Available online at www.sciencedirect.com
SCIENCE DIRECT
Theoretical Population Biology 69 (2006) 323–331
Theoretical Population Biology
www.elsevier.com/locate/tpb

The evolutionary consequences of plasticity in host–pathogen interactions

Peter D. Taylor*, Troy Day, Daniel Nagy, Geoff Wild, Jean-Baptiste André, Andy Gardner

Department of Mathematics and Statistics, Queen's University, Kingston, ON, Canada K7L 3N6

Received 13 May 2005
Available online 15 February 2006

Abstract

Interactions between individuals such as hosts and pathogens are often characterized by substantial phenotypic plasticity. Pathogens sometimes alter their exploitation strategies in response to defensive strategies adopted by their host and vice versa. Nevertheless, most game-theoretic models developed to explain the evolution of pathogen and host characteristics assume that no such plasticity occurs. Allowing for phenotypic plasticity in these models is difficult because one must focus on the evolution of pathogen and host reaction norms, and then allow for the potentially indefinite reciprocal changes in pathogen and host behaviour that occur during an infection as a result of their interacting reaction norms. Here, we begin to address these issues for a simple host–pathogen system in which the pathogen exhibits a level of virulence and the host exhibits a level of immune clearance. We find, quite generally, that plasticity promotes the evolution of higher levels of cooperation, in this case leading to reduced levels of both virulence and clearance.

Keywords: Host–pathogen; Evolution; Negotiation; Stackelberg; Clearance; Virulence

1. Introduction

Most models of evolutionary games do not explicitly account for the possibility of players altering their actions in response to information about their opponent's actions before the overall outcome is decided. Rather, game theory models in biology typically constrain players to decide upon a course of action *in ignorance* of that chosen by a partner or opponent, and then to use these chosen levels from that point onward. In nature, however, such interactions are often plastic (Agrawal, 2001). Although 'situation-dependent' strategies, such as the Retaliator of Maynard Smith and Price (1973), have been allowed for ever since the inception of evolutionary game theory, it is only in the last decade that these are receiving proper attention. For example, Doebeli and Knowlton (1998) studied the evolution of interspecific mutualism using "reactive" strategies in a continuous version of the Prisoner's Dilemma game. Their approach has since been adopted by other studies of a similar kind (e.g. Roberts and Sherratt, 1998; Wahl and Nowak, 1999a, b; see also review by Doebeli and Hauert, 2005).

One simple form of plasticity is what might be called *precedence*: one of the two players must necessarily "act first" in the sense that the other has knowledge of its strategy, and can respond to this information (Kokko, 1999; McNamara et al., 2003; Pen and Taylor, 2005). In this case, we would expect that the second player evolves to choose a *best response*, i.e. a behavioural optimum, conditional on its opponent's strategy. The first player might then maximize its fitness conditional upon the best response of the second. In this case, the strategy pair converges to what has become known as a *Stackelberg equilibrium* (Kokko, 1999; Fudenberg and Tirole, 1991; McNamara et al., 2003).

We can carry this idea further by supposing that both players can respond plastically to information about the strategy adopted by the other. But now some new questions arise. If both players have reaction norms that specify their action in response to that of their opponent, then this will result in an indefinite cycle of reciprocal

*Corresponding author.
E-mail address: taylorp@post.queensu.ca (P.D. Taylor).

0040-5809/$ - see front matter © 2006 Elsevier Inc. All rights reserved.
doi:10.1016/j.tpb.2005.09.004

plastic change in each player. We find ourselves in a potentially infinite loop of mutual manipulation, and to handle this we therefore need to restrict attention to reaction norms that will result in the asymptotic convergence of the actions of each player to a steady state in a relatively short period of time. This has been referred to as negotiation in the recent literature (McNamara et al., 1999, 2003; Taylor and Day, 2004).

In current theoretical treatments of negotiation, the traditional strategies (level of investment, level of aggression) are replaced by *negotiation rules*, i.e. functional rules for responding to a level of behaviour adopted by ones opponent. Negotiated games consider the fitness of a negotiation rule, given the negotiation rule of an opponent and the evolutionary outcome in this context is again expected to be a Nash equilibrium, but in the space of negotiation rules rather than in the space of behaviours that are adopted. This new perspective leads to new solutions and, ultimately, to predictions that are qualitatively different from those arrived at under precedence or other non-negotiated means (McNamara et al., 1999, 2003; Taylor and Day, 2004). In particular, the incorporation of negotiation into evolutionary games can lead to novel explanations for the evolution of such problematic behaviours as cooperation (Taylor and Day, 2004) and some forms of altruism (Pen and Taylor, 2005).

In the spirit of this issue, *Game Theory Now*, we explore the consequences of these new ideas in evolutionary game theory, in the biological context of host–pathogen coevolution. Game theoretic treatments of host–pathogen coevolution continue to be an active area of research (van Baalen, 1998; Day and Burns, 2003; Restif and Koella, 2003), and in light of recent applications of evolutionary theory to epidemiology (Gandon et al., 2001; Dieckmann et al., 2002) this area is likely to continue to flourish. The goal of this paper is to understand how the allowance for plasticity in host–pathogen interactions alters the predictions concerning the coevolution of pathogen virulence and host immune response. Previous work has demonstrated that host–pathogen coevolution can have two distinct outcomes: (i) commensalism, and (ii) increased pathogen-induced mortality rate (virulence) along with increased clearance rates (van Baalen, 1998). Does the occurrence of plasticity in hosts and/or pathogens tend to favour either of (i) or (ii)?

