



HAL
open science

Using Markov Decision Process to define an adaptive strategy to control the spread of an animal disease

Anne-France Viet, Laurent Jeanpierre, Maroua Bouzid, Abdel-Allah Mouaddib

► **To cite this version:**

Anne-France Viet, Laurent Jeanpierre, Maroua Bouzid, Abdel-Allah Mouaddib. Using Markov Decision Process to define an adaptive strategy to control the spread of an animal disease. *Computers and Electronics in Agriculture*, 2012, 80, pp.71-79. hal-00965368

HAL Id: hal-00965368

<https://hal.science/hal-00965368>

Submitted on 25 Mar 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Using Markov Decision Processes to define an adaptive strategy to control the spread of an animal disease

A.-F. Viet ^{a,b,c,d,e,*}, L. Jeanpierre ^{a,b,c}, M. Bouzid ^{a,b,c}, A.-I. Mouaddib ^{a,b,c}

^a Université Caen Basse-Normandie, UMR 6072 GREYC, BP 5186, F-14032 Caen, France

^b CNRS, UMR 6072 GREYC, BP 5186, F-14032 Caen, France

^c ENSICAen, UMR 6072 GREYC, BP 5186, F-14032 Caen, France

^d INRA, UMR 1300 Bioagression, Epidémiologie et Analyse de Risque, BP 40706, F-44307 Nantes Cedex 03, France

^e LUNAM Université, Oniris, BP 40706, F-44307 Nantes Cedex 03, France

A B S T R A C T

A strategy to control an animal disease within an area is often based on one or more actions systematically implemented. In this paper, we illustrate how to use a Markov Decision Process (MDP) to compute an adaptive strategy depending on the pathogen spread within a group of farmers with only one decision-maker for the group. The objective at the group level is to decrease the cost of the disease and its control. Status for each farm is assumed to be exactly known each year by the decision-maker. Possible actions each year are Doing nothing or Vaccinating. The computed MDP policy results in a non-systematic vaccination. Although the objective is only based on the total costs, the computed MDP policy reduces the prevalence, that is the amount of infected herds, compared to a systematically Doing nothing strategy.

Keywords:
Artificial intelligence
Epidemiology
Policy
Vaccination
Animal

1. Introduction

The decision to control the spread of a pathogen within a livestock population is related to various actors (farmers, veterinarians, associations...) who consider different criteria. For regulated animal diseases such as Foot-and-Mouth Disease (FMD), strategy for control is then defined at the country level and is mandatory. For other animal diseases, the farmer often decides whether to control them or not, on a voluntary basis. Their decisions are based on criteria, often economical ones (Klein et al., 2007). For a disease that can be transmitted by animal purchases, direct contact or by air, such a decision in a single farm is rarely sufficient as the pathogen spread within a farm is influenced by its spread in neighbouring farms. Since some farmers are grouped in associations, decision making for such disease control can be defined at the level of a group of farms to enhance its efficiency. The group proposes which control action should be implemented over time and farmers apply it.

To define a control strategy at a group level, the group of farmers should have an objective and should take into account both

uncertainty about the future pathogen spread and that some farmers may not always apply the proposed action. Example objective can be the limitation of the cost of the disease spread or the number of infected farms. Pathogen spread is not a deterministic process as transmission is not systematic. It results in uncertainty about the pathogen spread; mainly on the possible extent of infection or clearance (Dangerfield et al., 2008). To make the right decision, the variability of the pathogen spread should be taken into account. Moreover, the group should assume that some farmers may refuse or wrongly implement the retained policy. The fact that not all farmers take an action may decrease the efficiency of a strategy (Vonk Noordegraaf et al., 1998).

Modelling approaches are often used to explore the impact of decisions on pathogen spread in order to inform a leader for its decision. Classically, strategy consists of only one action applied every time but very few approaches consider different application times of the action. Often, studies concern the efficiency of a strategy consisting in applying an action either every year (O'Callaghan et al., 1999) or at each new case detection (Backer et al., 2009; Tildesley et al., 2006; LeMenach et al., 2005; Keeling et al., 2003; Ferguson et al., 2001). For pathogen spread within a livestock population, the definition of rules to adapt a strategy according to the current prevalence of the infection in an area is often based on expert knowledge (such as in Seegers et al., 2006). The best action at a time can also be determined by considering a decision tree based on the expected utility of actions (Rat-Aspert and Fourichon, 2010) or

* Corresponding author at: UMR 1300 BioEpAR, Oniris, BP 40706, F-44307 Nantes Cedex 3, France. Tel.: +33 (0) 2 40 68 78 54; fax: +33 (0) 2 40 68 77 68.

E-mail addresses: anne-france.viet@oniris-nantes.fr (A.-F. Viet), laurent.jeanpierre@unicaen.fr (L. Jeanpierre), maroua.bouzid-mouaddib@unicaen.fr (M. Bouzid), abdel-illah.mouaddib@unicaen.fr (A.-I. Mouaddib).

by simulating possible behaviours in the future (Merl et al., 2009) or even by other optimisation methods (Kristensen, 2003; Toft et al., 2005; Kobayashi et al., 2007; Ge et al., in press, 2010).

Markov Decision Processes (MDP) are classically used to define a policy optimising an objective function under uncertainty in various domains but rarely in animal epidemiology. MDP can formalise a problem with a choice of actions at each time step, actions that influence the future evolution of the system (Puterman, 2005). They are often used for planning in robotics and for managing resources. For example, Forsell et al. (2011) proposed a graph-based MDP to manage forests considering wind damage. For the study of pathogen spread within a livestock population, few studies used MDP-based approaches. One model concerning one herd used a MDP-based approach to define actions to manage the herd when there is one pathogen spread (Toft et al., 2005). For decision to control FMD in Ge et al. (in press, 2010), an MDP coupled with a Bayesian approach to estimate some epidemiological parameters of a Multi-level hierarchical MDP (Kristensen and Jorgensen, 2000) was used. In the case of an epidemic disease for which eradication should be achieved, the action “Do Nothing” is possible to represent flexibility before taking a specific action (Ge et al., 2007). For an endemic disease, control is not regulated and the action “Do Nothing” is a real option to consider. Decision-makers face the choice of controlling the disease or doing nothing. To our knowledge, no MDP approach was ever used to compute a policy for endemic diseases at the area level.

