundance are more complex. First, adults are subject to additional exogenous mortality from insecticide φ and second, pre-ovipositional adults may be recruited into either the wild-mated or sterile-mated stage. Pre-ovipositional females attain reproductive maturity after $\tau_{f^{pre}}$ days at which point they are instantaneously mated and, if mated to a wild-type male, immediately commence oviposition. While re-mating by female medfly is possible (Kraaijeveld & Chapman 2004) we assume that released and wild-type males are equivalent with respect to post-copulatory traits (Barclay 2005). However, insecticide-induced mortality complicates the determination of pre-ovipositional stage survivorship $\sigma_{f^{pre}}$ as mortality over the duration of this stage will not be constant in the presence of insecticide. Recruitment into the two mated stages must reflect this potential for variation in pre-ovipositional survivorship. Therefore, we let $M_a(t)$ denote the total instantaneous mortality rate experienced by all adult medfly at time t which is found as

$$M_a(t) = \mu_a + \varphi(t) \quad [5]$$

where μ_a is the baseline instantaneous natural mortality rate for all adult stages and $\varphi(t)$ is the instantaneous rate of insecticide-induced mortality at time t. We then use $M_a(t)$ to determine the change in pre-ovipositional survivorship over the duration of the stage

$$\frac{d\sigma_{f^{pre}}}{dt} = \sigma_{f^{pre}}(t)\left(M_a\left(t - \tau_{f^{pre}}\right) - M_a(t)\right) \quad [6].$$

Using equations [5] and [6] the instantaneous rates of change for pre-ovipositional females, wild-mated females and sterile-mated females are found as

$$\frac{df^{pre}}{dt} = \lambda f^{wild}(t - \tau_e - \tau_l - \tau_p)\sigma_e\sigma_l\sigma_p - M_a(t)f^{pre}(t)$$
$$- \lambda f^{wild}(t - \tau_e - \tau_l - \tau_p - \tau_{f^{pre}})\sigma_e\sigma_l\sigma_p\sigma_{f^{pre}}(t) \quad [7]$$

$$\frac{df^{wild}}{dt} = \left(\frac{m^{wild}(t)}{m^{wild}(t) + m^{ste}(t)}\right)\lambda f^{wild}(t - \tau_e - \tau_l - \tau_p - \tau_{f^{pre}})\sigma_e\sigma_l\sigma_p\sigma_{f^{pre}}(t)$$
$$- M_a(t)f^{wild}(t) \quad [8]$$

$$\frac{df^{ste}}{dt} = \left(\frac{m^{ste}(t)}{m^{wild}(t) + m^{ste}(t)}\right)\lambda f^{wild}(t - \tau_e - \tau_l - \tau_p - \tau_{f^{pre}})\sigma_e\sigma_l\sigma_p\sigma_{f^{pre}}(t)$$
$$- M_a(t)f^{ste}(t) \quad [9]$$

where m^{wild} denotes the number of wild-type males and m^{ste} is the number of self-limiting males. At any given instant, the number of wild type males is simply found as $m^{wild} = f^{pre} + f^{wild} + f^{ste}$.

The control variables, m^{ste} and φ, are not static between decision periods; insecticide toxicity declines with time and self-limiting insects are removed from the population by either natural mortality or by exposure to insecticide residues. We capture the change in insecticide mortality as a simple exponential decay

$$\frac{d\varphi}{dt} = -\alpha\varphi(t) \quad [10]$$

where α determines the persistence of the toxin. Self-limiting male abundance declines according to

$$\frac{dm^{ste}}{dt} = -M_{m^{ste}}(t)m^{ste}(t) \quad [11]$$

where $M_{m^{ste}}$ denotes the self-limiting equivalent of M_a and is found at a given instant in time as

$$M_{m^{ste}}(t) = \mu_{m^{ste}} + \varphi(t) \quad [12]$$

where $\mu_{m^{ste}}$ is the baseline self-limiting male mortality rate. Note that equation [12] assumes that self-limiting males do not differ from wild-type adults in their susceptibility to insecticide. From one week to the next larval medfly erode the value of the crop at rate

$$\frac{dc}{dt} = \gamma(2l(t)) \quad [13]$$

where γ is the per capita larval feeding rate and the doubling of the larval population accounts for damage caused by male larvae.

Equations [2-13] define the transition function for our model. That is they determine how the state of the system evolves over time and in response to management actions (Powell 2014; Powell & Meisel 2016). Given that we define the state of a system S_t to be the minimally dimensioned function of history required to model the transitions of that system from time t onwards, then the state for this model, at the start of any given week w is summarised as

$$S_w = \Big(e(t), l(t), p(t), f^{pre}(t), f^{wild}(t), f^{ste}(t), f^{wild}(t - \tau_e),\ f^{wild}(t - \tau_e - \tau_l), f^{wild}\big(t - \tau_e - \tau_l - \tau_p\big), f^{wild}\big(t - \tau_e - \tau_l - \tau_p - \tau_{f^{pre}}\big), M_a(t), M_a(t - \tau_e)\sigma_{f^{pre}}(t), \varphi(t), m^{ste}(t), M_s(t), c(t)\Big) \quad [14]$$

The manager’s objective is to minimise losses from larval medfly feeding and the cumulative cost of medfly suppression over the duration of the season. Thus, the optimal decision for any week a_w^* is that which satisfies

$$a_w^* = \arg\min_{a_w}\{C(S_w, a_w) + \sum_{w'=w+1}^{W} (C(S_{w'}, a_{w'}))\} \quad [15]$$

where $C(S_w, a_w)$ is a contribution function determining the full cost accrued between week w and week $w + 1$ which is found as

$$C(S_w, a_w) = c_a + \int_{t'=t_0}^{t'=t_w} \gamma(2l(t')) \quad [16]$$

where c_a denotes the point cost for the action implemented at the start of week w. The value of S_{w+1} is found via equations [2-13].

The Lookahead Policy

For a sufficiently small state space equation [15] could be solved exactly via backwards recursion (Clark & Mangel 2000; Marescot *et al.* 2013) however, the state space for this problem is complex and continuously valued. Thus, we approach this problem using a lookahead policy, a variant of approximate dynamic programming that uses a series of computationally tractable approximations of the model to suggest actions (Nicol *et al.* 2010; Powell 2011, 2014). We adopt the terminology of Powell (2014) and Powell & Meisel (2015) and define equation [15] to be the base model, the full problem that we seek to solve. Solution of the base model for all combinations of state, action and time would return a series of decisions which are optimal with respect to both their immediate and future costs but this is unattainable for a problem of this magnitude.

Rather than address the base model directly, lookahead policies proceed by constructing a series of truncated approximations of equation [15], termed lookahead models (Powell 2014; Powell & Meisel 2016). For example, to identify the action to be taken in week w given state S_w we construct the deterministic lookahead model

$$a_w = \arg\min_{a_w}\{C(S_w, a_w) + \sum_{w'=w+1}^{w'=w+H} (C(S_{\tilde{w}w'}, a_{\tilde{w}w'}))\} \quad [17]$$

where $H < W$ is a suitably short planning horizon chosen to capture important behaviours while maintaining computational tractability. Tildes distinguish variables from the base model from those in the lookahead model while the ww' subscript denotes a variable encountered in week w' within the lookahead model when solving for the action to be taken in week w (Powell 2014). Thus, $a_{\widetilde{ww'}}$ is the action implemented in week w' within the lookahead model when solving for a_w . This notation emphasises the distinction between the lookahead model (equation [17]) and the base model (equation [15]). Solving a lookahead model is not equivalent to solving the base model but by solving the lookahead model over a sufficiently long horizon, good, if sub-optimal, actions will be taken that will outperform those suggested by a purely myopic policy (Nicol *et al.* 2010; Powell 2011, 2014; Powell, Simao & Bouzaiene-Ayari 2012; Powell & Meisel 2016).

For the current application, we note that the immediate contribution from a given state-action pair is resolved over a one-week interval. Thus, a myopic policy applied to this problem would simulate the effects of each of the A available actions from S_w to S_{w+1} and produce a decision based upon which of the A alternative downstream states had the least cost. By repeating the process of iteratively applying each action to each state for up to H periods beyond $w + 1$, a lookahead model may attain a better representation of an actions value. In optimising over the interval $(w, \dots, w + H)$ a lookahead model generates a tree of decisions and states. By identifying the terminal state with the least cost and then solving back towards the original input state at the root of the tree the model isolates the action which is optimal with respect to the limited horizon of the lookahead model (Nicol *et al.* 2010; Powell 2011, 2014). The returned decision is then applied to state S_w and the model increments to week $w + 1$ where a new lookahead model is constructed and solved. While a lookahead model will return an optimal sequence of decisions over the full lookahead horizon H only the first of these decisions is implemented; the downstream decisions in the sequence were made only to improve the decision to be made now and are discarded (Powell 2011, 2014; Powell & Meisel 2016).

Parameterisation

The demographic model is parameterised using data presented in Diamantidis *et al.* 2009 (for adult parameters) and Diamantidis *et al.* 2011 (for juvenile parameters). Specifically, we utilise parameters derived from a Hawaiian medfly population which Diamantidis *et al* (2011) would classify as short-lived. Stage-specific instantaneous mortality rates μ_i for the pre-adult stages (eggs, larvae and pupae) were obtained from the stage duration and survivorship data presented by Diamantidis *et al (2011)* as

$$\mu_i = -\frac{\log(\sigma_i)}{\tau_i} \qquad [18].$$

The baseline natural instantaneous mortality rate for adult medfly μ_a was specified as the inverse of the longevity of an adult female where longevity refers to the average number of days a female medfly is expected to live post-eclosion. This was taken to be 50 days, based upon Figure 3 of Diamantidis *et al.* (2009). A constant value of 5 female eggs per female per day was taken to be a reasonable representation of daily egg production λ and the duration of the pre-ovipositional stage was taken to be 10 days (Diamantidis *et al.* 2009). Costs are unitless and intended only to provide a reasonable approximation of the relative costs of releases versus foliar sprays. We specify the point cost of spraying as $c^{spray} = 4000$ and the cost per sterile insect released as $c^{ste} = 0.003$. Thus, the cost of spraying is constant, irrespective of medfly abundance, while the absolute cost of a sterile release depends upon the total number of insects released which will vary with the chosen release ratio and the medfly population density. As this model only considers management over a single season, we do not consider the discounting of future costs (Zhang & Swinton 2009). We assume that the medfly population is highly susceptible to foliar spraying (Vontas *et al.* 2011) and so specify the maximum mortality rate from insecticide exposure as $\varphi_{max} = 0.9$. We assume that self-limiting males are equivalent to wild-type insects with respect to their natural mortality rate and their susceptibility to insecticide. The baseline values for all model parameters are summarised in Table 1.

Table 1*:* Summary of model parameters, their meanings and, where applicable, their origin. Instantaneous mortality rates for juvenile life stages may be derived using equation [18]. The mortality rate for adult stages is found as the inverse of adult female longevity.

Parameter	**Meaning**	**Point estimate/mean**	**Citation**
λ	Female eggs laid per female per day	5	Diamantidis et al 2009
τ_e	Egg stage duration (days)	2	Diamantidis et al 2011
σ_e	Egg stage survivorship	0.89	Diamantidis et al 2011
τ_l	Larval stage duration (days)	5.92	Diamantidis et al 2011
σ_l	Larval stage survivorship	0.865	Diamantidis et al 2011
τ_p	Pupal stage duration (days)	10.4	Diamantidis et al 2011
σ_p	Pupal stage survivorship	0.993	Diamantidis et al 2011
$\tau_{f^{pre}}$	Pre-ovipositional stage duration (days)	10	Diamantidis et al 2009
Adult female longevity (days)		50	Diamantidis et al 2009

c^{spray}	Point cost of insecticide application	4000	N/A
c^{ste}	Point cost per self-limiting male released	0.003	N/A
α	Instantaneous insecticide decay rate	0.175, 0.35, 0.7	N/A
φ_{max}	Maximum instantaneous insecticide mortality rate	0.9	N/A
γ	Instantaneous rate at which costs/losses from larval medfly accrue	0.001, 0.01	N/A
W	Season length (weeks)	16	N/A
t_w	Number of time increments comprising a week	700	N/A

Simulations

All simulations were performed using code written in R version 3.3.3 and were initiated with a population of 20,000 medfly; 10,000 mated females and 10,000 adult males. A season lasts for $W = 16$ weeks. The continuous time dynamics separating each decision period were approximated using an RK4 solver adapted to deal with time delays with a resolution of $t_w = 700$ time increments per week. The performance of three alternative lookahead policies with progressively longer planning horizons, $H = \{1,2,3\}$, were evaluated relative to one another and to a myopic policy, $H = 0$. We chose the progression $H = \{1,2,3\}$ to provide a simple illustration of the benefits and limitations of lookahead policies. In practice, the planning horizon would be determined by practical constraints such as the accuracy of available forecasts and the time available for computation. Costs from larval medfly infestation accumulated at one of two rates, $\gamma = 0.001$ or $\gamma = 0.01$, where the lower value represents a comparatively resilient crop in which large larval infestations are required before substantive losses occur and the larger value depicts a more susceptible crop for which the inverse is true. Once applied, insecticides decayed exponentially at one of three exponential rates, $\alpha = \{0.175,0.35,0.7\}$ where $\alpha = 0.35$ is taken to be the default value and depicts an insecticide which decays to low levels of toxicity ($\varphi < 0.1$) within approximately 1 week while $\alpha = 0.7$ and $\alpha = 0.175$, respectively reflect insecticides with lesser and greater persistence. We focus upon the parameters α and γ to demonstrate how the actions suggested by each policy change to accommodate variation in the urgency of medfly management (as determined by α, the susceptibility of the crop) and the extent to which each policy is able to plan actions in a way which effectively accounts for pesticide persistence (determined by γ).

Results

Effect of planning horizon (H) length under different damage regimes

For any given scenario, a myopic policy ($H = 0$) is always outperformed by a lookahead policy ($H \geq 1$). Even a minimal lookahead policy which evaluates actions over only one additional time period ($H = 1$) relative to a myopic policy confers substantial benefits in terms of cost reduction (Fig.1a). By evaluating an action over only a single week the myopic policy overemphasises short-term costs which biases it towards the application of insecticide and inaction (Fig.1b). Sterile releases are employed less frequently and only at low-mid release ratios (Fig.1b,e). The lookahead policies more frequently utilise self-limiting releases, often implementing larger release ratios than the myopic policy, and less frequently elect to perform no action (Fig.1c,d,f,g).

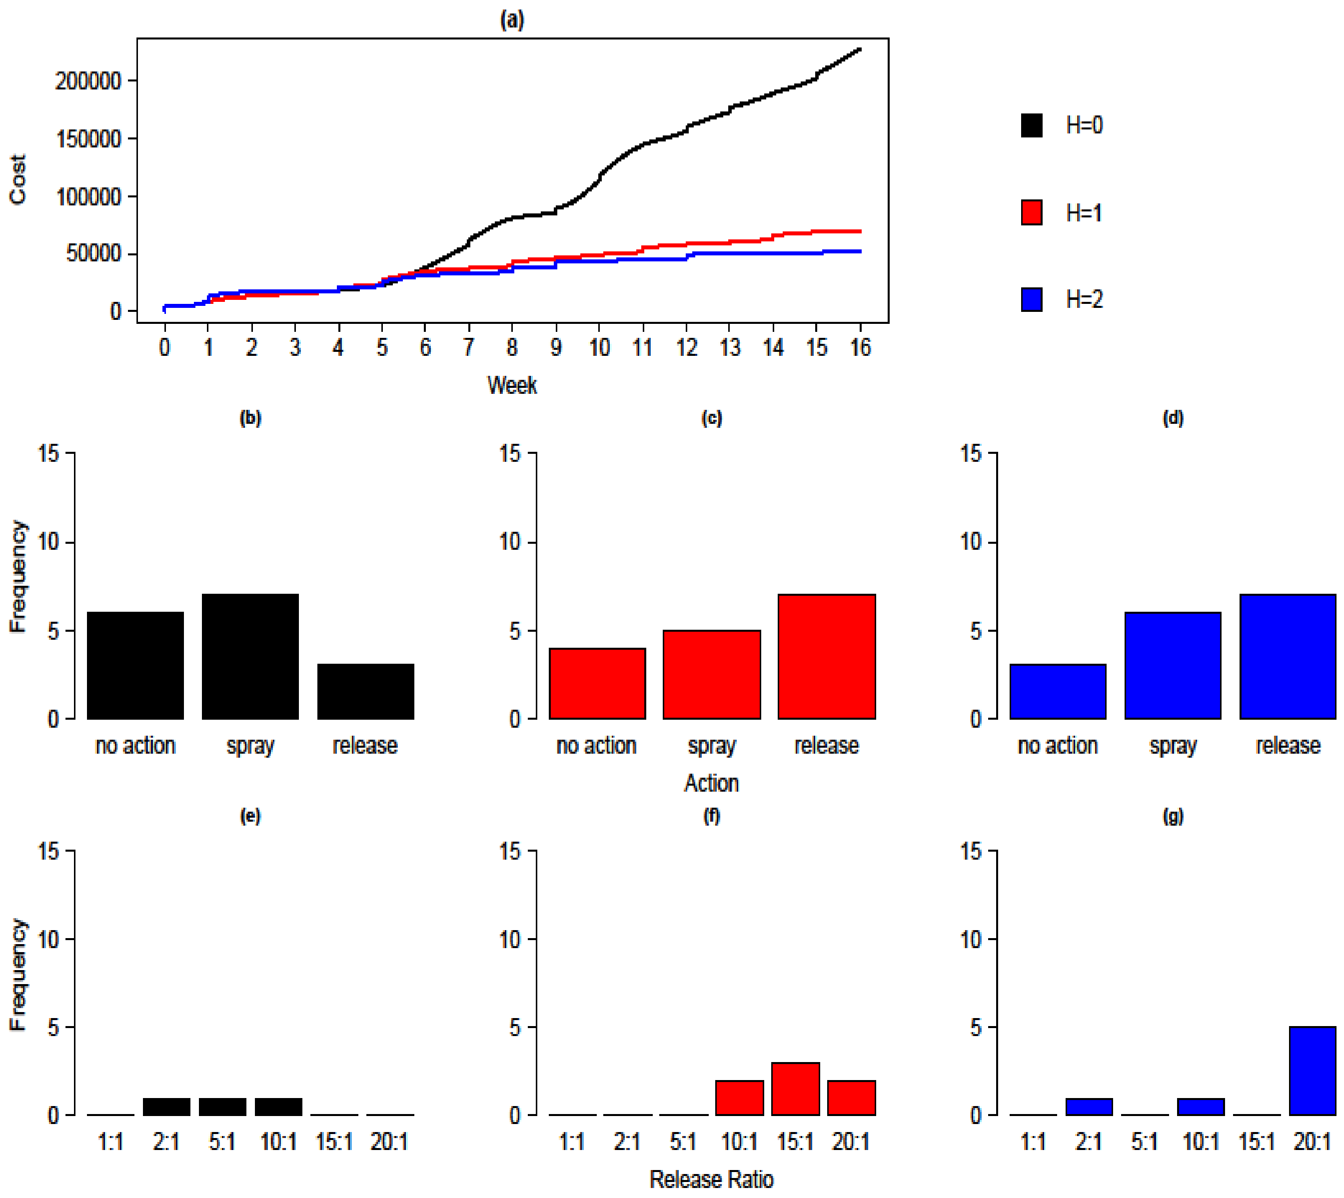

Figure 1*:* Comparisons of (a) the cumulative cost, (b-d) frequency of implemented actions and (e-g) sterile release ratios for three alternate policies attempting to manage a hypothetical medfly population) feeding upon a susceptible crop, $\gamma = 0.01$, over a 16-week season using insecticidal foliar sprays and sterile insect releases. Only a single action may be implemented in any given week. Black depicts a myopic policy, $H = 0$, while red and blue depict lookahead policies with planning horizons of $H = 1$ and $H = 2$, respectively where H is measured in weeks. Costs have no units and adult medfly mortality from insecticide exposure decays at instantaneous rate $\alpha = 0.35$. The

population was initiated with 10,000 mated adult female medfly and a 1:1 sex ratio was assumed.

The rate at which costs inflicted by the larval stage accrue interacts with the planning horizon H to influence action selection. Figure 2 illustrates cost trajectories, action frequencies and release ratios for a medfly population feeding on a resilient crop with a damage accrual rate of $\gamma = 0.001$. Counterintuitively, a reduction in the damage rate does not guarantee a substantive reduction in costs relative to a scenario with a more damaging pest (compare Fig.1a and Fig.2a). When pest-induced costs accumulate slowly, the negative effects of a poor decision are temporally delayed and a longer planning horizon is needed to incorporate these costs into decision-making. A myopic policy is now strongly biased against self-limiting releases and more frequently elects to perform no action, incurring significant costs in doing so (Fig.2a,b,e). A lookahead policy improves the quality of decisions, prompting a more proactive approach with fewer periods of inaction (Fig.2c,d), but releases are still employed less frequently relative to a scenario with a higher damage rate (Fig.1c,d) and smaller release ratios are more likely to be selected (Fig.2f,g).