The remainder of this article consists of four sections. In Section 2, we describe a simple model of host–pathogen coevolution. In Section 3, we consider the effect of plasticity. We look first at precedence, when either pathogen or host must "act first," and we obtain analytical results. Following that we consider negotiation in which each player can respond to the current level of activity of the other and we use simulations to obtain our results. Our results are discussed in Section 4. Interestingly, in all cases examined we find that the allowance for plasticity results in the evolution of both reduced pathogen virulence and reduced host clearance.

2. A simple host–pathogen system

Most models for the evolutionary dynamics of pathogen virulence and host defense mechanisms are based on some form of underlying epidemiological model. Such models are typically of the susceptible-infected variety (i.e., *S–I* models; Hethcote, 2000; Diekmann and Heesterbeek, 2000). A particularly simple example is the following:

$$\frac{dS}{dt} = F - \beta SI + cI,$$
$$\frac{dI}{dt} = \beta SI - vI - cI. \tag{2.1}$$

Here S and I are the sizes of the susceptible and infected host population, F is a function that specifies the dynamics of the susceptible population in the absence of the pathogen, β is the transmission rate coefficient, v is the pathogen-induced mortality rate (i.e., virulence), and c is the per capita rate at which the infected hosts recover from infection (referred to as the host clearance rate) and (in this model) become susceptible again.

To study evolutionary change in pathogen virulence in a game-theoretic context using the above model, we first suppose that the epidemiological dynamics reach a stable endemic equilibrium (i.e., an equilibrium where the pathogen has not gone extinct) and we then imagine introducing a mutant pathogen strain into the population (Day and Proulx, 2004). We can then derive an expression for the growth rate of this mutant strain. This growth rate expression is a measure of the mutant's fitness (Metz et al., 1992) and it can be used to obtain the evolutionarily stable (ES) pathogen strain. Under a wide variety of models, it has been shown that the ES pathogen strain is one that maximizes the expected number of new infections produced by a single infected host (per available susceptible host) (Bremermann and Thieme, 1989; Frank, 1996). For the above model, this quantity can be calculated as $\beta/(v+c)$; β is the rate at which new infections are produced, and $1/(v+c)$ is the expected duration of an infection.

From the expression for pathogen fitness, $\beta/(v+c)$, we can see that the ESS strain is one for which virulence is zero. Clearly not all pathogens are avirulent, however, and this has led many researchers to suppose that there are tradeoffs between various parameters in the above fitness expression. The tradeoff that has received the most attention is one between the transmission rate, β, and virulence, v. In particular, it is often supposed that pathogen strains cannot have a high transmission rate without also inducing a high mortality rate on their hosts. This can be incorporated into the above model by treating transmission rate as an increasing function of virulence. A particularly simple form that we will use below is $\beta = v^n$, where n is a shape parameter. Since we want the transmission rate β to exhibit diminishing returns against the virulence v, we take n to be between zero and one.

Similarly, we can study evolutionary change in host clearance rate in a game-theoretic context by introducing a

mutant host genotype at the endemic equilibrium of model (2.1) that codes for an altered level of clearance (van Baalen, 1998; Day and Burns, 2003). Again we can then derive an expression for the growth rate of this mutant host genotype. This is a measure of host fitness and it can be used to calculate the ES host clearance rate. The expression for the mutant host growth rate that is obtained from models such as (2.1) is a bit complicated, however, owing to the fact that there are two kinds of host individuals that might have the mutant genotype (susceptible hosts and infected hosts; see van Baalen, 1998; Day and Burns, 2003).

Given the two above-mentioned fitness expressions for mutant pathogens and for mutant hosts, we can then also analyse coevolutionary models. Specifically, we can use these two fitness expressions to determine the joint, coevolutionarily stable level of pathogen virulence and host clearance (van Baalen, 1998; Day and Burns, 2003). Generally speaking, virulence is selectively advantageous for the pathogen owing to its effect on transmission rate, but it is also selected against because infections last longer (and therefore can transmit pathogen propagules for longer) when virulence is low. The ESS virulence strikes a balance between these benefits and costs, and just where this balance is struck depends on the host clearance rate (van Baalen, 1998). Host clearance also has costs and benefits. Clearance is obviously beneficial to the host since it allows the host to rid itself of the infection. Clearance is also assumed to have costs, however, because energy devoted to defense mechanisms is energy taken away from other components of host fitness. For example, high clearance rates might come at the cost of reduced host fecundity (e.g., Demas et al., 1997; Moret and Schmid-Hempel, 2000). The ESS clearance rate then strikes a balance between these benefits and costs, and where this balance is struck depends on the pathogen's level of virulence.

Our goal here is to consider the outcome of such coevolutionary interactions when there is some plasticity of behaviour—either a possibility for the pathogen to adjust the level of virulence in response to the clearance rate exhibited by the host, or a possibility for the host to adjust the level of clearance in response to the level of virulence exhibited by the pathogen, or both. When both can respond, we will introduce "response rules" for both players that specify the virulence that the pathogen exhibits in response to the clearance rate of the host and vice versa. The interaction of a host and pathogen response rule during an infection results in some steady state level of virulence and clearance during the infection. These steady state values can then be used in the above-mentioned fitness expressions for host and pathogen to calculate the resulting fitness of each party. Our goal will be to find the coevolutionarily stable *response rules* of host and pathogen.

Note that this is a subtle but important change in perspective from previous host–pathogen models. The objects of the analysis here are no longer the virulence and clearance rates per se, but rather the response rules of the host and the pathogen. Of course, these response rules will generate some level of virulence and clearance during any given infection, but it is the response rule itself that is considered to be the evolutionary object of interest. We would like to know whether the level of virulence and clearance that results from the coevolutionarily stable response rules differs from those that are obtained in the absence of this type of negotiation.