Defining a policy at the area level without considering explicitly each decision at farm level avoids the problem of the model dimension. For endemic diseases, decisions are taken by each farmer. When considering a global objective at the area level, the interaction of all individual decisions should be considered. We face then a multi-component problem as in (Forsell et al., 2011; Ben-Ari and Gal, 1986; Kristensen, 1992). For such a problem, a multi-agent Markov Decision Process (MMDP, Boutilier, 1999) or a decentralized Markov Decision Process (Dec-MDP, Bernstein et al., 2002) may be considered. The complexity of these models depends on the number of agents and on the number of individual actions. In a given region, there can be much more than 100 farms, making these models impossible to use. For example, a MMDP with 10 farms, two actions and only three individual statuses would require more than $3^{10} * 2^{10}$ (more than 60 millions) tests for each considered time-step. For a multi-component problem, approximations were obtained either using a specific approach (parameter iteration, Ben-Ari and Gal, 1986; Kristensen, 1992) or considering a low number of relations between components to decompose the problem in smaller ones (Forsell et al., 2011). For our problem, we consider here a simplified situation at the group level to avoid approximate solution. To take into account that individual decisions may be different from the one given at the group level, we consider that some farmers would not comply with the chosen policy of the decision-maker at the group level. This induces lower transition probabilities resulting from the action than expected.

In this paper we show how to use an MDP in order to propose a global adaptive strategy to control a theoretical endemic pathogen spread within a group of farms. First, we describe the model assumptions and we give the formulation as an MDP. Results for the model applied to a theoretical pathogen spread are then presented. Finally, the potential use of this model for a real pathogen spread is discussed.

2. Methods

2.1. Problem description

In this paper, we consider a theoretical pathogen spread. We assume that the pathogen can be transmitted by direct contacts and

by purchase of animals. Once introduced into a farm, the transmission of the pathogen is assumed to be high in order to have only one infectious status. In few months, all animals are infected. We assumed that unspecified actions are taken to limit the infection within the herd. So, after a few years, infected herds can become free of infection and then susceptible again to a new infection by the pathogen. There is no immunity resulting from an infection. For prevention, we assumed that a vaccine is available to protect the susceptible farms of infection, but some farmers may be reluctant to use it. The vaccination induces a temporary immunity, and subsequently a vaccinated herd will become susceptible again. The vaccination is assumed here of no value for the infected farms.

We consider a group of N farms in an area. Each farm can buy animals from other farms in the group and from farms not in the group. Farmers, except for a few ones, follow the directive of the leader of the group. The decision-maker in this study is then the leader of the group.

Each year, the decision-maker should propose an action that would be applied by all the susceptible farmers. Two alternatives are available each year: Doing nothing or Vaccinating. To make his decision, the leader performs tests in each farm to determine its status regarding the pathogen spread: susceptible to an infection versus infected versus vaccinated. To simplify, we assumed that these tests are 100% specific and sensitive. Relaxing this hypothesis will be discussed in the last part of this paper.

The decision-maker makes his decision considering the economic consequences of the pathogen spread and the cost induced by vaccination.

2.2. Formal framework

We begin by a short introduction to Markov Decision Process (MDP) models, before presenting the MDP model used to model our problem and the associated epidemiological model. The MDP framework provides a formal description for modelling a large variety of stochastic, sequential decision problem in an uncertain environment (Puterman, 2005). It is a well defined framework with on-line and off-line algorithms for determining optimal behaviour.

2.2.1. Preliminaries

A MDP is defined by a tuple $\langle Ss, A, T, R \rangle$ such as

- Ss is the finite set of all states of the model, all the values which can be reached;
- A is the finite set of all available actions.
- $T : Ss \times A \times Ss \rightarrow [0,1]$ is the transition function. $T(s,a,s')$ is the probability to be in s' applying the action a in the state s . This probability represents the uncertainty on the outcome of the action a .
- $R : Ss \times A \rightarrow R$ is the reward function. $R(s,a)$ is the reward after applying the action a when in the state s .

A policy π is a function from Ss to A assigning an action to each possible state in Ss . Solving an MDP over an infinite horizon results in deriving an optimal stationary policy π^* . This policy is the one which maximises the expected sum of rewards, i.e. $E\left[\sum_{t=0}^h \gamma^t R(s_t, a_t) | \pi\right]$ with γ the discount factor ($0 < \gamma < 1$). Algorithms are available to exactly compute this optimal policy π^* , in particular Value Iteration and Policy Iteration (Puterman, 2005).

2.2.2. Our model

The system is represented by the number of farms in the status S (susceptible to be infected), I (infected) and V (vaccinated). It evolves by 1-month time-steps. The decision of vaccination is available only at the beginning of each year. In order to be consis-

tent with the Markov assumption, the state of the system includes an indication of the actual month (t_0 – t_{11}). To avoid vaccination at other time step than the first month, the cost of the vaccination is prohibitive except the first month (t_0).

2.2.2.1. *Model.* Our Markov Decision Process is defined by $\langle Ss, A, T, R \rangle$ with

- Ss , the set of states: a state of the system has two components:
 - The current month t_i
 - The state of the population regarding the pathogen spread $\{[x, y, z] | x + y + z = N\}$ corresponding to the repartition of the N farms into the three individual statuses S (susceptible), I (infected) and V (vaccinated).
- A , the available actions: {"Vaccinating"; "Doing nothing"}
- T , the transitions: they are computed based on a stochastic epidemiological model representing the pathogen spread between farms (given in the following subsection).
- R , the rewards: For a state (h, i, j) , the rewards are given by

$$\begin{cases} R((h, i, j), \text{Doing Nothing}) = -Cd \cdot i \\ R((h, i, j), \text{Vaccinating}) = -Cd \cdot i - Cv \cdot (h + j) \end{cases}$$

with Cd the cost of the disease for each farm in the individual status I and Cv the cost of the vaccination for each farm in the statuses S and V (booster shot).

2.2.2.2. *Epidemiological details.* For modelling the pathogen spread within the group, we consider a compartmental epidemiological model. At each time-step, we consider the number of farms in each status. The epidemiological model is based on a classical SIS model used for disease without immunity after an infection (Hethcote, 2000). We add the status V for representing the vaccinated status. In this epidemiological model, individual transitions concern a farm. They are illustrated Fig. 1.

In our model, the set of model states for a given month t is then all the repartition of farms among the three individual statuses (S, I, V). With three farms, $Ss = \{(3, 0, 0), (2, 1, 0), (2, 0, 1), (1, 2, 0), (1, 0, 2), (1, 1, 1), (0, 3, 0), (0, 2, 1), (0, 1, 2), (0, 0, 3)\}$ where (x, y, z) are respectively the number of farms in the status S, I and V . The state space size is then given by

$$|Ss| = 12 * C_{N+2}^2 = 12 * \frac{(N+2)(N+1)}{2} \quad (1)$$

where C_n^p is the binomial coefficient giving the number of combinations of p elements among n .

