Playing games with transmissible animal disease

Jonathan Cave
Research Interest Group

6 May 2008

THE UNIVERSITY OF WARWICK

Outline

- The nexus of game theory and epidemiology
- Some simple disease control games
  - A vaccination game with perceived risk
  - A game of interdependent risks
- Towards an elaborated structure
  - Dynamics
  - Differentiated interaction structures - (layered) networks
  - The evolution of conventions
  - Coevolution of structure and behaviour
- Different, differentiated diseases

The connection…

- Game theory is concerned with strategic behaviour - reasoned (rational) choices made by interdependent agents:
  - Players – those who make conscious choices
  - Strategies – what the players choose
  - Payoffs –players' preferences over combined choices (note: sometimes explicit 'rules' translate choices into outcomes over which players have preferences)
  - Information – what players know about these things
- Epidemiology provides various ways to formalise dynamic interdependence
- Basis of a game-theoretic analysis can be supplied by an epidemiological model
  - *Payoffs* affected by disease prevalence, incidence and (e.g.) market and welfare impact
  - *Strategies* (for controlling disease, risks, impacts, etc.) determined by disease characteristics
  - *Information* influenced by observed disease progress, choices (e.g. to notify, call in vets, etc.)
- Strategic behaviour in turn affects epidemiology
  - Animal movements, contact
  - Vaccination, culling, etc.
- This talk describes some simple models and their elaboration
- It tries to find common ground by using semi-mathematical language
- Hope is to get feedback on what's already old hat, what results are interesting, what extensions are promising…

Game theory basics

- Player i's payoffs denoted $U_i(\sigma_i,\sigma_{-i},\omega)$, where σ_i (σ_{-i}) are the strategies of I (and others) and ω is the *state* (not used in what follows)
- σ^* *is a Nash equilibrium at* ω iff for all i and all $s_i \neq \sigma_i^*$, $U_i(\sigma^*,\omega) \geq U_i(s_i,\sigma_{-i}^*,\omega) = U_i(\sigma^*|s_i,\omega)$ (mutual best replies)
- Game is:
  - *symmetric* if the strategy spaces and payoffs for each player are the same
  - *aggregate* if each player's payoff depends on its own strategy and the distribution of other players' strategies across the strategy set (the numbers playing each other strategy)
  - *Potential* if there is a real-valued function P of the strategies whose joint maxima identify the Nash equilibria (Formally, for each I, σ and s_i: $P(\sigma)-P(\sigma|s_i) = U_i(\sigma)-U_i(\sigma|s_i)$)
  - Example 1: a *network* of players playing 2-person games; i gets the average (or total) payoff from all his pairwise interactions
  - Example 2: a *market game* where the payoff to player i depends on his output and the aggregate of others' output
- Other solution concepts defined in terms of stability under specified dynamics:
  - Evolutionary stability: no sufficiently good deviation will be copied
  - Convergent stability: if many players adopt Q as an alternative to an equilibrium P and if payoff increases as players move closer to P than Q
  - Replicator dynamics: the prevalence of strategies that do best among those currently played increases
  - Players chosen at random select best replies to others' strategies with high (but < 1) probability

A simple vaccination game

- In deciding whether to vaccinate, farmers consider
  - (perceived) risk/cost of morbidity from vaccination (r_V)
  - (perceived) probability of infection ($\pi_{p,}$ which depends on the uptake level p)
  - (perceived) risk/cost of morbidity from infection (r_I)
- Decisions are indirectly influenced by others because the sum of others' decisions determines vaccine coverage
- This simple model shows how risk/cost perception influences expected vaccine uptake and coverage and the role played by pathogens' epidemiological characteristics
- All individuals have the same herd size, information and way to assess risk/costs