In order to make the exposition as simple as possible we have elected to simplify the underlying epidemiological model on which the fitness expressions are based. Specifically, we focus on a host–pathogen system in which all hosts are infected as juveniles. If the host's clearance rate is c during this stage and the pathogen's virulence is v, then the host will survive infection with probability $c/(c+v)$. If the host survives infection during the juvenile stage, then it can expect a total reproductive output of $b/(m_0+c)$ during the adult stage, where b is the rate of offspring production as an adult, and (m_0+c) is the mortality rate of an adult (which is assumed to increase linearly with the investment made in clearance during the juvenile stage from a baseline value of m_0). With these specifications, we have the following pathogen and host fitness expressions:

$$\text{Pathogen}: \quad P(v,c)=\frac{\beta}{v+c}=\frac{v^n}{v+c},$$
$$\text{Host}: \quad H(v,c)=\frac{c}{(v+c)}\,\frac{b}{(m_0+c)}. \tag{2.2}$$

Of course fitnesses will also depend on the population-wide average values $\hat{v}$ and $\hat{c}$, but we will usually suppress this dependence in the notation. To keep things simple we will assume that population-wide behaviour acts simply to normalize P and H, so that each of the expressions in (2.2) gets multiplied by a "constant" that depends on $\hat{v}$ and $\hat{c}$. These constants won't affect the equilibrium conditions nor the stability conditions (ES and CS) and we will omit them. We remark that in this simple case, both stability conditions (ES and CS) reduce to the two fitness maximum conditions $P_{vv}<0$ and $H_{cc}<0$. Expressions (2.2) will be used in the remainder of this article.

3. The effect of plasticity on the levels of investment

In this section we analyse the coevolutionary outcome of host–pathogen interactions under three different assumptions about the occurrence of phenotypic plasticity: (i) no plasticity in host or pathogen, (ii) plasticity in either the host or the pathogen (but not both), and (iii) plasticity in both the host and the pathogen.

3.1. No plasticity

First we consider a host–pathogen system in which the pathogen exhibits a fixed virulence v and the host exhibits a fixed clearance rate c during an infection. The fitness of both individuals, $P(v,c)$ for the pathogen and $H(v,c)$ for

the host, will depend upon both strategies as defined by (2.2). Given this, we expect a monomorphic population to reside at a Nash equilibrium determined by

$$P_v(\hat{v},\hat{c}) = H_c(\hat{v},\hat{c}) = 0, \tag{3.1}$$

where subscripts denote partial derivatives. We assume that this equilibrium is stable in the two fundamental ways, ES (Maynard Smith, 1974) and convergence stable (CS) (Eshel, 1983; Christiansen, 1991). The above results will serve as a baseline against which we can compare the evolutionary outcomes when there is plasticity.

3.2. Plasticity in either the host or the pathogen

We investigate what happens to the above equilibrium if one of the players can respond plastically to information about the action exhibited by the other player (Fudenberg and Tirole, 1991; Pen and Weissing, 2002; Abe et al., 2003; Pen and Taylor, 2005). We expect that this will change the equilibrium levels of virulence $\hat{v}$ and/or clearance rate $\hat{c}$, and our objective is to determine whether these will increase or decrease. If the pathogen knows the clearance rate c of the host, it will choose v to maximize its fitness P and the derivative condition for this is $P_v(v,c) = 0$. If the host knows the virulence level v of the pathogen, it will choose c to maximize its fitness H and the derivative condition for this is $H_c(v,c) = 0$. These "best response" curves are plotted in Fig. 1 for the functions of (2.2). They intersect at the Nash equilibrium (3.1).

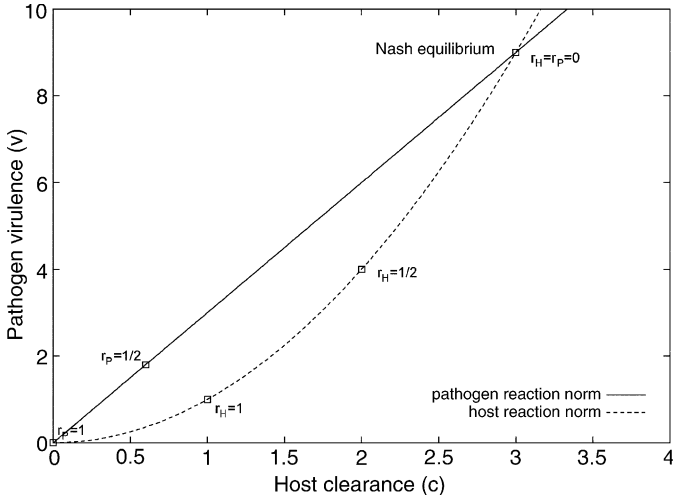

Fig. 1. Host and pathogen reaction norms and a number of equilibrium points of various types using the fitness functions (2.2) with $n = 3/4$, $m_0 = b = 1$. If the pathogen knows the host's clearance rate c (host precedes), the equilibrium is on the curve defined by $P_v = 0$, and this is the line $v = 3c$. If the host knows the pathogen's virulence v (pathogen precedes), the equilibrium is on the curve defined by $H_c = 0$, and this is the curve $v = c^2$. The Nash equilibrium (3.1) is at $c = 3$, $v = 9$ where these two intersect. Under precedence, the Stackelberg equilibria are marked for reliability $r = 1/2$ and $r = 1$. We use r_H when the host responds to the pathogen and r_P when the pathogen responds to the host.