At the group level, a transition from the model state (h, i, j) to the model state (l, m, n) corresponds to different sets of individual transitions. It depends on the number of farms moving from one status to another. For example, the transition from $(1, 2, 0)$ to $(2, 1, 0)$ without vaccination can result either (i) from a movement of one farm from I to S as the other farms kept the same individual status




Fig. 1. Schematic representation of the epidemiological model applied to farms in the group. Farms are separated according to their status regarding the pathogen spread (S : susceptible, I : infected, V : vaccinated). For each available action, arrows represent possible transitions between statuses.

and $1 = C_2^2$ for the situation (ii). We have then with $a = \text{"Doing Nothing"}$:

$$T((1, 2, 0), a, (2, 1, 0)) = (C_1^2 P(I \rightarrow S) P(I \rightarrow I)) (C_1^1 P(S \rightarrow S)) + (C_2^2 P(I \rightarrow S)^2) (C_1^1 P(S \rightarrow I)).$$

Following this reasoning, transitions at the group level are then computed considering all possible sets of individual transitions compatible with the related change of states and using multinomial coefficients. When the action a is "Doing nothing", the transition $T((h, i, j), a, (l, m, n))$ from (h, i, j) to (l, m, n) is given by

$$\sum_{u=0}^h \sum_{v=0}^i \sum_{w=0}^j \delta \left\{ \begin{array}{l} h + v + w - u = l \\ i + u - v = m \\ j - w = n \end{array} \right\} \left[\begin{array}{l} (C_h^u P(S \rightarrow I)^u P(S \rightarrow S)^{h-u}) \\ * (C_i^v P(I \rightarrow S)^v P(I \rightarrow I)^{i-v}) \\ * (C_j^w P(V \rightarrow S)^w P(V \rightarrow V)^{j-w}) \end{array} \right] \quad (2)$$

where $\delta \left\{ \begin{array}{l} \text{cond}_1 \\ \dots \\ \text{cond}_n \end{array} \right\}$ is a Dirac function, equal to 1 when the conditions $\text{cond}_1, \dots, \text{cond}_n$ are all satisfied and 0 elsewhere, C_n^p is the binomial coefficient for p elements among n , u the number of S farms becoming I , v the number of I farms becoming S and w the number of V farms becoming S .

Considering only the states satisfying conditions of the Dirac function, equation (Eq. (2)) can be simplified to:

Considering only the states satisfying conditions of the Dirac function, equation (Eq. (2)) can be simplified to:

$$T((h, i, j), a, (l, m, n)) = \begin{cases} B & \text{if } \begin{cases} n \leq j \\ m \leq i + h \\ l \geq j - n \\ l \leq j - n + h + i \end{cases} \\ 0 & \text{else} \end{cases}$$

with

$$B = \sum_{u=\max(0, h+(j-n)-l)}^{\min(h, h+i+(j-n)-l)} \left[\begin{array}{l} (C_h^u P(S \rightarrow I)^u P(S \rightarrow S)^{h-u}) \\ * (C_i^{u+i-m} P(I \rightarrow S)^{u+i-m} P(I \rightarrow I)^{m-u}) \\ * (C_j^{j-n} P(V \rightarrow S)^{j-n} P(V \rightarrow V)^n) \end{array} \right] \quad (3)$$

When the action a is "Vaccinating", the transition $T((h, i, j), a, (l, m, n))$ is given by

$$\sum_{u1=0}^h \sum_{u2=0}^{h-u1} \sum_{v=0}^i \sum_{w=0}^j \delta \left\{ \begin{array}{l} h + v + w - u1 - u2 = l \\ i + u1 - v = m \\ j - w + u2 = n \end{array} \right\} \left[\begin{array}{l} \left(\frac{h!}{u1!u2!(h-u1-u2)!} P(S \rightarrow I)^{u1} P(S \rightarrow V)^{u2} P(S \rightarrow S)^{h1-u1-u2} \right) \\ * (C_i^v P(I \rightarrow S)^v P(I \rightarrow I)^{i-v}) * (C_j^w P(V \rightarrow S)^w P(V \rightarrow V)^{j-w}) \end{array} \right] \quad (4)$$

or (ii) from both a movement of one farm from S to I and a movement of the two farms from I to S . As we do not differentiate between farms, there are then three possible movements: $2 = C_2^2$ for the situation (i) depending of which farm in the I status moves

where $\delta \left\{ \begin{array}{l} \text{cond}_1 \\ \dots \\ \text{cond}_n \end{array} \right\}$ is a Dirac function, equal to 1 when the conditions $\text{cond}_1, \dots, \text{cond}_n$ are all satisfied and 0 elsewhere, C_n^p is the

binomial coefficient for p elements among n , $u1$ the number of S farms becoming I , $u2$ the number of S farms becoming V , v the number of I farms becoming S and w the number of V farms becoming S .

Considering only the states satisfying conditions of the Dirac function, equation (Eq. (4)) can be simplified to:

$$T((h, i, j), a, (l, m, n)) = \begin{cases} B & \text{if } \begin{cases} m \leq i + h \\ n \leq h + j \\ l \leq j - n + h + i \end{cases} \\ 0 & \text{else} \end{cases}$$

with

$$B = \sum_{u2=\max(0, n-j)}^{\min(n, h, h-(m-i))} \left[\sum_{u1=\max(0, h+(j-n)-l)}^{\min(h-u2, h+i+(j-n)-l)} \left[\left(\frac{h!}{u1!u2!(h-u1-u2)!} P(S \rightarrow I)^{u1} P(S \rightarrow V)^{u2} P(S \rightarrow S)^{h1-u1-u2} \right) \right. \right. \\ \left. \left. * (C_i^{u1+i-m} P(I \rightarrow S)^{u1+i-m} P(I \rightarrow I)^{m-u1}) \right. \right. \\ \left. \left. * (C_j^{j+u2-n} P(V \rightarrow S)^{j+u2-n} P(V \rightarrow V)^{n-u2}) \right] \right] \quad (5)$$

To compute the transitions T at the group level, we need to define the transitions $P(I \rightarrow S)$, $P(V \rightarrow S)$, $P(S \rightarrow V)$ and $P(S \rightarrow I)$. Transitions from one state X to the same one are deduced from the other probabilities as we have $\sum_Y P(X \rightarrow Y) = 1$. The probabilities for the transitions from I to S and from V to S are assumed to be constant corresponding to the inverse of the respective mean sojourn time in the status I and V . The transition from S to V depends on the action: with the action ‘‘Doing nothing’’ $P(S \rightarrow V) = 0$, and with the action ‘‘Vaccinating’’ $P(S \rightarrow V) = pEffV$ with $pEffV$ the probability that the vaccination is correctly implemented in the farms and results in a full protection during nearly 1 year. The transition $P(S \rightarrow I)$ from S to I depends on the number of infected farms in the group, on a constant external risk and on the action (vaccination or not). The transition probability used has a frequency-dependent form (Begon et al., 2002). With the action ‘‘Doing nothing’’, the transition depends only on the number of infected herds and an external risk. The probability is given by $P(S \rightarrow I) = (Beta \cdot N_I / N_T + pOut)$ with $Beta$ the transmission rate, N_I the number of infected farms in the group, N_T the total size of the group and $pOut$ the probability to be infected from a farm outside of the group. With the action ‘‘Vaccinating’’, a part of susceptible farms would be vaccinated and then protected. Only those not protected by the vaccination may be infected. The probability is then given by $P(S \rightarrow I) = (1 - pEffV) \cdot (Beta \cdot N_I / N_T + pOut)$.

