

Applying optimal control theory to complex epidemiological models to inform real-world disease management

E.H. Bussell¹, C.E. Dangerfield¹, C.A. Gilligan¹ & N.J. Cunniffe¹

¹ Department of Plant Sciences, University of Cambridge, Cambridge CB2 3EA, United Kingdom

Keywords: optimal control, feedback, model predictive control, disease management

Corresponding Author: Dr. Nik Cunniffe, njc1001@cam.ac.uk

Twitter: [@nikcunniffe](https://twitter.com/nikcunniffe)

Summary

Mathematical models provide a rational basis to inform how, where and when to control disease. Assuming an accurate spatially-explicit simulation model can be fitted to spread data, it is straightforward to use it to test the performance of a range of management strategies. However, the typical complexity of simulation models and the vast set of possible controls mean that only a small subset of all possible strategies can ever be tested. An alternative approach – optimal control theory – allows the very best control to be identified unambiguously. However, the complexity of the underpinning mathematics means that disease models used to identify this optimum must be very simple. We highlight two frameworks for bridging the gap between detailed epidemic simulations and optimal control theory: open-loop and model predictive control. Both these frameworks approximate a simulation model with a simpler model more amenable to mathematical analysis. Using an illustrative example model we show the benefits of using feedback control, in which the approximation and control are updated as the epidemic progresses. Our work illustrates a new methodology to allow the insights of optimal control theory to inform practical disease management strategies, with the potential for application to diseases of plants, animals and humans.

1 Introduction

Mathematical modelling plays an increasingly important role in informing policy and management decisions concerning invading diseases [1, 2]. However, model-based identification of effective and cost-efficient controls can be difficult, particularly when models include highly detailed representations of disease transmission processes. There is a variety of mathematical tools for designing optimal strategies, but no standard for putting the results from mathematically motivated simplifications into practice. An open question is how to incorporate enough realism into a model to allow accurate predictions of the impact of control measures, whilst ensuring that the truly optimal strategy can still be identified [3]. In this paper we identify the difficulties – as well as

31 potential solutions – in achieving a practically useful optimal strategy, highlighting the potential roles of open
32 loop and model predictive control by way of a simple example.

33 **Realistic landscape-scale modelling**

34 The optimisation of epidemic interventions involves determining the management strategy that minimises
35 impacts of disease. This minimisation can be difficult when resources are limited and there are economic costs
36 associated with both control measures and disease. Methods that simulate the expected course of an epidemic
37 and explicitly model effects of interventions allow us to quantify the potential impact of a given strategy [4].
38 These simulation models accurately capture the dynamics of the real system and so have become important
39 tools in designing intervention strategies to inform policy decisions. Examples include vaccination policies for
40 human papillomavirus in the UK [5, 6], livestock culling policies for foot-and-mouth disease [7, 8] including
41 post-hoc vaccination optimisation [9, 10], and optimal host removal radii for tree diseases of citrus [11–13] and
42 sudden oak death [14].

43 Various complexities of disease dynamics, for example spatial heterogeneities and inherent individual
44 differences in susceptibility and pathogen transmission (risk structure), have been shown to be important
45 determinants of patterns and rates of epidemic spread [15–17]. To ensure accurate epidemic predictions, these
46 factors must be included in simulation models designed to aid decision making. However, inclusion of these
47 factors typically results in highly complex models with many possible control measures, making optimisation
48 computationally infeasible when interventions can be combined, and particularly when control measures can
49 also vary over time, in space or according to disease risk [18]. For most simulation models the only viable
50 option is then to use the model to evaluate a small subset of plausible strategies that remain fixed during the
51 epidemic, potentially scanning over a single parameter such as a culling radius. We shall refer to this approach
52 as ‘Strategy Testing’. Using this approach makes it difficult to have high confidence in the best-performing
53 strategy, since with no framework for choosing it, the set of strategies under test is likely to be biased. Further
54 to this, as the set to test cannot span the entire space of control options, it is unlikely that the true optimum will
55 be found.

56 **Optimal control of epidemiological models**

57 Many mathematical techniques exist for characterising the true optimal control for a disease, such as equilibrium
58 or final size analysis, depending on the system being analysed [15]. We here focus on optimising time-varying
59 control of dynamical systems, for which optimal control theory (OCT) is widely used [19]. The controls
60 considered in this paper are mathematical representations of possible disease interventions. By analysing a set
61 of equations describing the disease dynamics, OCT can mathematically characterise the optimal control and
62 provide insight into the underlying dynamics, without the repeated simulation required to optimise simulation
63 models. However, because of the underlying mathematical complexity, little progress can be made with OCT
64 unless the underpinning models for disease spread are highly simplified. Early work in OCT focussed on
65 optimal levels of vaccination and treatment [20], with extensions to consider further interventions including
66 quarantine, screening, and health-promotion campaigns appearing later [21]. Disease models can also be
67 coupled with economic effects [22–24], and within OCT this has been used to balance multiple costs, such as
68 surveillance and control [25], or prophylactic versus reactive treatment [26].