Here, we allow the possibility that information received from the partner may be incomplete, and thus we model fitness as

$$\text{Pathogen fitness (host goes first)}\quad rP(v,c) + (1-r)P(v,\hat{c}), \tag{3.2}$$

$$\text{Host fitness (pathogen goes first)}\quad rH(v,c) + (1-r)H(\hat{v},c), \tag{3.3}$$

where r denotes the *reliability* of the information received or in another interpretation it could represent the probability that the pathogen or host is able to respond to the information. For example, suppose the host goes first. When $r = 1$ the pathogen receives perfect information about the host's c-value and can respond precisely by choosing v to maximize $P(v,c)$. When $r = 0$, there is no effective information transfer and the pathogen treats the host as a random member of the population, with clearance rate $\hat{c}$, and responds by choosing v to maximize $P(v,\hat{c})$. For intermediate r, we take fitness (3.2) as an average of these two cases.

Host goes first: The first derivative condition for the maximization of (3.2) as a function of v is

$$rP_v(v,c) + (1-r)P_v(v,\hat{c}) = 0. \tag{3.4}$$

Eq. (3.4) can be regarded as defining v implicitly as a function of c: $v = \tilde{v}(c)$. The host, on the other hand, will be expected to exhibit a value of c that maximizes its fitness (which depends on what the pathogen will do):

$$\tilde{H}(c) = H(\tilde{v}(c), c). \tag{3.5}$$

The first derivative condition for that is

$$\frac{\mathrm{d}\tilde{H}}{\mathrm{d}c} = H_v\frac{\mathrm{d}\tilde{v}}{\mathrm{d}c} + H_c = 0. \tag{3.6}$$

We find the derivative of $\tilde{v}$ by differentiating Eq. (3.4) with respect to c, with $v = \tilde{v}(c)$ and then evaluating everything at the Stackelberg equilibrium $\hat{c}$:

$$\frac{\mathrm{d}}{\mathrm{d}c}(rP_v(\tilde{v}(c),c) + (1-r)P_v(\tilde{v}(c),\hat{c}))\big|_{c=\hat{c}} = P_{vv}\frac{\mathrm{d}\tilde{v}}{\mathrm{d}c} + rP_{vc} = 0,$$

$$\frac{\mathrm{d}\tilde{v}}{\mathrm{d}c} = -r\frac{P_{vc}}{P_{vv}}. \tag{3.7}$$

Putting this into (3.6) we get the host's condition

$$\frac{H_c}{H_v} = r\frac{P_{vc}}{P_{vv}}. \tag{3.8}$$

Eqs. (3.4) and (3.8) together, evaluated at the population-wide average values $\hat{v}$ and $\hat{c}$, give us the equilibrium. In the economics literature, typically with $r = 1$, this is known as a Stackelberg equilibrium (Fudenberg and Tirole, 1991). In case $r = 0$, there is no useful information transfer and we obtain, as expected, the Nash equilibrium (3.1).

Pathogen goes first: Here we suppose that the host, upon infection, can assess, with some uncertainty, the level of virulence of the pathogen and is able to respond by investing more or less resources into clearance. Then, exactly as above, with host fitness (3.3), the equilibrium

conditions are:

$$H_c(v,c)=0, \tag{3.9}$$

$$\frac{P_v}{P_c}=r\frac{H_{cv}}{H_{cc}}, \tag{3.10}$$

where r in this case denotes the reliability of the information received by the pathogen.

3.2.1. The effect of precedence on the equilibrium levels of virulence and clearance

To assess the effect of precedence on the equilibrium levels of v and c we will need the signs of the various first and second derivatives. We make the assumption that, generally, near equilibrium, increased virulence will decrease the host's fitness and prompt the host to increase clearance rate. Similarly, increased clearance will decrease the pathogen's fitness and prompt an increase in virulence. That is, higher levels of investment from one player will decrease the fitness of the other and promote a higher level of investment from the other. We now show that these assumptions, together with that of the ES stability of the Nash equilibrium, imply

$$H_v<0,\ \ P_c<0,\ \ P_{vc}>0,\ \ H_{cv}>0,\ \ H_{cc}<0,\ \ P_{vv}<0. \tag{3.11}$$

The first two inequalities follow directly from our assumption on the fitness effects. Now we argue that P_{vc} should be positive. Suppose the pathogen is at equilibrium so that $P_v=0$. Now suppose the host increases clearance c. Our assumption is that this should cause the pathogen to increase v. Since the pathogen will be moving to a new fitness maximum, that means the derivative of P at the old equilibrium has become positive, that is, the derivative P_v has shifted from being 0 to being positive. Thus, an increase in c causes an increase in P_v. That tells us that $P_{vc}>0$. A similar argument shows that $H_{cv}>0$. Finally the last two follow from the fact that the Nash equilibrium is a fitness maximum for each player when the strategy of the other is fixed (the ES condition). The particular fitness functions of (2.2) have all these properties.

Now we suppose the host precedes the pathogen and examine the effect of this on the host's clearance rate c. The Nash equilibrium condition for the host is

$$H_c=0. \tag{3.12}$$

If the host goes first, the Stackelberg equilibrium condition is

$$H_c=rH_v\frac{P_{vc}}{P_{vv}} \tag{3.13}$$

and conditions (3.11) tell us that the right-hand side of this equation is positive when evaluated at the Nash equilibrium. That means that in the shift from Nash to Stackelberg, H_c changes from zero to positive. Since H_{cc} is negative, this implies that c decreases. Thus, in going first, the host employs a lower clearance rate. Since $\mathrm{d}\tilde{v}/\mathrm{d}c$ is positive, the pathogen will also employ a lower virulence. A parallel argument shows that when the pathogen goes first the resulting levels of investment are also lower. Thus the ability of either the host or the pathogen to respond plastically to information gained through precedence generally lowers both virulence and clearance. Sample results are plotted on Fig. 1 for $r=1$ and $1/2$. For the functions (2.2) used there, the decrease (from the Nash equilibrium) in both virulence and clearance is greater when the host precedes than when the pathogen precedes.

