Accounting for self-interest in the public-health management of infectious diseases (the epidemiological game theory menagerie)

Timothy Reluga, Department of Mathematics, PSU

August 5, 2009
Makerere University, Kampala, Uganda

Policy Resistance[Sterman, 2006]

policy resistance – the tendency for interventions to be defeated by the system's response to the intervention itself. Examples of Policy Resistance

- Low tar and nicotine cigarettes increase intake of carcinogens because smokers compensate by smoking more cigarettes.
- Antibiotics have stimulated the evolution of drug-resistant pathogens.
- HAART treatment has dramatically reduced mortality among those living with HIV, but has increased risky behaviors, including unprotected sex and substance abuse, among youth and other groups, causing a rebound in incidence.

Failures of SIR Theory

None of our standard epidemic models can explain problems of policy resistance.

So, why is our theory incomplete? Because our theory treats a biological system like a physics system. Atoms and molecules, for the sake of many problems, are points with a relatively limited state-space.

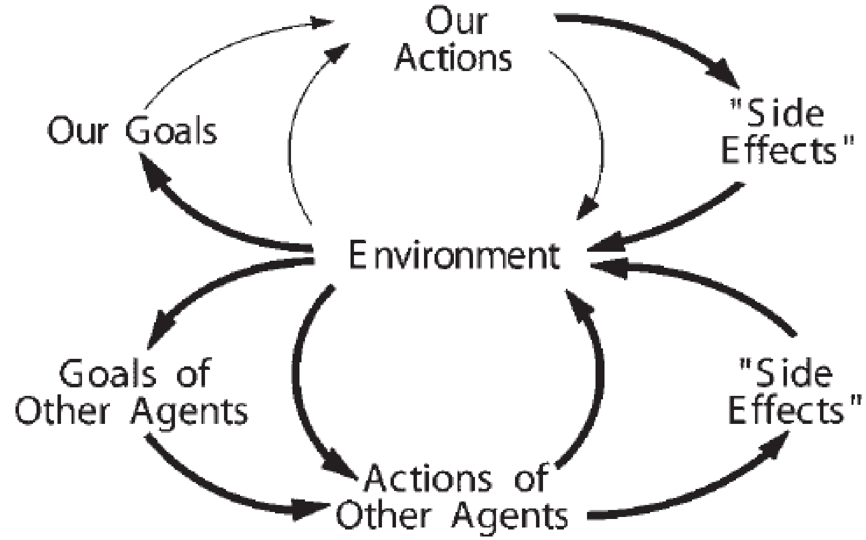

But people have complicated internal state-spaces such that their behaviors are functions of their environment in more complicated ways.

The new problem - theories that describe people with free-will

Challenges - we need a bottom-up approach that accounts for the behaviors of people and institutions as independent but interconnected actors.

Game theory is the study of interacting optimizing agents. Classical game theory has focused on ideal behaviors:

- Independent actors
- Perfect information
- Unlimited computing power
- Logical objectives

Classic game theory provides a baseline.

A model of public and private investments in health

$$\dot{S} = -\sigma(\overline{c}_s, c_t)\beta I S + \gamma I,$$
$$\dot{I} = \sigma(\overline{c}_s, c_t)\beta I S - \gamma I,$$
$$1 = S + I.$$

Here, individuals are investing $\overline{c}_s$ per day in their own prevention, while paying c_t in taxes per day, compared to a baseline income of u dollars per day.

If we let $p_S(t)$ be the probability that an individual is in the susceptible state S at time t and $p_I(t)$ be the probability that and individual is infected at time t, then individual's futures are governed by a Markov process...

$$\dot{p}_S = -\sigma(c_s, c_t)\beta I p_S + \gamma p_I,$$
$$\dot{p}_I = \sigma(c_s, c_t)\beta I p_S - \gamma p_I,$$

Note that $c_s \neq \overline{c}_s$ generally.

Outline of the mathematical approach

- Predict the population dynamics based on people's "average" behavior in the large-population limit.
- Use Markov decision process theory or dynamical systems theory to calculate the value of a particular strategy/behavior in a given environment created by the average behavior.
- Study the payoff. Try to identify game "solutions" with properties like Nash, ESS, invasion potential, convergently stable:
  - If everybody is doing the same thing, do I have any reason to be different?
  - Am I doing well, no matter what other people do?
- Determine the best reachable policy.

Determining future values for risks and costs

Suppose $u_i(t)$ is the expected future value of being in state i at time t, while $Q_{ji}dt$ is the probability of changing state for i to j. Then

$$u_i(t-dt) = v_i dt + (1 - hdt)\left(\sum_{j\neq i} u_j(t) Q_{ji} dt + (1 - \sum_j Q_{ji} dt) u_i(t)\right)$$

$$\frac{u_i(t-dt) - u_i(t)}{dt} = v_i - hu_i(t) + \sum_j u_j(t) Q_{ji}$$

$$\frac{d\mathbf{u}^T}{dt} = \mathbf{u}^T(h\mathbf{I} - \mathbf{Q}) - \mathbf{v}^T$$

Taking into account the initial distribution of for an individual's state, the expected value $U = \mathbf{u}^T\mathbf{p}_0$. This is a classic result of Markov decision process theory [Howard, 1960].

Payoffs to individuals and communities

For an individual, the payoff of investing c_s is exposure reduction while everybody else is investing $\overline{c}_s$

$$U(c_s;\overline{c}_s,c_t)=\frac{u-c_t}{h}-\frac{(h+\gamma)c_s+\beta I^*\sigma\left(c_s,c_t\right)c_i}{h\left[h+\gamma+\beta I^*\sigma\left(c_s,c_t\right)\right]}, \tag{3}$$

$$I^*=1-\frac{\gamma}{\sigma(\overline{c}_s,c_t)\beta} \tag{4}$$

where h is the rate of discounting of future returns and βI^* is the stationary force of infection.

In a Platonic Republic, where the government's sole objective is to maximize the welfare of it's citizens*, public policy should be to chose the tax-rate c_t to maximize $U(\overline{c}_s;\overline{c}_s,c_t)$.

*ignoring population growth effects

So what's the payoff-surface look like?

The best payoff is for $c_t = .05$, $\overline{c}_s = .07$.