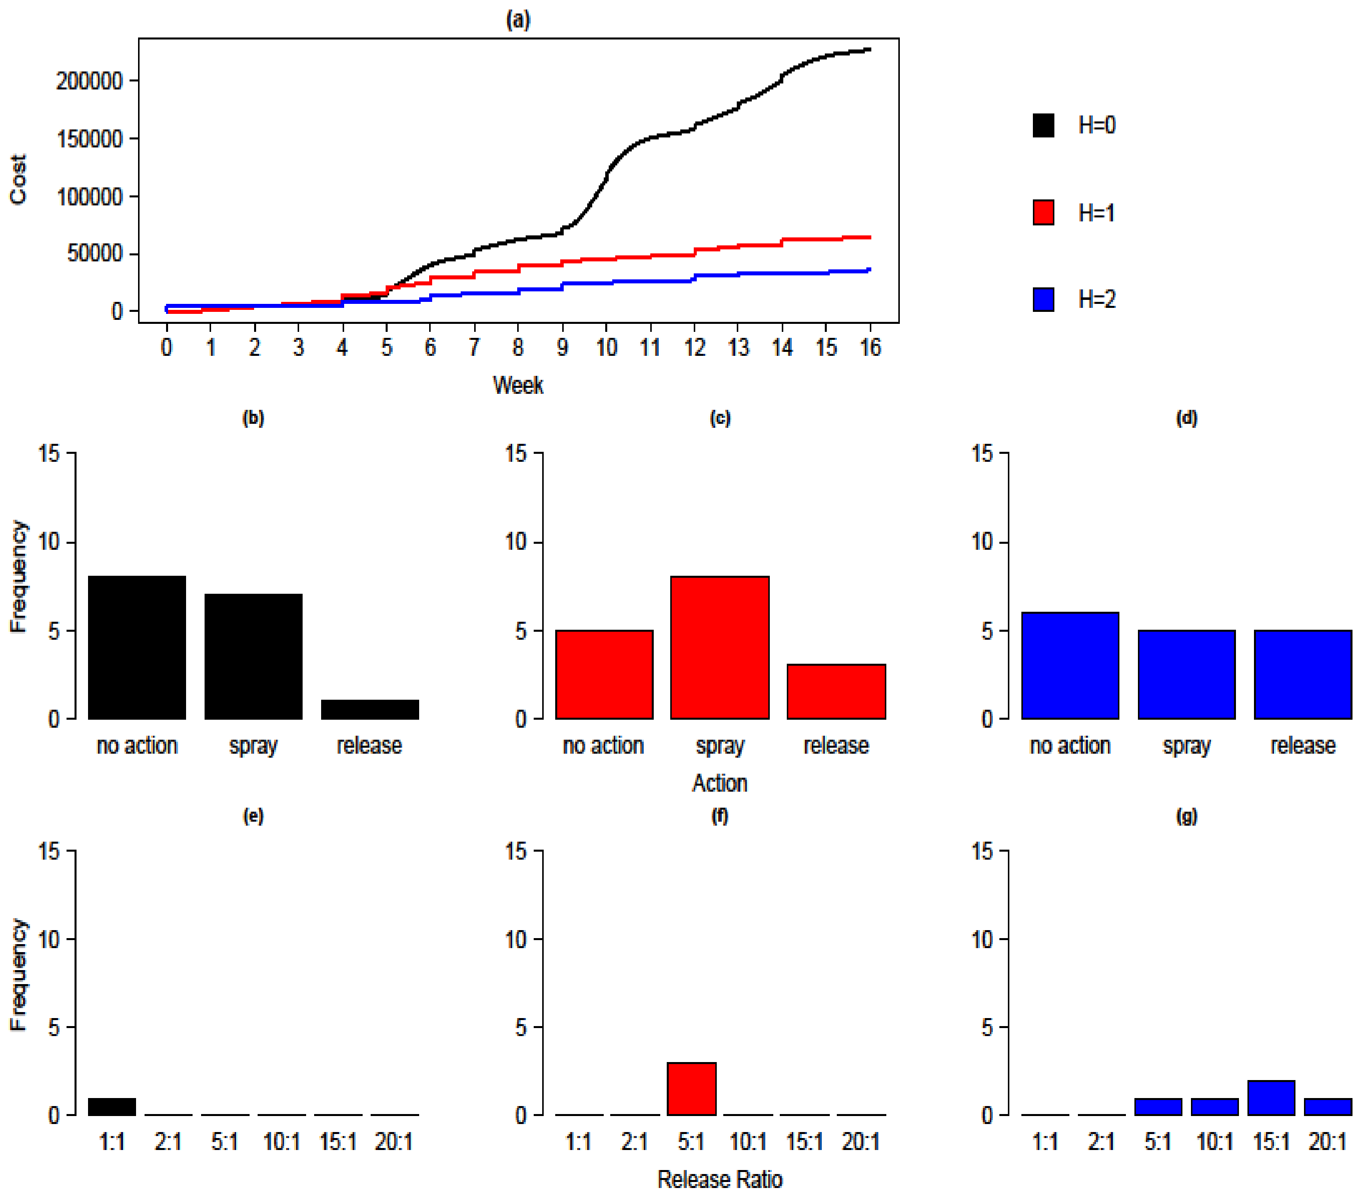

Figure 2*:* Comparisons of (a) the cumulative cost, (b-d) frequency of implemented actions and (e-g) sterile release ratios for three alternate policies attempting to manage a hypothetical medfly population) feeding upon a resilient crop, $\gamma = 0.001$, over a 16-week season using insecticidal foliar sprays and sterile insect releases. Only a single action may be implemented in any given week. Black depicts a myopic policy, $H = 0$, while red and blue depict lookahead policies with planning horizons of $H = 1$ and $H = 2$, respectively where H is measured in weeks. Costs have no units and adult medfly mortality from insecticide exposure decays at instantaneous rate $\alpha = 0.35$. The

population was initiated with 10,000 mated adult female medfly and a 1:1 sex ratio was assumed.

The interaction between the extent of the planning horizon and the threat posed by the pest influences the temporal sequence of actions in addition to their frequency. Figure 3 depicts the chronological sequence of actions implemented over the course of the 16-week season for the policies illustrated in Figure 1 ($\gamma = 0.01$, the right-hand column) and Figure 2 ($\gamma = 0.001$, the left-hand column). The set of alternate actions is listed on the y-axis. For $\gamma = 0.01$ (Fig.2d-f) the initial action is always to spray, irrespective of the policy. Spraying provides rapid knockdown of the founding population (which is comprised entirely of reproductive adults), reducing both the initial rate of population growth and the cost of larger release ratios. In contrast, when $\gamma = 0.001$, this initial spray is not identified as cost-effective until the planning horizon of the lookahead policy is extended to $H = 2$ (Fig.2c).

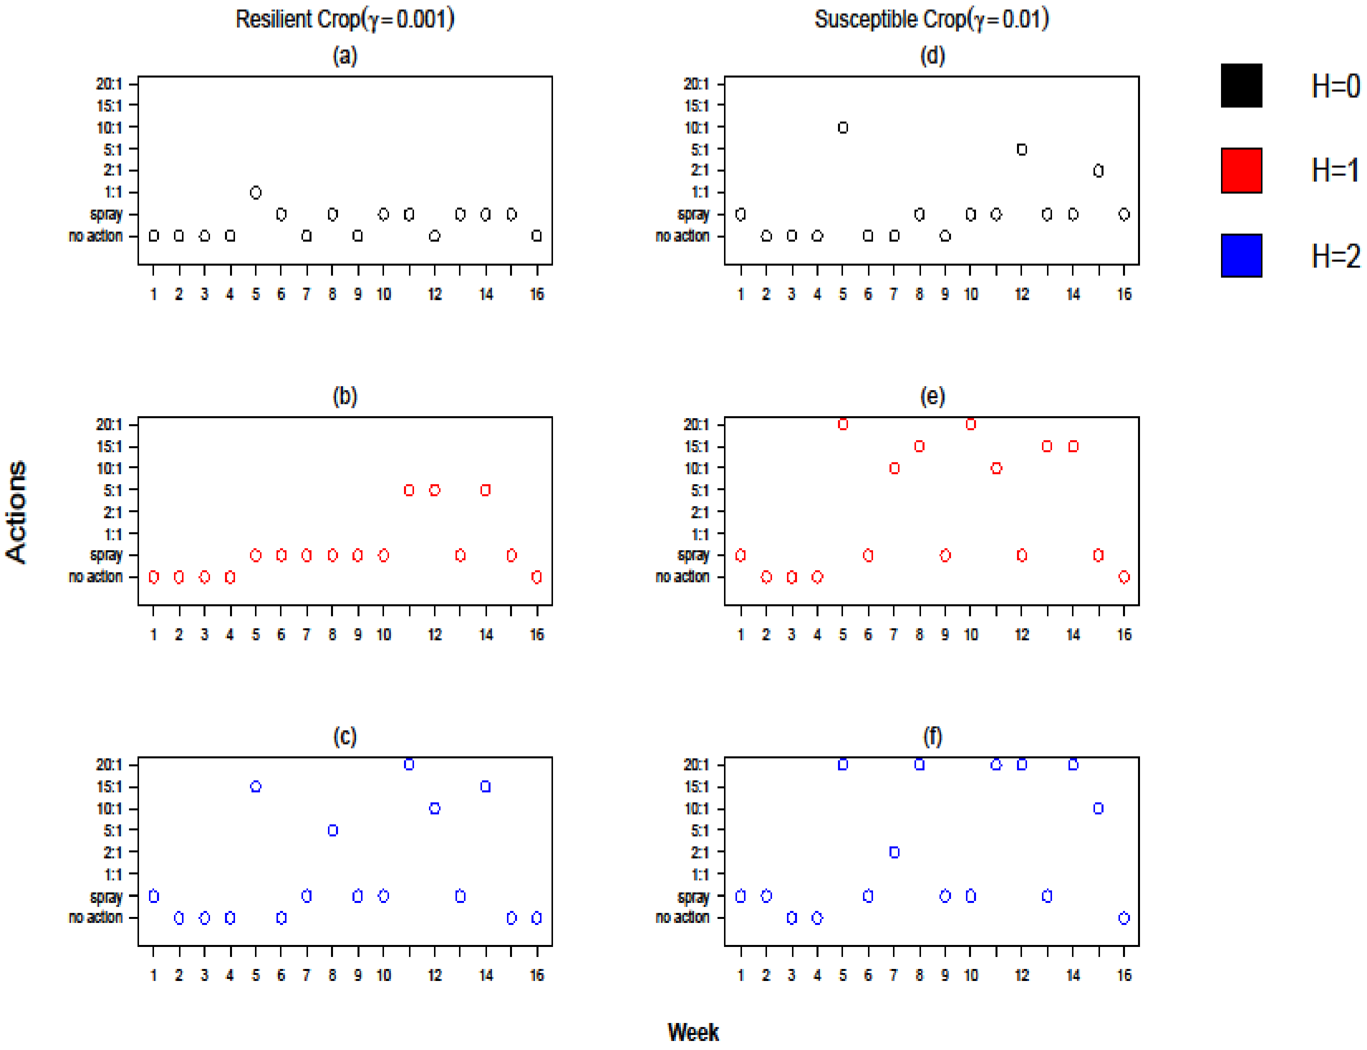

Figure 3*:* Sequences of implemented actions for three alternative policies for the management of a simulated medfly population over 16 weeks. For all panels the full range of alternative actions available in each decision period are listed on the y-axis and adult medfly mortality from insecticide exposure decays at instantaneous rate $\alpha = 0.35$. Costs from larval feeding accrue at instantaneous rate $\gamma = 0.001$ (a resilient crop) in the left column (a-c) and $\gamma = 0.01$ (a susceptible crop) in the right column (d-f). Black points indicate actions implemented by a myopic policy ($H = 0$) while red and blue points illustrate actions implemented by lookahead policies with planning horizons of

$H = 1$ and $H = 2$, respectively where H is measured in weeks. The population was initiated with 10,000 mated adult female medfly and a 1:1 sex ratio was assumed.

Taken together, the results shown in Figures 1-3 illustrate how lookahead policies reduce costs by employing proactive and diverse action sequences relative to a myopic policy. Thus, while still inevitably sub-optimal, lookahead policies can make interventions that prevent the medfly population from growing to highly damaging densities while a myopic policy will only intervene once the pest has attained a density for which significant damage is imminent. The extent to which this is possible depends upon the value of γ, the rate at which larvae induced costs accumulate relative to H, the length of the planning horizon. As the value of γ is increased, costs resulting from a poor decision are increasingly experienced in the near-term future. This enables even a myopic policy to suggest actions which will, to an extent, pre-empt costs before they occur (Fig.3a,d). However, lookahead policies still provide significant cost reductions (Fig.1a) by accounting for the possibility of damage over a longer time span. When the value of γ is low the inverse applies and the consequences of a poor decision may lie outside the planning horizon and so are not factored into decision making. This promotes policies that favour spraying and inaction (Fig.3a,b) unless a longer planning horizon is used (Fig.3c).

Effect of spray persistence upon the integration of spraying and sterile releases

Varying the persistence of sprays, as proxied by α the instantaneous rate at which adult medfly mortality from spray exposure declines, influences the ease with which lookahead policies are able to effectively integrate spraying with sterile releases. Figures 1-3 utilised a decay rate of $\alpha = 0.35$ such that spray toxicity decays to trivial values within the space of one simulated week, minimising the residual killing of sterile insects at the point of release. Figure 4 illustrates the cumulative costs and implemented action frequencies for the management of a medfly population with either a myopic policy or one of two lookahead policies when the insecticide decay rate is doubled ($\alpha = 0.7$, parameters are otherwise identical to those of Figure 1). Reducing spray persistence substantially increases costs for all policies (Fig4.a) but especially for a myopic policy where costs accumulate rapidly. All policies adopt spraying as their most frequent action (Fig.4 b-d) but only the lookahead policies make multiple uses of sterile releases, which contribute towards reducing management costs relative to the myopic policy which utilises only a single sterile release. The increase in spraying frequency with reduced spray persistence appears contradictory but emphasises the role of an initial knockdown in rendering self-limiting releases cost-effective over short time scales. Fewer adults are killed by each spraying event and so more will survive to enter the mated stage and produce offspring which themselves now have a greater likelihood of surviving to reproductive maturity. Thus, greater effort is required to reduce the adult population to a level for which, relative to the current planning horizon, releases appear cost-effective. The preference for larger release ratios may facilitate the persistence of the sterile population at a non-trivial level in the face of recurrent episodes of insecticide application.

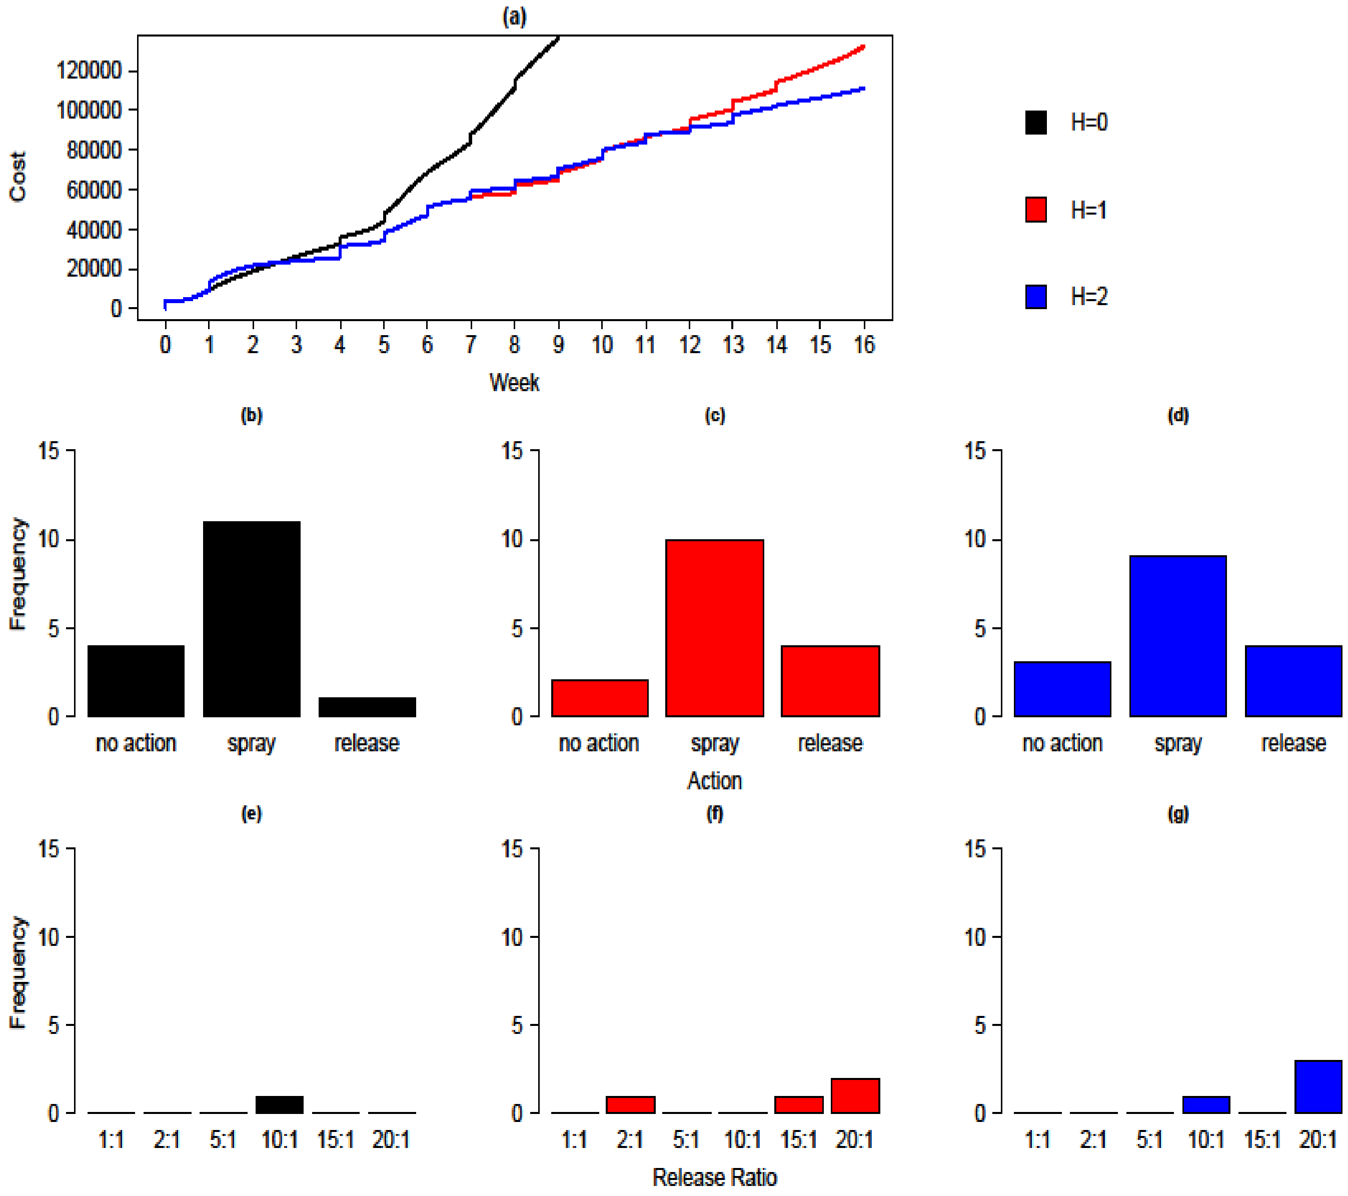

Figure 4*:* Comparisons of (a) the cumulative cost, (b-d) frequency of implemented actions and (e-g) sterile release ratios for three alternate policies attempting to manage a hypothetical medfly population) feeding upon a susceptible crop, $\gamma = 0.01$, over a 16-week season using transient insecticidal foliar sprays and sterile insect releases. Only a single action may be implemented in any given week. Black depicts a myopic policy, $H = 0$, while red and blue depict lookahead policies with planning horizons of $H = 1$ and $H = 2$, respectively where H is measured in weeks. Costs have no units and adult medfly mortality from insecticide exposure decays at instantaneous rate $\alpha = 0.7$. The

population was initiated with 10,000 mated adult female medfly and a 1:1 sex ratio was assumed.

Figure 5 illustrates cost trajectories and action frequencies when $\alpha = 0.175$ and all other parameters are identical to those in Figure 4. Reducing the insecticide decay rate reduces overall management costs for all policies, while still permitting effective integration of sprays with sterile releases. The reduction in costs is intuitive, if each insecticide spray persists at non-trivial concentrations for longer, fewer sprays are required for effective knockdown (Fig.5,b-c). All three policies, including the myopic policy now make more frequent use of sterile releases than foliar sprays but lookahead policies achieve significant cost reductions by favouring larger release ratios (Fig.5, e-g). Given that the persistence of the spray necessarily entails that some sterile insects will be lost to residue exposure subsequent to their release, it is striking that sterile releases are still identified as effective actions irrespective of the potential for overlap between releases and insecticide residues. The utility of releases in conjunction with a persistent spray is likely due to the improved knockdown potential of the spray rendering releases cost-effective even to a myopic policy. The preference of the lookahead policies for larger release ratios is likely promoted by the increased affordability of such ratios when knockdown from spraying is effective, and by the need to offset the loss of sterile insects to spray residues.