Static results

- Generally get stable convergence to homogeneous Nash equilibrium P*
- Expected variation in behaviour is here replaced by uniform mixed strategies: consider a 'combination' σ of strategies P* and Q
  - In σ, fractions μ and 1-μ play P* and an alternative Q
  - (Uptake/coverage) $p = \mu P^* + (1-\mu)Q$
  - Payoff to playing P* is $U(P^*, \sigma, \mu) = V(P^*, \mu P^* + (1-\mu)Q)$
  - Payoff to playing Q is $U(Q, \sigma, \mu) = V(Q, \mu P^* + (1-\mu)Q)$
  - Advantage of playing P* rather than Q is $A(P^*,Q) = (\pi_p - \rho)(P^*-Q)$
- *Lemma*: For any given ρ, there is a unique P* s.t. $A(P^*,Q) > 0$ for all $Q \neq P^*$ and all $\mu > 0$.
  - Letting $\mu \rightarrow 0$ shows that $P^*(\rho)$ is a Nash equilibrium
  - If P and Q are not Nash, but $|P^*-P| < |P^*-Q|$ then $A(P,Q)>0$ (stability)
- *Theorem*: if $\rho \geq \pi_0$ then the best reply to p = 0 is 0. Because higher p means lower π_p, the best reply to any p > 0 is also $P_i = 0$, and the unique equilibrium is P* = 0. By the same token if if $\rho \geq \pi_1$ then the unique equilibrium is P* = 1. Otherwise, there is a unique internal solution where all players use a strategy P* such that $\pi_{P^*} = \rho$.

Adding the SIR model

- We add a standard SIR dynamic model:

$$\begin{aligned}\dot{S} &= \delta\left(1-p\right)-\beta SI-\delta S\\ \dot{I} &= \beta SI-\gamma I-\delta I\\ \dot{R} &= \delta p+\gamma I-\delta R\end{aligned}$$

- δ = mean birth/death rate, β = mean transmission rate, γ = 1/(infectious period), p = uptake.
  - Assumes symmetrical mortality, no infection before (not) being vaccinated, etc.
  - Steady-state uptake = coverage.
  - Third equation is redundant (population balance).
  - Rescale to $\tau = t/\gamma$ (time in mean infectious period units), $\phi = \delta/\gamma$ (infectious period in mean lifetimes) and $R_0 = \beta/(\gamma+\delta)$ (2^o cases spawned by each 1^o case):

$$\begin{aligned}\frac{\partial S}{\partial \tau} &= \phi\left(1-p\right)-R_0\left(1+\phi\right)SI-\phi S\\ \frac{\partial I}{\partial \tau} &= R_0\left(1+\phi\right)SI-\left(1+\phi\right)I\end{aligned}$$

Long-term behaviour

- Whether the disease becomes endemic or disappears depends on the coverage relative to a critical threshold:

$$\hat{p} = \frac{\max\{R_0 - 1, 0\}}{R_0}$$

If $p \geq \hat{p}$, the system converges to $S^* = 1$; otherwise, it converges to the endemic steady state

$$S^E = 1 - \hat{p};\ I^E = \frac{\phi(\hat{p} - p)}{1+\phi}$$ so henceforth we assume $R_0 < 1$ and $p < \frac{\beta - \gamma - \delta}{\beta}$

- At coverage p, the long-term probability of infection for an unvaccinated animal depends on the relative rate at which it dies or becomes infected

$$\pi_p = \frac{R_0(1+\phi)S^E I^E}{R_0(1+\phi)S^E I^E + \phi S^E} = 1 - \frac{1}{R_0(1-p)}$$

- This is independent of ϕ and thus of the birth/death rate and the infectious period.
- There is a mixed strategy (imperfect uptake) equilibrium if $\pi_1 < \rho < \pi_0$, or

$$\rho < \frac{R_0 - 1}{R_0}$$, so the equilibrium $P^* = 1 - \frac{1}{R_0(1-\rho)}$

An illustration: impact of increasing R_0 (2º cases per 1º case)

- The LHS shows equilibrium uptake as a function of relative risk. Horizontal lines are ‘elimination thresholds’ – limit is step function at $\rho = 1$.
- The RHS shows the impact of an upward shift in risk perception (from <1 to the new value ρ). The upper part is the incentive to switch vaccination practice; lower part is corresponding change in uptake (from old to new equilibrium P) as functions of new risk level.
- This shows that behaviour is more responsive as $\beta/(\delta+\gamma)$ increases; and that recovery is slower than collapse