3.3. Plasticity in both the host and the pathogen

Here we look at *negotiation*, in which there is an opportunity for the two players to repeatedly respond to one another plastically. We use a negotiation framework introduced by McNamara et al. (1999) and extended by Taylor and Day (2004) in which each player responds to an action from the other with an action specified by its response rule, with this process going back and forth converging hopefully to a final pair of strategies $v^{\#}$ and $c^{\#}$ which determine the fitness of each player. The negotiation is effected by a pair of linear response rules

$$\begin{aligned} v&=\rho_v-\lambda_v c,\\ c&=\rho_c-\lambda_c v, \end{aligned} \tag{3.14}$$

where the *responsiveness* λ measures the degree to which an individual responds to a change in the offer of its partner. Note that we are using subscripts here, on the variables ρ and λ, not as partial derivatives, but as identifiers. In this framework, the pathogen chooses ρ_v and λ_v, and the host chooses ρ_c and λ_c, and the limit points of the recursive equations are

$$\begin{aligned} v^{\#}&=\frac{\rho_v-\rho_c\lambda_v}{1-\lambda_v\lambda_c},\\ c^{\#}&=\frac{\rho_c-\rho_v\lambda_c}{1-\lambda_v\lambda_c}, \end{aligned} \tag{3.15}$$

where convergence requires $|\lambda_v\lambda_c|<1$. With these, the fitnesses become functions of the four negotiation strategies:

$$\begin{aligned} P^{\#}(\rho_v,\lambda_v,\rho_c,\lambda_c)&=P(v^{\#},c^{\#}),\\ H^{\#}(\rho_v,\lambda_v,\rho_c,\lambda_c)&=H(v^{\#},c^{\#}). \end{aligned} \tag{3.16}$$

The Nash equilibrium for this system consists of the equations

$$\frac{\partial P^{\#}}{\partial\rho_v}=0,\ \ \frac{\partial P^{\#}}{\partial\lambda_v}=0,\ \ \frac{\partial H^{\#}}{\partial\rho_c}=0,\ \ \frac{\partial H^{\#}}{\partial\lambda_c}=0. \tag{3.17}$$

However, as pointed out by Taylor and Day (2004), this system is degenerate and only provides two independent equations. If we set λ_v and λ_c as parameters, we can solve these four equations for the ρ-variables:

$$\rho_v=\frac{(3+\lambda_v-4\lambda_v\lambda_c)^2}{(1-\lambda_v)(1-4\lambda_c)}, \tag{3.18}$$

$$\rho_c = \frac{(3+\lambda_v-4\lambda_v\lambda_c)(1-\lambda_c)}{(1-\lambda_v)(1-4\lambda_c)^2}. \tag{3.19}$$

This system is equivalent to the system of four equations above. These equations define a 2-dimensional manifold of neutrally stable equilibrium points in four-dimensional negotiation phase space, and it is mathematically unclear where on this surface we expect a model population to settle. If we project this system into (v,c) space, we get

$$v = \frac{3(3+\lambda_v-4\lambda_v\lambda_c)}{(1-\lambda_v)(1-4\lambda_c)^2}, \tag{3.20}$$

$$c = \frac{(3+\lambda_v-4\lambda_v\lambda_c)}{(1-\lambda_v)(1-4\lambda_c)}. \tag{3.21}$$

We ran simulations of this game with individual variation in the parameters of the two response rules, and with an individual's fitness (and thus the representation of its response rule in the next generation) given by a particular realization of Eq. (2.2). The results are displayed in Figs. 1 and 2 and show an evolution of negotiation rules that yields lower values of both virulence and clearance.

4. Discussion

We have developed a simple evolutionary model of host–pathogen interaction. Our objective is not to detail the complex interactions that may arise between a host and its pathogen, but to illustrate how different degrees of flexibility in an individual's ability to respond to its opponent's behaviour can significantly alter the evolutionary outcome of biological interactions, with reference to a real-life example. In this evolutionary game, the host is described by a single variable called clearance, measuring the rate at which it recovers from infection. Clearance is the outcome of up-regulation of immune function, which we assume also results in a down-regulation of host fecundity. The pathogen is described by a single variable called virulence, measuring its impact on host mortality. Virulence is assumed to be a by-product of pathogen investment into reproduction (e.g., Anderson and May, 1982; Frank, 1996), and is therefore positively correlated with the pathogen transmission rate.

We examined the co-evolution of virulence and clearance in three distinct contexts. First, we considered the case where the two opponents are unable to respond to one another's actions; i.e. they both must choose a single clearance or virulence level and then retain these values for the duration of the infection. This context is typical of such coevolutionary analyses and we present it here simply as a benchmark for comparison. The end-point of this evolutionary process is the familiar Nash equilibrium levels of virulence and clearance (Fig 1).

Secondly, we considered a situation in which only one player can respond to the action taken by the other. In other words, one of the players (player 1) commits to an action before its opponent (player 2) and then must retain this action for the duration of the infection. The second player then uses a "response rule" or reaction norm that specifies its action as a function of the action chosen by player 1. In this case, it is the action of player 1 and the response rule of player 2 that are the objects of selection, and we expect the response rule of player 2 to evolve so as to maximize its fitness for each possible action chosen by player 1 (the so-called best response rule). The action of player 1 is expected to evolve so as to maximize its fitness conditional on player 2 adopting the best response rule. The resulting joint outcome is referred to as the Stackelberg equilibrium. We found that whether the host or the pathogen takes precedence influences the results quantitatively but not qualitatively: in both cases, both virulence and clearance are lower than in the Nash equilibrium.