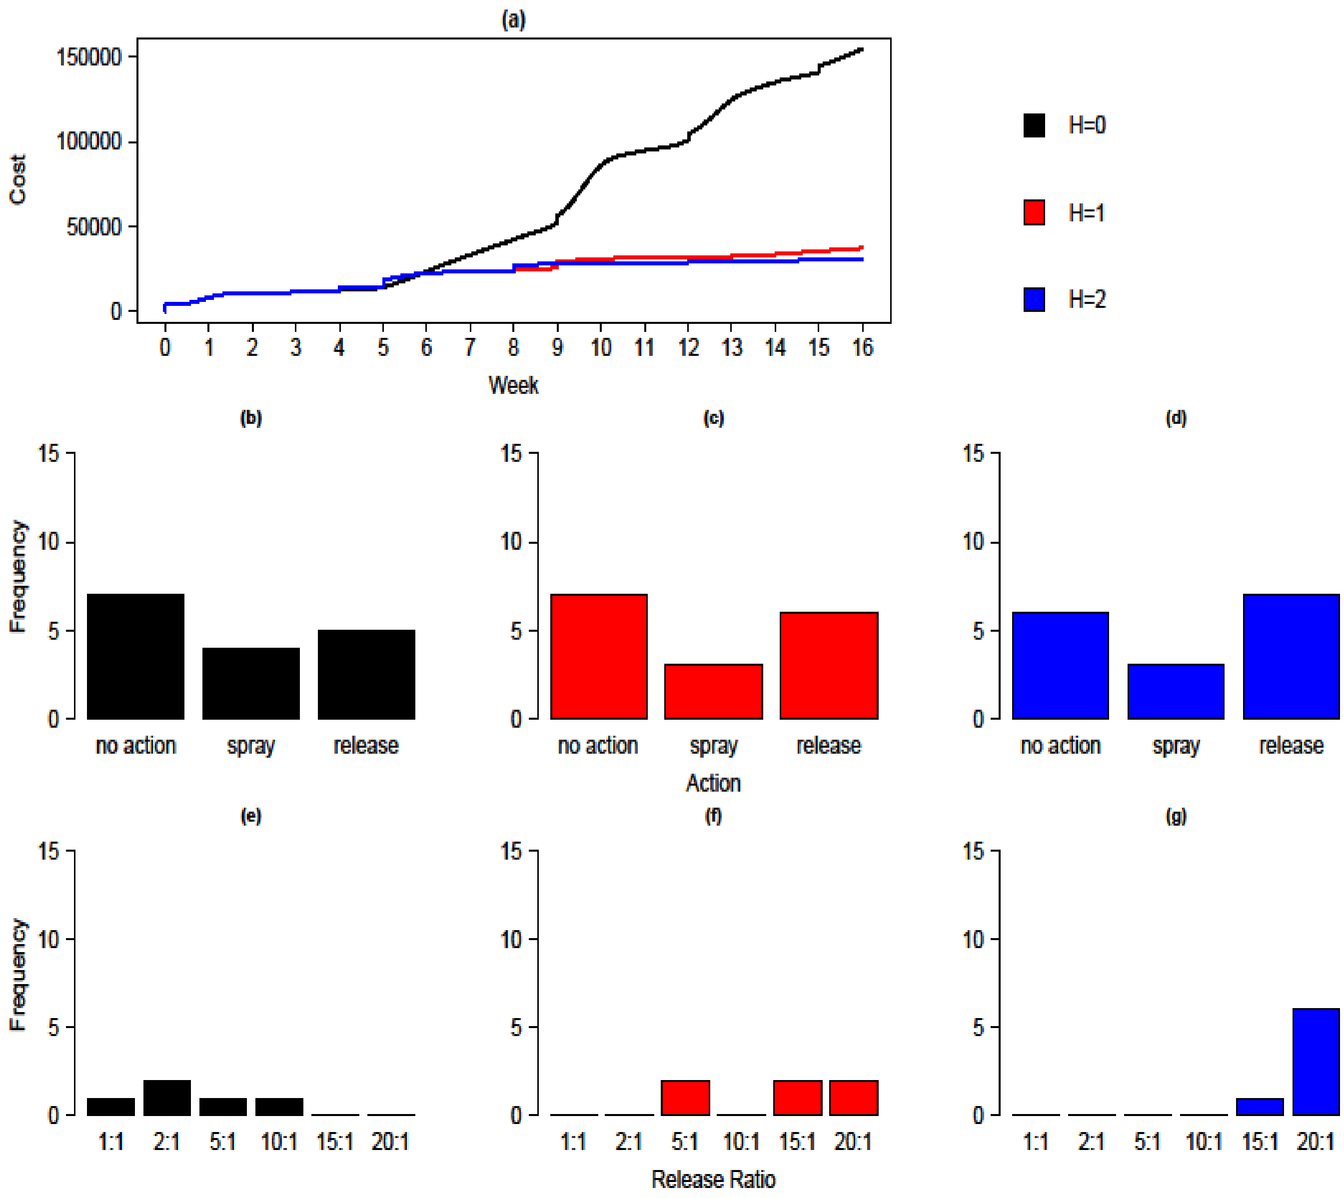

Figure 5: Comparisons of (a) the cumulative cost, (b-d) frequency of implemented actions and (e-g) sterile release ratios for three alternate policies attempting to manage a hypothetical medfly population) feeding upon a susceptible crop, $\gamma = 0.01$, over a 16-week season using persistent insecticidal foliar sprays and sterile insect releases. Only a single action may be implemented in any given week. Black depicts a myopic policy, $H = 0$, while red and blue depict lookahead policies with planning horizons of $H = 1$ and $H = 2$, respectively where H is measured in weeks. Costs have no units and adult medfly mortality from insecticide exposure decays at instantaneous rate $\alpha = 0.175$. The

population was initiated with 10,000 mated adult female medfly and a 1:1 sex ratio was assumed.

The contrast between persistent ($\alpha = 0.175$) and transient ($\alpha = 0.7$) insecticides is further explored in Figure 6, which illustrates the action sequences for each policy shown in Figures 4 and 5. For a persistent spray, the sequential application of sprays does not occur: spraying events are spread across the season and may be interspersed with sequential sterile releases (Fig.6a-c). The inverse is true of a transient spray: releases are periodic and, when using a lookahead policy, spread across the season, while sprays are applied sequentially (Fig.6d-f).

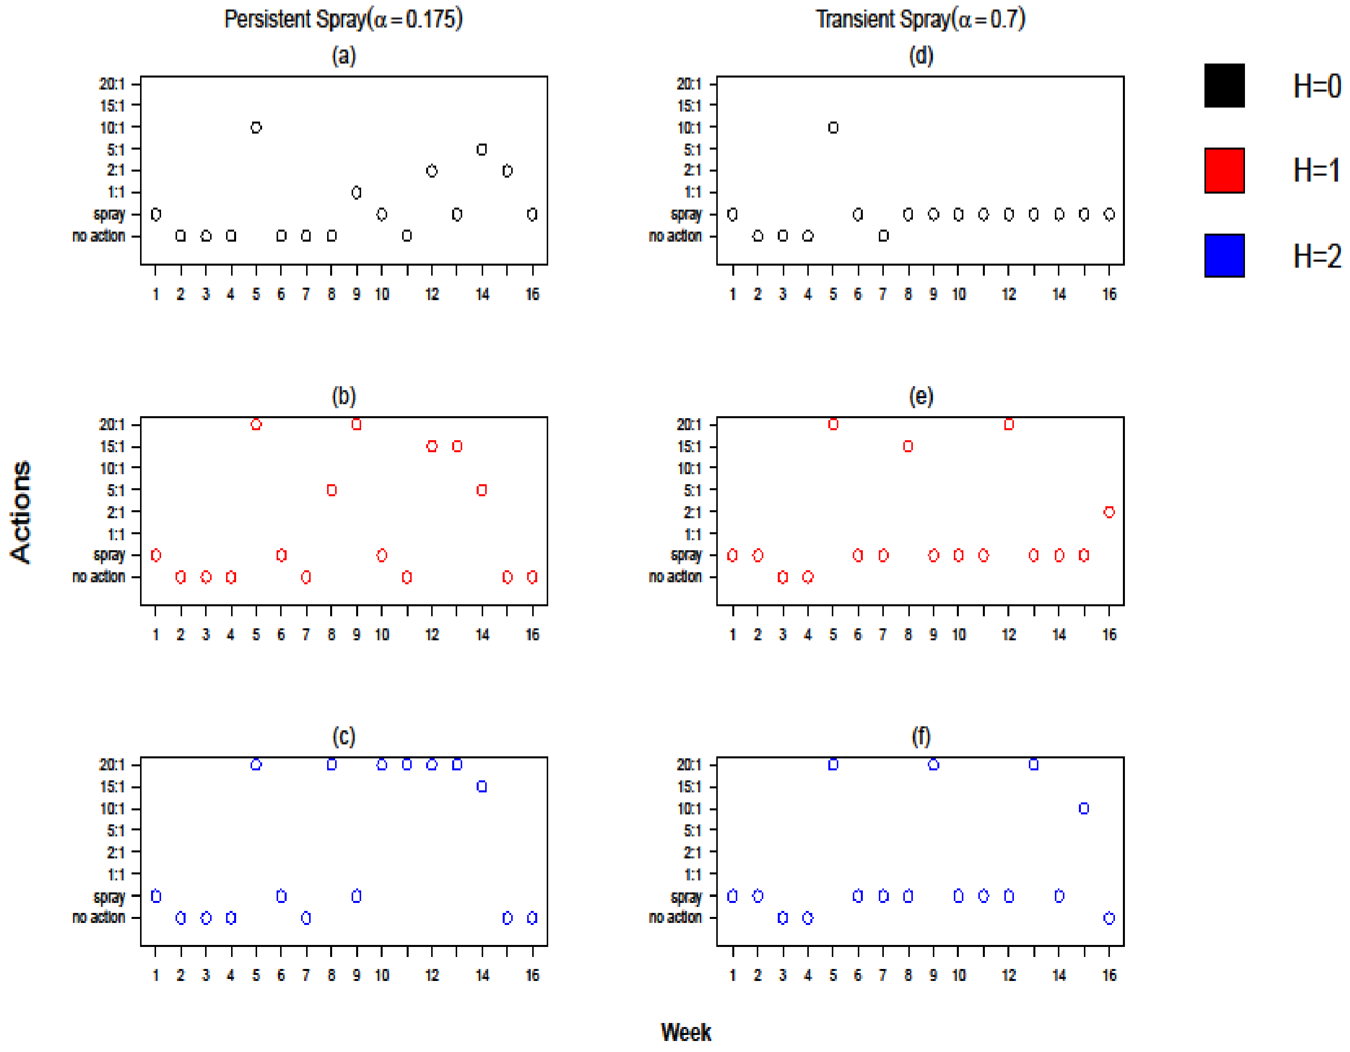

Figure 6*:* Sequences of implemented actions for three alternative policies for the management of a simulated medfly population over 16 weeks. For all panels the full range of alternative actions available in each decision period are listed on the y-axis and costs from larval feeding accrue at instantaneous rate $\gamma = 0.01$. Adult medfly mortality from insecticide exposure decays at instantaneous rate $\alpha = 0.175$ (a persistent insecticide) in the left column (a-c) and $\alpha = 0.7$ (a transient insecticide) in the right column (d-f). Black points indicate actions implemented by a myopic policy ($H = 0$) while red and blue points illustrate actions implemented by lookahead policies with

planning horizons of $H = 1$ and $H = 2$, respectively where H is measured in weeks. The population was initiated with 10,000 mated adult female medfly and a 1:1 sex ratio was assumed.

Lookahead policies and sub-optimality

For the pest management problem considered here a lookahead policy always outperforms a myopic policy but actions suggested by a lookahead policy are optimal only with respect to the lookahead model which derived them and are not guaranteed to constitute optimal actions with respect to the base model. Figure 7 shows the accumulation of costs over time for three lookahead policies ($H = \{1,2,3\}$) and three levels of insecticide persistence ($\alpha = \{0.175,0.35,0.7\}$) applied to medfly populations with instantaneous damage rates of $\gamma = \{0.001,0.01\}$. For any given set of conditions, the most basic lookahead policy, $H = 1$, is always outperformed by an alternative lookahead policy with a longer planning horizon. However, extending the planning horizon is not a consistent guarantee of improved performance: a lookahead policy with a planning horizon of $H = 3$ underperforms relative to a policy with a shorter horizon of $H = 2$ in Fig.7 a,d and e while it provides additional cost reductions in Fig. 7b,c and f. This discrepancy can be understood in terms of the distribution of costs over time for each scenario and the implementation of the lookahead policy itself. When the costs of an error are concentrated in the near-term and insecticides offer effective knockdown (Fig.7d,e) such that sterile releases can be effectively evaluated over short time windows (Fig.7a) then extending the planning horizon to $H = 3$ produces action sequences less effective than those for $H = 2$. Actions taken in solving a lookahead model are intended to improve the decision made in the current time period. Extending the time horizon too far when costs/benefits are concentrated in the near-term future may, therefore, lead to decisions being based upon information that is less relevant to the current period and actions that are optimal with respect to a future that never comes to pass. However, for instances where the costs of a poor decision are concentrated in the future (Fig.7b) or sprays provide ineffective knockdown then extending the planning horizon is beneficial.

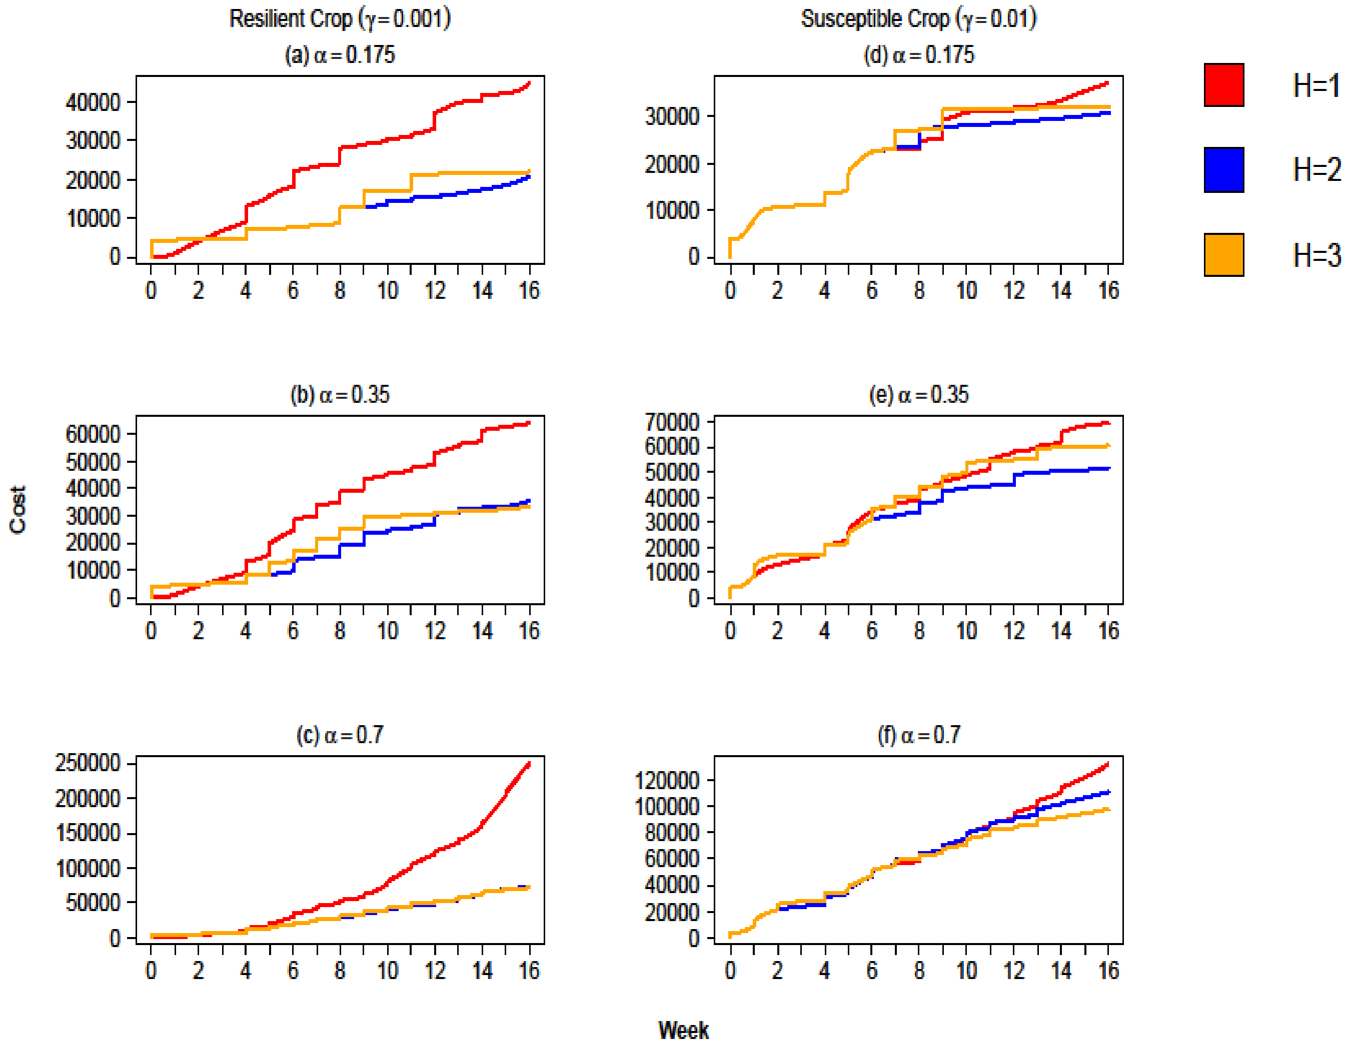

Figure 7*:* Accumulation of costs over time for three alternate lookahead policies attempting to manage a hypothetical medfly population (initiated with 10,000 mated

adult females) over a 16-week season using insecticidal foliar sprays and sterile insect releases. Line colour indicates the extent of the planning horizon in weeks utilised by a given lookahead policy: red for $H = 1$, blue for $H = 2$ and orange for $H = 3$. Costs from larval feeding accrue at instantaneous rate $\gamma = 0.001$ (a resilient crop) in the left column (a-c) and $\gamma = 0.01$ (a susceptible crop) in the right column (d-f). Adult medfly mortality from insecticide exposure decays at instantaneous rate $\alpha = 0.175$ (a persistent insecticide) in the top row (a,d), $\alpha = 0.35$ in the middle row (b,e) and $\alpha = 0.7$ (a transient insecticide) in the bottom row (c,f).

Discussion

We approached the management of a hypothetical continuously reproducing, stage-structured pest population inspired by the Mediterranean fruit fly using a lookahead policy that identifies appropriate actions by solving a temporally truncated approximation of the full problem. A lookahead policy consistently returned action sequences with superior integration of both self-limiting insect releases and foliar sprays than those of a myopic policy (*e.g.* Fig.1), even in instances where effective knockdown was difficult to achieve (Fig.4) or insecticide residues were persistent and potentially threatened self-limiting insects (Fig.5). This is congruent with the findings of similar models (Nicol *et al.* 2010; Nicol & Chadès 2011) and a promising indicator that there is scope for lookahead policies to provide a valuable tool for the management of populations with state spaces that are difficult to capture within a traditional dynamic programming framework.

While lookahead policies always outperformed a myopic policy, the cost-efficiency of a given lookahead policy was determined by the distribution of costs over time relative to the extent of the planning horizon H (Fig.1, Fig.2, Fig.4, Fig.5, & Fig.7). Recall that costs arriving outside of the planning horizon cannot be accounted for by an unmodified lookahead policy (Powell 2011, 2014). When the costs resulting from a poor decision are mostly deferred into the future (γ is low) extending the planning horizon was consistently beneficial. In contrast, extending the planning horizon too far when costs were concentrated in the near-term future (γ is high) can be counter-productive (Fig.7d, e). That is, incorporating the values of more distant states into the decision process when costs are imminent can produce lower-quality decisions. This could likely be remediated by the introduction of a discount rate. Discounting of future costs would devalue the terminal states of a given lookahead model, placing greater emphasis upon the short-term behaviour of the system while still potentially enabling any benefits to a longer horizon to persist (Nicol *et al.* 2010). The effects of insecticide persistence can be understood in similar terms. Inefficient knockdown effectively emphasises the costs of control in the short-term relative to the long-term costs accruing from larvae which is equivalent to displacing the costs of errors into the future. Thus, transient sprays require longer planning horizons to effectively evaluate the benefits of an action, especially when larvae increase costs rapidly (Fig.7f). The opposite reasoning applies when spray persistence is increased; the costs of control are deemphasised and costs from the pest dominate decision-making.

The sensitivity of lookahead policy performance to the degree of overlap between the chosen planning horizon and the distribution of costs over time has implications for the further application of lookahead policies to the management of insect pest populations. The extent of the planning horizon will need to be chosen with reference to the rate at which costs from the damaging life stage(s) accrue and the life history of the pest. The former determines how many insects will be deemed too many, and the latter determines both the duration of the damaging stage(s) and the rapidity with which additional individuals are recruited into the damaging stage. Furthermore, consideration will need to be given to how the negative effects of the damaging life stage accumulate. For example, the timescale for the negative effects of an agricultural pest will be intimately connected to the duration of the crop-damaging life stage but, if the pest also vectors pathogens then negative effects may persist beyond the lifetime of any one pest. Similar issues apply to the management of insect vectors of human pathogens, such as mosquitoes, where the benefits of management are frequently expressed in terms of the number of disease cases averted (Alphey, Alphey & Bonsall 2011). For such problems, a longer planning horizon may be required to ensure that decisions at least partially account for the benefits of averting additional disease cases or, where the costs of the disease and the rate of spread are well-understood and predictable, a hybrid lookahead policy (Powell 2014; Powell & Meisel 2016) could be employed using a cost function approximation to summarize the costs of cases arising beyond the planning horizon.

For clarity, we assumed that sterile releases comprised self-limiting insects with an early-acting bi-sex lethal effect: the simplest form of sterile release to model. However, self-limiting methods are diverse with respect to both their mode of action and the timing of any lethal effects (Alphey *et al.* 2010; Bourtzis *et al.* 2016; Leftwich, Bolton & Chapman 2016; Harvey-Samuel, Ant & Alphey 2017). For example, male-selecting self-limiting strains for the suppression of medfly have been developed and are more likely to be available for medfly management in the near future than a bi-sex lethal strain (Gong *et al.* 2005; Morrison *et al.* 2009; Leftwich *et al.* 2014). For such male-selecting strains, given that the male offspring from self-limiting matings survive, the suppression potential of a male-selecting strain is reduced relative to that of a bi-sex self-limiting strain (Alphey, Bonsall & Alphey 2009). Furthermore, females inheriting the lethal allele are not killed until the pupal stage (Gong *et al.* 2005; Leftwich *et al.* 2014). This may increase the length of the planning horizon required for a lookahead policy to usefully evaluate such a release. However, while we did not here consider the potential for the evolution of resistance to insecticides (Vontas *et al.* 2011), male-selecting insect releases may offer additional utility in instances where the evolution of resistance to other pest management tools is considered a threat to future management success (Alphey *et al.* 2007; Alphey, Bonsall & Alphey 2009; Harvey-Samuel *et al.* 2015).