69 The optimal strategies identified by OCT can be very complex, often specifying controls that switch strategies
70 at specific times during the course of an epidemic. The added complexity of these switching controls can
71 significantly improve disease management when tested on a spatially explicit model, but can perform poorly
72 if the exact time of the switch is not known [27], for example when parameters are not known precisely. This
73 demonstrates that numerical results from OCT are not always directly applicable to the real world because
74 of additional complexities in real systems. It is also unclear however, how insight from OCT alone could
75 be translated into practical advice. To move towards robust strategies that could be used practically, more
76 recent work has focussed on including additional features and heterogeneities into the models used in OCT,
77 in particular spatial dynamics. Space is usually only included to a limited extent, for example by using
78 metapopulation models (e.g. [28, 29]), or partial differential equations (e.g. [30]) to optimise spatial strategies,
79 so whether the heterogeneities added are sufficient to identify robust and practical control strategies remains
80 an open question.

81 **Moving towards practical control**

82 Despite finding the mathematically optimal control strategy, major simplifications to the system as modelled
83 are required to allow progress to be made using OCT. It is therefore often unclear how these strategies would
84 perform if adopted by policy makers. On the other hand, models with sufficient realism to inform policy
85 directly are often impossible to fully optimise. Therefore, a framework is needed to couple the results of OCT
86 with the realism required in policy making and in simulation type models. The question is then how should
87 we make practical use of OCT?

88 In §2 we describe two methods from control systems engineering for applying OCT results, and compare
89 these versus Strategy Testing using a simple illustrative model in §3. We seek to answer how, under current
90 computational constraints, results from OCT can be applied whilst maintaining the realism required for practical
91 application.

92 **2 Applying optimal control to realistic systems**

93 Outside of epidemiology, OCT has had wider use on approximate models of complex systems. A recent study
94 reviews the use of OCT for agent-based models (ABMs) [31], a type of model that simulates the individual
95 behaviour of autonomous agents. An *et al.* [31] suggest the use of a model that approximates the dynamics
96 of the ABM, designed to be simple enough to allow mathematical analysis of the optimal control. A suitable
97 approximate model is chosen and fitted either to real data, or to synthetic data from the ABM. The OCT results
98 from the approximating model are then mapped onto the ABM to be tested: a process referred to as 'lifting',
99 which could equally well apply to the detailed epidemic simulation models considered in this paper. We
100 now describe two possible frameworks from control systems engineering for making use of this control lifting
101 approach.

102 **Open-loop control**

103 The first method is the simplest application of control lifting, and the framework implicitly suggested by An *et al.*
104 [31]. Control is optimised on the approximate model once using the initial conditions of the simulation model,

105 and this control is lifted to the simulator and used for the full simulation run time (figure 1). Crucially, any time
106 variation in this control is fixed at the very start of the simulation, and so does not incorporate any feedback. It
107 is therefore referred to as ‘open-loop’ control, as it is fully specified by the simulation initial conditions and the
108 trajectory predicted by the approximate model. Use in epidemiology is uncommon, although Clarke *et al.* [32]
109 use OCT in an approximate model to find optimal levels of Chlamydia screening and contact tracing which are
110 then mapped onto a network simulation.

111 **Model predictive control**

112 Open-loop control requires the approximate model to remain accurate over a long time scale. Since simulation
113 models typically include many heterogeneities that must be omitted from the approximate model, systematic
114 deviations between the simulation and approximate predicted trajectories for disease progress are highly likely.
115 Model predictive control (MPC) is an optimisation technique incorporating system feedback [33, 34]. At regular
116 update times the values of the state variables in the approximate model are reset to match the conditions in the
117 simulation at that time. The control is then re-optimised and the new control strategy is used going forward in
118 the simulation to the next update time. This corresponds to a series of open-loop problems solved at regular
119 update steps (figure 1). Model predictive control can therefore take into account perturbations in the disease
120 progress trajectory caused by heterogeneities omitted from the approximate model.

121 Model predictive control has had some use within the epidemiological literature, the majority being for
122 control of drug applications for single individuals rather than control of epidemics at the population level.
123 Examples include finding management strategies for HIV that are robust to measurement noise and modelling
124 errors [35, 36], and control of insulin delivery in patients with diabetes [37]. These studies highlight the benefits
125 of MPC for robust control, i.e. control that remains effective despite system perturbations. However, only one
126 study concentrates on epidemic management [38], and that does not explicitly test the feedback control on
127 simulations.

128 **3 Optimising strategies on an illustrative network model**

129 **Methods**

130 To demonstrate open-loop and MPC for epidemic management we use a stochastic SIR network model including
131 host demography and risk structure. The model is deliberately kept simple to show how the underpinning
132 idea is broadly applicable across human, animal and plant diseases. Whilst the model and its parameters are
133 arbitrary we use it to represent a scenario in which a simulation model has already been fitted to a real disease
134 system; the network model is therefore used here as a proxy for a potentially very detailed simulation model.