These ideas are discussed in the economics literature (e.g. Fudenberg and Tirole, 1991) but here one makes an assumption of rational behaviour. Note that we do not need this here but can expect reaction norms to reflect a complex hard-wired response. There is empirical evidence that organisms are less likely to adopt fitness-maximizing behaviours within contexts that are less frequently encountered (Herre, 1987). This may be interpreted as there having been insufficient selection for the appropriate response in such contexts, suggesting that behavioural reaction norms will tend to be less optimal in homogeneous populations. However, the result need not imply suboptimality if individuals are responding to incomplete information, and this has been allowed for within our analysis. In general, the assumption of optimality is one of the most powerful tools we have for making predictions in evolutionary biology (Parker and Maynard Smith, 1990; Sutherland, 2005), so we have proceeded on this assumption.

Third, we considered the situation where both players can sequentially respond to one another's actions. In this case both players have response rules or reaction norms, and it is these rules that are the objects of selection (McNamara et al., 1999; Taylor and Day, 2004). During an interaction (e.g., an infection) one player first chooses its action, and the second player then responds with its action, which is determined by its reaction norm and the observed action of the first player. Then the first player responds to the second and so on. This 'negotiation' proceeds indefinitely, typically eventually converging to a steady state pair of actions. In this case, again, we find that this results in a more "peaceful" outcome relative to the Nash equilibrium (lower virulence and clearance).

Why does plasticity have this effect? In a population at the Nash equilibrium, a pathogen that reduces its virulence incurs a fitness decrement. Similarly, a host reducing its clearance below the Nash equilibrium also reduces its fitness. How can natural selection promote such sacrifice? One way for this to occur is for the phenotypes of interacting partners to be positively correlated. That is to say, less aggressive individuals (i.e., pathogens with lower virulence, or hosts with lower clearance) are associated

with less aggressive partners. In this case, when both partners are less aggressive, both will benefit. The most studied cause of such correlation is genealogical closeness (kin selection; Hamilton, 1963, 1964; Maynard Smith, 1964). However, phenotypic plasticity provides another mechanism whereby such correlations may arise (Nee, 1989; Taylor and Day, 2004; Gardner and West, 2004). Specifically, if reaction norms evolve in such a way as to result in a positive correlation, then such cooperative outcomes will occur. Our results demonstrate that such reaction norms do evolve, and therefore the question at hand can be recast as: why do reaction norms result in such positive correlations, and thus more cooperative outcomes, evolve?

In the standard game theoretic analysis host and pathogen associate at random and express hard-wired behaviours. Thus, at the Nash equilibrium, there is no scope for phenotypic correlation between partners. In precedence games, however, phenotypic plasticity of player 2 potentially leads to a correlation between players' strategies. First, consider the case where player 1 is the host. If clearance is large, then the best pathogen response is one with a large virulence as well, because the optimal level of virulence increases with the risk of immune clearance (van Baalen, 1998). Second, consider the case where player 1 is the pathogen. If virulence is large, then the best host response is one with large clearance too, because the optimal level of clearance increases with virulence (van Baalen, 1998). Therefore, when only one player can respond plastically, we expect its reaction norm to evolve to have a positive slope, causing a positive correlation between the aggressiveness expressed by each partner. As a result, this leads to a more peaceful settlement than the Nash equilibrium.

The situation in which both players respond plastically, presents more of a challenge, both conceptually and analytically. When your partner is plastic, then your own phenotype determines his, and you have thus gained some control over him, but when this control is bidirectional it is not clear how it should operate and some form of reaction norm needs to be assumed. Our use of negotiation rules leads to a neutrally stable and therefore uncertain behavioural state and we have had to resort to simulations to determine the evolutionary outcome. These show that