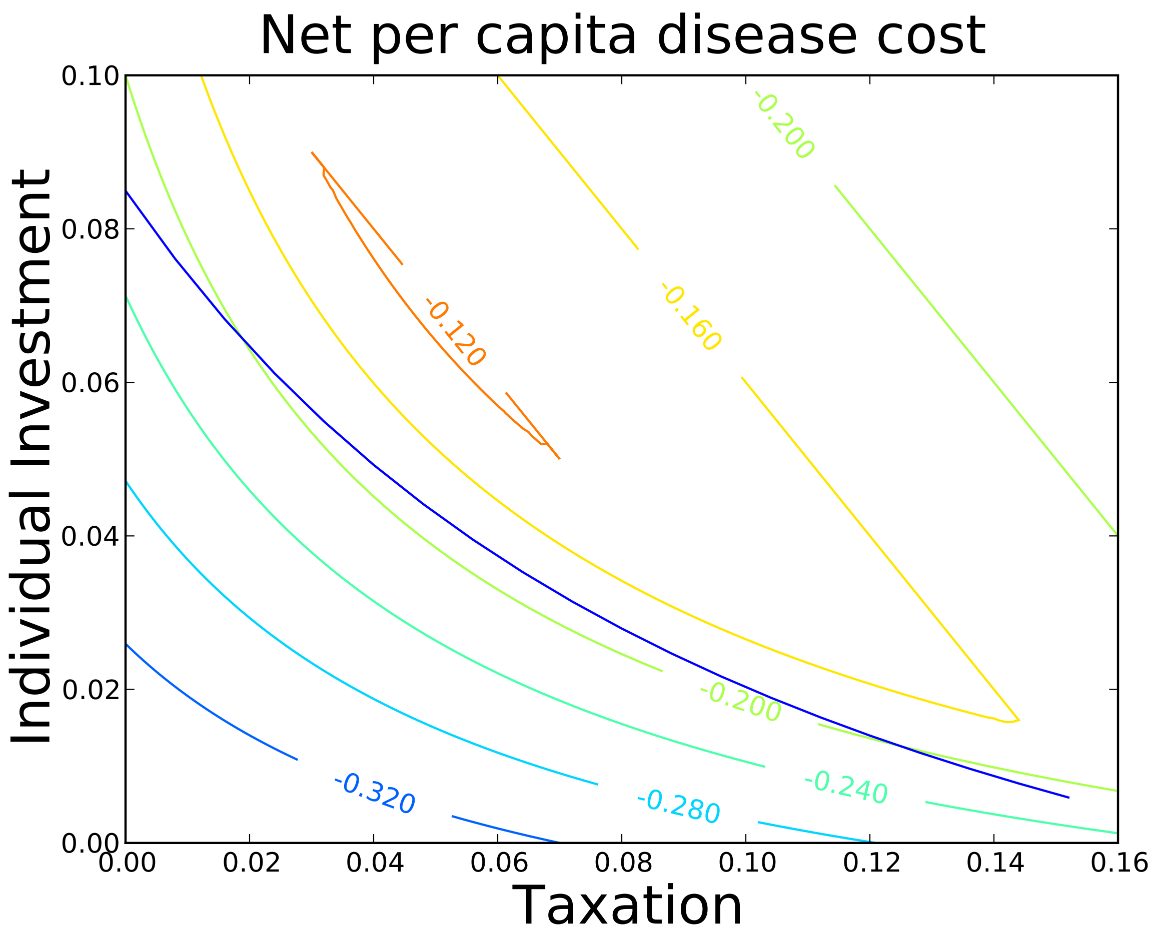

However, rather than maximizing the community's payoff, individuals should choose c_s to maximize their own payoff $U(c_s; \overline{c}_s, c_t)$ for a given environment $(\overline{c}_s, c_t)$.

Relative payoffs $U(c_s; \overline{c}_s) - U(\overline{c}_s; \overline{c}_s)$

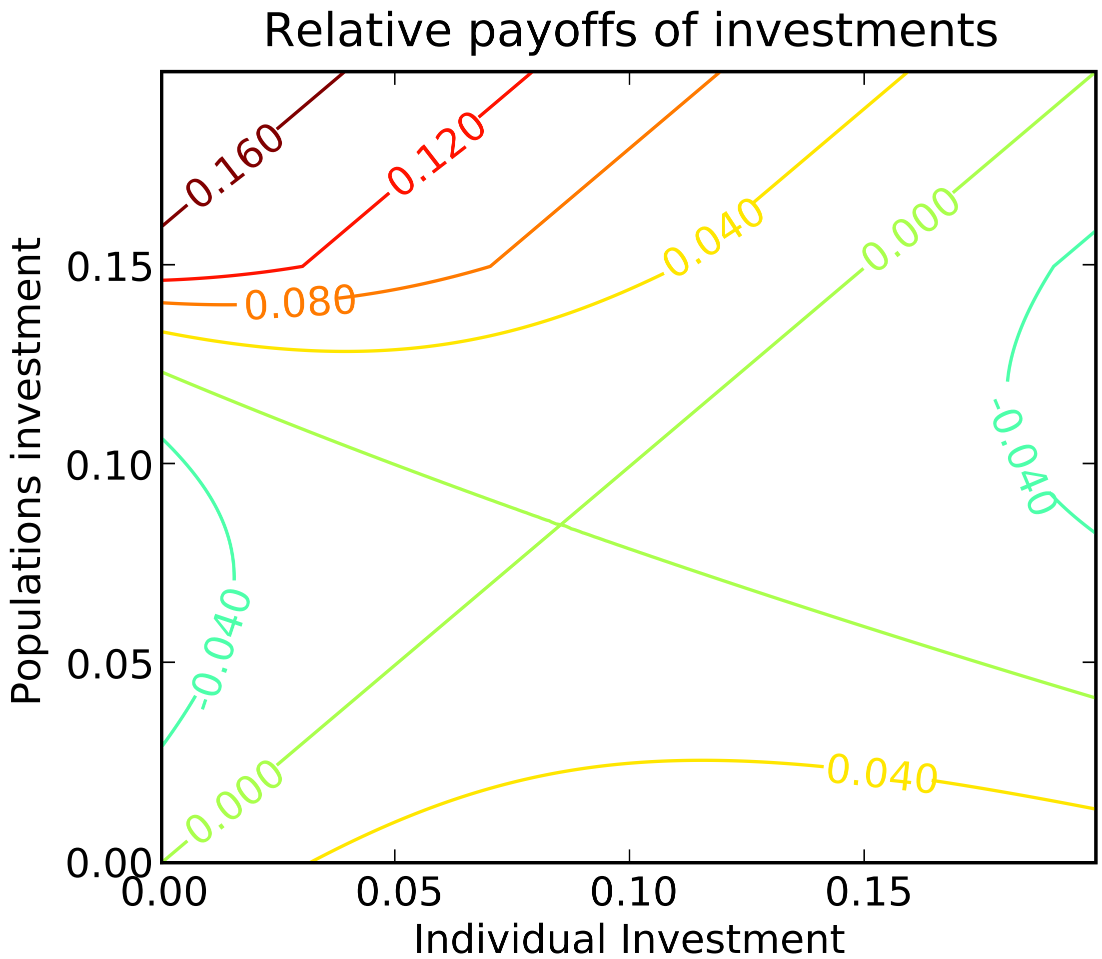

Contour plot of relative payoffs showing how individuals can improve their payoffs compared to the average.

Solution concepts

Calculating the you future payoff isn't enough, though, because your future isn't predetermined. Your future depends on both your decisions and the decisions of everybody else.

The question, then is how does your payoff depend on your choices c_s and the choices of others $\overline{c}_s$. What is the shape of

$$U(c_s, \overline{c}_s)? \tag{5}$$

There may be decisions c_s^*

- ...that are best for you, no matter what everybody else does.
- ...where you never do worse than anybody else.
- ...such that nobody would want to change their mind if given a second chance.

Partial orderings of the strategy space

One way to formally describe the properties of the various solution concepts that is valid for arbitrary strategy-spaces in a game describing choice in a single population is in terms of two partial orderings

The Nash ordering based on what strategies beat resident's

$$N_{>}(\bar{c}_s) = \{c_s : U(c_s, \bar{c}_s) > U(\bar{c}_s, \bar{c}_s)\}, \tag{6}$$

and a Smith ordering describing what resident strategies are beatable.

$$S_{>}(c_s) = \{\bar{c}_s : U(c_s, \bar{c}_s) > U(\bar{c}_s, \bar{c}_s)\}, \tag{7}$$

These conditions provide an axiomatic approach to existence of fitness landscapes, ESS's, recursively dominate strategies, and weak equilibria in local and global senses.