135 **Simulation Model**

136 In our model, infection spreads stochastically across a network of nodes that are clustered into three distinct
137 regions (figure 2a). Each node contains a host population stratified into high and low risk groups. The infection
138 can spread between individuals within nodes and between connected nodes. The net rate of infection of risk

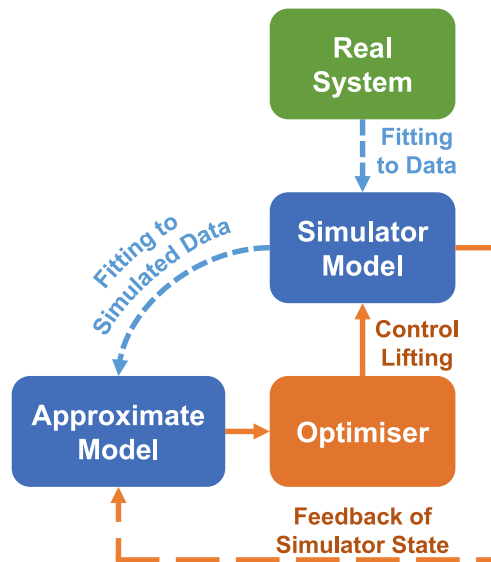


Figure 1: Open-loop and model predictive control (MPC). The model hierarchy is shown, with optimised controls from the approximate model directly lifted to the simulation model. The real system is in green, the models and fitting processes are in blue, and the control framework is in orange. Without the orange dashed feedback loop, this is open-loop control. MPC resets the state of the approximate model at regular update steps, before re-optimising and lifting controls to the simulation model until the next update time.

139 group r in node i is given by:

$$140 \quad S_i^r \sum_j \sigma_{ij} \left(\rho^{rH} I_j^H + \rho^{rL} I_j^L \right), \quad (1)$$

141 where S and I are numbers of susceptible and infected hosts respectively, subscripts identify the node, and
 142 superscripts specify high (H) or low (L) risk group. The sum is over all connected nodes including the focal
 143 node itself, with pairwise connectivities given by σ_{ij} , and risk structure given by ρ . Full details of the model
 144 are given in the supplementary material. Although not limited to these applications, the model in Equation 1
 145 could represent crop or livestock diseases spreading through farms, or sexually transmitted infections spreading
 146 through towns, cities or countries.

147 Mass vaccination is the only intervention we consider, with the potential to target based on both risk group
 148 and region but randomised across host infection status (i.e. susceptibles cannot be preferentially vaccinated).
 149 Logistical and economic constraints are included through a maximum total vaccination rate (η_{\max}) that can be
 150 divided between risk groups and regions. Within each group susceptibles are vaccinated at rate: $f \eta_{\max} S/N$,
 151 where f is the proportion of control allocated to that group, and N is the total group population.

152 Optimal allocation of the vaccination resources minimises an objective function J representing the disease
 153 burden of the epidemic across all infected hosts over the simulation time (T): $J = \int_{t=0}^T I(t) dt$. In common with
 154 the particular control we consider and the risk and spatial structures, this simple choice of objective function
 155 was made merely to illustrate our methods, but the framework generalises immediately to more complex
 156 settings.

157 **Approximate Models**

158 Exhaustive optimisation of control using the simulation model, across space, risk group and time, is clearly
 159 infeasible. We consider two different deterministic approximate models of the simulator to assess the best level