In addition to their ability to at least partially capture the downstream effects of contemporary decisions, there are four additional reasons to recommend lookahead policies for ecological management problems of high dimensionality. First, they do not depend upon functional approximations and so readily accommodate time-dependent processes and non-stationary parameters or functions (Powell 2014). For example, the current model could be extended to model changes in the vulnerability of fruits to attack over the course of a season (Zhang & Swinton 2012). Second, by explicitly modelling future state transitions lookahead policies naturally account for non-separable interactions between state variables which would otherwise be difficult to capture with an approximation based upon a parametric function (Powell 2014). Third, they readily incorporate forecasts of future events, even where the forecasts are updated or revised over time (Powell 2014). Taken together with the first property, an indifference to non-stationarity, this renders the combined online application of lookahead policies and adaptive management to information-dense problems a possibility. Lookahead policies are agnostic to the particular set of equations and parameters used to model the focal population and so can accommodate the updates required to account for *e.g.* structural uncertainty (Boettiger, Mangel & Munch 2015;

Chadès *et al.* 2017). Finally, the time required to solve a problem using a lookahead policy is independent of the magnitude or complexity of the state space: state spaces of arbitrary complexity can theoretically be accommodated (Nicol *et al.* 2010; Powell, Simao & Bouzaiene-Ayari 2012). The runtime is instead determined by the extent of the planning horizon, the number of alternative actions and the range of random outcomes. While we considered a deterministic problem, lookahead policies can account for stochasticity (Nicol *et al.* 2010; Nicol & Chadès 2011; Powell 2011; Powell & Meisel 2016) and, utilising methods from stochastic programming, can accommodate continuously valued random variables (Powell 2014). The structural agnosticism of lookahead policies entails that it would also be possible to extend this framework to include delays between the selection and implementation of an action.

With respect to our chosen example of a medfly suppression problem, we have presented a relatively basic representation for ease of understanding. However, the flexibility of lookahead policies offers the potential for substantive refinements which would augment the utility of the model to pest management professionals. Given a specific crop and locale, reliable estimates for both the costs of actions and potential costs of medfly induced losses could be utilised. Furthermore, restrictions could be placed upon the timing and frequency of insecticide applications, particularly in relation to the timing of the harvest. The susceptibility of the crop to medfly could be allowed to vary over the course of the season, potentially resulting in periods in which medfly are present but cannot yet attack the crop. Additionally, forecasts of pertinent environmental variables such as temperature or precipitation and how they interact with the ecology of the pest, the growth of the crop and the utility of available actions could be incorporated. Alternatively, where such detailed models are unwarranted or the data for their formulation is lacking then the basic framework provided here could instead be built upon to provide intuition regarding how, for example, different forms of transgenic release (see above) could be utilised in conjunction with other management tools.

However, lookahead policies also carry several key disadvantages. In its basic form, a lookahead policy is a brute-force method, proceeding by rote application of the same procedure to each decision period, which is indifferent to the structure of the problem and so cannot utilise advantageous structural features to accelerate computation or improve decision quality as, for example, a value function approximation could (Powell, Simao & Bouzaiene-Ayari 2012; Powell 2014). The requirement that a new lookahead model be solved for each time period may also slow computation relative to alternative methods. Lookahead policies will only explore states that are directly connected to the input state and will not generalise beyond these states without augmentation. Thus, for infinite horizon problems where the global value of a given state is pertinent it may be necessary to augment a lookahead policy with a means of more fully enumerating the state space (for example, by incorporating machine learning methods (Nicol & Chadès 2011)). Finally, we reiterate that a decision which is optimal with respect to a lookahead model is unlikely to be optimal with respect to the base model (Powell 2014).

Lookahead policies are imperfect solutions to imperfect problems. They are not a replacement for exact methods or myopic heuristics but an extension: they represent an additional tool to be applied when we are required to address management questions beyond the scope of exact methods. They are also not an invitation to complicate a problem for its own sake: a given problems state should remain a “minimally dimensioned function of history” (Powell 2014) and ADP methods should be invoked where this minimally dimensioned function exceeds the bounds of exact methods. If we remain aware of their shortcomings, ADP methods such as lookahead policies broaden the range of questions which can be asked and could represent a valuable tool for ecologists seeking to work with managers to produce good, if sub-optimal, solutions to detail-rich problems.

Data Availability

The R code generated by this project has been uploaded to the Open Science Framework: https://osf.io/h6cwa/.

Chapter 4: Insect pest control, approximate dynamic programming and the management of the evolution of resistance

Sean C.Hackett[1], Michael B.Bonsall[1,2]

[1] *Mathematical Ecology Research Group, Department of Zoology, South Parks Road, Oxford, OX1 3PS, UK*

[2] *St. Peter's College, New Inn Hall Street, Oxford, OX1 2DL, UK*

Abstract

Ecological decision problems, such as those encountered in agriculture, often require managing conflicts between short-term costs and long-term benefits. Dynamic programming is an ideal method for optimally solving such problems but agricultural problems are often subject to additional complexities which produce state spaces intractable to exact solutions. In contrast, lookahead policies, a class of approximate dynamic programming (ADP) algorithm, may attempt to solve problems of arbitrary magnitude. However, these algorithms focus on a temporally truncated caricature of the full decision problem over a defined planning horizon and as such are not guaranteed to suggest optimal actions. Thus, lookahead policies may offer promising means of addressing detail-rich ecological decision problems but may not be capable of fully utilising the information available to them, especially in scenarios where the best short and long-term solutions may differ. We constructed and applied lookahead policies to the management of a hypothetical, stage-structured, continually reproducing agricultural insect pest. The management objective was to minimise the combined costs of management actions and crop damage over a 16-week growing season. The manager could elect to utilise insecticidal sprays or one of six release ratios of male-selecting transgenic insects. Complicating matters was the expression of insecticide resistance at non-trivial frequencies in the pest population. We assessed the extent to which lookahead policies were able to recognise the potential threat of insecticide resistance and successfully integrate insecticides and transgenic releases to capitalise upon their respective benefits. Lookahead policies were competent at anticipating and responding to ecological and economic information. Policies with longer planning horizons made fewer, better-timed insecticidal sprays and made more frequent transgenic releases which consequently facilitated lower resistance allele frequencies. However, lookahead policies were ultimately inefficient resistance managers, directly responded to resistance only when it was dominant and prevalent. Effective long-term agricultural management requires the capacity to anticipate and respond to the evolution of resistance. Lookahead policies can accommodate all the information pertinent to making the best long-run decision but may lack the perspective to actually do so. However, combining lookahead policies with other ADP methods may remedy this by providing cost-effective frameworks for integrated pest management problems.

Introduction

Insect pests pose substantial challenges to both agriculture and public health through the direct consumption of agricultural produce and the vectoring of plant, livestock and human pathogens. In addressing these challenges, economic and ecological considerations coincide: cost effective management practices must account for the ecological idiosyncrasies of the pest. This constrains the types of intervention which will be the most useful and how the timing of those interventions should be planned. However, the selection pressure imposed by management interventions favours the survival of individuals expressing the capacity to tolerate or avoid control, facilitating the evolution of resistance. This adds an evolutionary component to questions of pest management which complicates the ease with which cost effective long-term suppression can be achieved. There is an implicit discounting between our capacity to manage a pest today and our capacity to manage it in the future. Susceptibility to control is effectively a renewable resource which is depleted when we fail to moderate the strength of selection for resistance by accepting short-term losses in exchange for guaranteed longer-term benefits (REX Consortium 2013; Mitchell & Onstad 2014).

In responding to the economic and ecological pressures of insect pest management it is rarely optimal to confront the focal pest with only a single management tool. For example, sterile insect releases can complement insecticidal sprays to achieve greater overall levels of suppression than when either control is used in isolation: in general, sprays are most effective when employed against a large population while the efficiency of sterile releases can increase as the target population declines (Barclay 2005). However, there has been a diversification of sterile insect methods in recent years, driven by developments in genetic engineering, resulting in novel approaches such as gene-drive systems, male-selecting strains and sex-ratio distortion (Bourtzis *et al.* 2016; Leftwich, Bolton & Chapman 2016; Harvey-Samuel, Ant & Alphey 2017; Alphey & Bonsall 2018). Male-selecting strains are of particular relevance to resistance management. These insects carry a dominant lethal mutation which is benign when expressed in a male insect but deleterious when expressed in a female unless they are fed on a supplemented diet which compensates for the transgene (Alphey & Bonsall 2018). While male-selecting releases have a reduced capacity for population suppression relative to a bi-sex lethal release the survival of the male offspring enables them to act as a source of alleles conferring susceptibility to other control measures such as insecticidal toxins. Thus, there exists the exciting potential for male-selecting releases to complement existing control programmes by permitting a proactive approach to resistance management whilst still contributing to population suppression (Alphey *et al.* 2007; Alphey, Bonsall & Alphey 2009; Harvey-Samuel *et al.* 2015; Zhou *et al.* 2018). However, the near-term benefits of male-selecting releases with respect to the prevention of damage to the crop are sensitive to the timing of lethality. If females are killed early then some damage will be averted relative to a scenario where no insects were released, but when lethality occurs late there will effectively be no immediate reductions in damage (Leftwich, Bolton & Chapman 2016; Alphey & Bonsall 2018). Thus, effective integration of male-selecting releases into pest management programmes necessitates that the potential future benefits of a release can be assessed and weighted against more immediate costs.

The integration of distinct control options for the management of an insect pest population can be conceptualised as a dynamic program, a sequential decision problem where decisions are made recurrently and actions implemented now can influence the options and rewards available in the future (Powell 2014; Hackett & Bonsall 2018). For a given objective and set of actions, the iterative algorithmic technique of dynamic programming will identify the optimal mapping between possible states and actions. The algorithm proceeds by recursively identifying the action which maximises the sum of both current and future benefits for each combination of time and state (Clark & Mangel 2000; Marescot *et al.* 2013). Thus, dynamic programming automatically accounts for the downstream consequences of decisions and has the desirable property of easily accommodating both stochasticity and non-linearity (Marescot *et al.* 2013). However, the range of questions which dynamic programming can address is fundamentally limited by the curse of dimensionality. To operate, the algorithm must be able to compile the solution for each combination of state and time in a lookup table which renders many complex, highly dimensioned problems intractable to this method (Powell 2014). This is unfortunate given the complexity inherent to many contemporary questions in insect pest management, particularly with regards to the delay or prevention of resistance evolution and the use of transgenic insect releases. While methods for reducing the dimensionality of a problem exist (Clark & Mangel 2000) such simplifications come at the expense of realism which may be undesirable when it is the granular details of a problem, such as the effects of insect dispersal (Yakob & Bonsall 2009; Miller & Sappington 2017), that are of greatest interest.

This curse of dimensionality can be circumvented, and the range of questions available for exploration greatly expanded, by utilising approximate dynamic programming (ADP) which provide a toolset for approaching problems of arbitrary complexity at the expense of guaranteed optimality (Powell 2011). That is, where dynamic programming would identify the best action, ADP will suggest a good action. Given that ecological decision problems frequently require the incorporation of numerous environmental, economic and ecological processes this trade-off may be justifiable (Nicol *et al.* 2010; Nicol & Chadès 2011). ADP encompasses four broad classes of policy, where the term policy simply refers to any rule used to pair actions and states, and each policy class sidesteps the limitations of dynamic programming by approximating a distinct aspect of the focal problem (Powell 2014). Cost function approximations (CFAs) append a correction term, which may contain tuneable parameters, to the function being optimised which has been designed to promote actions that are considered favourable from the perspective of the manager (Powell 2014; Powell & Meisel 2016). Value function approximations (VFAs) approximate the downstream benefits of a decision made now using a function which may be either parametric or non-parametric (Powell 2014; Powell & Meisel 2016). A strategy which is optimal with respect to the chosen

VFA is then identified. Policy function approximations do not directly address the focal optimisation problem. Instead, a parametric function depicting a fixed decision rule is specified and the parameter values which return the best valued action sequences are then identified (Powell 2014; Powell & Meisel 2016).

The fourth policy class is the lookahead policy. Lookahead policies are distinguished from the other classes by their lack of reliance upon functional approximations. Instead lookahead policies proceed by sequentially constructing and then explicitly solving a series of temporally truncated approximations of the full problem (Nicol *et al.* 2010; Powell 2014; Powell & Meisel 2016). The actions suggested by a lookahead policy are optimal only with respect to the truncated horizon (termed the planning horizon) over which they are solved and so are unlikely to be optimal with respect to the time horizon of the full problem. However, where the scale of the full problem renders it intractable it may still be possible to identify effective actions by employing a lookahead policy with a sufficiently long planning horizon. While a comparatively brute-force method compared to VFAs and PFAs, lookahead policies have several extremely valuable features: their runtime is independent of the magnitude of the state space, they are indifferent to non-stationary parameters and they can be applied to questions where accurately approximating the downstream value of an action is problematic (Nicol *et al.* 2010; Powell 2011, 2014; Powell & Meisel 2016; Hackett & Bonsall 2018).

Lookahead policies potentially have much to offer to the study of ecological decision problems such as those in insect pest management. Their agnosticism to the complexity of the state space allows for the specification of models as complex as required, enabling them to represent granular details which might otherwise need to be summarised, streamlined or omitted (Nicol *et al.* 2010; Powell 2014). For example, lookahead policies can accommodate details such as how both wild-type and released insects distribute themselves in space and how management actions unfold across this space (Yakob & Bonsall 2009; Garcia *et al.* 2016), the timing and nature of insect dispersal and mating (Sudo *et al.* 2018), variation in crop susceptibility to insect pests over time (Zhang & Swinton 2012) and the influence of environmental variables such as temperature and precipitation upon the population dynamics of the pest and the success of management actions (Shoemaker 1982). However, in addition to their near-term consequences for pest population dynamics and the performance of management actions, these factors also have longer-term implications for how resistance to control will evolve and spread (Téllez-Rodríguez *et al.* 2014; Garcia *et al.* 2016; Miller & Sappington 2017; Watkinson-Powell & Alphey 2017; Sudo *et al.* 2018) which may not be properly accounted for by a basic lookahead policy. Furthermore, for controls such as male-selecting releases which may require the acceptance of a short-term cost in exchange for longer-term gains lookahead policies may be unduly biased against utilising such an option by their limited temporal scope. The abridged nature of lookahead policies may also limit their capacity to adequately delay or reverse the evolution of resistance.

Here we apply lookahead policies to the management of a hypothetical agricultural insect pest with two life stages (juvenile and adult) and continuous reproduction. We consider the extent to which lookahead policies with planning horizons of different lengths are able to combine the use of insecticidal sprays and either early or late acting male-selecting transgenic releases to achieve population suppression. Additionally, we explore the capacity for lookahead policies to anticipate and respond to the threat of resistance evolution by initialising the pest population using insecticide resistance allele frequencies which would be considered problematic in practice. We consider the influence of demographic factors on how the lookahead policies integrate actions by repeating model analysis for two levels of natural juvenile mortality. We find that simple lookahead policies can ably anticipate and plan for the dynamics of the pest population and that they are able to recognise the benefits of both late and early acting male-selecting releases. However, lookahead policies tended to respond to insecticide resistance only after insecticide efficacy was significantly impaired.

Methods

Overview

We explore the utility of lookahead policies for the integration of transgenic insect releases (comprising exclusively male insects carrying a dominant female lethal mutation, also referred to as a male-selecting allele) and insecticidal sprays in the suppression of a hypothetical agricultural insect pest exhibiting non-trivial levels of insecticide resistance. Furthermore, we highlight the implications of a cost-minimising approach for effective resistance management. That is, the presented model is not a resistance management model but a pest management model with pesticide resistance. Outcomes are contrasted for both early and late acting lethal transgenes. The possible influence of pest demography upon the selection of actions is considered by repeating all simulations for a pest population with a greatly reduced juvenile mortality rate.

Pest demography and genetics

We consider a hypothetical continuously reproducing agricultural insect pest population with a 1:1 sex ratio and a life cycle that can be broadly sub-divided into two distinct stages: a non-reproductive juvenile stage which directly damages the crop and a benign reproductive adult stage. Only adult insects are affected by management actions;

juveniles cannot be mated to transgenic males and are rendered intractable to insecticides by being sequestered within, for example, the fruit or stem of the crop. An insect's genotype is determined by two independent genetic loci, each with two alternative alleles. One locus determines susceptibility to insecticide with alleles S (conferring susceptibility) and r (conferring resistance). The second locus represents the target for the lethal transgene with alleles w (denoting wild type) and L (the dominant lethal transgene). The transgene is assumed to function perfectly- no female insects inheriting the L allele survive to adulthood. Thus, male insects may possess any one of the following genotypes $g = \{SSww, Srww, rrww, SSwL, SrwL, rrwL\}$. Adult females may only possess the genotypes $g = \{SSww, Srww, rrww\}$. When the L allele is late acting, juvenile females with the genotypes $g = \{SSwL, SrwL, rrwL\}$ are viable but are killed before they attain maturity. Irrespective of the timing of lethality, released males always possess the genotype $SSLL$. Therefore, there are seven male genotypes to which adult females may be mated resulting in 21 distinct adult female stages f^{i*j} where i denotes the maternal genotype and j the paternal genotype. Mating is instantaneous at the point of maturation and occurs uniformly and at random. All male genotypes are assumed to be equally competitive. Once mated, females do not mate again and produce juveniles at instantaneous rate λ.

Equal sex ratio allows explicit tracking of only female numbers for all three ww genotypes and only male numbers for wL genotypes; the abundance of ww males is equivalent to the abundance of ww females and vice versa for female wL juveniles when considering late acting releases. Irrespective of their genotype, juveniles mature to adults at instantaneous rate σ and are subject to natural mortality at instantaneous rate

μ_j. Thus, changes in the abundance of each juvenile stage j_{a^g} (where a denotes sex and g indicates genotype) are determined as:

$$\begin{aligned}\frac{dj_{f^{SSww}}}{dt} = {} & \lambda(f^{SSww*SSww}(t) + 0.5f^{SSww*Srww}(t) + 0.5f^{SSww*SSwL}(t) \\ & + 0.25f^{SSww*SrwL}(t) + 0.5f^{Srww*SSww}(t) + 0.25f^{Srww*Srww}(t) \\ & + 0.25f^{Srww*SSwL}(t) + 0.125f^{Srww*SrwL}(t)) - \mu_j j_{f^{SSww}}(t) \\ & - \sigma_j j_{f^{SSww}}(t) \qquad [1]\end{aligned}$$

$$\begin{aligned}\frac{dj_{f^{Srww}}}{dt} = {} & \lambda(0.5f^{SSww*Srww}(t) + f^{SSww*rrww}(t) + 0.25f^{SSww*SrwL}(t) \\ & + 0.5f^{SSww*rrwL}(t) + 0.5f^{Srww*SSww}(t) + 0.5f^{Srww*Srww}(t) \\ & + 0.5f^{Srww*rrww}(t) + 0.25f^{Srww*SSwL}(t) + 0.25f^{Srww*SrwL}(t) \\ & + 0.25f^{Srww*rrwL}(t) + f^{rrww*SSww}(t) + 0.5f^{rrww*Srww}(t) \\ & + 0.5f^{rrww*SSwL}(t) + 0.25f^{rrww*SrwL}(t)) - \mu_j j_{f^{Srww}}(t) \\ & - \sigma_j j_{f^{Srww}}(t) \qquad [2]\end{aligned}$$

$$\begin{aligned}\frac{dj_{f^{rrww}}}{dt} = {} & \lambda(0.25f^{Srww*Srww}(t) + 0.5f^{Srww*rrww}(t) + 0.125f^{Srww*SrwL}(t) \\ & + 0.25f^{Srww*rrwL}(t) + 0.5f^{rrww*Srww}(t) + f^{rrww*rrww}(t) \\ & + 0.25f^{rrww*SrwL}(t) + 0.5f^{rrww*rrwL}(t)) - \mu_j j_{f^{rrww}}(t) \\ & - \sigma_j j_{f^{rrww}}(t) \qquad [3]\end{aligned}$$

$$\begin{aligned}\frac{dj_{m^{SSwL}}}{dt} = {} & \lambda(0.5f^{SSww*SSwL}(t) + 0.25f^{SSww*SrwL}(t) + f^{SSww*SSLL}(t) \\ & + 0.25f^{Srww*SSwL}(t) + 0.125f^{Srww*SrwL}(t) + 0.5f^{Srww*SSLL}(t)) \\ & - \mu_j j_{m^{SSwL}}(t) - \sigma_j j_{m^{SSwL}}(t) \qquad [4]\end{aligned}$$

$$\begin{aligned}\frac{dj_{m^{SrwL}}}{dt} = {} & \lambda(0.25f^{SSww*SrwL}(t) + 0.5f^{SSww*rrwL}(t) + 0.25f^{Srww*SSLL}(t) \\ & + 0.25f^{Srww*SrwL}(t) + 0.25f^{Srww*rrwL}(t) + 0.5f^{Srww*SSLL}(t) \\ & + 0.5f^{rrww*SSwL}(t) + 0.25f^{rrww*SrwL}(t) + f^{rrww*SSLL}(t)) - \mu_j j_{m^{SrwL}}(t) \\ & - \sigma_j j_{m^{SrwL}}(t) \qquad [5]\end{aligned}$$

$$\begin{aligned}\frac{dj_{m^{rrwL}}}{dt} = \lambda(&0.125 f^{Srww*SrwL}(t) + 0.25 f^{Srww*rrwL}(t) + 0.25 f^{rrww*SrwL}(t) \\ &+ 0.5 f^{rrww*rrwL}(t)) - \mu_j j_{m^{rrwL}}(t) \\ &- \sigma_j j_{m^{rrwL}}(t) \qquad [6]\end{aligned}$$

All adult stages are subject to both natural mortality μ_a and mortality from insecticide residues φ. The extent to which a given genotype (at the S/r locus) g tolerates exposure to insecticides is determined by δ^g where $0 \leq \delta^g \leq 1$ with $\delta^g = 0$ indicating complete resistance (insects of genotype g are unaffected by exposure to insecticides) and $\delta^g = 1$ indicating complete susceptibility (insects of genotype g experience the full effect of insecticicde exposure). For homozygotes, the value of δ^g is entirely determined by the insect's genotype. For Sr heterozygotes, δ^g is derived from the tolerances of homozygotes and the dominance of resistance h as

$$\delta^{Sr} = (1-h)\delta^{SS} + h\delta^{rr} \qquad [7]$$

where $0 \leq h \leq 1$ with $h = 0$ indicating fully recessive resistance and $h = 1$ indicating dominant resistance. We assume that resistance is not associated with any pertinent fitness costs.