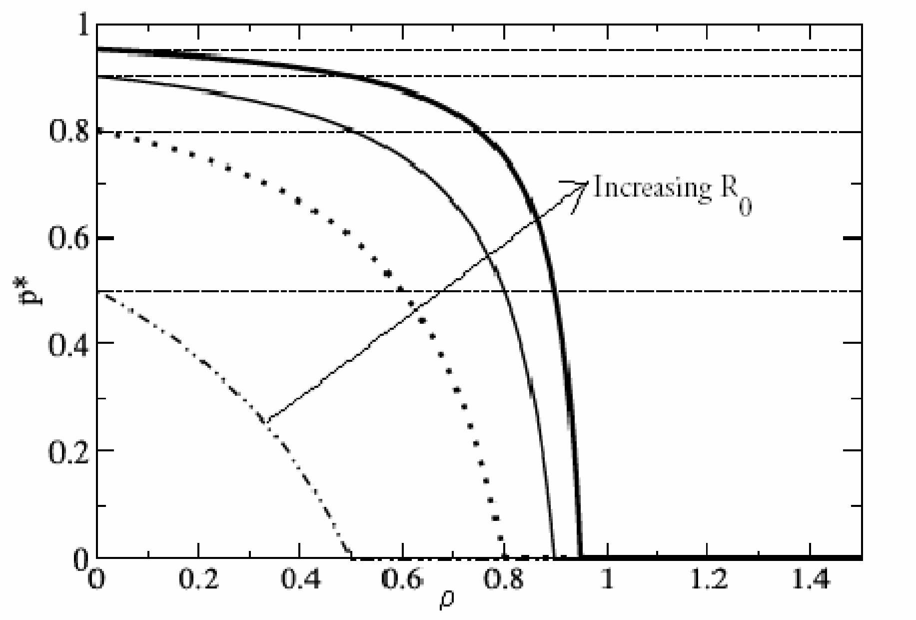

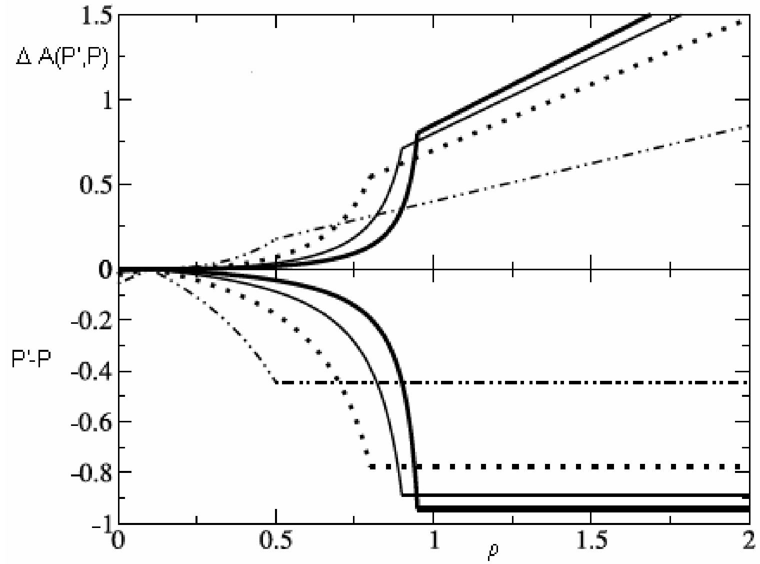

Implications

- For any *positive* perceived relative risk ($\rho>0$), equilibrium uptake falls below the critical threshold and disease will become endemic unless there are additional compulsions or incentives to vaccinate.
- If vaccination is seen as *riskier* than infection ($\rho>1$) no farmers will vaccinate in equilibrium. The minimal perceived risk above which there will be no vaccination is $1-1/R_0$.
- This abstracts from heterogeneity, impact of actual course of disease and political/media responses on risk perceptions, risk aversion, etc..
- During crises, perceived risks will rise; increased risk of infection "morbidity" have similar effects; if they cross the π_0 threshold, the impacts can be

A second model: interactive risk

- This model is based on the notion that precautions have spill over effects, which affect incentives to take care.
- Results depend on the direction of externalities (does A’s precaution increase or decrease B’s risk), the effectiveness of A’s precaution for A’s risk and the ‘aggregation technology’
- A player's risk depends (in this simple model) on his own precaution and a function of everyone else’s; positive spillovers may be ‘best effort’ (max), ‘weakest link’ (min) or anything in between.
- The analysis connects two strands in the literature
  - ‘Tipping equilibrium’ - if failure to take precautions reduces others’ incentives, safety may collapse; if taking precautions increases others’ incentives, high-security cascade may result. Allows ‘leadership’
  - Supermodularity (strategic complements) and submodularity (strategic substitutes) – affects equilibrium existence, uniqueness, optimality