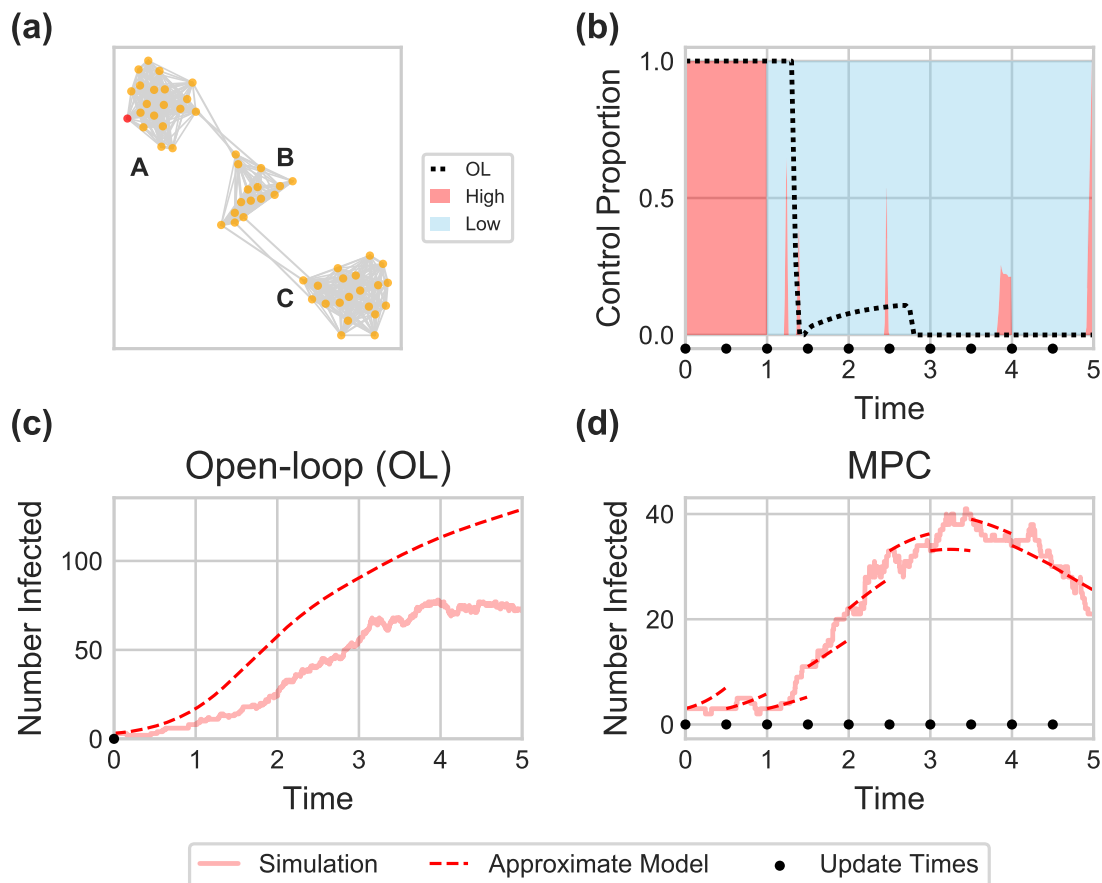


Figure 2: (a) shows the network used for the illustrative simulation model, including region labels. The epidemic is seeded in the red node in region A, and can spread between connected nodes (grey lines). In (b) the control allocation is shown for a single space based MPC run, with the corresponding open-loop allocation indicated by the black dotted line. (c) shows the total number of infected individuals under a single run of space based open-loop control. Control is based on the prediction of the approximate model starting from the initial conditions. (d) shows the number of infected individuals in the simulation and space based approximate model corresponding to the MPC control carried out in (b). Here the prediction is reset to match the simulation at every update step (0.5 time units) and the control is re-optimised. By taking account of differences in the number of infected individuals compared with those predicted at the initial time, MPC gives better predictions of the simulation state as well as improved control when compared to open-loop control (note different y axis scales).

160 of approximation. The first model is purely risk structured, factoring out all spatial information and leaving
 161 one high risk and one low risk population group. The second approximate model is more complex, in as
 162 much as it is also risk structured, but additionally includes a first approximation to the host spatial structure by
 163 dividing the host population into three large regions. Spatial dynamics are included between but not within
 164 the three regions, maintaining enough simplicity to obtain optimal control results. This could represent, for
 165 example, optimising control at the country level, but not at the regional level. We will refer to this model as
 166 the spatial approximate model. A single set of parameters is fitted for each model to data from an ensemble
 167 of simulation model runs. We then test which of the two approximate models is the more useful for control
 168 optimisation. Full details of the approximate models, fitting and optimisation procedures are given in the
 169 supplementary material.

170 **Control Scenarios**

171 We test six different control scenarios, which compare Strategy Testing (scenarios 1 and 2) with open-loop and
172 MPC applied to both of our approximate models (scenarios 3 to 6):

- 173 1. 'High': exclusively vaccinate high risk individuals
- 174 2. 'Split': partition control resources between high and low risk groups based on an optimisation performed
175 in advance
- 176 3. 'Risk OL': open-loop control on the risk based approximate model
- 177 4. 'Risk MPC': MPC on the risk based approximate model
- 178 5. 'Space OL': open-loop control on the spatial approximate model
- 179 6. 'Space MPC': MPC on the spatial approximate model

180 The optimal constant allocation for the 'Split' strategy was found by selecting the proportionate allocation to
181 each risk group which minimised the mean disease cost in a set of simulation realisations. (supplementary
182 material text and figure S8). Figures 2c and 2d illustrate the open-loop and MPC strategies, showing the model
183 resetting at regular intervals in MPC.

184 **Results**

185 The risk based optimisation results in a strategy that initially vaccinates high risk individuals, before switching
186 priorities and treating the more populous low risk group almost exclusively (figure 2b). This same switch is
187 seen in the optimal controls from the spatial approximate model, but a number of spatial switches are also seen
188 (figure S9). The spatial strategies are therefore much more complex than the risk based controls. Comparing
189 disease costs shows that incorporating greater realism, through a more complex approximate model as well
190 as by using MPC, allows for improved management of the epidemic (figures 3 and S10). Of the constant
191 'user-defined' strategies, splitting control between risk groups is more effective than just vaccinating the high
192 risk group. The optimal allocation to the high risk group is 63% of vaccination resources, with the rest used
193 to vaccinate low risk individuals. Using the risk-based approximate model gives an improvement over either
194 of these strategies, although there is little difference between open-loop and MPC (see below). Adding space
195 into the approximate model improves control further, leading to the smallest epidemic costs when the MPC
196 framework is used.

197 The illustrative model demonstrates the management improvements that can be achieved by combining
198 OCT with both open-loop and MPC. The key results of the OCT analyses are the control switching times. Using
199 the switching controls from either approximate model with open-loop control gives lower epidemic costs than
200 the naively chosen 'user-defined' strategies. Feedback allows the MPC controllers to re-evaluate the timing of
201 these switches during the epidemic, and potentially include additional switches, to match the exact trajectory
202 of the current realisation (figures 2b-d).