The abundance of adult males heterozygous for the L allele is determined by iterative application of the following equation to each of the three possible transgenic heterozygote genotypes $g = \{SSwL, SrwL, rrwL\}$:

$$\frac{dm^g}{dt} = \sigma_j j_{m^g}(t) - \left(\mu_a + \delta^g \varphi(t)\right) m^g(t). \qquad [8]$$

Given the equal sex ratio, the total abundances of $SSww$, $Srww$ and $rrww$ adult males are derived from the abundance of adult females with the corresponding genotype as $m^j = \sum_i f^{i*j}$. The total number of males available for mating M is found as the sum of all male genotypes excluding released males, m^{SSLL}. Thus, adult females are allocated between mated stages according to

$$\frac{df^{i*j}}{dt} = \left(\frac{m^j(t)}{M(t)}\right)\sigma_j j_{f^a}(t) - \left(\mu_a + \delta^i \varphi(t)\right) f^i(t) \quad [9]$$

where the indices i and j denote male and female genotype respectively.

Actions and costs

The management objective is to minimise the cumulative costs of both pest induced losses and pest management over the duration of a single season. The season is comprised of $W = 16$ discrete weeks (a broad representation of a temperate crop season) and a single action may be implemented at the start of each week w. Actions are selected with reference to the current state of the system S_w, which is defined as the minimally dimensioned function of history required to model the transitions of the system from a given point in time onwards (Powell 2016). Once implemented, an action instantaneously updates the pertinent control variable and the manager experiences a point cost. Insecticide application is uniform across the landscape and sets φ, the instantaneous rate of mortality from insecticide, to its highest value φ_{max} while incurring the cost c_{spray}. Subsequent to application, pesticide toxicity decays exponentially

$$\frac{d\varphi}{dt} = -\alpha\varphi(t) \qquad [10]$$

where α determines the persistence of the pesticide. For transgenic releases, six alternative release ratios are available (1:1, 2:1, 5:1, 10:1, 15:1 and 20:1). Thus, the cost of a release c_{rel} varies with the abundance of non-transgenic males and the release ratio. Once released, transgenic male abundance declines according to

$$\frac{dm^{SSLL}}{dt} = -\left(\mu_a + \delta^{SS}\varphi(t)\right)m^{SSLL}(t) \quad [11]$$

which assumes that the transgenic males are equivalent to wild-type males and experience no ill effects from carrying two copies of the lethal construct.

Costs inflicted by juvenile insects accrue as the product of total juvenile abundance (across both sexes and all genotypes) and the rate parameter γ which proxies the susceptibility of the crop to feeding damage, with greater values implying greater susceptibility. Thus, the increase in costs caused by larval feeding for late acting releases is found as

$$\frac{dc}{dt} = \gamma\left(2\sum_i {j_f}^i + 2\sum_j j_{m^j}\right) \quad [12]$$

where j_f refers to the abundance of female ww homozygotes ($i = \{SSww, Srww, rrww\}$) and j_m denotes the abundance of male wL heterozygotes ($j = \{SSwL, SrwL, rrwL\}$). These values are doubled for late-acting lethality to account for the abundance of the opposite sex. For early acting lethality, there are no female wL juveniles and it is necessary to double only j_f.

Equations [1]-[12] define the transition function for the model and how the state of the system changes over time in response to management actions (Powell 2014; Powell & Meisel 2016). Given equations [1-12] and the definition of the state of the system as the minimally dimensioned function of history required to model a system from a given point in time onwards, the state of the current model at the onset of any given week w may be formally summarised as

$$S_w = \{j_{f^{SSww}}(t), j_{f^{Srww}}(t), j_{f^{rrww}}(t), j_{m^{SSwL}}(t), j_{m^{SrwL}}(t), j_{m^{rrwL}}(t), f^{SSww*SSww}(t), f^{SSww*Srww}(t), f^{SSww*rrww}(t), f^{SSww*SSwL}(t), f^{SSww*SrwL}(t), f^{SSww*rrwL}(t), f^{SSww*SSLL}(t), f^{Srww*SSww}(t), f^{Srww*Srww}(t), f^{Srww*rrww}(t), f^{Srww*SSwL}(t), f^{Srww*SrwL}(t), f^{Srww*rrwL}(t), f^{Srww*SSLL}(t), f^{rrww*SSww}(t), f^{rrww*Srww}(t), f^{rrww*rrww}(t), f^{rrww*SSwL}(t), f^{rrww*SrwL}(t), f^{rrww*rrwL}(t), f^{rrww*SSLL}(t), m^{SSwL}(t), m^{SrwL}(t), m^{rrwL}(t), m^{SSLL}(t), \varphi(t), c(t)\}. \quad [13]$$

Given the stated objective of cost minimisation, for any given week the optimal action a_w^* is that which satisfies

$$a_w^* = \arg\min_{a_w}\{C(S_w, a_w) + \sum_{w'=w+1}^{W} (C(S_{w'}, a_{w'}))\} \quad [14]$$

where $C(S_w, a_w)$ is the contribution function which determines the full cost accruing between week w and week $w + 1$ as

$$C(S_w, a_w) = c_a + \int_{t'=0}^{t'=t_w} 2\gamma l(t') \quad [15]$$

where c_a denotes the point cost for the action implemented (c_{spray}, c_{rel}, or 0 in the case of inaction) at the start of week w and $2\gamma l(t')$ is equivalent to equation [12]. For a sufficiently small state space equation [14] could be solved exactly using dynamic programming (Marescot *et al.* 2013). The solution would be compiled in a lookup table detailing the optimal combinations of action and state for each permissible instance of the state for each decision period (Clark & Mangel 2000; Marescot *et al.* 2013).

However, the state space for the current problem is both large and continuously valued (equation [13]) necessitating the use of alternative methods.

Decision making

We choose to approach this problem using a lookahead policy a class of approximate dynamic programming algorithm which is indifferent to the dimensionality of the state space and forgoes attempts to solve equation [14], instead solving a series of temporally truncated approximations of equation [14] to suggest actions which, while sub-optimal, may be "good enough" (Nicol *et al.* 2010; Powell 2011, 2014). The problem we seek to solve, equation [14] is termed the base model while the truncated approximations solved in its place are termed lookahead models (Powell 2014; Powell & Meisel 2016). A lookahead policy proceeds forwards in time, from the current week w up to the time horizon W. For each week, a best action a_w is identified by solving a lookahead model over an interval of time spanning from w, the current week, up to some $H < W$, where H is a planning horizon which determines the number of decision periods into the future the model will "look ahead", as follows

$$a_w = \arg\min_{a_w}\{C(S_w, a_w) + \sum_{w'=w+1}^{w'=w+H} (C(\tilde{S}_{ww'}, \tilde{a}_{ww'}))\}. \quad [16]$$

where tildes distinguish values pertaining to the lookahead model from those in the base model. Thus, $\tilde{S}_{ww'}$ is the state experienced during week w' when solving a lookahead model to identify an action for week w. Equation [16] states that the best action, with respect to the lookahead model, is that which minimises cumulative costs over the interval $(w, \dots, w + H)$. This is not guaranteed to be the same action which would minimise costs over the full interval of the base model $(w, .., W)$ given the lookahead models incomplete knowledge of the downstream effects of an action.

Solving equation [16], a lookahead model, is not equivalent to solving equation [14], the base model. However, by solving equation [16] for a planning horizon sufficient length to capture important behaviours, actions can be suggested that will outperform those recommended by a purely myopic policy which minimises only the immediate costs of an action (Nicol *et al.* 2010; Powell 2011, 2014; Powell & Meisel 2016; Hackett & Bonsall 2018). Actions taken within the lookahead model are effectively exploratory, serving only to improve the quality of the current decision. Thus, once a_w, the action to be taken now, has been identified states and actions experienced as part of the lookahead model are discarded. Action a_w is then implemented and the model transitions to S_{w+1} where a new lookahead model is constructed and solved over the interval $(w+1, \dots, w+1+H)$ to return action a_{w+1}.

Simulations

Simulations were performed using code written in R version 3.3.3. The pest population is initiated with 10,000 mated adult females, 10,000 wild-type adult males and one of three insecticide resistance allele frequency of $r_0 = \{0.3, 0.5, 0.7\}$. The initial frequency of the resistance allele is then used to divide the founding females between the *SSww*, *Srww*, and *rrww* genotypes. The lethal transgene is assumed to be absent at start of the season. Given a 1:1 sex ratio, females are subsequently allocated to mated classes based upon the relative abundance of each genotype at the S/r locus. Performance was evaluated for a myopic policy ($H = 0$) and three lookahead policies $H = \{1,2,3\}$. For each policy, simulations were repeated for three levels of genetic dominance of resistance $h = \{0, 0.5, 1.0\}$ (describing recessive, additive and dominant resistance respectively), two levels of natural juvenile mortality $\mu_j = \{0.2, 0.05\}$ (which we term the default and robust populations, respectively) and for both early and late acting releases. Continuous time dynamics were approximated using a 4th-order Runge-Kutta (RK4) solver with a resolution of $t_w = 700$ time increments. Values for all model parameters are summarised in Table 1.

Table 1*:* Summary of model parameters and their meanings.

Parameter	**Meaning**	**Default value**
λ	Female eggs laid per adult female per day	1.0
μ_j	Instantaneous juvenile natural mortality rate.	Default =0.2, Robust =0.05
μ_a	Instantaneous adult natural mortality rate	0.05
σ_j	Instantaneous juvenile maturation rate	0.1
h	Dominance of resistance	0,0.5,1.0
δ^{SS}	Susceptibility of a *SS* adult to insecticide	1.0
δ^{Sr}	Susceptibility of a *Sr* adult to insecticide	$(1-h)\delta^{SS}+h\delta^{rr}$
δ^{rr}	Susceptibility of a *rr* adult to insecticide	0
c^{spray}	Point cost of insecticide application	4000
c^{ste}	Point cost per self-limiting male released	0.003
α	Instantaneous insecticide decay rate	0.35
φ_{max}	Maximum Instantaneous insecticide mortality rate	0.9

γ	Instantaneous rate at which costs/losses from juvenile feeding accrue	0.01
W	Season length (weeks)	16
t_w	Resolution of the 4th order RK4 solver	700

Results

Default Population

Myopic policies ($H = 0$) were consistently outperformed by lookahead policies ($H \geq 1$) across all scenarios with respect to cost minimisation. Additionally, for the default pest population ($\mu_j = 0.2$) lookahead policies with longer planning horizons outperformed those with shorter ones irrespective of the initial resistance allele frequency r_0 and the dominance of resistance h. For the same initial conditions lookahead policies with longer planning horizons suggest actions which provide better cost reductions and lower resistance allele frequencies. This is illustrated in Figure 1 which shows the change in the frequency of the insecticide resistance allele over the duration of the season for all four policies.

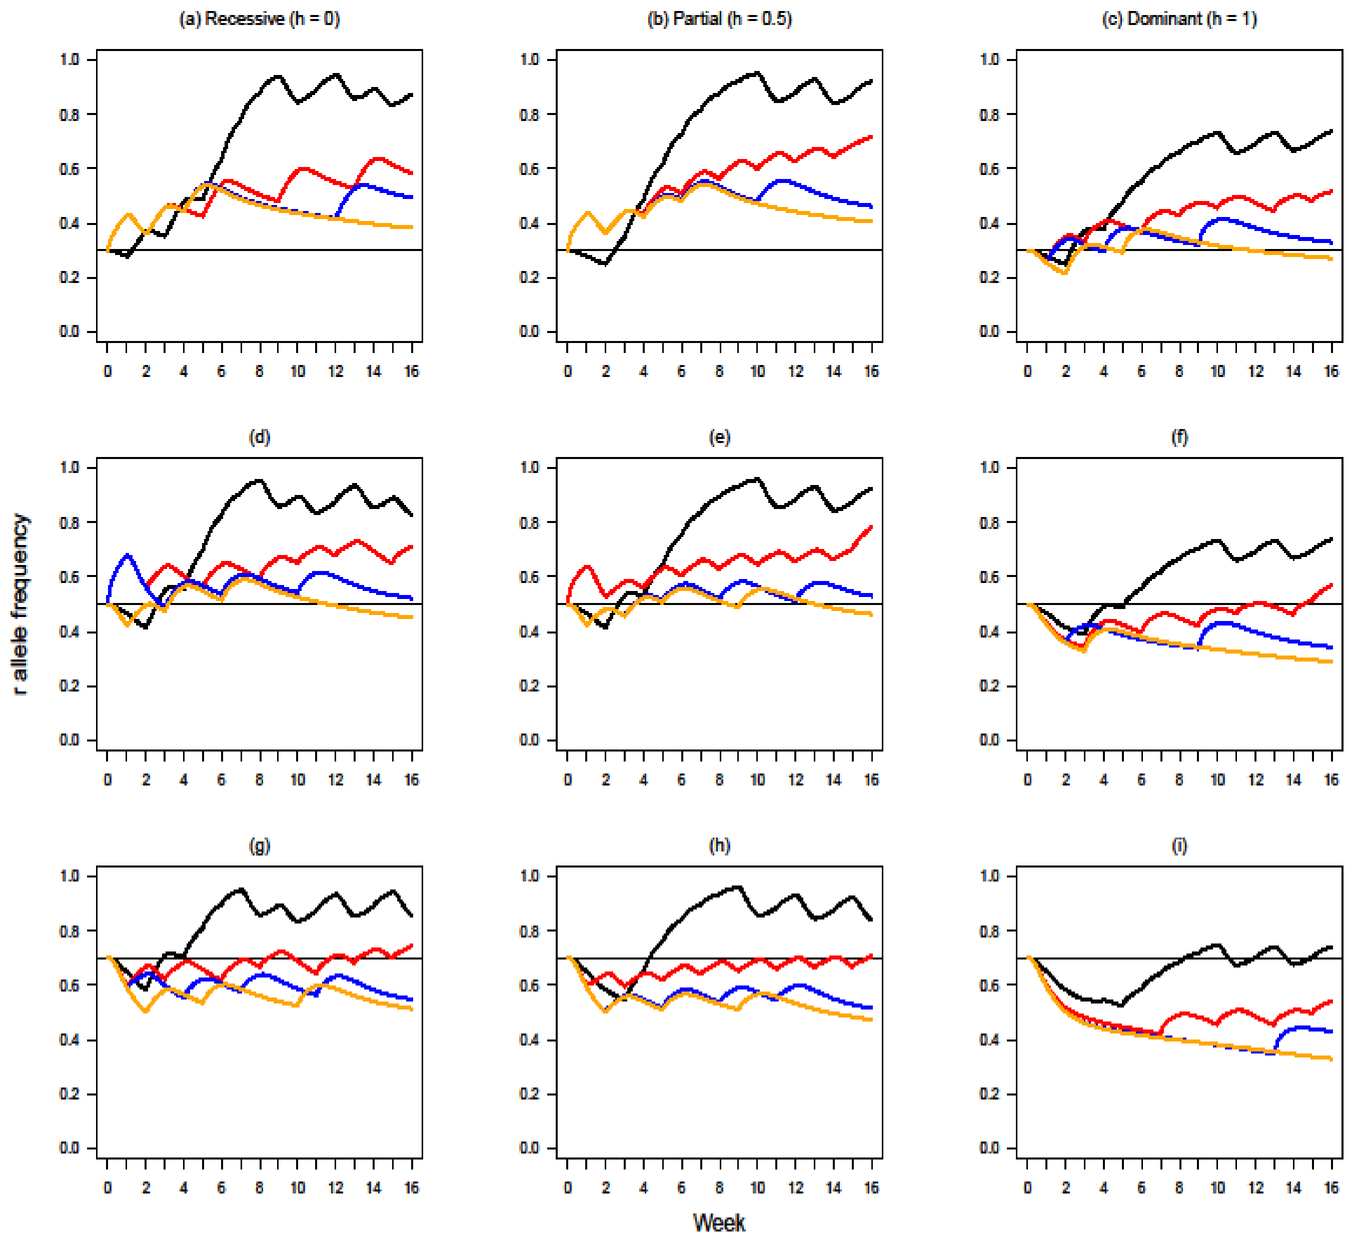

Figure 1*:* Insecticide resistance allele trajectories for four alternative policies attempting to manage a hypothetical insect pest over a 16-week season using insecticidal sprays and releases of transgenic male insects carrying a late-acting male-selecting transgene. Black lines depict trajectories for a myopic policy with a planning horizon of $H = 0$ while red, blue and orange lines portray lookahead policies with planning horizons of $H = 1, H = 2$ and $H = 3$ respectively. For all policies, H is measured in weeks. All panels within the same row share the same initial resistance allele frequency r_0 (indicated by a horizontal black line: from $r_0 = 0.3$ (a, b, c), to $r_0 = 0.5(d, e, f)$ to $r_0 = 0.7(g, h, i)$.), and all panels within the same column share the same

dominance of resistance h (completely recessive resistance: $h = 0$ (a, d, g), additive resistance: $h = 0.5$ (b, e, h) , dominant resistance: $h = 1$ (c, f, i).)

Instances of spraying can be discerned in Figure 1 from the increases in the frequency of resistance alleles that they induce. Similarly, transgenic releases can be identified by declines in resistance allele frequency. Notably, for all presented combinations of r_0 and h, policies with longer planning horizons make less frequent use of sprays than shorter ones. Policies with longer planning horizons have an improved capacity to weight the long-term benefits of a well-timed release against the short-term costs from feeding by both male and female transgenic heterozygote juveniles. Intuitively, the frequency of sprays is also sensitive to the initial resistance allele frequency and the dominance of the resistance allele. In general, partially dominant resistance incentivises spraying, particularly for the myopic policy and the shortest lookahead policy ($H = 1$). Conversely, dominant resistance discourages spraying and promotes additional releases, especially for lookahead policies with longer planning horizons of $H = \{2,3\}$ and this effect increases with the initial frequency of the resistance allele (*e.g.* compare Fig.1c and Fig.1i). The increase in the frequency of spraying when resistance is partially dominant highlights the value of insecticidal sprays to the model: sprays provide an immediate decrease in adult abundance and thus a reduction in the production of juveniles. In contrast, late acting transgenic releases do not reduce the extant number of reproductive adult females, they instead restrict the future recruitment of adult females. Furthermore, when releases are late acting, females inheriting the lethal allele will still cause feeding damage. Thus, given the stated objective of cost minimisation, the model tends to respond to reductions in spray efficacy with an increase in spraying effort and will only make a concerted switch towards transgenic releases as the sole method of suppression when spray efficacy degrades significantly.

The general trends observed for late acting releases also hold true for early acting releases which produce qualitatively similar outcomes to late-acting releases with respect to resistance allele frequency (Figure 2). However, while the end point is similar, for any given combination of h and r_0 a lesser cost is incurred and the action sequences employed are distinct from those utilised for late-acting releases. By removing juvenile females from the population before they feed, early-acting releases confer a partial reduction in feeding damage at the point of use in addition to reducing future reproduction. Consequently, sprays are used less frequently when releases are early-acting. This is illustrated in Figure 3 which shows the frequency with which each action category (no action, spray or release of transgenics) was selected for the lookahead policies with planning horizons of $H = 3$ depicted in Figure 1 (late-acting, blue) and Figure 2 (early-acting, red).