2.3. Computational study

We assume that the group consists of 200 farms. Since a theoretical pathogen spread is assumed, values of parameters (Table 1) are independent of any disease and are chosen only to have an endemic situation; i.e. a constant prevalence of infected farms in the group over time without any action. The reproduction number (R_0) is here of 5 which corresponds to an endemic situation of 80% of infected farms at equilibrium.

We compute the MDP policy with the Value Iteration algorithm, assuming an infinite horizon and a discounted rate of 0.9975398 corresponding to an annual discounted rate of 97% classically used in economic studies. We compare the obtained policy to two commonly used ones: Do Nothing and Systematic Vaccination, using each year the respective actions ‘‘Doing nothing’’ and ‘‘Vaccinating’’. The effect of these policies was tested over 50 simulated years. For policy comparison, we consider the rate of vaccination in the area and the variability of both the prevalence and the total

Table 1

Parameters used for simulation (values given per month).

Parameter	Notation	Value
Cost for each infected herd		1
Cost for the vaccination of a herd at t_0		7
Cost for the vaccination of a herd at other than t_0		999999
Transmission rate between herds in the group	$Beta$	0.1
External risk of infection	$pOut$	0.001
Recovered rate (Infected to Susceptible)	$P(I \rightarrow S)$	0.02
Vaccination efficiency	$pEffV$	0.95
Rate of end of protection by vaccination	$P(V \rightarrow S)$	0.1

costs in the area. For each policy, we run 100,000 simulations with the same initial population consisting in 80% of infected farms and 20% of susceptible ones. It corresponds to a plausible endemic situation in an area.

As the cost of the vaccine would influence the computed policy, we do a sensitivity analysis of the model to this parameter. We run five sets of simulations with various vaccination costs: 1, 3, 6, 9 and 12.

2.4. Computation time

The computation times are evaluated for different group sizes (Table 2). To evaluate the transition matrix density, we mention the size of the state space and the amount of non-zero values for each line. The matrix is relatively sparse, mainly due to the month indicator step. Nevertheless, the computation time and the memory used are still high although we took great caution when programming to lower them. First, we implemented the model with a multi-threads approach both to compute the matrix for different lines in parallel and to run Value Iteration. For the matrix creation, the formulas (Eqs. (3) and (5)) are difficult to implement directly even for their simplified version. Instead, we decided to build the matrix line by line, simulating possible moves of groups: for a given initial state, we explore the movements’ tree for farms considering the individual transitions probability. When all herds have been considered once (one branch is fully explored), the resulting probability is added to the corresponding final state’s probability. When all branches are done, the probability distribution of all the final states is the expected outcome of the considered action and initial state. Data structures were also optimised to use the least possible memory.

3. Results

The policy computed with the MDP leads to a high number of model states with the action ‘‘Doing Nothing’’ (Fig. 2). When the number of farms in the S and I statuses are low, no vaccination is retained which is not surprising as a high number of farms are protected by vaccination (in V status).

When simulated, the MDP computed policy induces a non-systematic vaccination (Fig. 3). Nevertheless, in the first year, the vaccination was used in all replications as the initial model state is assumed to be in the area of vaccination policy. Over time, the

Table 2

Computation results for each group size (N) with the corresponding number of states ($|Ss|$), density indicators, computation times (in seconds) for building the matrix and solving the model, cpu time and elapse time (in hours:minutes:seconds). The program was run (A) on a laptop (Core i5) using four threads and (B) a shared server machine (48 processors, 260 GB of memory) using 30 threads.

N	Ss	Mean density (%)	Maximal density (%)	Matrix building		Model solving		Cpu time (B)	Elapse time (B)
				(A)	(B)	(A)	(B)		
50	15,912	0.157	1.464	1.70	0.51	25.62	12.40	00:12:41	00:00:14
100	61,812	0.074	0.676	19.52	3.15	131.00	31.15	00:12:25	00:00:35
150	137,712	0.047	0.434	99.74	15.47	406.66	66.62	00:36:41	00:01:24
200	243,612	0.035	0.319	na	50.67	na	131.08	01:29:02	00:03:03
250	379,512	0.027	0.252	na	127.59	na	367.01	03:36:48	00:08:18
300	545,412	0.022	0.207	na	273.74	na	817.99	07:22:03	00:18:18
350	741,312	0.019	0.176	na	592.11	na	985.46	11:18:11	00:26:22
400	967,212	0.016	0.152	na	951.39	na	1253.80	17:31:45	00:36:55

na: not applicable.




Fig. 2. Policy computed by the MDP: each model state is given by (S, I) with S and I the proportion of herds in the health status susceptible and infected, respectively.

vaccination is less used depending on the current model state at the time when the decision should be taken.

The MDP policy allows us to propose a strategy reducing the prevalence with a total cost nearly equivalent to the Do Nothing strategy after 10 years (Figs. 4 and 5). Compared to the Systematic Vaccination strategy, the total cost is lower but the prevalence is less reduced (Figs. 4 and 5). Due to the non-systematic vaccination, the cost of the strategy is lower than for the Systematic Vaccination strategy (Fig. 5). We note that for the epidemiological param-

eters used here, the prevalence is not reduced to zero even after 50 years of Systematic Vaccination (Fig. 4). This result is not in contradiction with the minimal coverage of vaccination that can be computed using R_0 and is equal to 80%. Although this minimal coverage is lower than p_{EffV} (Table 1), the clearance is not achieved. It can be explained by the fact that vaccination is used only by susceptible farms for preventing infection and not by all farms. Moreover, vaccination is done only once a year (1 over 12 times-steps) with an assumption of waning immunity.

As expected, the cost of the vaccination influences the computed policy. For the values of the cost vaccination tested, the computed policy corresponds to a Systematic Vaccination for vaccination costs below 6, inducing the same prevalence and the same total cost for the two strategies (MDP Policy and Systematic Vaccination, Figs. 6 and 7). For vaccination costs higher than 6, the vaccination is too expensive regarding its efficiency and its use is limited (for value 9) or forbidden in nearly all states (for value 12) by the computed policy.

4. Discussion

In this paper, we used a classical approach for solving sequential decision problems also used for populations management. In our study, we consider a group of herds with a unique decision-marker. The dynamics of the transmission between herds may be represented by a meta-population approach with movements of animals between herds (such as in Le Menach et al., 2006). The




Fig. 3. Yearly proportion of replications where all uninfected farms are vaccinated using the strategy based on the MDP policy (in light grey) and the strategy Systematic Vaccination (in black) (100,000 replications for each strategy). The Do Nothing strategy is not represented here as for that strategy, by definition, vaccination is never used in any of replications.