203 In the risk based strategies there is little difference between open-loop and MPC. This is because the precise
204 timing of the switch from high to low risk group vaccination does not significantly affect the epidemic cost
205 (supplementary figure S11). The timings of disease introduction into regions B and C are highly variable

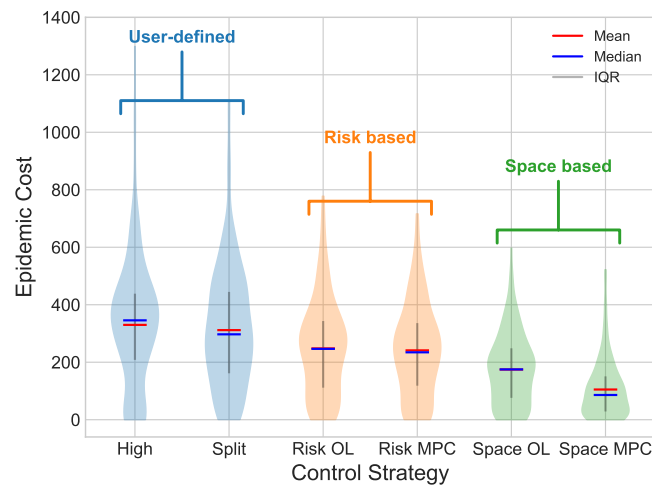


Figure 3: Results of different control optimisation schemes on the example simulation model. Spatial MPC performs best, showing an improvement over both open-loop and user-defined strategies.

206 between simulation runs (supplementary figure S2). The potential for additional switches in the spatial
207 approximate model gives more flexibility for the MPC controller to respond to this variability, and so spatial
208 MPC shows a significant improvement over open-loop. Extending the objective to include a cost associated
209 with switching strategy would allow a more detailed assessment of the practicality of implementing these
210 complex strategies, although to provide a simple demonstration of our key ideas we have not pursued this
211 here.

212 In the spatial case, open-loop performs worse than MPC as it is unable to adapt to perturbations in the
213 dynamics. The control must then rely on the model predicted spread, necessarily an approximation to the
214 actual trajectory. The performance of the control is closely linked to the accuracy of this approximation,
215 raising the important issue of choosing the approximate model. In our example, spatial dynamics are clearly
216 important because of the timing of spread between regions, and so the more informed controls of the spatial
217 model outperform the risk based strategies.

218 4 Discussion and outstanding questions

219 Our results show that the choice of approximate model affects the performance of both open-loop and MPC
220 strategies. Here we have found a suitable approximate model in an ad hoc manner, but a key challenge for the
221 future is to develop a more formal method for choosing the most appropriate approximate model. A more
222 accurate model will give better predictions, and hence control that is closer to the true optimum, but this must
223 be balanced with added complexity and optimisation constraints. One difficulty in doing this is that it is not
224 always clear where the boundary of mathematical or computational feasibility is, and so how complex the
225 model can be made in practice. It is also difficult to mathematically determine, in a systematic way, which
226 aspects of the dynamics are important to capture accurately. This key issue must be considered though, since
227 the implications relate directly to applications in the real world.

228 Practical disease control requires surveys of the real system to assess the state of the epidemic. Both open-
229 loop and MPC optimise control using predictions of the future dynamics, making them both feed-forward
230 controllers. The approximate model underlying these frameworks allows more informed decisions between

231 surveys, resulting in control that is closer to the true optimum. Accurate predictions can avoid continuous
232 or very frequent surveys which may be expensive or logistically challenging. The repeated updates of MPC
233 improve these predictions but will each require surveillance of the real system, so the frequency of updates
234 must be chosen so as to balance improved knowledge of the system with any surveillance constraints.

235 In addition to this feed-forward control, MPC incorporates a feedback loop not included in open-loop
236 strategies. The feedback mechanism, included through the update steps, allows for control that copes better
237 with uncertainty, i.e. it is more robust. As shown by [35], MPC can still successfully control a system when
238 measurements of that system are inaccurate, or the parameters in the simple approximate model are incorrect.
239 In our case this allowed it to perform better on the complex simulation model, where additional heterogeneities
240 were present. Open-loop strategies are not robust to these perturbations so may not perform well in practice.
241 For certain systems and parameters the open-loop strategy may in fact perform worse than naive constant
242 control [28]. We also see this for certain parameter sets in the illustrative model. It would be easy to dismiss
243 the OCT results as not useful for policy, but our results show that by using feedback, these results can still be
244 beneficial.

245 In this paper we have focussed on a top-down approach, finding robust, practically-applicable strategies by
246 making use of OCT to optimise simulation models. Equally, many studies analyse dynamical systems using
247 OCT without simulation models. These studies rarely consider practical application of the resulting optimal
248 controls, and so may benefit from a framework for testing the results. With this bottom-up approach, a system
249 for testing the results on realistic systems is vital to ensure that these results are robust to additional realism.
250 Using an MPC framework as considered here could be one way in which OCT researchers could demonstrate
251 the potential impact of their work to a wider audience.