Nash equilibria and ESS's

A Nash equilibrium in a population game is a strategy c_s^*, such that, when everybody plays c_s^*, nobody wants to change their mind.

$$N_>(c_s^*) = \emptyset \tag{8}$$

A Nash equilibrium has global invasion potential if it never does worse than the resident strategy.

$$S_\geq(c_s^*) = \text{everything.} \tag{9}$$

Such a strategy is usually called an Evolutionary Stable Strategy (ESS)

Theorem 1. *If $\sigma(\bar{c}_s, c_t)$ is decreasing and convex in $\bar{c}_s$, then there is a unique Nash equilibrium $c_s^*(c_t)$ for every taxation rate $c_t \geq 0$.*

Policy resistance - illusions of opportunity

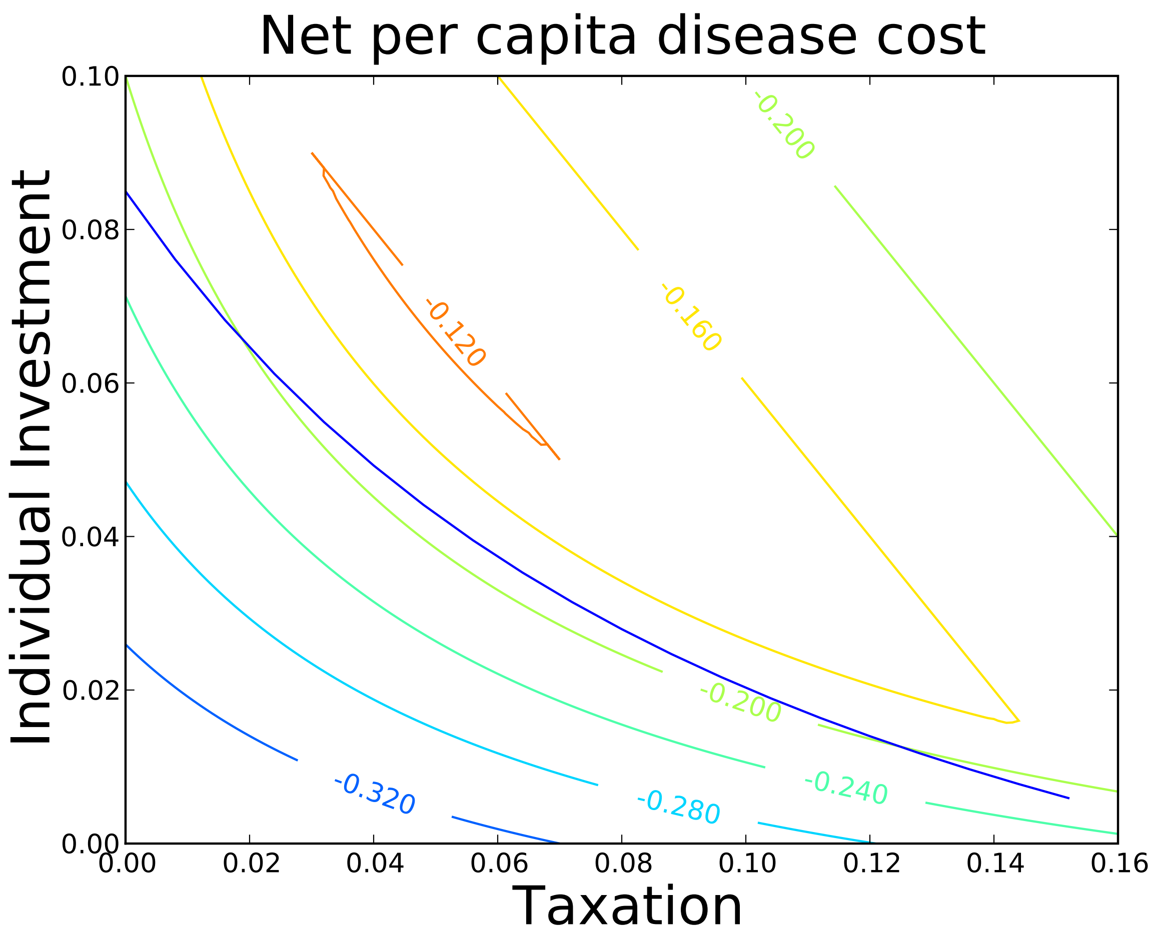

Theorem 2. *If the effects of government and individual interventions are independent, such that* $\sigma(c_s, c_t) = \sigma_s(c_s)\sigma_t(c_t)$*, then increased taxation decreases equilibrium individual investment in self-protection (* $\mathrm{d}c_s^*/\mathrm{d}c_t \leq 0$*).*

So what does it mean?

- Results depend greatly on how individual and community investment interact to reduce risk ($\sigma(c_s, c_t)$'s shape). If taxation is inefficient, it does more harm than good. But scale effects on capital potentially allow taxation to reduce risk much more than is possible by individuals alone.

- The ideal outcome depends on perspective. What's ideal for the community is not necessarily ideal for individuals.

- Large, sudden changes in policy can have unforeseen consequences if not designed to account for behavioral feedbacks.

Simple models can provide very satisfying just-so stories, but how much can we trust them?

The fundamental ingredient in these analyses are the calculation of Nash equilibria, but I don't have very much intuition yet for what the Nash equilibria look like in different epidemiological models yet.

- What if individuals have strategies that can change depending on the age of the individual?
- What if different subpopulations have different costs and risks? How does that effect the analysis?

Allowing Age-dependent behaviors in an SI model

$$\dot{S}_y = r - \overline{\sigma}_y \lambda S_y - f S_y - m_y S_y,$$
$$\dot{I}_y = \overline{\sigma}_y \lambda S_y - (\gamma_y + f) I_y,$$
$$\dot{S}_o = f S_y - \overline{\sigma}_o \lambda S_o - m_o S_o,$$
$$\dot{I}_o = f I_y + \overline{\sigma}_o \lambda S_o - \gamma_o I_o.$$

with $\lambda = \beta_y I_y + \beta_o I_o$.