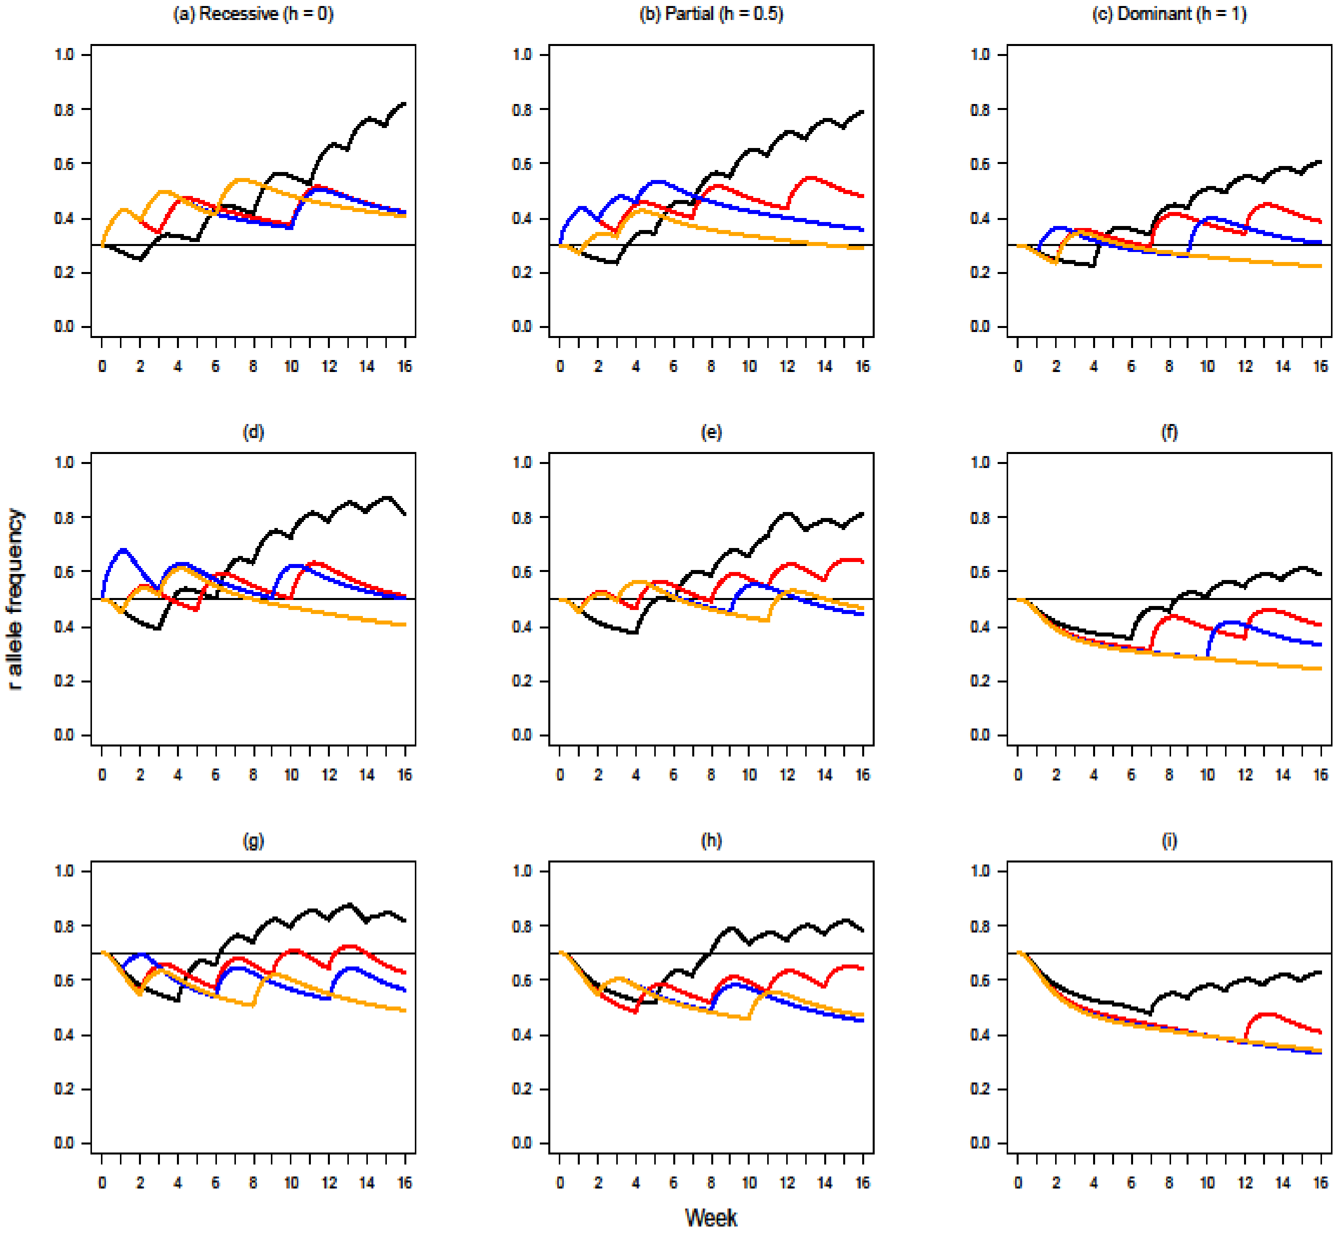

Figure 2*:* Insecticide resistance allele trajectories for four alternative policies attempting to manage a hypothetical insect pest over a 16-week season using insecticidal sprays and releases of transgenic male insects carrying an early-acting male-lethal transgene. Black lines depict trajectories for a myopic policy with a planning horizon of $H = 0$ while red, blue and orange lines portray lookahead policies with planning horizons of $H = 1, H = 2$ and $H = 3$ respectively. For all policies, H is measured in weeks. All panels within the same row share the same initial resistance allele frequency r_0 (indicated by a horizontal black line: from $r_0 = 0.3$ (a, b, c), to $r_0 = 0.5(d, e, f)$ to $r_0 = 0.7(g, h, i)$.), and all panels within the same column share the same

dominance of resistance h (completely recessive resistance: $h = 0$ (a, d, g), additive resistance: $h = 0.5$ (b, e, h), dominant resistance: $h = 1$ (c, f, i).)

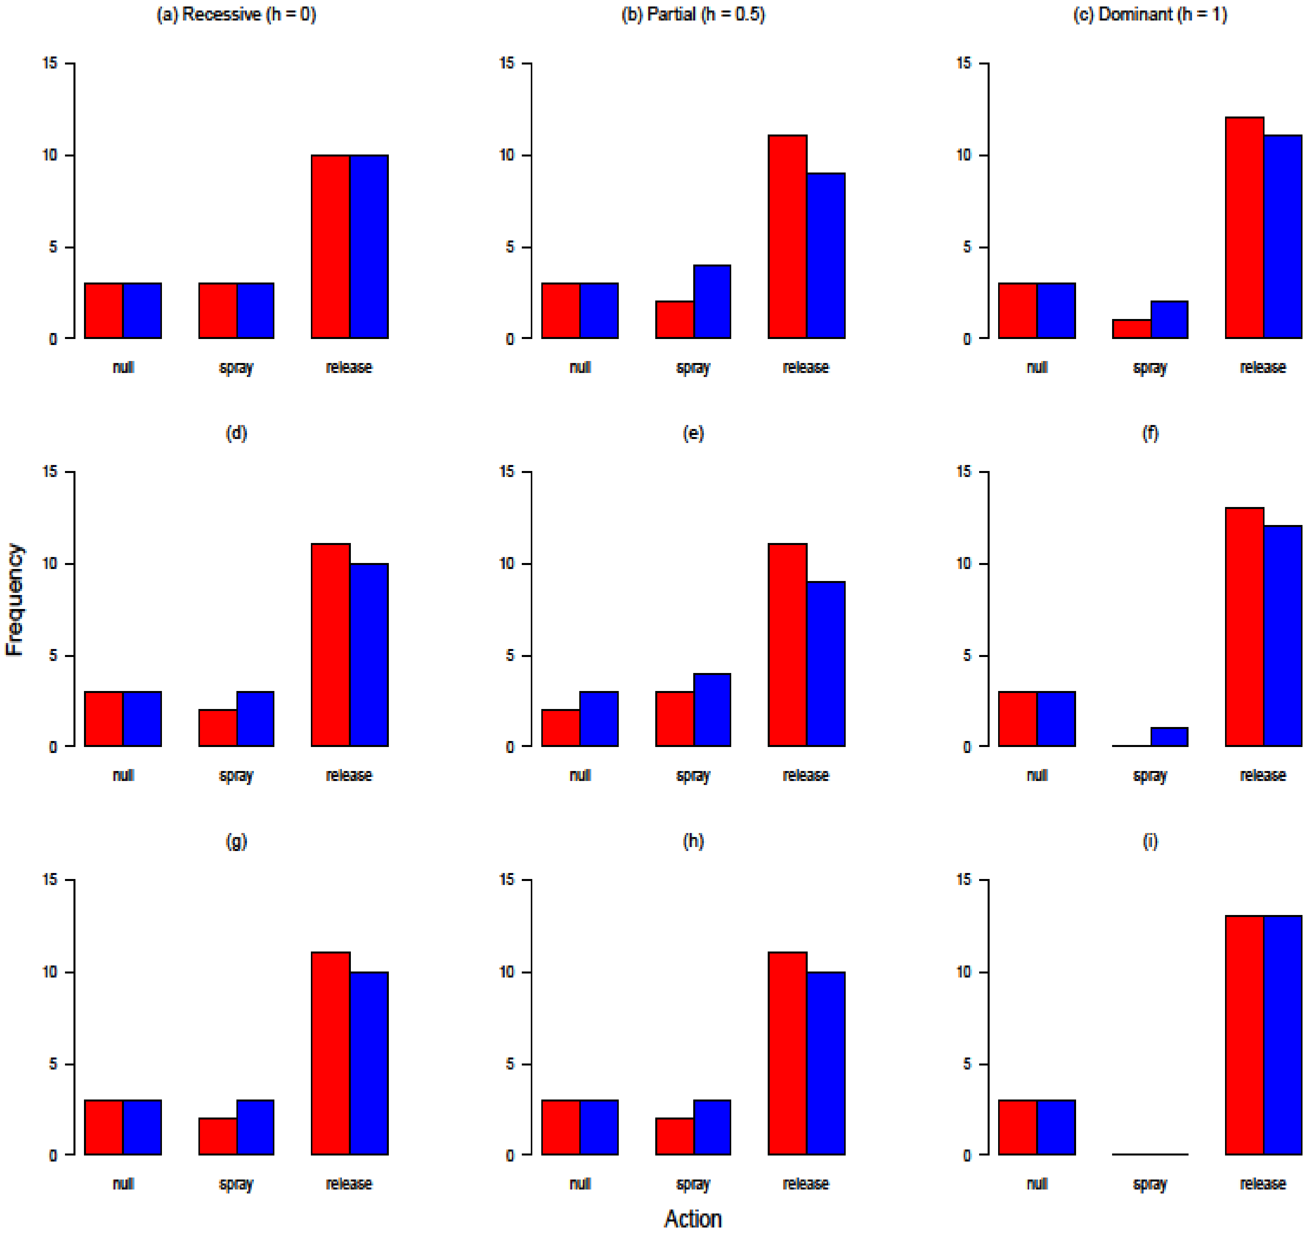

Figure 3*:* Frequencies of implemented actions over a 16-week season for a lookahead policy with a planning horizon of $H = 3$ weeks. In a given week, the model may select to do nothing (null), apply foliar insecticide (spray) or choose from one of six transgenic male insect release ratios (here aggregated together as one category: release). Red bars: transgenic males carry an early acting female lethal construct, blue bars: construct is late acting. All panels within the same row share the same initial resistance allele frequency r_0(indicated by a horizontal black

line), and all panels within the same column share the same dominance of resistance h. Moving from the top row to the bottom row the initial resistance allele frequency increases from $r_0 = 0.3$ (a, b, c), to $r_0 = 0.5(d, e, f)$ to $r_0 = 0.7(g, h, i)$. Panels in the lefthand column illustrate results for completely recessive resistance $h = 0$ (a, d, g), panels in the central column show results for additive resistance $h = 0.5$ (b, e, h) while panels in the righthand column show results for dominant resistance $h = 1$ (c, f, i).

Irrespective of the type of release, and in spite of the reductions in resistance allele frequency observed for longer planning horizons, it is notable that no policy makes an intensive effort to eradicate resistance. The actions selected by a given policy result in a decline in pest abundance over the season (albeit via different means and over different time scales) which is a desirable outcome with respect to the cost-minimising objective function as fewer pests entails less feeding damage. However, given that actions carry a cost and the damage is inflicted upon the crop in proportion to juvenile abundance it follows that the impetus for the model to act decreases as the pest population declines. Furthermore, the focus upon a single season entails that there is no penalty to ending the season with a high resistance allele frequency as the value of insecticides in the subsequent season is not accounted for. Thus if, due to actions taken during the early weeks of the season, the pest population has been significantly reduced but resistance has become prevalent then no action will be taken to mitigate this: the effects of a largely resistant population in the following season are not pertinent to decision making. The policies presented here are only incentivised to concentrate upon reducing the resistance allele frequency when the potential for damage is high and susceptibility to sprays is low. This can be observed in Figure 1 and Figure 2: when resistance is initially prevalent (*e.g.* $r_0 \geq 0.5$) and particularly when it is non-recessive (so some heterozygotes may survive insecticide exposure) then most policies will reduce the resistance allele frequency to some extent prior to utilising sprays. Thus, while lookahead policies with longer planning horizons tend to produce lower resistance allele frequencies they are not expressly managing resistance. Instead, they are utilising their improved capacity to resolve how the consequences of actions unfold over time to select actions such that costs are minimised, an approach which incidentally favours fewer instances of spraying and more transgenic releases.

Robust population

When juveniles are robust ($\mu_j = 0.05$) there is a reduction in the frequency of transgenic releases and a greater emphasis on insecticide application. Consequently, resistance allele frequencies tend to be greater than those observed for the default population as illustrated by Figure 4 which presents the resistance allele trajectories observed for late acting releases. The increased preference for spraying applies independently of the type of transgenic release and affects all considered policies, even those with longer planning horizons (*e.g.* $H = 3$, Figure 5). However, the general trends detailed for the default population ($\mu_j = 0.2$) still hold true for the robust population. For example, while sprays are applied more frequently by all policies, those with longer planning horizons still deploy sprays less frequently than policies with shorter horizons for given values of h and r_0. This shift in action preference is attributable to the sensitivity of the model to juvenile abundance and persistence. Juveniles inflict damage upon the crop but cannot be directly removed. This biases policies with short planning horizons against transgenic releases as they cannot adequately perceive the potential future benefits that may follow from temporarily accepting the presence of heterozygous larvae within the crop. Instead, they focus on spraying which yields immediate population suppression. In contrast, policies with longer horizons are better able to integrate releases and sprays and tend to favour few well-timed sprays interspersed with transgenic releases (*e.g.*Fig.1 & Fig.2). However, increasing juvenile survival indirectly increases the amount of damage any given juvenile population can inflict. Thus, the threat posed by even a small juvenile population is increased which promotes spraying, the action which produces the greatest immediate reduction in juvenile recruitment, whilst narrowing the range of circumstances under which releases are selected.

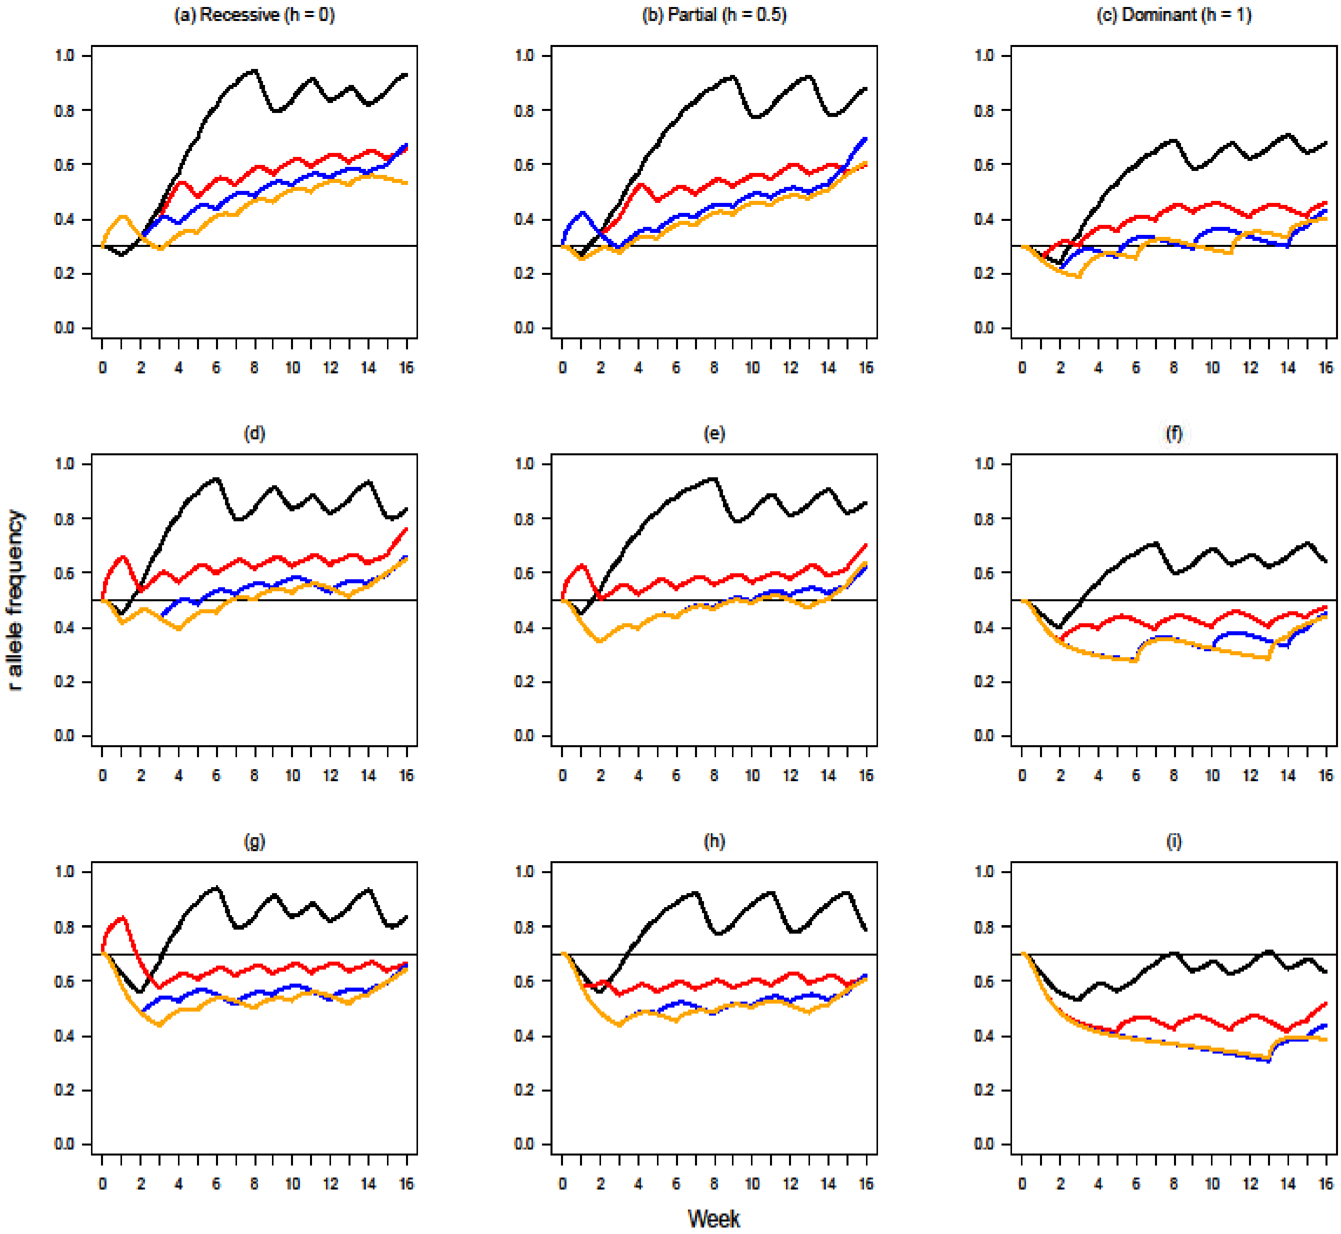

Figure 4*:* As for Figure 1 but with a reduced instantaneous juvenile mortality rate, $\mu_j = 0.05$.