252 We have assumed throughout that an accurate simulation model of the real system in question can be
253 built, and that a single set of parameters can be fitted for the chosen deterministic approximate model. In
254 reality there may be considerable uncertainty in parameters for the simulator so fitting a single deterministic
255 model may be challenging. A question for future study would be how to handle these uncertainties, perhaps
256 also incorporating improved knowledge of parameters as the simulation proceeds [39]. A final challenge
257 for optimal control more generally relates to the choice of objective function. The strategies found by OCT
258 are highly dependent on the exact form of the objective function, and so more research is needed into how
259 to quantify the balancing of very different costs, for example treatment costs and disease burden [28]. In
260 human disease, cost-effectiveness analyses are usually based on quality adjusted life years (QALYs) [40], and
261 a maximum economic cost per QALY gained. A similar concept could perhaps be used for plant and animal
262 diseases, including calculations of yield losses [41] as well as effects on welfare, biodiversity and tourism for
263 example [42].

264 **5 Conclusions**

265 Simulation models that capture the complex dynamics and uncertainties necessary for informing policy have
266 significant limitations when designing optimal time-dependent management strategies. Optimal control theory
267 is a powerful mathematical approach for characterising these interventions, but results have proved difficult to
268 translate into effective policy advice. Increased focus on how optimal control can be used and translated to policy
269 makers would allow for more robust decisions on disease management, incorporating greater epidemiological

270 understanding.

271 Lifting of controls from OCT to realistic simulations as in [31] goes some way to addressing this problem,
272 but the differences in model structure between the simulation and approximate models is likely to limit the
273 practical performance of the control measures. In this paper we show that incorporating feedback can alleviate
274 this problem and help lead to robust and practical control strategies. Whilst these techniques may be able
275 to transfer optimal control results to more realistic simulations and so to practical application, it does raise
276 the issue of communicability of results. With complex feedback strategies between two models, one complex
277 in structure and the other mathematically complex, the overall result is no longer simple to explain. How
278 applicable these strategies can ever be, will then be determined by the reliability of, and trust in, the simulation
279 model, and its ability to establish the benefit of a more complex intervention strategy.

280 Overall, there is clearly a benefit to policy and management decisions from making use of OCT results. We
281 suggest that potential policies, including OCT results, can and should be tested on realistic simulations, and
282 that MPC is an effective platform for practical applications of OCT.

283 **Additional Information**

284 **Acknowledgements**

285 We thank Andrew Craig, Eleftherios Avramidis and Hola Adrakey for helpful discussions.

286 **Data Accessibility**

287 All code and generated data are available at <https://github.com/ehbussell/Bussell2018Model>.

288 **Authors' Contributions**

289 E.H.B., C.E.D. and N.J.C. designed the study, E.H.B. conducted the analysis and wrote the initial draft of the
290 manuscript. All authors contributed to data interpretation, manuscript editing and discussion.

291 **Competing Interests**

292 We have no competing interests.

293 **Funding**

294 E.H.B. acknowledges the Biotechnology and Biological Sciences Research Council of the United Kingdom
295 (BBSRC) for support via a University of Cambridge DTP Ph.D. studentship.

296 **References**

- 297 [1] Heesterbeek H, et al. 2015. Modeling infectious disease dynamics in the complex landscape of global health. *Science* **347**, aaa4339
298 doi:[10.1126/science.aaa4339](https://doi.org/10.1126/science.aaa4339)
- 299 [2] Metcalf CJE, Edmunds WJ, Lessler J. 2015. Six challenges in modelling for public health policy. *Epidemics* **10**, 93–96 doi:[10.1016/j.epidem.2014.08.008](https://doi.org/10.1016/j.epidem.2014.08.008)
- 300 [3] Cunniffe NJ, Koskella B, E Metcalf CJ, Parnell S, Gottwald TR, Gilligan CA. 2015. Thirteen challenges in modelling plant diseases. *Epidemics* **10**, 6–10
301 doi:[10.1016/j.epidem.2014.06.002](https://doi.org/10.1016/j.epidem.2014.06.002)