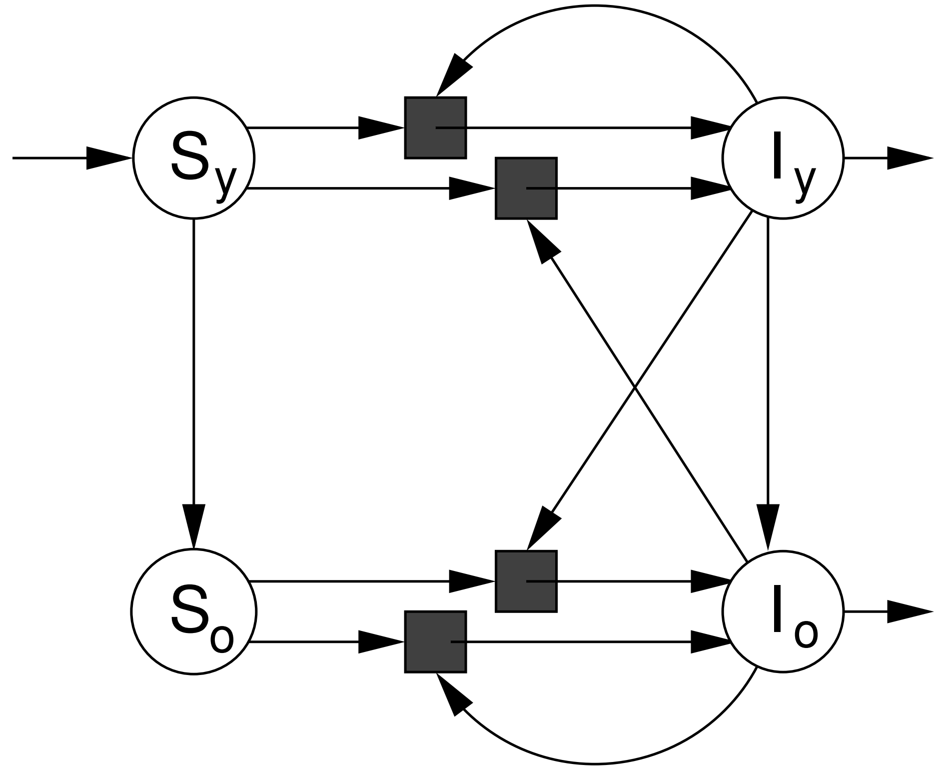

$$\boldsymbol{Q} = \begin{bmatrix} -\sigma_y \lambda - f - m_y & 0 & 0 & 0 \\ \sigma_y \lambda & -\gamma_y - f & 0 & 0 \\ f & 0 & -\sigma_o \lambda - m_o & 0 \\ 0 & f & \sigma_o \lambda & -\gamma_o \end{bmatrix}, \quad \boldsymbol{p}(0) = \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \end{bmatrix}.$$

Costs and benefits (felicities)

To describe differences in decisions, we associate values with each state. We assume young and old states have values u_y and u_o per day, that infection costs c_{yi} and c_{oi}, and that total self-protection costs c_v per day.

The felicities of each state then are represented by the vector

$$\boldsymbol{v} = \begin{bmatrix} u_y - (1-\sigma_y)c_v \\ u_y - c_{yi} \\ u_o - (1-\sigma_o)c_v \\ u_o - c_{oi} \end{bmatrix}, \tag{10}$$

The Utility Function

$$U(\sigma, \lambda^*) = \frac{u_{yi}\,\lambda^*\,\sigma_y + u_y - c_v + \sigma_y c_v}{\lambda^*\,\sigma_y + f + m_y + h} + \frac{f\left[u_{oi}\lambda^*\sigma_o + u_o - (1-\sigma_o)\,c_v\right]}{(\lambda^*\,\sigma_y + f + m_y)\,(\lambda^*\,\sigma_o + m_o)}$$

where λ^* is the equilibrium force of infection depending on the population's aggregate behavior $\overline{\sigma}$,

$$u_{yi} = \frac{1}{\gamma_y + f}\,(u_y - c_{yi} + f u_{oi}) \tag{11}$$

and

$$u_{oi} = \frac{1}{\gamma_o}\,(u_o - c_{oi})\,. \tag{12}$$

For simplicity, we assume discounting is built into the removal rates m_y, m_o, γ_y, and γ_o.

Best Response

The best response correspondence for stationary force of infection λ^* is

$$\sigma^B = \text{argmax}_\sigma U(\sigma, \lambda^*) =$$

$$H\left(\frac{u_y + \lambda^* u_{yi} + f E_o}{m_y + f + \lambda^*} - \frac{u_y - c_v + f E_o}{m_y + f}\right) \times H\left(\frac{u_o + \lambda^* u_{oi}}{m_o + \lambda^*} - \frac{u_o - c_v}{m_o}\right)$$

where

$$E_o = \max_{\sigma_o \in [0,1]} \frac{u_{oi} \lambda^* \sigma_o + u_o - (1 - \sigma_o)\, c_v}{(\lambda^*\, \sigma_o + m_o)}$$

and the Heaviside correspondence

$$H(x) = \begin{cases} 0 & \text{if } x < 0, \\ [0, 1] & \text{if } x = 0, \\ 1 & \text{if } x > 0. \end{cases} \tag{13}$$

Determining Nash Equilibria

The Nash Equilibria are strategies which are best-responses to themselves.

$$\sigma^B(\lambda^*(\sigma)) \ni \sigma$$

which can be rewritten in the scalar necessary condition

$$\lambda^*(\sigma^B(\lambda)) \ni \lambda$$

Theorem 3. *If $\mathcal{R}_0 > 1$ and the stationary force of infection λ^* is a strictly increasing function of both $\overline{\sigma}_y$ and $\overline{\sigma}_o$, $c_v > 0$, then there is a unique Nash equilibrium strategy σ^* as long as*

$$\frac{u_o - c_v}{m_o} \neq \frac{u_y - c_v - (m_y + f)\,(u_{yi} - u_{oi})}{m_y},$$