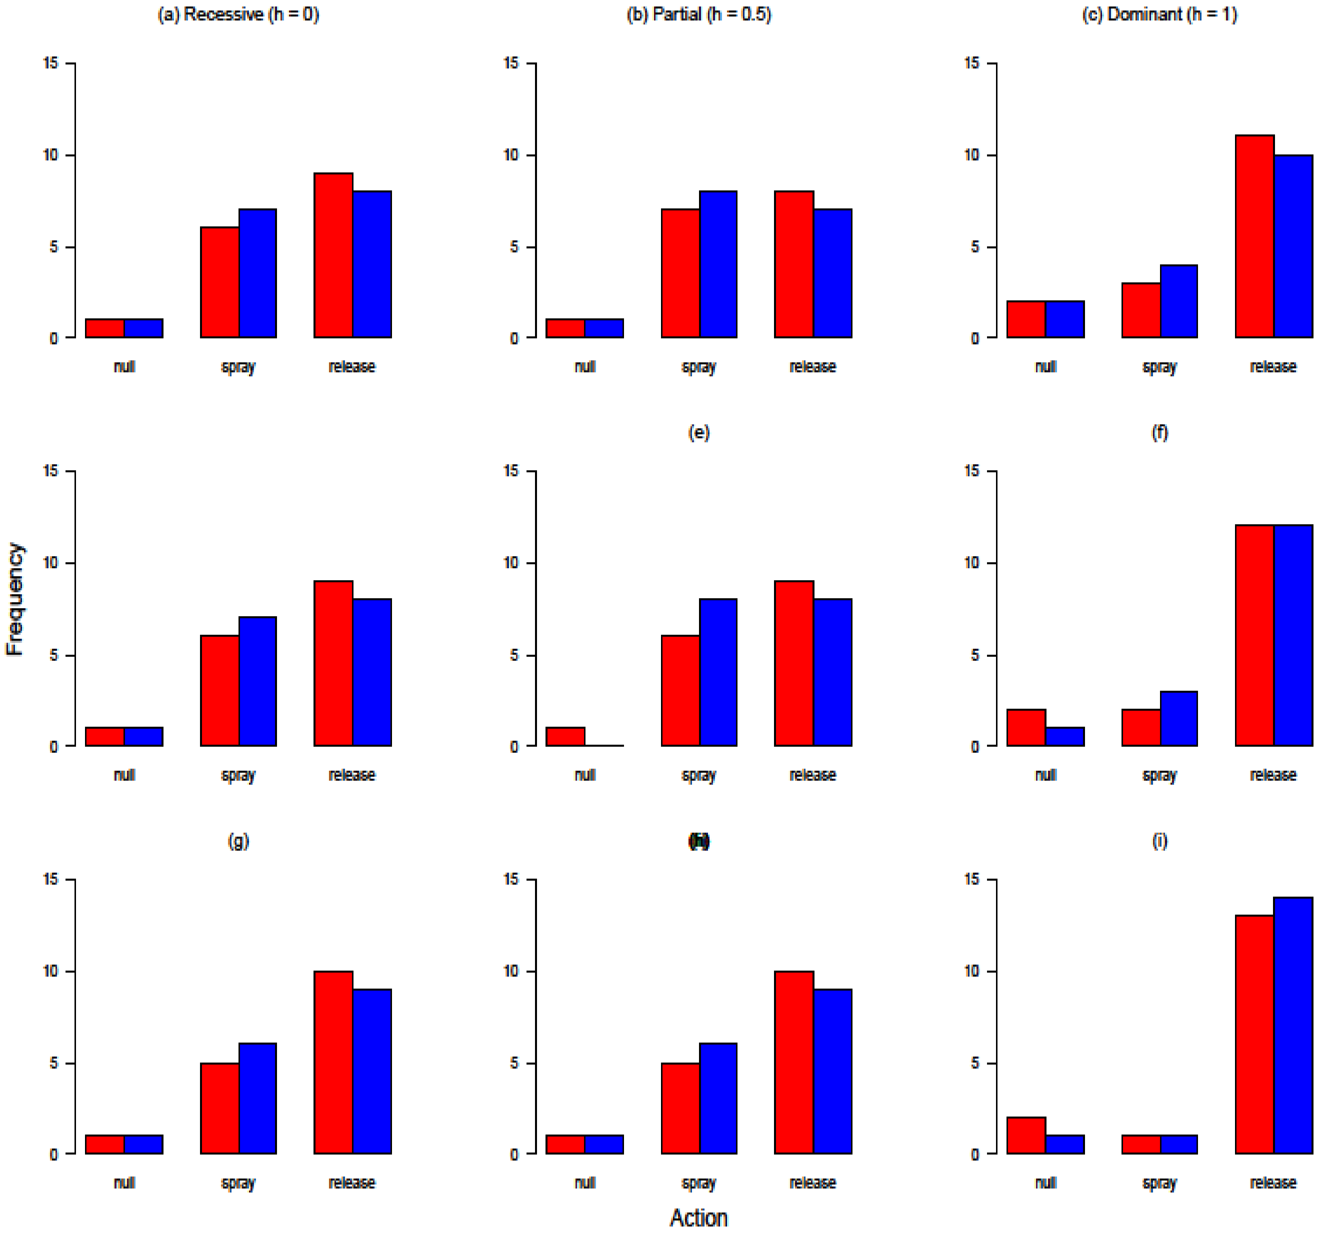

Figure 5*:* As for Figure 3 but with a reduced instantaneous juvenile mortality rate, $\mu_j = 0.05$.

Discussion

We applied a class of approximate dynamic programming algorithms, the lookahead policy, which identifies useful actions by solving a truncated version of the full optimisation problem, to the management of a simple representation of a continuously reproducing, stage-structured agricultural insect pest. Lookahead policies consistently suggested sequences of decisions which provided lower costs and resistance allele frequencies than a myopic policy (which optimises only with respect to a single time period). These benefits were made available by the superior capacity of lookahead policies to integrate insecticidal sprays and male-selecting transgenic insect releases and could be improved upon by extending the planning horizon, H (Fig.1, Fig.2 & Fig.4). This was true for both early and late-acting transgenic releases although, intuitively, the lowest costs and resistance allele frequencies were reported for early-acting releases which also made fewer applications of insecticide (Fig.3, Fig.5). Lookahead policies were also shown to be able to adjust their decisions to accommodate significant changes to the demography of the pest (*e.g.* a substantive reduction in the juvenile mortality rate). However, while these findings indicate that approximate optimisation methods such as the lookahead policy may be a valuable tool in addressing complex ecological decision problems (Nicol *et al.* 2010; Nicol & Chadès 2011) there is significant room for improvement if these approaches are to be gainfully applied to problems with an evolutionary component, where the longer term responses of the target population must be anticipated (Ashley *et al.* 2003; Brown & Staňková 2017). Given an adequate planning horizon, unmodified lookahead policies as presented here show themselves to be competent ecologically enlightened managers, able to anticipate and plan for the ecological responses of the pest to management actions, but if they are to anticipate the effects of evolution upon the focal population a more refined policy may be required (Ashley *et al.* 2003; Brown & Staňková 2017).

Lookahead policies combined releases with insecticidal sprays to achieve suppression in spite of the damage incurred by the heterozygous offspring resulting from matings between wild type females and transgenic males. Indeed, policies with longer planning horizons $H = \{2,3\}$ tended to make greater use of transgenic releases than insecticide. That is, by making fewer, better timed applications of insecticide, lookahead policies with longer planning horizons were able to attain a better outcome than equivalent policies with shorter horizons. Such policies are better able to resolve the downstream effects of a given action which is especially useful with regards to male-selecting transgenic releases, the benefits of which accrue across two temporally staggered phases. First, there is the initial suppression benefit received in the form of the removal of females heterozygous for the male-selecting allele from the juvenile population. Second, males heterozygous for the male-selecting allele can survive to maturity and potentially confer the allele upon their own offspring. This further reduces the future reproductive potential of the population, introgressing alleles for insecticide susceptibility and increasing the efficacy of a subsequent insecticide application (Alphey *et al.* 2007; Alphey, Bonsall & Alphey 2009; Harvey-Samuel *et al.* 2015; Leftwich, Bolton & Chapman 2016; Zhou *et al.* 2018; Alphey & Bonsall 2018). By at least partially recognising these benefits lookahead policies were able to incorporate releases into their action sequences, reducing their dependency upon insecticide, an ecologically desirable outcome given the controversy surrounding pesticides and the environmental consequences of their overuse (Guedes *et al.* 2016).

Lookahead policies were adept at incorporating differences between early and late acting male-selecting releases and changes in juvenile survivorship into their decision making as these were factors with demographic consequences that could be observed within the bounds of the planning horizon. In contrast, they were mediocre resistance managers in spite of even the lowest initial resistance allele frequency used in the simulations ($r_0 = 0.3$) being great enough to be considered problematic if observed under field conditions (Tabashnik *et al.* 2008, 2014; Tabashnik & Carrière 2017) . As such, these lookahead policies were reactionary managers, acting to reduce resistance only when it reduced the efficacy of sprays to the extent that it prevented adequate population suppression. Two major features of the current model framework impede the scope for lookahead policies to manage resistance. First, the emphasis upon cost minimisation entails that "best practice" with respect to the objective function is likely to feature a non-trivial level of resistance (Brown & Staňková 2017) . Second, even the longest planning horizons used here are likely too short to offer sufficient information to the model regarding the future costs of resistance (Ashley *et al.* 2003; Vander Wal *et al.* 2014; Brown & Staňková 2017).

If lookahead policies are to be useful as a predictive tool in exploring the implications of management scenarios for problems with an evolutionary component then their aptitude for accommodating detail-rich problems will need to be linked to a capacity to anticipate consequences which lie beyond the planning horizon. The policies possess the valuable capacity to include information relevant to the evolution of resistance such as the presence of fitness costs (to both transgenes and resistance alleles) (Gassmann, Carrière & Tabashnik 2009; Zhou *et al.* 2018), how fitness costs manifest (which may have demographic consequences as well as evolutionary ones) (Guedes *et al.* 2016; Hackett & Bonsall 2016) or the extent to which the pest may be pre-adapted to toxins (ffrench-Constant & Bass 2017). However, their limited ability to reconcile the short term ecological effects and long term evolutionary effects of their actions restricts the extent to which this information can be used effectively and biases them towards management practices which anticipate demographic responses but simply react to evolutionary responses (Ashley *et al.* 2003; Brown & Staňková 2017).

One route to augmenting the long-term perspective of lookahead policies is through the adoption of a hybrid policy-based approach. A virtue of approximate dynamic programming is that each policy type it encompasses circumvents the curses of dimensionality by distinct means (Powell 2011, 2014). Different policy types are not mutually exclusive and can be combined in a complementary fashion to create a hybrid policy when appropriate (Powell 2014; Powell & Meisel 2016). For example, a cost function approximation (CFA) could be added as a third term to equation [16] where it would act as a penalty term devaluing actions which returned resistance allele frequencies above a defined threshold, which could be informed by existing policy or field data. Alternatively, the lookahead policy could operate within the bounds of a policy function approximation (PFA), which could, among other things, place limits upon the maximum permissible number of sprays in the season, reject actions resulting in a resistance allele frequency exceeding a threshold value or enforce spatial or temporal refuge requirements (Carrière & Tabashnik 2001; Sudo *et al.* 2018). These are exemplar cases by which the long-term perspective of a lookahead policy could be enhanced but the full creative scope for hybrid policy development is significant.

Data Availability

The R code generated by this project has been uploaded to the Open Science Framework: https://osf.io/ja52e/.

Chapter 5: Discussion

This project has considered the implications of potentially confounding factors such as the evolution of insecticide resistance, the realisation of fitness costs to resistance and stage-structured population dynamics for the optimal management of agricultural insect pests using dynamic programming methods. A recurrent theme has been the extent to which results we wish to achieve now conflict with those we would like to see in the future, particularly with respect to the application of approximate dynamic programming algorithms. Another theme has been how the biological idiosyncrasies of a pest can shift the definition of what constitutes a good decision. Both themes arise from the nature of the manager-pest interaction. Insect pest populations are not a static target, instead they and those seeking to manage them are locked into a "leader-follower" game in which the manager assumes the role of the leader and the pest reacts to the conditions imposed upon it (Brown & Staňková 2017). Furthermore, while the insect cannot anticipate the decisions of the manager, its ecology and demography define constraints on the actions that can be usefully applied. The best decisions are made when managers are able to make full use of their capacity to anticipate the potential responses of the pest while decisions which fail to use this information are unlikely to lead to the best long-run outcome (Ashley *et al.* 2003; Vander Wal *et al.* 2014; Brown & Staňková 2017; Tabashnik & Carrière 2017).

The friction between current and future benefits is most apparent when considering the evolution of resistance to insecticidal toxins. Preserving susceptibility to control is essentially an economic harvesting problem: to what extent can we exploit susceptibility without depleting it such that we cannot exploit it in the future? If our interest is in maximising the lifespan of control agents then a long-term evolutionary perspective that considers likely resistance mechanisms is warranted. The combined implementation of the high-dose/refuge strategy for transgenic crops expressing Bt toxins represented a promising acknowledgement that the lifespan of any toxic control agent is potentially finite (Tabashnik 2004). While theoretically congruent with an evolutionarily enlightened approach to resistance management, the efficacy of refugia is limited in practice by the need to accommodate the near-term economic needs of growers. This was highlighted in chapter 2, in which a dynamic programming model consistently selected refuges larger than those which would be implemented in practice (Hackett & Bonsall 2016) and evidence suggests that the refugia currently deployed against several major pests are insufficient to delay resistance (Tabashnik & Carrière 2017). In practice refugia are often selected as if the assumption of a high-dose toxin holds true but factors such as delays in the implementation of a new toxin or pyramid, improper placement of the crop relative to other products expressing the same or similar toxins, cross resistance and pre-existing tolerances can see this assumption violated before a crop becomes commercially available (Carrière, Fabrick & Tabashnik 2016; Tabashnik & Carrière 2017). Resistance can still be delayed when crops fail to express a high-dose but to do so requires the use of significantly larger refuges than would otherwise be used (Tabashnik, Brévault & Carrière 2013; Tabashnik & Carrière 2017). For example, a 70% refuge requirement was mandated for Cry1Ac cotton in Australia from 1996 to 2003 once it became apparent that the high-dose assumption was not satisfied for *Helicoverpa armigera*. Susceptibility to Cry1Ac was maintained in the *H. armigera* population until pyramided cotton expressing Cry1Ac and Cry2Ab could be introduced, allowing the switch to a minimum refuge size of only 10% (Downes *et al.* 2010). In contrast, the related pest *H. zea* has acquired resistance to both toxins in the

US where minimum refuge requirements were less stringent (Tabashnik & Carrière 2017).

It is also frequently assumed that resistance will carry a fitness cost which disadvantages resistant insects in the absence of toxins, delaying the spread of the allele when refugia are present (Tabashnik 2004; Alphey *et al.* 2008; Gassmann, Carrière & Tabashnik 2009). However, chapter 2 illustrated that the presence of a fitness cost does not guarantee synergy with the high-dose/refuge strategy by showcasing the effect of two different types of fitness cost upon the decisions of a dynamic programming model (Hackett & Bonsall 2016). Fitness costs which apply only under a narrow range of circumstances may not be sufficient to impede resistance evolution; it is therefore crucial that the manifestation of fitness costs be understood under field conditions (ffrench-Constant & Bass 2017). Understanding the molecular mechanisms by which a susceptible pest might attain resistance may allow for the anticipation of potential fitness costs (Coustau, Chevillon & ffrench-Constant 2000) but this does not preclude the possibility that costless resistance will evolve by another means (ffrench-Constant & Bass 2017).

Whether the evolution of resistance is delayed or facilitated is also subject to a range of complicating factors which are idiosyncratic to the target pest and its environment including the type and timing of pest dispersal, the demography of the pest, oviposition preferences, the spatial arrangement of refugia and seasonal variation in the toxicity of transgenic crops (Téllez-Rodríguez *et al.* 2014; Brévault, Tabashnik & Carrière 2015; Carrière, Fabrick & Tabashnik 2016; Garcia *et al.* 2016; Sudo *et al.* 2018; Carrière *et al.* 2018). In identifying the best long-term response to a particular management scenario, we would ideally be able to utilise a framework which can account for such details whilst also maintaining a far-sighted perspective of how the actions taken today will influence the scenario we experience tomorrow. Dynamic programming can achieve this but only while the scale of the problem is suitably constrained. However, at the expense of guarantees of optimality, approximate dynamic programming algorithms (ADP) can accommodate problems of arbitrary magnitude.

Chapter 3 demonstrated how a class of approximate dynamic programming algorithm, the lookahead policy, could effectively combine the application of foliar insecticides and bi-sex lethal sterile insect releases to manage a stage-structured, continuously reproducing insect pest (Hackett & Bonsall 2018). This framework was then modified in chapter 4 to allow for the evolution of insecticide resistance and to include male-selecting transgenic releases to assess the extent to which a basic lookahead policy was able to anticipate and respond to the risk of resistance evolution. Ultimately, lookahead policies, when implemented with an appropriate planning horizon, were able to effectively anticipate and respond to the population dynamics of the pest but were less able to actively manage resistance. However, this does not entail that ADP methods are of limited utility to questions of insect pest management and ecological decision problems more generally. These methods have been gainfully applied to problems of extremely high dimensionality and complexity across a diverse range of fields (Powell 2011; Powell, Simao & Bouzaiene-Ayari 2012; Powell & Meisel 2016) and there is as yet no reason to suggest that they could not bring similar benefits to questions in ecology. For example, as highlighted in chapter 4, a bespoke hybrid policy could be developed that would be capable of accounting for the granular details of a specific problem, such as pest dispersal, whilst retaining the longer-term perspective required to preserve susceptibility to control.

In choosing to utilise ADP methods we are accepting that sub-optimal decisions will be made. However, the loss of optimality may be compensated for by the gains in versatility. Is this trade-off acceptable? Ecological decision problems, whether for conservation or the management of a pest or pathogen, are frequently complex and ultimately messy – encompassing a breadth of competing economic, ecological and evolutionary demands, and almost always requiring the careful allocation of limited resources (Wilson *et al.* 2006; Nicol *et al.* 2010; Marescot *et al.* 2013; Vander Wal *et al.* 2014; Chadès *et al.* 2017). Classical dynamic programming methods are fundamentally constrained in the level of detail which they may incorporate before accommodations must be made and variables are either excluded or simplified. The resultant model can be solved optimally but the solution is only optimal with respect to the model from which it was derived (Powell 2011). Thus, if we cannot be confident that the factors which have been abridged or removed are not important to the system we seek to manage how useful is such an answer in practice? Approximate dynamic programming allows us to consider a problem as it is, not as we might wish it to be. It offers the capacity to carefully construct bespoke policies which can more fully represent the complexities of the population we wish to control and the tools we wish to use.

As an example of the types of bespoke policy which are possible, consider the long-term management of an insect pest feeding upon an insecticidal crop. Ideally, any policy we use would be capable of attaining sufficient suppression within a given growing season to meet the near-term economic needs of growers whilst also preserving susceptibility to the insecticidal action of the crop into the following season and beyond. Suppose also that male-selecting transgenic releases are available for the pest. What kind of ADP policy could be developed to manage this population? In this instance, there are two control options available: (i) the determination of the proportion of the landscape allocated to refuge plants at the start of the season and (ii) the frequency and magnitude of male-selecting releases. The basic lookahead policies applied in chapters 3 and 4 would struggle to accommodate this problem given the conflicting temporal resolutions of the two action types. To effectively specify refugia requires that the policy resolve time in terms of seasons while selecting and implementing releases necessitates resolving time in terms of weeks or days. To circumvent this, a hybrid policy could be utilised in which the release strategy within each season is controlled by a policy function approximation (PFA) and the refuge allocation for each season is decided by a lookahead policy. The PFA would detail a set of rules, perhaps derived from field trials and previous management programs, governing when and how releases are to be deployed. The rules could be fixed or could contain tuneable parameters. Meanwhile, the specification of refugia for each season would be controlled by a lookahead policy which, given the release program specified by the PFA, seeks to identify a refuge allocation which preserves susceptibility to the insecticidal crop and offers sufficient suppression. This policy formulation allows the lookahead policy to plan over multiple seasons (*e.g.* with a planning horizon of three years as opposed to three weeks), potentially providing it with the perspective necessary to actively manage resistance. Given that practical levels of resistance to insecticidal crops can emerge within five years of deployment (Tabashnik & Carrière 2017), even a policy which can only foresee potential outcomes over a small number of years could be advantageous.

Approaches which can plan for both near and long-term outcomes and effectively integrate complementary technologies, or simply do as well as they can with whatever tools are available, stand to benefit not only agricultural ecology but challenges in disciplines such as vector control and the management of antibiotic resistance (Alphey, Alphey & Bonsall 2011; Brown & Staňková 2017). The introduction of novel technologies and chemistries such as crops expressing vegetative insecticidal proteins will not escape the need for resistance management strategies (Tabashnik & Carrière 2017): unregulated killing today will still threaten our capacity to kill tomorrow. Given the benefits that have already been derived from technologies such as transgenic crops (Dively *et al.* 2018), the adoption of long-sighted management strategies which aim to anticipate resistance could help sustain these benefits into the future. In developing such strategies, both dynamic programming and approximate dynamic programming methods provide a versatile toolbox for evaluating how we might approach such problems and for exploring how the evolutionary and ecological complexities of particular pests might impact upon management efforts.

Bibliography

Alphey, N., Alphey, L. & Bonsall, M.B. (2011) A Model Framework to Estimate Impact and Cost of Genetics-Based Sterile Insect Methods for Dengue Vector Control ed P.L. Oliveira. *PLoS ONE*, **6**, e25384.

Alphey, L., Benedict, M., Bellini, R., Clark, G.G., Dame, D. a, Service, M.W. & Dobson, S.L. (2010) Sterile-Insect Methods for Control of Mosquito-Borne Diseases: An Analysis. *Vector-Borne and Zoonotic Diseases*, **10**, 295–311.

Alphey, N. & Bonsall, M.B. (2018) Genetics-based methods for agricultural insect pest management. *Agricultural and Forest Entomology*, **20**, 131–140.

Alphey, N., Bonsall, M.B. & Alphey, L. (2009) Combining pest control and resistance management: synergy of engineered insects with Bt crops. *Journal of Economic Entomology*, **102**, 717–32.

Alphey, N., Coleman, P.G., Bonsall, M.B. & Alphey, L. (2008) Proportions of different habitat types are critical to the fate of a resistance allele. *Theoretical Ecology*, **1**, 103–115.

Alphey, N., Coleman, P.G., Donnelly, C. a & Alphey, L. (2007) Managing insecticide resistance by mass release of engineered insects. *Journal of Economic Entomology*, **100**, 1642–1649.

Ashley, M. V., Willson, M.F., Pergams, O.R.W., O'Dowd, D.J., Gende, S.M. & Brown, J.S. (2003) Evolutionarily enlightened management. *Biological Conservation*, **111**, 115–123.

Baltzegar, J., Cavin Barnes, J., Elsensohn, J.E., Gutzmann, N., Jones, M.S., King, S. &

Sudweeks, J. (2018) Anticipating complexity in the deployment of gene drive insects in agriculture. *Journal of Responsible Innovation*, **5**, S81–S97.

Barclay, H. (2005) Mathematical models for the use of sterile insects. *Sterile Insect Technique. Principles and Practice in Area-wide Integrated Pest Management*, 1st ed (eds V. Dyck, J. Hendrichs, & A.S. Robinson), pp. 147–174. Springer Netherlands, Dordrecht.

Barclay, H.J. (2016) Determination of the sterile release rate for stopping growing age-structured populations. *International Journal of Pest Management*, **62**, 40–54.

Becker, J.M. & Liess, M. (2015) Biotic interactions govern genetic adaptation to toxicants. *Proceedings of the Royal Society B: Biological Sciences*, **282**, 20150071.

Becker, J.M. & Liess, M. (2017) Species Diversity Hinders Adaptation to Toxicants. *Environmental Science & Technology*, **51**, 10195–10202.

Bellows, T. (1981) The descriptive properties of some models for density dependence. *Journal of Animal Ecology*, **50**, 139–156.