Fig. 4. Percentiles over time of the prevalence of infected farms in the group using the three strategies: Do Nothing (heavy solid lines), MDP policy (dotted lines) and Systematic Vaccination (thin solid lines) (100,000 replications for each strategy).

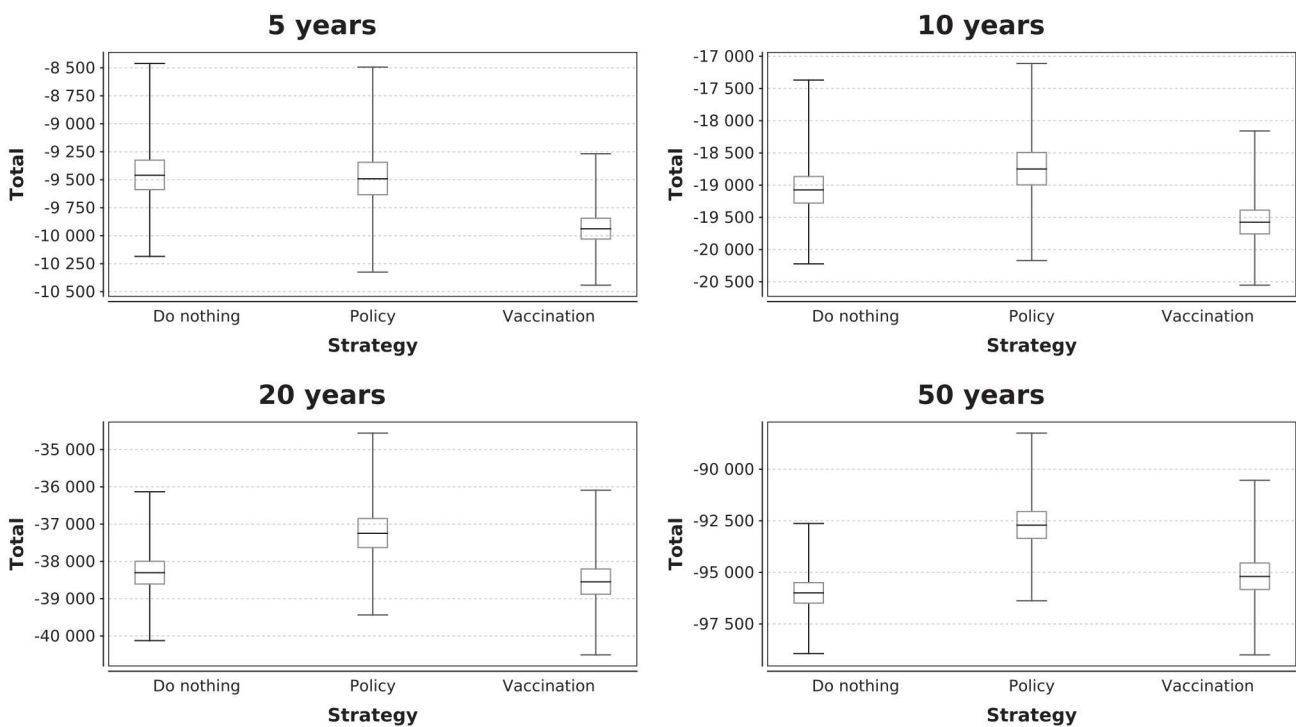


Fig. 5. Boxplot of the total cost of the pathogen spread and its control within the group using the three strategies: Do Nothing, MDP policy and Systematic Vaccination at different time steps: 5, 10, 20 and 50 years (100,000 replications for each strategy).

use of an MDP to propose a sequence of actions in order to reach an objective in a meta-population is common in ecology, mainly for questions related to conservation and biodiversity (Possingham, 1996; Chadès et al., 2008; Nicol et al., 2009). As in McDonald-Madden et al. (2008), we do not consider explicitly the movement of animals between populations. Nevertheless, in our model, the farms are in interactions due to the pathogen spread between them but without given the network of animal movements. Modelling the pathogen spread influencing the choice of actions is already considered for human population and in agriculture. For human populations, there is one decision-maker per individual and no decision-maker for a group of individuals. The game theory can then be used (such as in Reluga et al., 2006). In agriculture, a more complex approach related to MDP was proposed to consider decision regarding pathogen transmission including spatial location (Peyrard et al., 2007).

We consider here a simplification of a multi-component problem in order to be able to compute the exact optimal policy. We consider then no individual decision by farmer. If we wanted to represent the individual decision, a Graph MDP (Forsell et al., 2011; Peyrard et al., 2007) could be considered for diseases transmitted from one farm to another by closed contact between animals and by purchases. For parametrisation of the model, data on animal movements between farms (as in: Ezanno et al., 2006; Ribbens et al., 2009) should be available to define a network. To represent decisions both at individual level and at group level, interdependency between them should be considered.

Our computed MDP policy is optimal and adaptive and may not have been proposed by experts. Often in animal health modelling, an action is systematically used whatever the current system state. As in our approach, adaptation on the current system state is