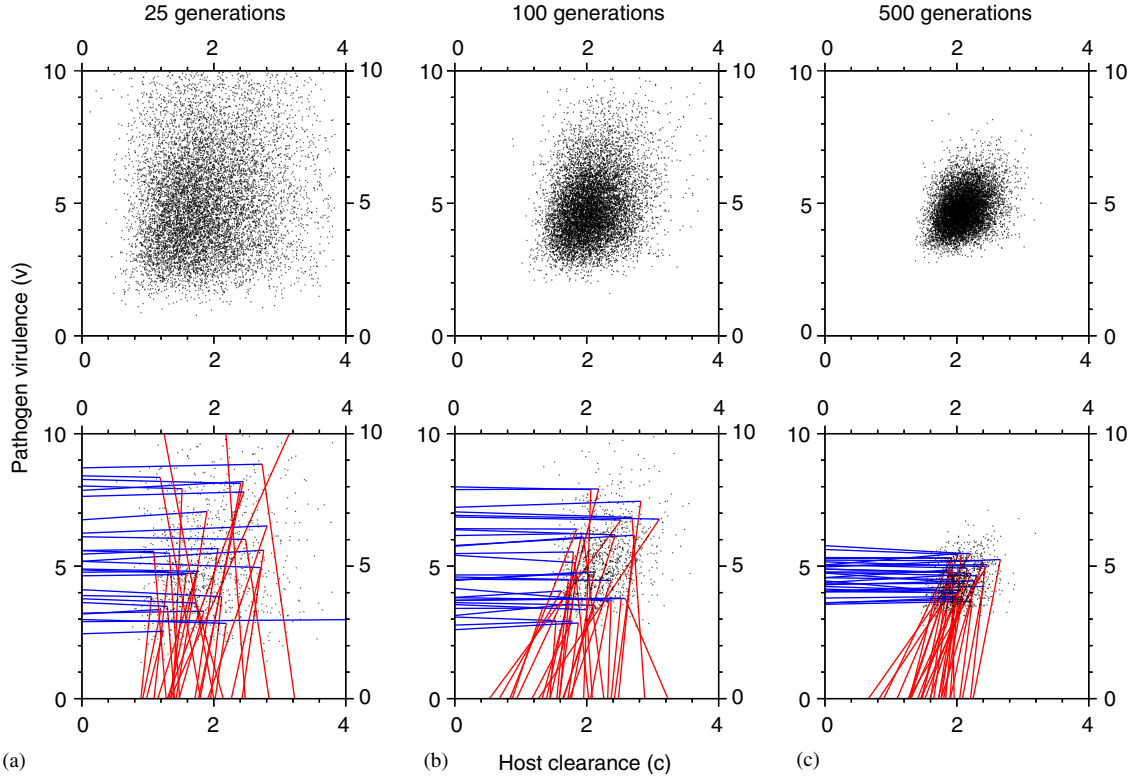

Fig. 2. Change over time from simulation studies of the negotiation game under the joint forces of selection and mutation. The top figures show the projection of the population into (v, c) space. The bottom figures show sample response lines. The lines that slope up from the horizontal axis are host responses and the lines that are nearly horizontal are pathogen responses. The dots in the top figure are plotted at the intersection of a corresponding pair of lines. The population size is 10 million individuals (host–pathogen pairs); the top figures provide a sample of 1000 and the bottom figures provide a sample of 30. In negotiation parameter space we begin (generation 0) with λ normally distributed with mean 0 and variance 0.01 and ρ exponentially distributed with mean 5. We calculate the host and pathogen fitness for each individual (Eq. (2.2) with $n = 3/4$, $m_0 = b = 1$) and use these to form the next generation as follows: for each of the 10 million "slots" a host and a pathogen are picked independently from the previous generation with a probability proportional to their respective fitness values. At that point, the ρ and λ parameters for both host and pathogen are allowed to mutate by being multiplied by $\exp(\varepsilon)$ where ε is normally distributed with mean 0 and standard deviation 0.01: (a) after 25 generations, (b) after 100 generations, (c) after 500 generations.

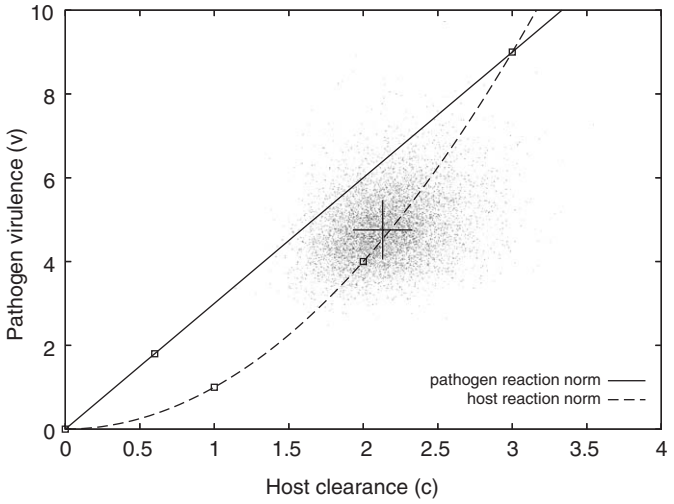

Fig. 3. Fig. 1(c) is superimposed on Fig. 1, the cross marking the mean of the distribution.

the host reaction norms tend to evolve to have positive slope on average (Fig. 2), while the pathogen reaction norms are fairly flat, and as a result the evolutionary outcome of the system is more cooperative than the Nash equilibrium.

Negotiation is therefore a relevant question each time two partners with plastic abilities interact. This will frequently arise in the natural world. For example, host–pathogen interactions are intimate and lasting relationships, during which numerous signals may be exchanged. On the one hand, hosts clearly receive information from their pathogens and respond by regulating immune function. On the other hand, numerous pathogens might also be able to assess the physiological status of their host and adjust their behaviour accordingly. Indeed, as in any species with restricted habitat choice, phenotypic plasticity could generally be a key determinant of the strategy of pathogens (Agrawal, 2001; Thomas et al., 2002). Nevertheless, plasticity in host–pathogen systems has usually been interpreted as a mechanism of optimization in varying environments. Here we suggest that the understanding of these mechanisms could benefit a lot from a novel point of view based on negotiation. For instance, the development of benevolent relationships between immune system and digestive bacteria involves the reciprocal exchange of information regarding the respective 'actions' of each protagonist (Mowat, 2003; Aldridge et al., 2005). More generally, the majority of host–bacterial interactions that end up without violent escalation (see Merrell and Falkow, 2004) could be interpreted as the results of negotiation processes (Fig. 3).

Acknowledgments

We are grateful for many excellent comments from four referees. This work has been supported by the Natural Sciences and Engineering Research Council of Canada (to PDT and TD), by Queen's University Advisory Research Council and by the Canada Research Chairs program.

References

Abe, J., Kamimura, Y., Ito, H., Matsuda, H., Shimada, M., 2003. Local mate competition with lethal male combat: effects of competitive asymmetry and information availability on a sex ratio game. J. Evol. Biol. 16, 607–613.

Agrawal, A.A., 2001. Phenotypic plasticity in the interactions and evolution of species. Science 294, 321–326.

Aldridge, P.D., Gray, M.A., Hirst, B.H., Khan, C.M.A., 2005. Who's talking to whom? Epithelial–bacterial pathogen interactions. Molec. Microbiol. 55, 655–663.

Anderson, R.M., May, R.M., 1982. Coevolution of hosts and pathogen. Parasitology 85, 411–426.

Bremermann, H.J., Thieme, H.R., 1989. A competitive exclusion principle for pathogen virulence. J. Math. Biol. 27, 179–190.