Bird, L.J. & Akhurst, R.J. (2007) Effects of host plant species on fitness costs of Bt resistance in Helicoverpa armigera (Lepidoptera: Noctuidae). *Biological Control*, **40**, 196–203.

Boettiger, C., Mangel, M. & Munch, S. (2015) Avoiding tipping points in fisheries management through Gaussian process dynamic programming. *Proceedings of the Royal Society B*, **282**, 8–11.

Bourtzis, K., Lees, R.S., Hendrichs, J. & Vreysen, M.J.B. (2016) More than one rabbit out of the hat: Radiation, transgenic and symbiont-based approaches for sustainable management of mosquito and tsetse fly populations. *Acta Tropica*, **157**, 115–130.

Brévault, T., Tabashnik, B.E. & Carrière, Y. (2015) A seed mixture increases dominance of resistance to Bt cotton in Helicoverpa zea. *Scientific reports*, **5**, 9807.

Brown, J.S. & Staňková, K. (2017) Game theory as a conceptual framework for managing insect pests. *Current Opinion in Insect Science*, **21**, 26–32.

Carey, J.R. (2011) Biodemography of the Mediterranean fruit fly: Aging, longevity and adaptation in the wild. *Experimental Gerontology*, **46**, 404–411.

Carrière, Y., Crowder, D.W. & Tabashnik, B.E. (2010) Evolutionary ecology of insect adaptation to Bt crops. *Evolutionary Applications*, **3**, 561–573.

Carrière, Y., Degain, B.A., Unnithan, G.C., Harpold, V.S., Heuberger, S., Li, X. & Tabashnik, B.E. (2018) Effects of seasonal changes in cotton plants on the evolution of resistance to pyramided cotton producing the Bt toxins Cry1Ac and Cry1F in Helicoverpa zea. *Pest Management Science*, **74**, 627–637.

Carrière, Y., Fabrick, J.A. & Tabashnik, B.E. (2016) Can Pyramids and Seed Mixtures Delay Resistance to Bt Crops? *Trends in Biotechnology*, **34**, 291–302.

Carrière, Y. & Tabashnik, B.E. (2001) Reversing insect adaptation to transgenic insecticidal plants. *Proceedings of the Royal Society B: Biological Sciences*, **268**, 1475–80.

Chadès, I., Nicol, S., Rout, T.M., Péron, M., Dujardin, Y., Pichancourt, J.B., Hastings, A. & Hauser, C.E. (2017) Optimization methods to solve adaptive management problems. *Theoretical Ecology*, **10**, 1–20.

Clark, C.W. & Mangel, M. (2000) *Dynamic State Variable Models in Ecology: Methods and Applications*. Oxford University Press.

Consortium, R. (2013) Heterogeneity of selection and the evolution of resistance. *Trends in Ecology & Evolution*, **28**, 110–118.

Cooper, J. & Dobson, H. (2007) The benefits of pesticides to mankind and the environment. *Crop Protection*, **26**, 1337–1348.

Coustau, C., Chevillon, C. & ffrench-Constant, R. (2000) Resistance to xenobiotics and parasites: can we count the cost? *Trends in Ecology & Evolution*, **15**, 378–383.

Dangal, V. & Huang, F. (2015) Fitness costs of Cry1F resistance in two populations of fall armyworm, Spodoptera frugiperda (J.E. Smith), collected from Puerto Rico and Florida. *Journal of Invertebrate Pathology*, **127**, 81–86.

Diamantidis, A.D., Carey, J.R., Nakas, C.T. & Papadopoulos, N.T. (2011) Population-specific demography and invasion potential in medfly. *Ecol Evol*, **1**, 479–488.

Diamantidis, A.D., Papadopoulos, N.T., Nakas, C.T., Wu, S., MÜller, H.G. & Carey, J.R. (2009) Life history evolution in a globally invading tephritid: Patterns of survival and reproduction in medflies from six world regions. *Biological Journal of the Linnean Society*, **97**, 106–117.

Dively, G.P., Venugopal, P.D., Bean, D., Whalen, J., Holmstrom, K., Kuhar, T.P., Doughty, H.B., Patton, T., Cissel, W. & Hutchison, W.D. (2018) Regional pest suppression associated with widespread Bt maize adoption benefits vegetable growers. *Proceedings of the National Academy of Sciences*, 201720692.

Downes, S., Mahon, R.J., Rossiter, L., Kauter, G., Leven, T., Fitt, G. & Baker, G. (2010) Adaptive management of pest resistance by Helicoverpa species (Noctuidae) in Australia to the Cry2Ab Bt toxin in Bollgard II®cotton. *Evolutionary Applications*, **3**, 574–584.

Dyck, V., Henrichs, J. & Robinson, A. (eds). (2005) *Sterile Insect Technique, Principles and Practice in Area-Wide Integrated Pest Management.* Springer Netherlands.

Enkerlin, W.R. (2005) Impact of fruit fly control programmes using the sterile insect technique. *Sterile Insect Technique. Principles and Practice in Area-wide Integrated Pest Management* (eds V. Dyck, J. Hendrichs, & A.S. Robinson), pp. 651–676. Springer Netherlands, Dordrecht.

ffrench-Constant, R.H. & Bass, C. (2017) Does resistance really carry a fitness cost? *Current Opinion in Insect Science*, **21**, 39–46.

Furlong, M.J., Wright, D.J. & Dosdall, L.M. (2013) Diamondback moth ecology and management: problems, progress, and prospects. *Annual review of entomology*, **58**, 517–41.

Garcia, A.G., Ferreira, C.P., Cônsoli, F.L. & Godoy, W.A.C. (2016) Predicting evolution of insect resistance to transgenic crops in within-field refuge configurations, based on larval movement. *Ecological Complexity*, **28**, 94–103.

García, M., Ortego, F., Hernández-Crespo, P., Farinós, G.P. & Castañera, P. (2015) Inheritance, fitness costs, incomplete resistance and feeding preferences in a laboratory-selected MON810-resistant strain of the true armyworm Mythimna unipuncta. *Pest Management Science*.

Gassmann, A.J., Carrière, Y. & Tabashnik, B.E. (2009) Fitness costs of insect resistance to Bacillus thuringiensis. *Annual Review of Entomology*, **54**, 147–63.

Gassmann, A.J., Fabrick, J.A., Sisterson, M.S., Hannon, E.R., Stock, S.P. & Carrie, Y. (2009) Effects of Pink Bollworm Resistance to Bacillus thuringiensis on Phenoloxidase Activity and Susceptibility to Entomopathogenic Nematodes. *Journal of Economic*

Entomology, **102**, 1224–1232.

Gassmann, A.J., Onstad, D.W. & Pittendrigh, B.R. (2009) Evolutionary analysis of herbivorous insects in natural and agricultural environments. *Pest Management Science*, **65**, 1174–81.

Glaum, P.R., Ives, A.R. & Andow, D. a. (2011) Contamination and management of resistance evolution to high-dose transgenic insecticidal crops. *Theoretical Ecology*, **5**, 195–209.

Gong, P., Epton, M.J., Fu, G., Scaife, S., Hiscox, A., Condon, K.C., Condon, G.C., Morrison, N.I., Kelly, D.W., Dafa'alla, T., Coleman, P.G. & Alphey, L. (2005) A dominant lethal genetic system for autocidal control of the Mediterranean fruitfly. *Nature Biotechnology*, **23**, 453–456.

Gould, F. (1998) Sustainability of transgenic insecticidal cultivars: integrating pest genetics and ecology. *Annual Review of Entomology*, **43**, 701–26.

Guedes, R.N.C., Smagghe, G., Stark, J.D. & Desneux, N. (2016) Pesticide-Induced Stress in Arthropod Pests for Optimized Integrated Pest Management Programs. *Annual Review of Entomology*, **61**, 43–62.

Gurr, G.M. & Kvedaras, O.L. (2010) Synergizing biological control: Scope for sterile insect technique, induced plant defences and cultural techniques to enhance natural enemy impact. *Biological Control*, **52**, 198–207.

Hackett, S.C. & Bonsall, M.B. (2016) Type of fitness cost influences the rate of evolution of resistance to transgenic Bt crops ed S. Diamond. *Journal of Applied Ecology*, **53**, 1391–1401.

Hackett, S.C. & Bonsall, M.B. (2018) Management of a stage-structured insect pest: an application of approximate optimization. *Ecological Applications*, 0–3.

Harvey-Samuel, T., Ant, T. & Alphey, L. (2017) Towards the genetic control of invasive species. *Biological Invasions*, **19**, 1683–1703.

Harvey-Samuel, T., Morrison, N.I., Walker, A.S., Marubbi, T., Yao, J., Collins, H.L., Gorman, K., Davies, T.G.E., Alphey, N., Warner, S., Shelton, A.M. & Alphey, L. (2015) Pest control and resistance management through release of insects carrying a male-selecting transgene. *BMC Biology*, **13**, 49.

Houston, A.I. & McNamara, J.M. (1999) *Models for Adaptive Behavior. An Approach Based on State*. Cambridge University Press, Cambridge.

Hutchison, W.D., Burkness, E.C., Mitchell, P.D., Moon, R.D., Leslie, T.W., Fleischer, S.J., Abrahamson, M., Hamilton, K.L., Steffey, K.L., Gray, M.E., Hellmich, R.L., Kaster, L. V, Hunt, T.E., Wright, R.J., Pecinovsky, K., Rabaey, T.L., Flood, B.R. & Raun, E.S. (2010) Areawide suppression of European corn borer with Bt maize reaps savings to non-Bt maize growers. *Science (New York, N.Y.)*, **330**, 222–5.

Ingber, D. a. & Gassmann, A.J. (2015) Inheritance and Fitness Costs of Resistance to Cry3Bb1 Corn by Western Corn Rootworm (Coleoptera: Chrysomelidae). *Journal of Economic Entomology*, **108**, 2421–2432.

Jakka, S.R.K., Knight, V.R. & Jurat-Fuentes, J.L. (2014) Fitness Costs Associated With Field-Evolved Resistance to Bt Maize in <I>Spodoptera frugiperda</I> (Lepidoptera: Noctuidae). *Journal of Economic Entomology*, **107**, 342–351.

Kraaijeveld, K. & Chapman, T. (2004) Effects of male sterility on female remating in the Mediterranean fruitfly, Ceratitis capitata. *Proceedings of the Royal Society B-Biological Sciences*, **271**, S209–S211.

Leftwich, P.T., Bolton, M. & Chapman, T. (2016) Evolutionary biology and genetic techniques for insect control. *Evolutionary Applications*, **9**, 212–230.

Leftwich, P.T., Koukidou, M., Rempoulakis, P., Gong, H.-F., Zacharopoulou, A., Fu, G., Chapman, T., Economopoulos, A., Vontas, J. & Alphey, L. (2014) Genetic elimination of field-cage populations of Mediterranean fruit flies. *Proceedings of the Royal Society B*, **281**, 20141372-.

Levine, E., Spencer, J., Isard, S., Onstad, D. & Gray, M. (2002) Adaptation of the western corn rootworm to crop rotation: evolution of a new strain in response to a management practice. *American Entomologist*, **48**, 94–107.

Lu, Y., Wu, K., Jiang, Y., Guo, Y. & Desneux, N. (2012) Widespread adoption of Bt cotton and insecticide decrease promotes biocontrol services. *Nature*, **487**, 362–365.

Macnair, M. (1991) Why the evolution of resistance to anthropogenic toxins normally involves major gene changes: the limits to natural selection. *Genetica*, **84**, 213–219.

Malacrida, A.R., Gomulski, L.M., Bonizzoni, M., Bertin, S., Gasperi, G. & Guglielmino, C.R. (2007) Globalization and fruitfly invasion and expansion: The medfly paradigm. *Genetica*, **131**, 1–9.

Mangel, M. (1987) Oviposition site selection and clutch sizes in insects. *Journal of Mathematical Biology*, **25**, 1–22.

Mangel, M. (2015) Stochastic Dynamic Programming Illuminates the Link Between Environment, Physiology, and Evolution. *Bulletin of Mathematical Biology*, **77**, 857–877.

Marescot, L., Chapron, G., Chades, I., Fackler, P.L., Duchamp, C., Marboutin, E. & Gimenez, O. (2013) Complex decisions made simple: A primer on stochastic dynamic programming. *Methods in Ecology and Evolution*, **4**, 872–884.

Maynard Smith, J. & Slatkin, M. (1973) The stability of predator-prey systems. *Ecology*, **54**, 384–391.

Miller, N.J. & Sappington, T.W. (2017) Role of dispersal in resistance evolution and spread. *Current Opinion in Insect Science*, **21**, 68–74.

Mitchell, P.D. & Onstad, D. (2014) Valuing Pest Susceptibility to Control. *Insect Resistance Management: Biology, Economics and Prediction*, Second edi (ed D. Onstad), pp. 25–53. Academic Press.

Morrison, N.I., Segura, D.F., Stainton, K.C., Fu, G., Donnelly, C.A. & Alphey, L.S. (2009) Sexual competitiveness of a transgenic sexing strain of the Mediterranean fruit fly, *Ceratitis capitata*. *Entomologia Experimentalis et Applicata*, **133**, 146–153.

Nicol, S. & Chadès, I. (2011) Beyond stochastic dynamic programming: A heuristic sampling method for optimizing conservation decisions in very large state spaces. *Methods in Ecology and Evolution*, **2**, 221–228.

Nicol, S.C., Chadès, I., Linke, S. & Possingham, H.P. (2010) Conservation decision-making in large state spaces. *Ecological Modelling*, **221**, 2531–2536.

Oerke, E.-C. (2006) Crop losses to pests. *The Journal of Agricultural Science*, **144**, 31.

Okuyama, T. & Hsu, J.-C. (2013) Larval competition within and between insecticide resistant and susceptible individuals in the oriental fruit fly, Bactrocera dorsalis. *Journal of Applied Entomology*, **137**, 289–295.

Onstad, D. (2006) Modeling larval survival and movement to evaluate seed mixtures of transgenic corn for control of western corn rootworm (Coleoptera:

Chrysomelidae). *Journal of economic entomology*, **99**, 1407–1414.

Onstad, D.W. (2014) Major Issues In Insect resistance Management. *Insect Resistance Management: Biology, Economics and Prediction*, 3rd ed (ed D.W. Onstad), pp. 1–20. Academic Press.

Onstad, D.W., Shelton, A. & Flexner, J.L. (2014) Insect Resistance, Natural Enemies and Density Dependent Processes. *Insect Resistance Management: Biology, Economics and Prediction*, Second (ed D. Onstad), pp. 403–419. Academic Press.

Plant, R. & Mangel, M. (1987) Modeling and simulation in agricultural pest management. *SIAM Review*, **29**.

Powell, W.B. (2011) *Approximate Dynamic Programming*, 2nd ed. John Wiley & Sons, Inc., Hoboken.

Powell, W.B. (2014) Clearing the Jungle of Stochastic Optimization. *Bridging Data and Decisions* (eds A. Newman & J. Leung), pp. 109–137. INFORMS, Cantonville.

Powell, W.B. (2016) A Unified Framework for Optimization under Uncertainty. *Optimization Challenges in Complex, Networked and Risky Systems* (eds A. Gupta & C. Agostino), pp. 45–83. INFORMS, Cantonville.

Powell, W.B. & Meisel, S. (2016) Tutorial on Stochastic Optimization in Energy—Part II: An Energy Storage Illustration. *IEEE Transactions on Power Systems*, **31**, 1468–1475.

Powell, W.B., Simao, H.P. & Bouzaiene-Ayari, B. (2012) Approximate dynamic programming in transportation and logistics: a unified framework. *EURO Journal on Transportation and Logistics*, **1**, 237–284.

Raymond, B., Sayyed, A.H., Hails, R.S. & Wright, D.J. (2007) Exploiting pathogens and their impact on fitness costs to manage the evolution of resistance to Bacillus thuringiensis. *Journal of Applied Ecology*, **44**, 768–780.

Romeu-Dalmau, C., Bonsall, M.B., Willis, K.J. & Dolan, L. (2015) Asiatic cotton can generate similar economic benefits to Bt cotton under rainfed conditions. *Nature Plants*, **1**, 15072.

Roughgarden, J. (1971) Density-dependent natural selection. *Ecology*, **52**, 453–468.

Shea, K. & Possingham, H. (2000) Optimal release strategies for biological control agents: an application of stochastic dynamic programming to population management. *Journal of Applied Ecology*, **37**, 77–86.

Shoemaker, C. (1981) Applications of dynamic programming and other optimization methods in pest management. *IEEE Transactions on Automatic Control*, **26**, 1125–1132.

Shoemaker, C. (1982) Optimal integrated control of univoltine pest populations with age structure. *Operations Research*, **30**, 40–61.

Sisterson, M., Antilla, L. & Carrière, Y. (2004) Effects of insect population size on evolution of resistance to transgenic crops. *Journal of Economic Entomology*, **97**, 1413–1424.

Sjöberg, P., Rämert, B., Thierfelder, T. & Hillbur, Y. (2015) Ban of a broad-spectrum insecticide in apple orchards: effects on tortricid populations, management strategies, and fruit damage. *Journal of Pest Science*, **88**, 767–775.

Sudo, M., Takahashi, D., Andow, D.A., Suzuki, Y. & Yamanaka, T. (2018) Optimal management strategy of insecticide resistance under various insect life histories:

Heterogeneous timing of selection and interpatch dispersal. *Evolutionary Applications*, **11**, 271–283.

Szyniszewska, A.M. & Tatem, A.J. (2014) Global Assessment of Seasonal Potential Distribution of Mediterranean Fruit Fly, Ceratitis capitata (Diptera: Tephritidae) ed K. Bourtzis. *PLoS ONE*, **9**, e111582.

Tabashnik, B. (2004) Delaying evolution of insect resistance to transgenic crops by decreasing dominance and heritability. *Journal of Evolutionary Biology*, **17**, 904–918.

Tabashnik, B.E., Brévault, T. & Carrière, Y. (2013) Insect resistance to Bt crops: lessons from the first billion acres. *Nature Biotechnology*, **31**, 510–21.

Tabashnik, B.E. & Carrière, Y. (2017) Surge in insect resistance to transgenic crops and prospects for sustainability. *Nature Biotechnology*, **35**, 926–935.

Tabashnik, B.E., Gassmann, A.J., Crowder, D.W. & Carriére, Y. (2008) Insect resistance to Bt crops: evidence versus theory. *Nature biotechnology*, **26**, 199–202.

Tabashnik, B.E., Mota-Sanchez, D., Whalon, M.E., Hollingworth, R.M. & Carrière, Y. (2014) Defining terms for proactive management of resistance to Bt crops and pesticides. *Journal of economic entomology*, **107**, 496–507.

Tabashnik, B.E., Sisterson, M.S., Ellsworth, P.C., Dennehy, T.J., Antilla, L., Liesner, L., Whitlow, M., Staten, R.T., Fabrick, J. a, Unnithan, G.C., Yelich, A.J., Ellers-Kirk, C., Harpold, V.S., Li, X. & Carrière, Y. (2010) Suppressing resistance to Bt cotton with sterile insect releases. *Nature biotechnology*, **28**, 1304–7.

Téllez-Rodríguez, P., Raymond, B., Morán-Bertot, I., Rodríguez-Cabrera, L., Wright, D.J., Borroto, C.G. & Ayra-Pardo, C. (2014) Strong oviposition preference for Bt over non-Bt maize in Spodoptera frugiperda and its implications for the evolution of resistance. *BMC biology*, **12**, 48.

Vargas-Teran, M., Hoffman, H.. & Tweddle, N.. (2005) Impact of screwworm eradication programmes using the sterile insect technique. *Sterile Insect Technique* (eds V. Dyck, J. Hendrichs, & A.S. Robinson), pp. 629–650. Springer Netherlands, Dordrecht.

Vontas, J., Hernández-Crespo, P., Margaritopoulos, J.T., Ortego, F., Feng, H.-T., Mathiopoulos, K.D. & Hsu, J.-C. (2011) Insecticide resistance in Tephritid flies. *Pesticide Biochemistry and Physiology*, **100**, 199–205.

Vander Wal, E., Garant, D., Calmé, S., Chapman, C.A., Festa-Bianchet, M., Millien, V., Rioux-Paquette, S. & Pelletier, F. (2014) Applying evolutionary concepts to wildlife disease ecology and management. *Evolutionary Applications*, **7**, 856–868.

Watkinson-Powell, B. & Alphey, N. (2017) Resistance to genetic insect control: Modelling the effects of space. *Journal of Theoretical Biology*, **413**, 72–85.

Wilson, K.A., McBride, M.F., Bode, M. & Possingham, H.P. (2006) Prioritizing global conservation efforts. *Nature*, **440**, 337–340.

Yakob, L. & Bonsall, M.B. (2009) Importance of space and competition in optimizing genetic control strategies. *Journal of economic entomology*, **102**, 50–57.

Zhang, W. & Swinton, S.M. (2009) Incorporating natural enemies in an economic threshold for dynamically optimal pest management. *Ecological Modelling*, **220**, 1315–1324.

Zhang, W. & Swinton, S.M. (2012) Optimal control of soybean aphid in the presence of natural enemies and the implied value of their ecosystem services. *Journal of*

Environmental Management, **96**, 7–16.

Zhang, L.J., Wu, Z.L., Wang, K.F., Liu, Q., Zhuang, H.M. & Wu, G. (2015) Trade-off between thermal tolerance and insecticide resistance in Plutella xylostella. *Ecology and evolution*, **5**, 515–30.

Zhou, L., Alphey, N., Walker, A.S., Travers, L.M., Hasan, F., Morrison, N.I., Bonsall, M.B. & Raymond, B. (2018) Combining the high-dose/refuge strategy and self-limiting transgenic insects in resistance management-A test in experimental mesocosms. *Evolutionary Applications*, 1–12.