- 302 [4] Lofgren ET, et al. 2014. Opinion: mathematical models: a key tool for outbreak response. *Proc. Natl. Acad. Sci. U.S.A.* **111**, 18095–18096
303 doi:[10.1073/pnas.1421551111](https://doi.org/10.1073/pnas.1421551111)
- 304 [5] Jit M, Choi YH, Edmunds WJ. 2008. Economic evaluation of human papillomavirus vaccination in the United Kingdom. *BMJ* **337**, a769
305 doi:[10.1136/bmj.a769](https://doi.org/10.1136/bmj.a769)
- 306 [6] Choi YH, Jit M, Gay N, Cox A, Garnett GP, Edmunds WJ. 2010. Transmission dynamic modelling of the impact of human papillomavirus vaccination in
307 the United Kingdom. *Vaccine* **28**, 4091–4102 doi:[10.1016/j.vaccine.2009.09.125](https://doi.org/10.1016/j.vaccine.2009.09.125)
- 308 [7] Keeling MJ, Woolhouse ME, Shaw DJ, Matthews L, Chase-Topping M, Haydon DT, Cornell SJ, Kappey J, Wilesmith J, Grenfell BT. 2001. Dynamics of the
309 2001 UK foot and mouth epidemic: stochastic dispersal in a heterogeneous landscape. *Science* **294**, 813–817 doi:[10.1126/science.1065973](https://doi.org/10.1126/science.1065973)
- 310 [8] Keeling MJ. 2005. Models of foot-and-mouth disease. *Proc. R. Soc. Lond. B Biol. Sci.* **272**, 1195–1202 doi:[10.1098/rspb.2004.3046](https://doi.org/10.1098/rspb.2004.3046)
- 311 [9] Keeling MJ, Woolhouse MEJ, May RM, Davies G, Grenfell BT. 2003. Modelling vaccination strategies against foot-and-mouth disease. *Nature* **421**, 136–42
312 doi:[10.1038/nature01343](https://doi.org/10.1038/nature01343)
- 313 [10] Tildesley MJ, Savill NJ, Shaw DJ, Deardon R, Brooks SP, Woolhouse ME, Grenfell BT, Keeling MJ. 2006. Optimal reactive vaccination strategies for a
314 foot-and-mouth outbreak in the UK. *Nature* **440**, 83–86 doi:[10.1038/nature04324](https://doi.org/10.1038/nature04324)
- 315 [11] Cunniffe NJ, Stutt RO, DeSimone RE, Gottwald TR, Gilligan CA. 2015. Optimising and communicating options for the control of invasive plant disease
316 when there is epidemiological uncertainty. *PLoS Comput. Biol.* **11**, e1004211 doi:[10.1371/journal.pcbi.1004211](https://doi.org/10.1371/journal.pcbi.1004211)
- 317 [12] Hyatt-Twynam SR, Parnell S, Stutt RO, Gottwald TR, Gilligan CA, Cunniffe NJ. 2017. Risk-based management of invading plant disease. *New Phytol.* **214**,
318 1317–1329 doi:[10.1111/nph.14488](https://doi.org/10.1111/nph.14488)
- 319 [13] Adrakey HK, Streftaris G, Cunniffe NJ, Gottwald TR, Gilligan CA, Gibson GJ. 2017. Evidence-based controls for epidemics using spatio-temporal
320 stochastic models in a Bayesian framework. *J. R. Soc. Interface* **14**, 20170386 doi:[10.1098/rsif.2017.0386](https://doi.org/10.1098/rsif.2017.0386)
- 321 [14] Cunniffe NJ, Cobb RC, Meentemeyer RK, Rizzo DM, Gilligan CA. 2016. Modeling when, where, and how to manage a forest epidemic, motivated by
322 sudden oak death in California. *Proc. Natl. Acad. Sci. U.S.A.* **113**, 5640–5645 doi:[10.1073/pnas.1602153113](https://doi.org/10.1073/pnas.1602153113)
- 323 [15] Keeling MJ, Rohani P, 2008. *Modeling Infectious Diseases in Humans and Animals*. Princeton University Press
- 324 [16] Anderson RM, Medley GF, May RM, Johnson AM. 1986. A preliminary study of the transmission dynamics of the human immunodeficiency virus (HIV),
325 the causative agent of AIDS. *Math. Med. Biol.* **3**, 229–263 doi:[10.1093/imammb/3.4.229](https://doi.org/10.1093/imammb/3.4.229)
- 326 [17] Smith DL, Lucey B, Waller LA, Childs JE, Real LA. 2002. Predicting the spatial dynamics of rabies epidemics on heterogeneous landscapes. *Proc. Natl.*
327 *Acad. Sci. U.S.A.* **99**, 3668–3672 doi:[10.1073/pnas.042400799](https://doi.org/10.1073/pnas.042400799)
- 328 [18] Bellman R, 2013. *Dynamic Programming*. Courier Corporation
- 329 [19] Lenhart S, Workman JT, 2007. *Optimal Control Applied to Biological Models*. CRC Press
- 330 [20] Sethi SP, Staats PW. 1978. Optimal control of some simple deterministic epidemic models. *J. Oper. Res. Soc.* **29**, 129–136 doi:[10.1057/jors.1978.27](https://doi.org/10.1057/jors.1978.27)
- 331 [21] Behncke H. 2000. Optimal control of deterministic epidemics. *Optim. Control Appl. Methods* **21**, 269–285 doi:[10.1002/oca.678](https://doi.org/10.1002/oca.678)
- 332 [22] Perrings C, et al. 2014. Merging economics and epidemiology to improve the prediction and management of infectious disease. *EcoHealth* **11**, 464–475
333 doi:[10.1007/s10393-014-0963-6](https://doi.org/10.1007/s10393-014-0963-6)
- 334 [23] Epanchin-Niell RS. 2017. Economics of invasive species policy and management. *Biol. Invasions*, 1–22 doi:[10.1007/s10530-017-1406-4](https://doi.org/10.1007/s10530-017-1406-4)
- 335 [24] Cunniffe NJ, Laranjeira FF, Neri FM, DeSimone RE, Gilligan CA. 2014. Cost-effective control of plant disease when epidemiological knowledge is
336 incomplete: modelling Bahia bark scaling of citrus. *PLOS Comput. Biol.* **10**, 1–14 doi:[10.1371/journal.pcbi.1003753](https://doi.org/10.1371/journal.pcbi.1003753)
- 337 [25] Ndeffo Mbah ML, Gilligan CA. 2010. Balancing detection and eradication for control of epidemics: Sudden oak death in mixed-species stands. *PLoS*
338 *ONE* **5**, 1–8 doi:[10.1371/journal.pone.0012317](https://doi.org/10.1371/journal.pone.0012317)
- 339 [26] Brown VL, White KAJ. 2011. The role of optimal control in assessing the most cost-effective implementation of a vaccination programme: HPV as a case
340 study. *Math. Biosci.* **231**, 126–134 doi:[10.1016/j.mbs.2011.02.009](https://doi.org/10.1016/j.mbs.2011.02.009)
- 341 [27] Forster G, Gilligan CA. 2007. Optimizing the control of disease infestations at the landscape scale. *Proc. Natl. Acad. Sci. U.S.A.* **104**, 4984–4989
342 doi:[10.1073/pnas.0607900104](https://doi.org/10.1073/pnas.0607900104)
- 343 [28] Rowthorn RE, Laxminarayan R, Gilligan CA. 2009. Optimal control of epidemics in metapopulations. *J. R. Soc. Interface* **6**, 1135–1144
344 doi:[10.1098/rsif.2008.0402](https://doi.org/10.1098/rsif.2008.0402)
- 345 [29] Ndeffo Mbah ML, Gilligan CA. 2011. Resource allocation for epidemic control in metapopulations. *PLoS ONE* **6**, doi:[10.1371/journal.pone.0024577](https://doi.org/10.1371/journal.pone.0024577)