A classification scheme and summary analysis

- **Case I:** Partial effectiveness, negative externalities – A’s precautions reduce everyone’s risk. The reduction is not complete, so A knows that others’ free-riding is costly to him.
  - Single or multiple (homogeneous) equilibria with tipping
  - One equilibrium dominates (high-precaution?), unique equilibrium may be optimal (e.g. if cost so low that each would want to take care even if no-one else did), but may not be (e.g. if costs so high that no-one wants to take care alone)
  - Number taking precautions $\leq$ socially optimal number
- **Case II:** Complete effectiveness, negative externalities – A’s precaution completely immunises him (and gives others some benefit).
  - Typically unique equilibrium (no tipping), but incentive to take care falls as others do (or follow suit)
  - Either full or no-precaution equilibrium could be efficient, but no guarantee
- **Case III:** Positive externalities – A’s investment increases others’ return and ‘crowds out’ their investment
  - Free-riding prevents multiple equilibrium
- Key is whether A’s precaution encourages or discourages others – and reciprocal impact on A

A more careful analysis – 2x2 case

- Game played by agents choosing one of two strategies. Payoff depends on individual, aggregate choice.
- Simple case is each agent playing ‘against’ others to whom it is linked – payoff is average based on

	No precaution	*Precaution*
No precaution	A, A	B, C
Precaution	C, B	D, D

- Payoff externalities:

Externality	*Impact of other's precaution on:*	
	Unprotected	*Protected*
α	Good	Good
β	Good	Bad
γ	Bad	Good
δ	Bad	Bad

- Substitutes if C-A < D-B
- Complements if C-A > D-B
- Precaution is *risk dominant* if A+B<C+D; no-precaution is risk-dominant if A+B>C+D
- Equilibrium regimes:

Equilibrium	*Description*
I	Unique no precaution
II	Pure partial compliance
III	Unique full precaution
IV	2 uniform conventions

Conventions – the 'local evolution' model

- Each farm is 'near' others as described by a *graph* Γ – a set of epidemiologically linked pairs (ij)
- Farm i's neighbourhood is $N_i(\Gamma) = \{j: ij \in \Gamma\}$
- i is chosen at random to rethink its behaviour: it chooses
  - A best reply to strategies of $N_i(\Gamma)$ with probability $1-\varepsilon > 0$
  - A 'mistake' with probability ε
- The resulting Markov process converges almost surely
  - To a risk-dominant equilibrium if there are two strategies per farm and all farms are linked to all other farms
  - To a generalised stable strategy if there are more than 2 strategies
  - To a (possibly) diverse allocation if the network has e.g. clusters
- Dynamics show tipping, cascades and (temporary) cycles

A classification of 2x2 case

Best	→		*Worst*	*Equil*	*Pareto*	*Risk Dom.*	*Payoff*	*Welfare*
A	B	C	D	I	1	Y	α	Y
A	B	D	C	I	1	Y	γ	Y
A	C	B	D	I	1	Y	α	Y
A	C	D	B	IV	1:2	?	α	1:2
A	D	B	C	IV	1:2	Y	γ	1:2
A	D	C	B	IV	1:2	?	γ	1:2
B	A	C	D	I	1	Y	β	?
B	A	D	C	I	1	Y	δ	?
B	C	A	D	II	2	na	β	Y
B	C	D	A	II	2	na	β	Y
B	D	A	C	I	0	Y	δ	N
B	D	C	A	II	2	na	δ	?
C	A	B	D	II	2	na	α	?
C	A	D	B	III	0	N	α	N
C	B	A	D	II	2	na	β	Y
C	B	D	A	II	2	na	β	Y
C	D	A	B	III	1	Y	α	?
C	D	B	A	III	1	Y	β	?
D	A	B	C	IV	1:2	?	γ	1:2
D	A	C	B	IV	1:2	Y	γ	1:2
D	B	A	C	IV	1:2	?	δ	1:2
D	B	C	A	III	1	Y	δ	Y
D	C	A	B	III	1	Y	γ	Y
D	C	B	A	III	1	Y	δ	Y