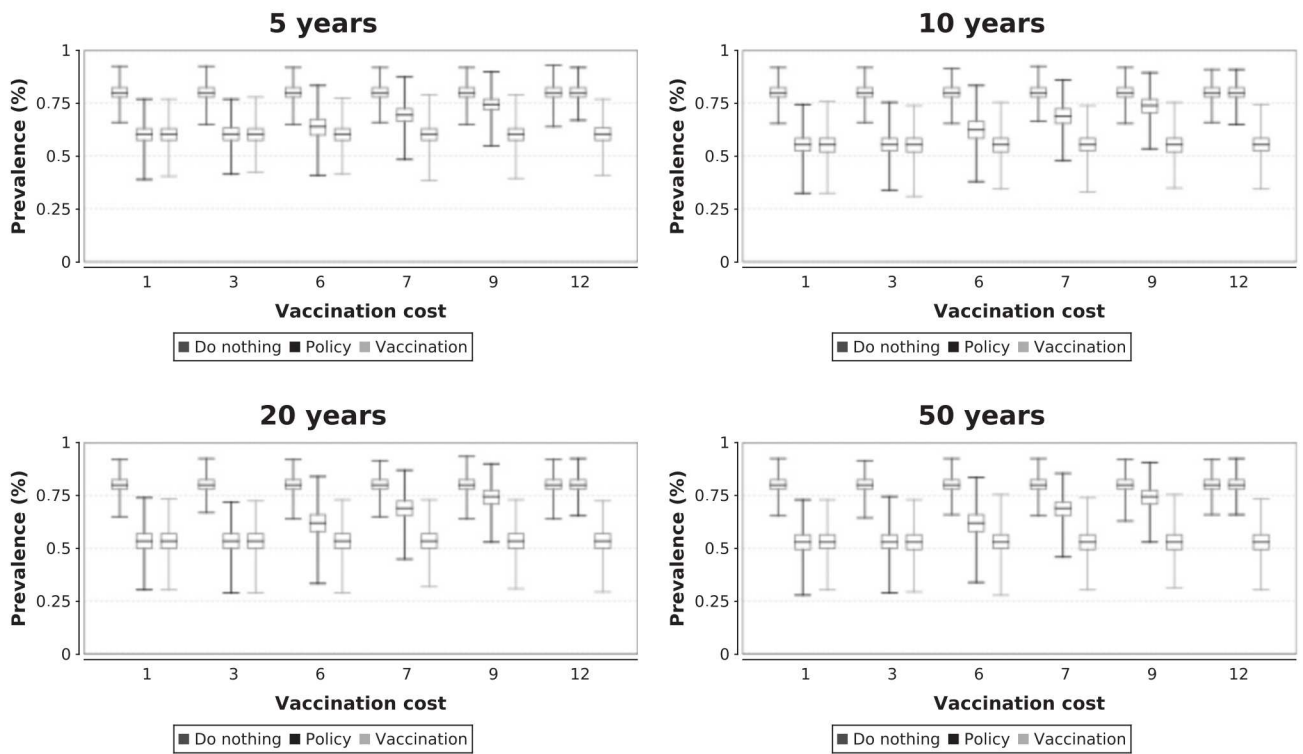


Fig. 6. For different costs of vaccination, boxplot of the prevalence of infected farms in the group with the three strategies: Do Nothing, MDP policy and Systematic Vaccination at different time steps: 5, 10, 20 and 50 years (100,000 replications for each strategy and each cost).

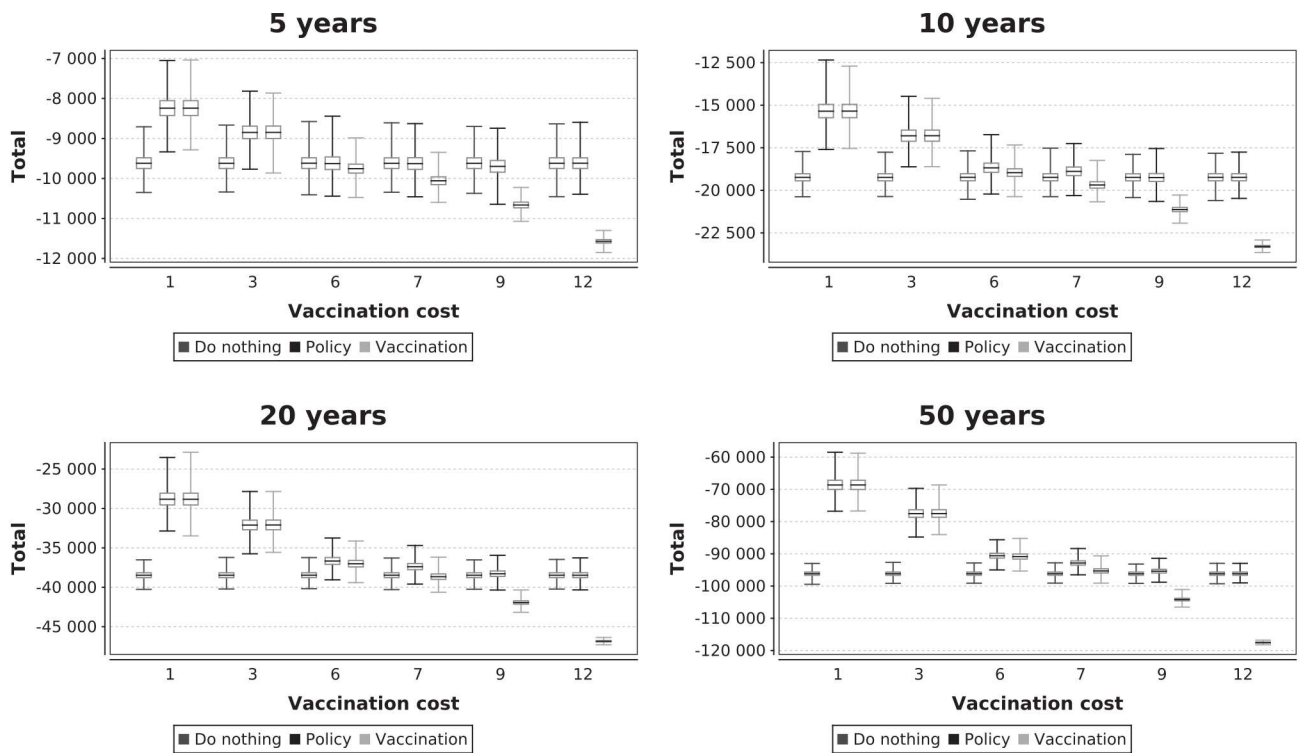


Fig. 7. For different costs of vaccination, boxplot of the total cost of the pathogen spread and its control within the group using the three strategies: Do Nothing, MDP policy and Systematic Vaccination at different time steps: 5, 10, 20 and 50 years (100,000 replications for each strategy and each cost).

considered in some models where actions are used under constraints based on expert knowledge (such as in Seegers et al., 2006). The optimality is evaluated on the simulation results for

specified policies. With our approach, we propose a policy both adaptive and optimal. The constraints are defined mainly based on the expected behaviour of the pathogen spread with and

without actions and the associated costs. These constraints are complex and may not have been easily identified by experts.

For model building and optimisation, the computation time is long although the matrix is sparse. Pre-computing transitions matrix and multi-threading both matrix building and Value Iteration reduces the computation time on computers with several processors. The counterpart is the memory needed to store the matrix. It explains why we did not provide computation time on the laptop for $N \geq 200$. In case of a lack of memory space, our program is able to compute parts of the matrix before Value Iteration and to compute the remaining matrix lines as needed. In this case, the computing time is highly increased. Limiting computing time is possible providing approximations. During the transition matrix building, transitions with very low probabilities can be avoided to reduce the density of the matrix. Such an approximation reduces the computing time for the optimisation. For example, cutting probabilities below 10^{-8} reduces the matrix density from 0.035% to 0.021% and the cpu time from 1H30 to 1H in the shared server machine. If only a few initial states are needed, the states can be pruned using simulated trajectories from these initial states. With three individual statuses, the model can be computed in a shared server machine even for a group size of 400. Nevertheless, as the number of individual statuses increases, the model becomes intractable. With four individual statuses and 200 herds, the space state size is then of $12 * C_{203}^3 = 12 * 1373701$, that is more than 16 millions states. Exact resolution of such a model is then not possible in a reasonable time. Approximate resolution of large MDP by a hierarchical approach was proposed in many works such as in (Givan et al., 2003; Barry et al., 2010). The use of such an approach for our model is currently evaluated.