Force of infection λ when $\beta_y \approx \beta_o$

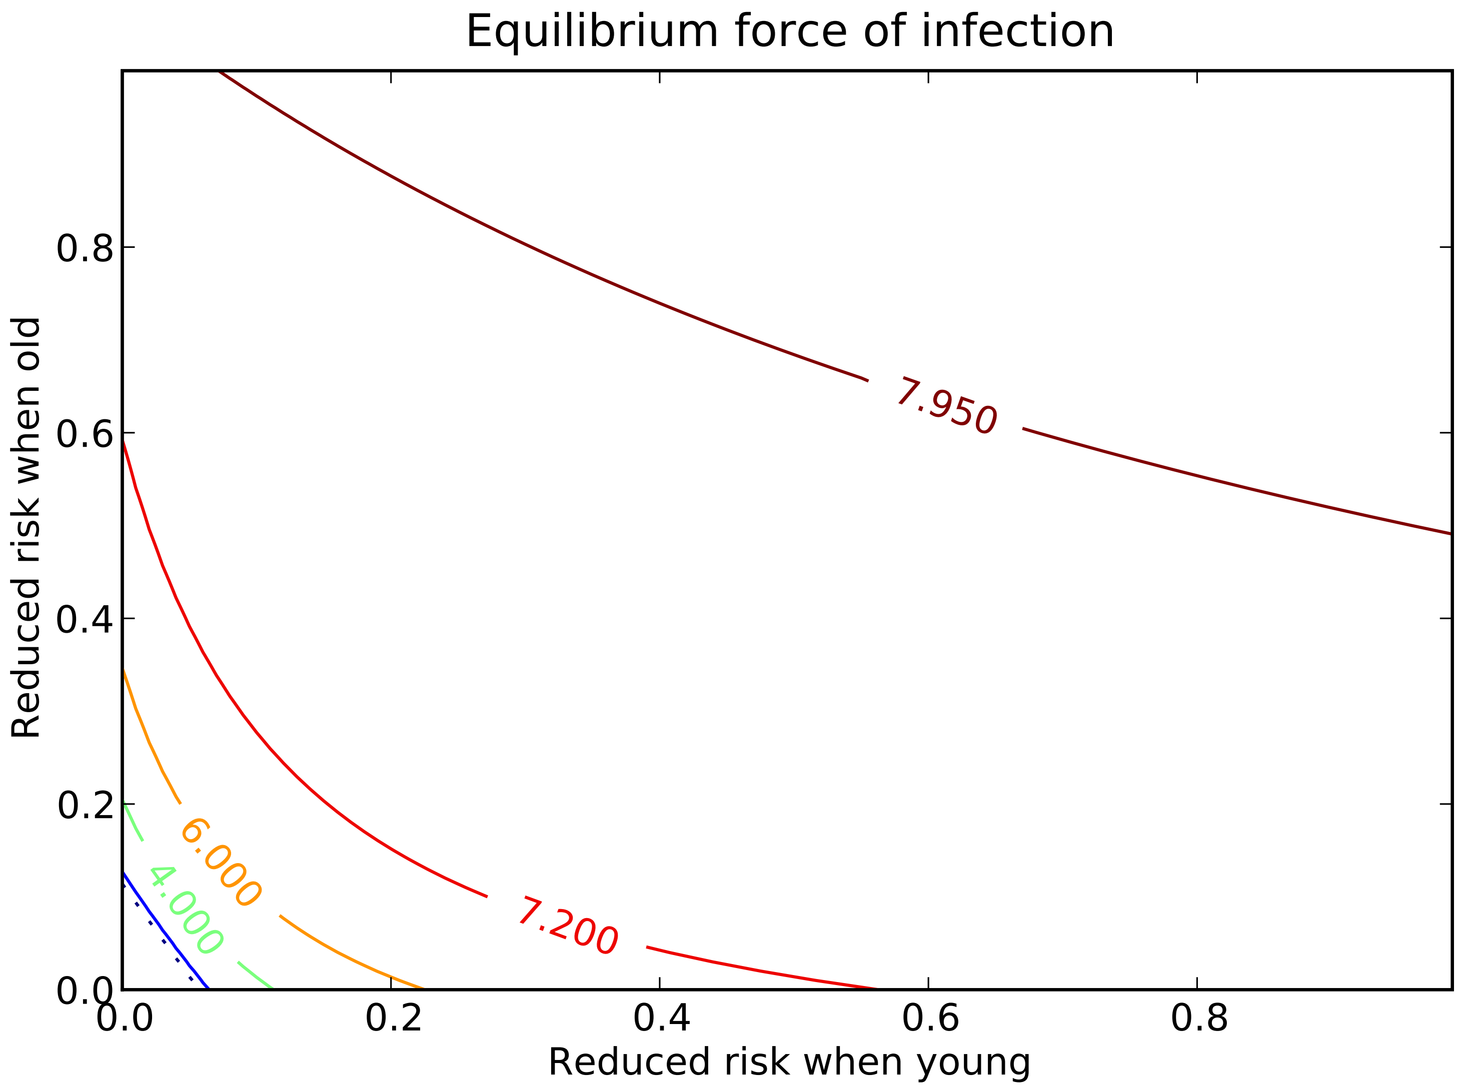

Solution to the fixed point equation

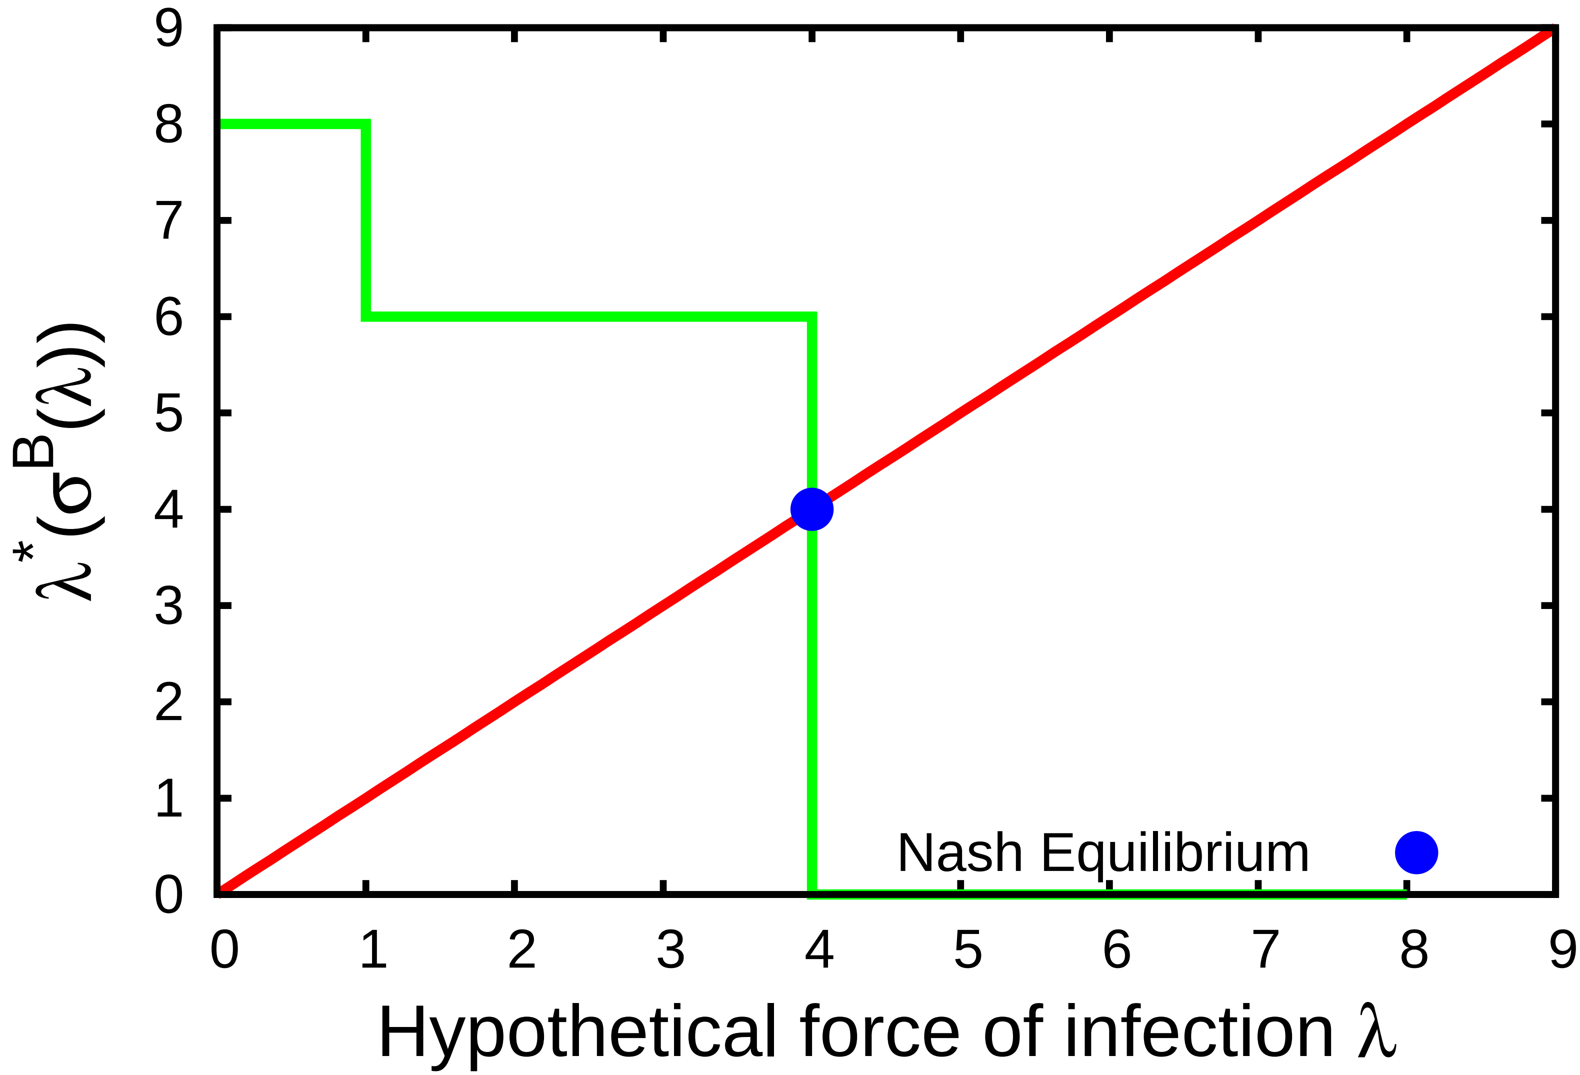

Non-Monotonicity of Stationary Solutions ($\beta_y \approx 0$)