- 346 [30] Neilan RM, Lenhart S. 2011. Optimal vaccine distribution in a spatiotemporal epidemic model with an application to rabies and raccoons. *J. Math. Anal.*
347 *Appl.* **378**, 603–619 doi:[10.1016/j.jmaa.2010.12.035](https://doi.org/10.1016/j.jmaa.2010.12.035)
- 348 [31] An G, Fitzpatrick BG, Christley S, Federico P, Kanarek A, Neilan RM, Oremland M, Salinas R, Laubenbacher R, Lenhart S. 2017. Optimization and control
349 of agent-based models in biology: a perspective. *Bull Math Biol* **79**, 63–87 doi:[10.1007/s11538-016-0225-6](https://doi.org/10.1007/s11538-016-0225-6)
- 350 [32] Clarke J, White KA, Turner K. 2013. Approximating optimal controls for networks when there are combinations of population-level and targeted
351 measures available: chlamydia infection as a case-study. *Bull. Math. Biol.* **75**, 1747–1777 doi:[10.1007/s11538-013-9867-9](https://doi.org/10.1007/s11538-013-9867-9)
- 352 [33] Camacho EF, Bordons C, 2012. *Model Predictive Control in the Process Industry*. Springer Science & Business Media
- 353 [34] Lee JH. 2011. Model predictive control: review of the three decades of development. *Int. J. Control Autom. Syst.* **9**, 415–424 doi:[10.1007/s12555-011-0300-6](https://doi.org/10.1007/s12555-011-0300-6)
- 354 [35] Zurawski R, Teel AR. 2006. A model predictive control based scheduling method for HIV therapy. *J. Theor. Biol.* **238**, 368–382 doi:[10.1016/j.jtbi.2005.05.004](https://doi.org/10.1016/j.jtbi.2005.05.004)
- 355 [36] David J, Tran H, Banks HT. 2011. Receding horizon control of HIV. *Optim. Control Appl. Meth.* **32**, 681–699 doi:[10.1002/oca.969](https://doi.org/10.1002/oca.969)
- 356 [37] Hovorka R, et al. 2004. Nonlinear model predictive control of glucose concentration in subjects with type 1 diabetes. *Physiol. Meas.* **25**, 905–920
357 doi:[10.1088/0967-3334/25/4/010](https://doi.org/10.1088/0967-3334/25/4/010)
- 358 [38] Sélley F, Besenyei Á, Kiss IZ, Simon PL. 2015. Dynamic control of modern, network-based epidemic models. *SIAM J. Appl. Dyn. Syst.* **14**, 168–187
359 doi:[10.1137/130947039](https://doi.org/10.1137/130947039)
- 360 [39] Thompson RN, Gilligan CA, Cunniffe NJ. 2018. Control fast or control smart: when should invading pathogens be controlled? *PLoS Comput. Biol.* **14**,
361 e1006014 doi:[10.1371/journal.pcbi.1006014](https://doi.org/10.1371/journal.pcbi.1006014)
- 362 [40] Whitehead SJ, Ali S. 2010. Health outcomes in economic evaluation: the QALY and utilities. *Br. Med. Bull.* **96**, 5–21 doi:[10.1093/bmb/ldq033](https://doi.org/10.1093/bmb/ldq033)
- 363 [41] Savary S, Ficke A, Aubertot JN, Hollier C. 2012. Crop losses due to diseases and their implications for global food production losses and food security.
364 *Food Secur.* **4**, 519–537 doi:[10.1007/s12571-012-0200-5](https://doi.org/10.1007/s12571-012-0200-5)
- 365 [42] Boyd IL, Freer-Smith PH, Gilligan CA, Godfray HCJ. 2013. The consequence of tree pests and diseases for ecosystem services. *Science* **342**, 1235773
366 doi:[10.1126/science.1235773](https://doi.org/10.1126/science.1235773)