With an MDP, the policy is computed once for all situations of the pathogen spread in the area, considering all possible behaviours from each model state. Computing the policy before the simulation reduces the simulation time. It produces a guideline that can be used whatever the situations. The decisions are available for situations with a fast pathogen spread and also for those with low ones. This pre-computing step is not possible when the parameter values change over time as in Merl et al. (2009). With our approach, it can be of interest to do a robustness analysis to explore the efficiency of the computed policy when some parameters are slightly different from the ones used for computing.

In the computational study, the policy was computed assuming an infinite horizon although the interest of decision makers can be defined on a finite one. For an infinite horizon, a discounted rate must be used for the algorithm to converge. The rate used here was chosen to correspond to an actualisation rate used in economic studies (3% per year). Due to the time step of 1 month, the economic impact of the 5th year is reduced by nearly 14%. Our objective was economical, so decision makers can accept the discounted rate. Nevertheless, if the objective is a number of herds, they may not accept such a decrease rate. In case of a finite horizon noted H , a policy by year are calculated. For an animal disease, it may be easier to have only one strategy that can be followed every year. Using the same policy each year may induce a non-respect of the objective. Indeed, choosing a finite horizon with the same discounted rate increases very slightly the number of states with a vaccination (results not shown). Using them whatever the year, the MDP policies may induce similar prevalence and cost but more cost than the true finite horizon policy. Moreover, with a finite horizon H , the future dynamics after the year H are not considered although decisions may continue after H . To avoid this drawback, moving or rolling horizon was sometimes considered (van Den Broek, 2002; Balakrishnan and Hung Cheng, 2009) but one strategy should be computed each year and there is no guarantee of optimality.

When considering only a vaccination every year, a 1 year time-step is sufficient. Nevertheless, we define this model in order to extend it later to a situation with more actions. Indeed, we consider that an action may have varying impact and cost over time due to herd management. For example, a biosecurity action can be considered each month. Depending on the current month, the herd can be in pasture or in a building. According to the localisation of animals (pasture or building), the biosecurity action can have different impact and costs. Introducing such a biosecurity action is possible in our model as the current month is known for each model state. To simplify the model presentation here, we consider only the vaccination action to be independent of herd management. The action is available every year at the same date as the boosting is often done each year. Even without considering another action, our model can be easily extended to a situation where the Vaccinating action should be separated of at least 12 months. In this case, the transition from t_0 to t_1 would only be possible in case of vaccination. In case of the action "Doing Nothing", the transition from t_0 would be only to t_0 . Vaccination occurs then at least after 12 months after the previous one but not exactly at the same date each year.

This model is theoretical and cannot be directly used for a real pathogen spread. Tests can be used to identify the health status of each herd but they are often nor 100% sensitive nor 100% specific (Christensen and Gardner, 2000). These uncertainties may be introduced in the model. A Partially Observable MDP (Kaelbling et al., 1998) can be used to formalise it, but with an increased size of the model. Moreover, the external risk may vary, influenced by the pathogen spread outside of the group. Indeed, individual farmers can be independent or outside of the group considered here. They may also try to limit the pathogen spread, decreasing the prevalence outside of the group and thus the external risk of infection. Variation of the external risk can be introduced into the model considering another component of the state (an external risk of pathogen introduction), increasing the model size. For a real pathogen spread, we should also take into account that the transmission parameters are often wrongly estimated. Impact of such uncertainty can be studied considering a robustness analysis for the computed policy. If this uncertainty influences a lot the results, approach with a Bayesian approach as in Ge et al. (in press, 2010) may be considered.

5. Conclusion

In this paper, we showed how to define an MDP to propose a preventive strategy based on the computed policy to control a pathogen spread at the level of a group of farms. The group's objective here consists in having the lowest total cost of pathogen spread for the farms of the group. Although the objective is economic only, the computed policy induces a decrease of the pathogen spread for a cost equals or lower than a Do Nothing strategy.

Moreover, vaccination is generally applied less than in a Systematic vaccination strategy, and that implies an increased prevalence. However, this loss turns into an increase of the disease cost that is compensated by the reduced cost of vaccination. The total cost is then lower when vaccinating parsimoniously. The analysis we conducted on the influence of the vaccination cost shows that the MDP produces automatically the Systematic Vaccination when its cost is cheap enough, the Do Nothing strategy when its cost is too expensive, and finds the right threshold otherwise.

Acknowledgements

The authors want to thank Régis Sabbadin for advices for improving this paper. Financial support for this research was provided by INRA, Cemagref and Basse-Normandie, Bretagne, Pays de

la Loire and Poitou-Charentes Regional Councils under SANCRE project, in the framework of “For and About Regional Development” programs.