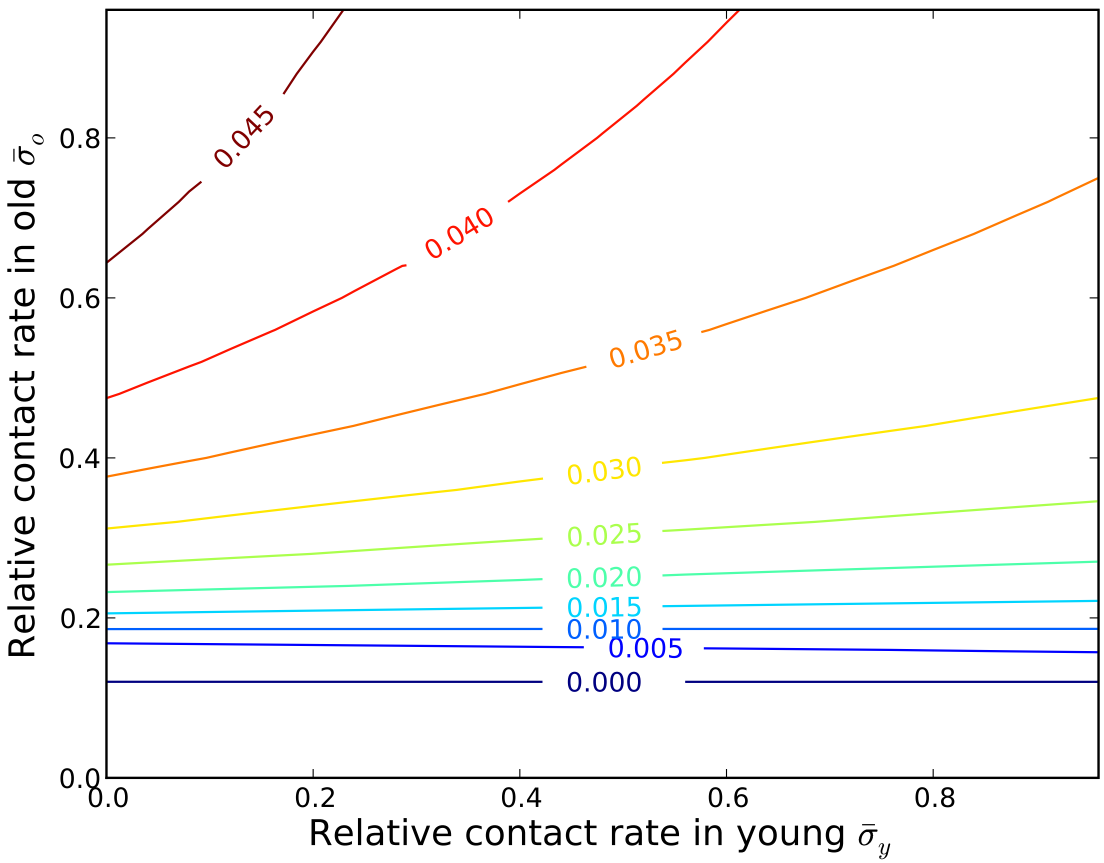

However, the stationary force of infection is not always monotone in the self-protection fractions. This can potentially create 3 Nash equilibria, of which 2 are "stable".