A more general model

- N interdependent agents (i)
  - p_i – risk faced by agent i
  - L_i – loss incurred if risk 'fires'
  - c_i – cost of precaution (prevents direct loss)
  - X_i – strategy (N, P (precaution))
  - $I_i(\{K\},X_i)$ – expected indirect cost to i when {K} choose P and i chooses X_i
  - Only direct losses to i affect others so P protects others perfectly
- Expected payoffs to i's choice:
  - P: $c_i + I_i(\{K\},P)$
  - N: $p_iL_i + (1-\alpha p_i)I_i(\{K\},N)$ – α is the non-additivity of harm, running from $\alpha = 0$ (suffer both direct and indirect damage) to $\alpha = 1$ (suffer either direct or indirect damage – only go bankrupt once☺)
  - Indifferent if $c_i = C^*(\{K\}) = p_iL_i+(1-\alpha p_i)I_i(\{K\},N) - I_i(\{K\},P)$ (take precaution if cost lower than $C^*(\{K\})$)
- Different situations
  - Case I: $I_i(\{K\},P) = I_i(\{K\},N) = I_i(\{K\})$ and $\alpha = 1$ so $C^*(\{K\}) = p_i[L_i - I_i(\{K\})]$.
    - I_i falls as {K} gets bigger – higher I_i means lower C*
    - C* rises, and tipping is possible.
  - Case II: $I_i(\{K\},P) = 0$ and $\alpha = 1$ so $C^* = p_iL_i+(1-p_i)I_i(\{K\},N)$
    - C* falls as {K} expands (I_i raises the critical cost)
  - Case III: $I_i(\{K\},P) = I_i(\{K\},N) = I_i(\{K\})$ so $C^*(\{K\}) = p_i[I_i(\{K\}) - \text{Investment}_{i]}$
    - C* again falls as {K} expands, but for a different reason (free-ride on others' investments

Herd behaviour

- Consider a Nash equilibrium in which Xi = N, all i (no precaution). A ‘critical mass’ is a coalition {K} such that if X_i = P for all i in K then C_j*({K}) $\geq$ c_j for all j not in K.
- [skipped for brevity – results on existence, characterisation of smallest minimal critical mass coalition]

A Case I example

- Let r_{ij} be the risk that infection from i transfers to j (r_{ii} is the direct risk at farm i) with (common) loss L
- (P,P) is Pareto optimal in an area that strictly includes the shaded region (so it is optimal whenever it is an equilibrium)
- In the central area, tipping is possible
- With more than three farms, cascades are possible (following the costs)

	No precaution	*Precaution*
No precaution	$-[r_{11}+(1-r_{11})r_{21}]L$, $-[r_{22}+(1-r_{22})r_{12}]L$	$-r_{11}L$, $-c_2-r_{12}L$
Precaution	$-c_1-r_{21}L$, $-r_{22}L$	$-c_1$, $-c_2$

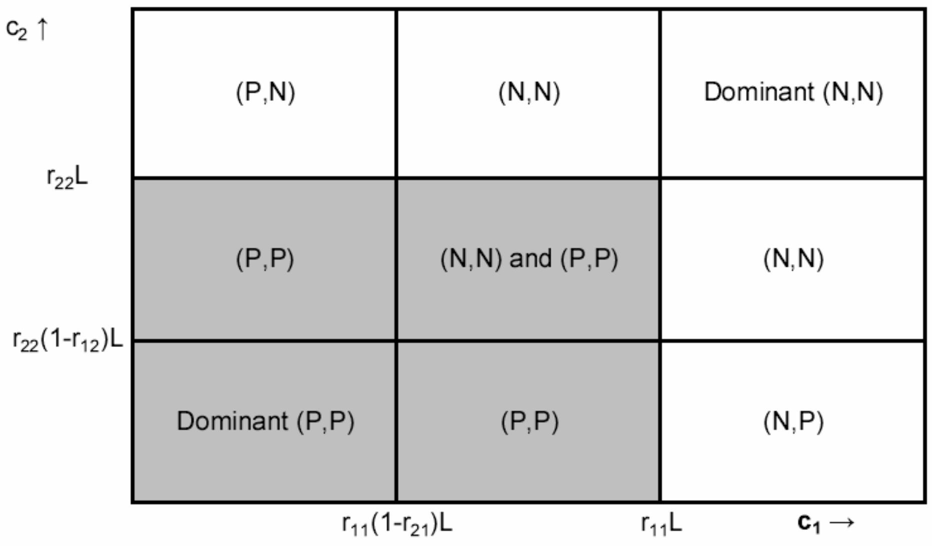

Future directions

- Coevolution of structure and behaviour
- Path-dependence
- Degrees of ‘public good’-ness (between the full group and binary network models)
- Etc.