References

- Backer, J.A., Hagenaars, T.J., van Roermund, H.J.W., de Jong, M.C.M., 2009. Modelling the effectiveness and risks of vaccination strategies to control classical swine fever epidemics. *J. R. Soc. Interface* 6, 849–861.
- Balakrishnan, J., Hung Cheng, C., 2009. The dynamic plant layout problem: incorporating rolling horizons and forecast uncertainty. *Omega* 37, 165–177.
- Barry, J., Kaelbling, L.P., Lozano-Perez, T., 2010. Hierarchical Solution of Large Markov Decision Processes. In: ICAPS Workshop on Planning and Scheduling in Uncertain Domains, Toronto, Canada.
- Begon, M., Bennett, M., Bower, R.G., French, N.P., Hazel, S.M., Turner, J., 2002. A clarification of transmission terms in host-microparasite models: numbers, densities and areas. *Epidemiol. Infect.* 129, 147–153.
- Ben-Ari, Y., Gal, S., 1986. Optimal replacement policy for multicomponent systems: an application to a dairy herd. *Eur. J. Oper. Res.* 23, 213–221.
- Bernstein, D.S., Givan, R., Immerman, N., Zilberstein, S., 2002. The complexity of decentralized control of Markov Decision Processes. *Math. Oper. Res.* 27, 819–840.
- Boutilier, C., 1999. Sequential optimality and coordination in multi-agent systems. In: Proceedings of the 16th International Joint Conference on Artificial Intelligence, pp. 478–485.
- Chadès, I., McDonald-Madden, E., McCarthy, M.A., Wintle, B., Linkie, M., Possingham, H.P., 2008. When to stop managing or surveying cryptic threatened species. *PNAS* 105, 13936–13940.
- Christensen, J., Gardner, I.A., 2000. Herd-level interpretation of test results for epidemiologic studies of animal diseases. *Prev. Vet. Med.* 45, 83–106.
- Dangerfield, C.E., Ross, J.V., Keeling, M.J., 2008. Integrating stochasticity and network structure into an epidemic model. *J. R. Soc. Interface* 38, 761–774.
- Ezanno, P., Fourichon, C., Beaudeau, F., Seegers, H., 2006. Between-herd movements of cattle as a tool for evaluating the risk of introducing infected animals. *Anim. Res.* 55, 189–208.
- Ferguson, N.M., Donnelly, C.A., Anderson, R.M., 2001. The Foot-and-Mouth epidemic in Great-Britain: pattern of spread and impact of interventions. *Science* 292, 1155–1160.
- Forsell, N., Wikström, P., Garcia, F., Sabbadin, R., Blennow, K., Eriksson, L.O., 2011. Management of the risk of wind damage in forestry: a graph-based Markov decision process approach. *Ann. Oper. Res.* 190, 57–74. doi:10.1007/s10479-009-0522-7.
- Ge, L., Mourits, M.C.M., Huirne, R.B.M., 2007. Towards flexible decision support in the control of animal epidemics. *Rev. Sci. Tech. Off. Int. Epizoot.* 26, 551–563.
- Ge, L., Mourits, M.C., Kristensen, A.R., Huirne, R.B., 2010. A modeling approach to support dynamic decision-making in the control of fmd epidemics. *Prev. Vet. Med.* 95, 167–174.
- Ge, L., Kristensen, A.R., Mourits, M.C., Huirne, R.B., in press. A new decision support framework for managing foot-and-mouth disease epidemics. *Ann. Oper. Res.* (doi: 10.1007/s10479-010-0774-2).
- Givan, R., Dean, T., Greig, M., 2003. Equivalence notions and model minimization in Markov Decision Processes. *Artif. Intell.* 142, 163–223.
- Hethcote, H.W., 2000. The mathematics of infectious diseases. *SIAM rev.* 42, 599–653.
- Kaelbling, L.P., Littman, M.L., Cassandra, A.R., 1998. Planning and acting in partially observable stochastic domains. *Artif. Intell.* 101, 99–134.
- Keeling, M.J., Woolhouse, M.E.J., May, R.M., Davies, G., Grenfell, B.T., 2003. Modelling vaccination strategies against foot-and-mouth disease. *Nature* 421, 136–142.
- Klein, E., Laxminaryan, R., Smith, D.L., Gilligan, C.A., 2007. Economic incentives and mathematical models of disease. *Environ. Dev. Econ.* 12, 707–732.
- Kobayashi, M., Carpenter, T.E., Dickey, B.F., Howitt, R.E., 2007. A dynamic, optimal disease control model for foot-and-mouth disease: I. Model description. *Prev. Vet. Med.* 79, 257–273.
- Kristensen, A.R., 1992. Optimal replacement in the dairy herd: a multi-component system. *Agr. Syst.* 39, 1–24.
- Kristensen, A.R., 2003. A general software system for Markov decision processes in herd management applications. *Comput. Electron. Agr.* 38, 199–215.
- Kristensen, A.R., Jørgensen, E., 2000. Multi-level hierarchic Markov processes as a framework for herd management support. *Ann. Oper. Res.* 94, 69–89.
- Le Menach, A., Vergu, E., Grais, R.F., Smith, D.L., Flahault, A., 2006. Key strategies for reducing spread of avian influenza among commercial poultry holdings: lessons for transmission to humans. *P. Roy. Soc. B – Biol. Sci.* 273, 2467–2475.
- LeMenach, A., Legrand, J., Grais, R.F., Viboud, C., Valleron, A.-J., Flahault, A., 2005. Modeling spatial and temporal transmission of foot-and-mouth disease in France: identification of high-risk areas. *Vet. Res.* 36, 699–712.
- McDonald-Madden, E., Baxter, P.W.J., Possingham, H.P., 2008. Subpopulation triage: how to allocate conservation effort among populations. *Conserv. Biol.* 22, 656–665.
- Merl, D., Johnson, L.R., Gramacy, R.B., Mangel, M.A., 2009. Statistical framework for the adaptive management of epidemiological interventions. *PLoS ONE* 4, e5807.
- Nicol, S., Chades, I., Linke, S., Possingham, H.P., 2009. Conservation Decision-Making in Large State Spaces. In: Proceedings of the 18th World IMACS Congress and MODSIM09 International Congress on Modelling and Simulation. Cairns, Australia.
- O’Callaghan, C.J., Medley, G.F., Peter, T.F., Mahan, S.M., Perry, B.D., 1999. Predicting the effect of vaccination on the transmission dynamics of heartwater (*Cowdria ruminantium* infection). *Prev. Vet. Med.* 42, 17–38.
- Peyrard, N., Sabbadin, R., Aubertot, J.N., 2007. A graph-based markov decision process framework for optimising integrated management of diseases in agriculture. In 17th Inter. Congress on Modelling and Simulation (MODSIM’07), Christchurch, New-Zealand.
- Possingham, H.P., 1996. Decision theory and biodiversity management: how to manage a metapopulation. In *Frontiers of Population Ecology*. CSIRO, Melbourne, pp. 391–398.
- Puterman, M.L., 2005. Markov Decision Processes: Discrete Stochastic Dynamic Programming. John Wiley & Sons Inc.
- Rat-Aspert, O., Fourichon, C., 2010. Modelling collective effectiveness of voluntary vaccination with and without incentives. *Prev. Vet. Med.* 93, 265–275.
- Reluga, T.C., Bauch, C.T., Galvani, A.P., 2006. Evolving public perceptions and stability in vaccine uptake. *Math. Biosci.* 204, 185–198.
- Ribbens, S., Dewulf, J., Koenen, F., Mintiens, K., de Kruif, A., Maes, D., 2009. Type and frequency of contacts between Belgian pig herds. *Prev. Vet. Med.* 88, 57–66.
- Seegers, H., Billon, D., Grimard, B., Beaudeau, F., 2006. A Simulation Model To Study Reproductive Disorders And Reproductive Management In Dairy Herds. In, 11th International Symposium on Veterinary Epidemiology and Economics, Cairns, Australia.
- Tildesley, M.J., Savill, N.J., Shaw, D.J., Deardon, R., Brooks, S.P., Woolhouse, M.E.J., Grenfell, B.T., Keeling, M.J., 2006. Optimal reactive vaccination strategies for a foot-and-mouth outbreak in the UK. *Nature* 440, 83–86.
- Toft, N., Kristensen, A.R., Jørgensen, E., 2005. A framework for decision support related to infectious diseases in slaughter pig fattening units. *Agr. Syst.* 85, 120–137.
- van den Broek, W., 2002. Moving horizon control in dynamic games. *J. Econ. Dyn. Control* 26, 937–961.
- Vonk Noordegraaf, A., Buijtel, J.A.A.M., Dijkhuizen, A.A., Franken, P., Stegeman, J.A., Verhooff, J., 1998. An epidemiological and economic simulation model to evaluate the spread and control of infectious bovine rhinotracheitis in the Netherlands. *Prev. Vet. Med.* 36, 219–238.