Possible Solution to the fixed point equation

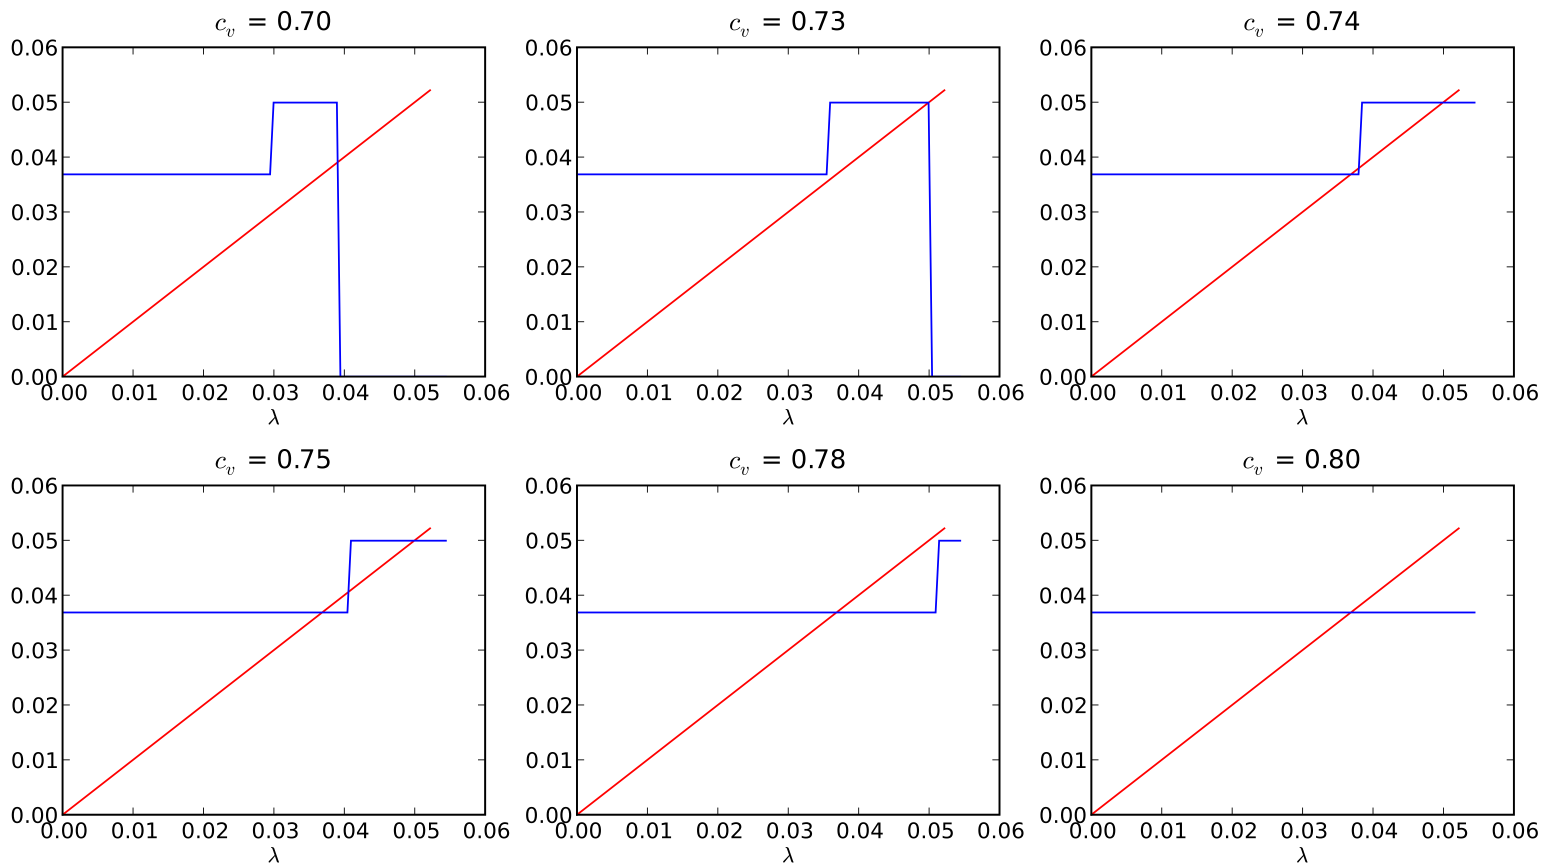

When there are 3 equilibria, neither extreme can invade the other

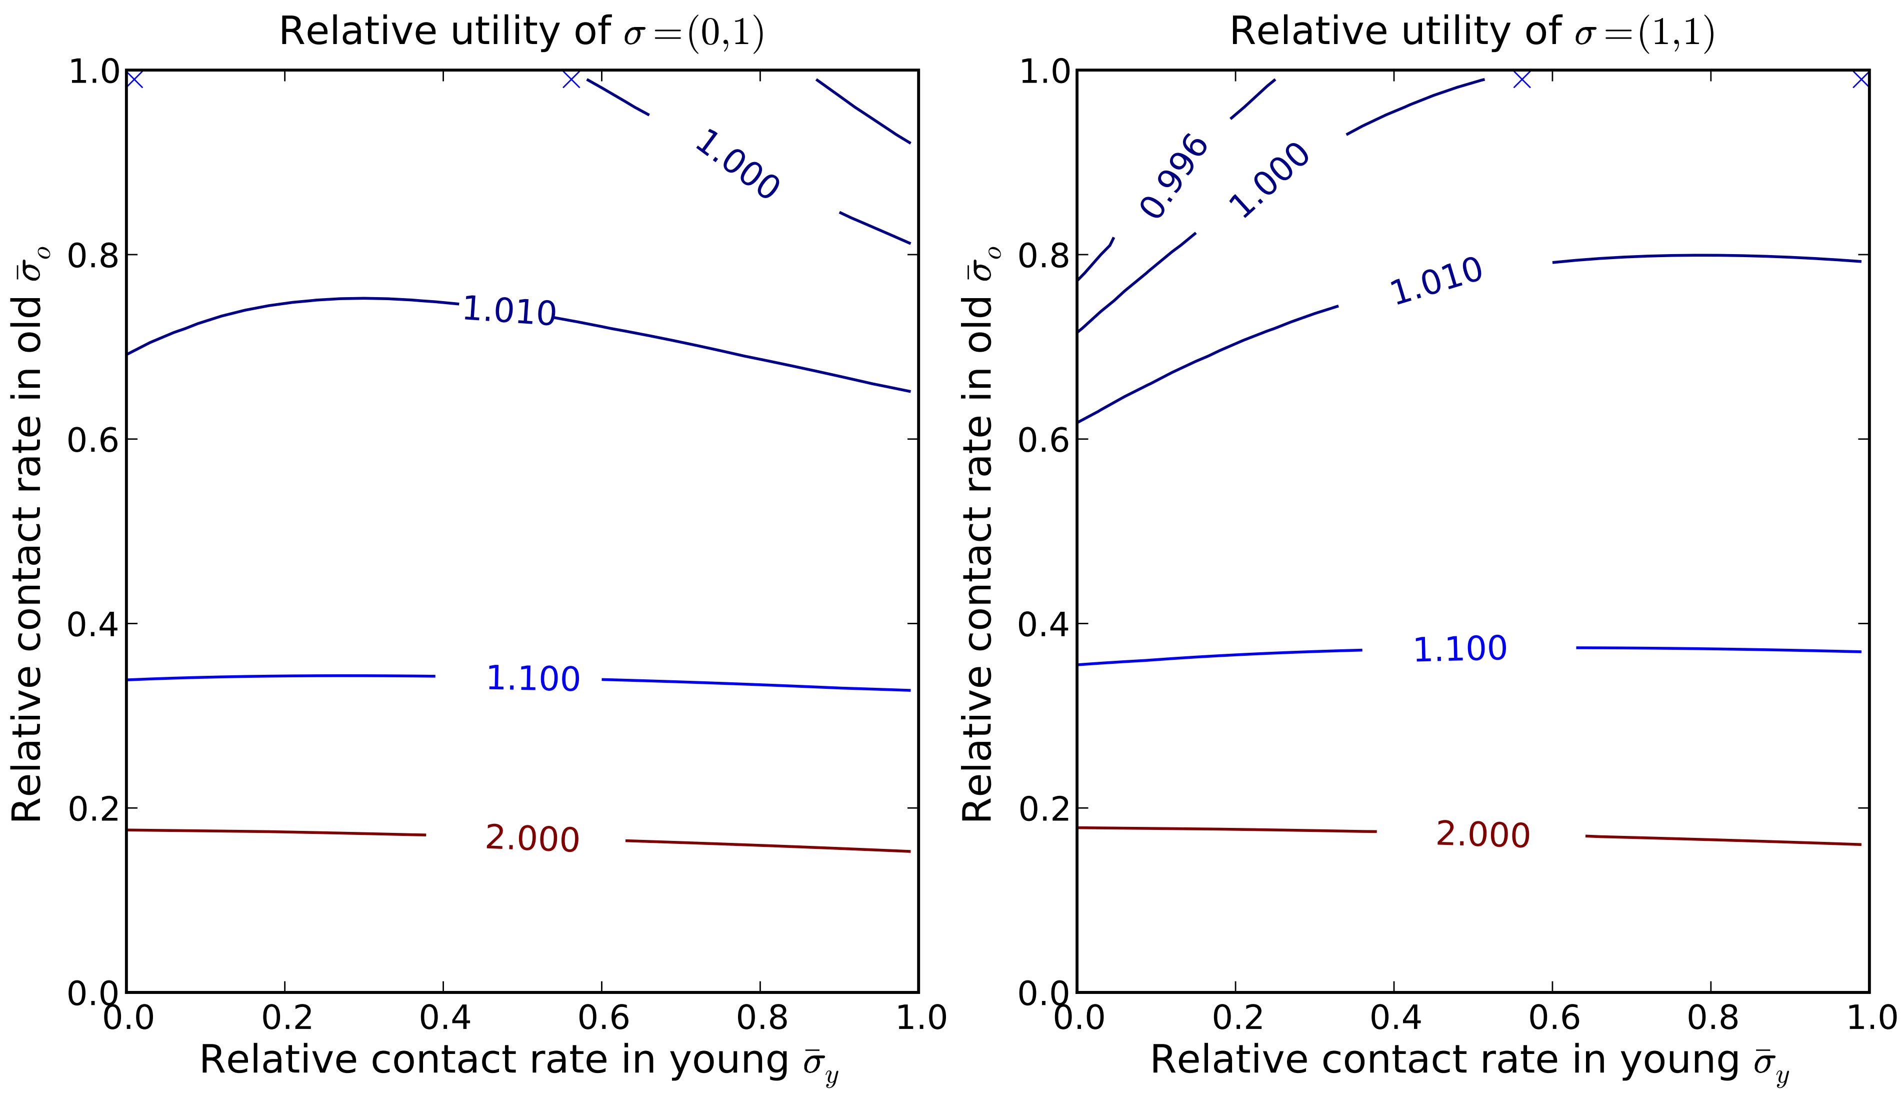

Interacting subpopulations with different interests

Rather than having a single homogeneous population, suppose we consider the problem of two seperate populations (for instance, men and women)