Christiansen, F.B., 1991. On conditions for evolutionary stability for a continuously varying character. Am. Nat. 138, 37–50.

Day, T., Burns, J.G., 2003. A consideration of patterns of virulence arising from host–pathogen coevolution. Evolution 57, 671–676.

Day, T., Proulx, S., 2004. A general theory for the evolutionary dynamics of virulence. Am. Nat. 163, E40–E63.

Demas, G.E., Chefer, V., Talan, M.I., Nelson, R.J., 1997. Metabolic costs of mounting an antigen-stimulated immune response in adult and aged C57BL/6J mice. Am. J. Physiol. Reg. Integr. Comp. Physiol. 42, R1631–R1637.

Diekmann, O., Heesterbeek, J.A.P., 2000. Mathematical Epidemiology of Infectious Disease. Wiley, New York.

Dieckmann, U., Metz, J.A.J., Sabelis, M.W., Sigmund, K. (Eds.), 2002. In Pursuit of Virulence Management. Cambridge University Press, Cambridge.

Doebeli, M., Hauert, C., 2005. Models of cooperators based on the Prisoner's dilemma and the snowdrift game. Ecol. Lett. 8, 748–766.

Doebeli, M., Knowlton, N., 1998. The evolution of interspecific mutualism. Proc. Natl. Acad. Sci. USA 95, 8676–8680.

Eshel, I., 1983. Evolutionary and continuous stability. J. Theor. Biol. 103, 99–111.

Frank, S.A., 1996. Models of pathogen virulence. Q. Rev. Biol. 71, 37–78.

Fudenberg, D., Tirole, J., 1991. Game Theory. The MIT Press, Cambridge, MA.

Gandon, S., Mackinnon, M.J., Nee, S., Read, A.F., 2001. Imperfect vaccines and the evolution of pathogen virulence. Nature 414, 751–756.

Gardner, A., West, S.A., 2004. Cooperation and punishment, especially in humans. Am. Nat. 164, 753–764.

Hamilton, W.D., 1963. The evolution of altruistic behavior. Am. Nat. 97, 354–356.

Hamilton, W.D., 1964. The genetical evolution of social behaviour I. J. Theor. Biol. 7, 1–16.

Hethcote, H.W., 2000. The mathematics of infectious diseases. SIAM Rev. 42, 599–653.

Herre, A.E., 1987. Optimality, plasticity and selective regime in fig wasp sex ratios. Nature 329, 627–629.

Kokko, H., 1999. Cuckoldry and the stability of biparental care. Ecol. Lett. 2, 247–255.

Maynard Smith, J., 1964. Group selection and kin selection. Nature 201, 1145–1147.

Maynard Smith, J., 1974. The theory of games and the evolution of animal conflicts. J. Theor. Biol. 47, 209–221.

Maynard Smith, J., Price, G.R., 1973. The logic of animal conflict. Nature. 246, 15–18.

McNamara, J.M., Gasson, C.E., Houston, A.I., 1999. Incorporating rules for responding into evolutionary games. Nature 401, 368–371.

McNamara, J.M., Houston, A.I., Barta, Z., Osorno, J.-L., 2003. Should the young ever be better off with one parent than with two? Behav. Ecol. 14, 301–310.

Merrell, D.S., Falkow, S., 2004. Frontal and stealth attack strategies in microbial pathogenesis. Nature 430, 250–256.

Metz, J.A.J., Nisbet, R.M., Geritz, S.A.H., 1992. How should we define 'fitness' for general ecological scenarios? Trends Ecol. Evol. 7, 198–202.

Moret, Y., Schmid-Hempel, P., 2000. Survival for immunity: The price of immune system activation for bumblebee workers. Science 290, 1166–1168.

Mowat, A.M., 2003. Anatomical basis of tolerance and immunity to intestinal antigens. Nat. Rev. Immunol. 3, 331–341.

Nee, S., 1989. Does Hamilton's rule describe the evolution of reciprocal altruism? J. Theor. Biol. 141, 81–91.

Parker, G.A., Maynard Smith, J., 1990. Optimality theory in evolutionary biology. Nature 348, 27–33.

Pen, I., Taylor, P.D., 2005. Modeling information exchange in worker-queen conflict over sex allocation. Proc. R. Soc. Lond. B 272, 2403–2408.

Pen, I., Weissing, F.J., 2002. Optimal sex allocation: steps towards a mechanistic theory. In: Hardy, I. (Ed.), The Sex Ratio Handbook. Cambridge University Press, Cambridge, pp. 26–45.

Restif, O., Koella, J.C., 2003. Shared control of epidemiological traits in a coevolutionary model of host–pathogen interactions. Am. Nat. 161, 827–836.

Roberts, G., Sherratt, T.N., 1998. Development of cooperative relationships through increasing investment. Nature 394, 175–179.

Sutherland, W.J., 2005. The best solution. Nature 435, 569.

Taylor, P.D., Day, T., 2004. Stability in negotiation games and the emergence of cooperation. Proc. R. Soc. B (London) 271, 669–674.

Thomas, F., Brown, S.P., Sukhdeo, M., Renaud, F., 2002. Understanding pathogen strategies: a state-dependent approach? Trends Parasitol. 18, 387–390.

van Baalen, M., 1998. Coevolution of recovery ability and virulence. Proc. R. Soc. B 265, 317–325.

Wahl, L., Nowak, M.A., 1999a. The continuous prisoner's dilemma: I. Linear reactive strategies. J. Theor. Biol. 200, 307–321.

Wahl, L., Nowak, M.A., 1999b. The continous prisoner's dilemma: II. Linear reactive strategies with noise. J. Theor. Biol. 200, 323–338.