$$\dot{S}_1 = \gamma I_1 - \lambda_1 \overline{\sigma}_1 S_1, \tag{14a}$$

$$\dot{I}_1 = \lambda_1 \overline{\sigma}_1 S_1 - \gamma I_1, \tag{14b}$$

$$\dot{S}_2 = \gamma I_2 - \lambda_2 \overline{\sigma}_2 S_2, \tag{14c}$$

$$\dot{I}_2 = \lambda_2 \overline{\sigma}_2 S_2 - \gamma I_2, \tag{14d}$$

$$\lambda_1 = \beta_{11} I_1 + \beta_{21} I_2, \tag{14e}$$

$$\lambda_2 = \beta_{12} I_1 + \beta_{22} I_2 \tag{14f}$$

with $S_1 + I_1 = 1$ and $S_2 + I_2 = 1$ [Hethcote and Yorke, 1984, Beretta and Capasso, 1986].

The game equilibria depend on the mixing pattern and the differences in costs †

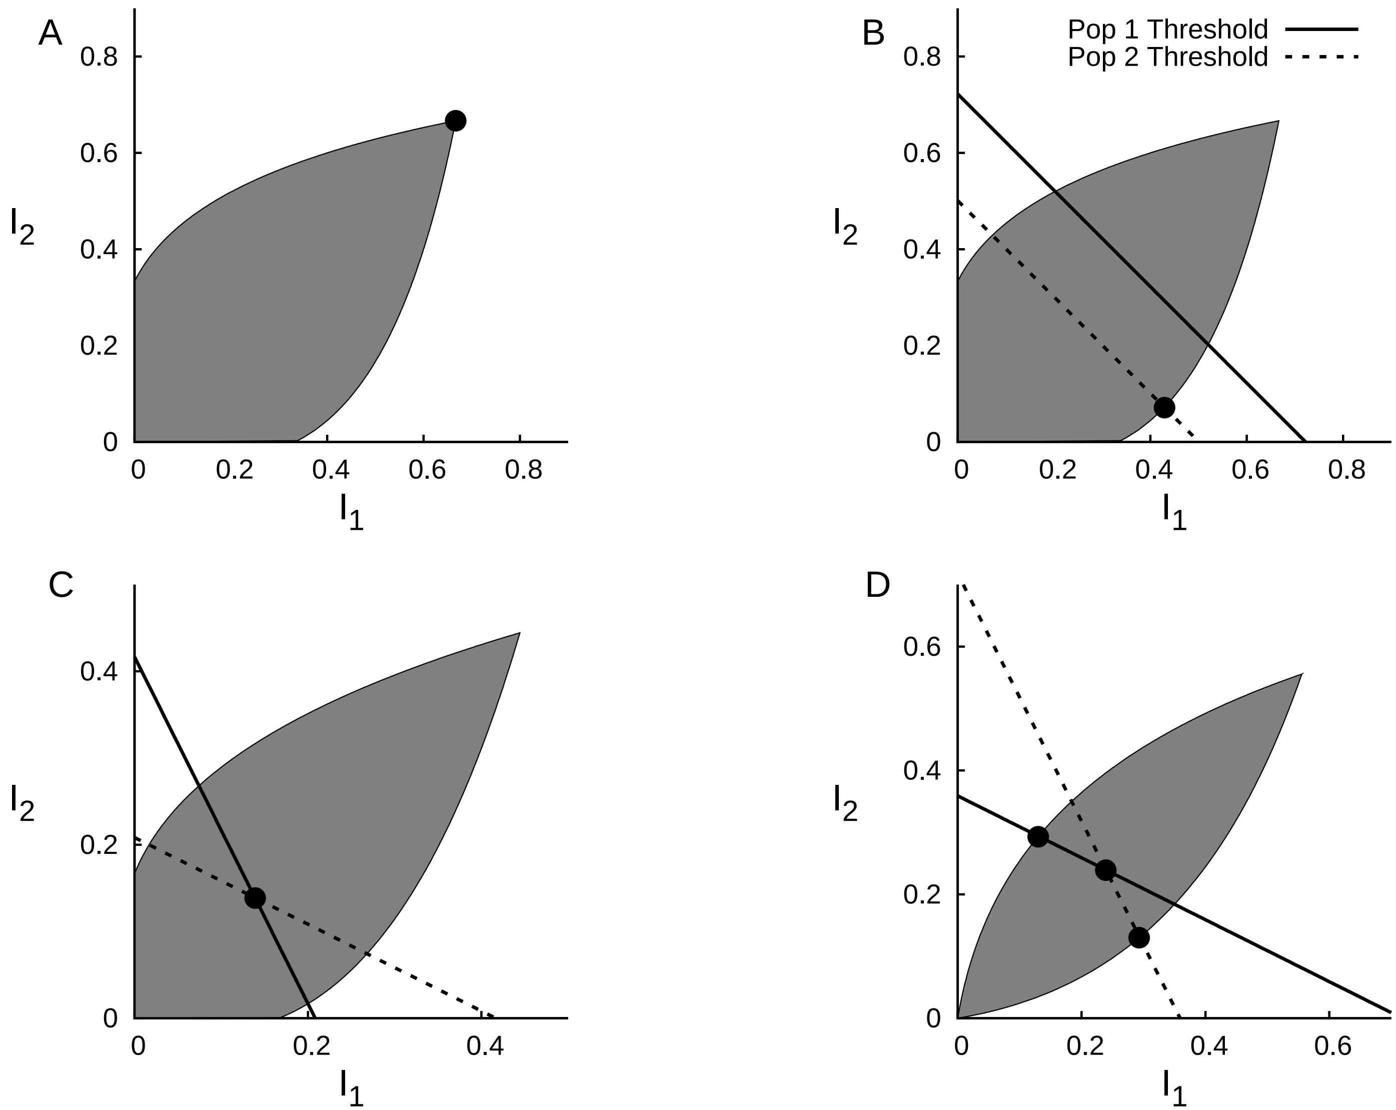

†The analysis of this game parallels similarities to the analysis of Lotka's competition model.

Summary

- Games of behavioral response to infectious diseases can be constructed by extending population-scale models to include Markov process descriptions of individual's lives.
- These games allow us to account for "policy resistance" when we make decisions.
- Infectious disease games can admit multiple solutions, each of which has it's own implications for community structure.
- Non-monotone effects can lead to non-unique game equilibria. (imperfect efficacy, disassortative mixing, age-dependent virulence)
- Policy models can incorporate game theory into their optimization objectives to anticipate policy resistance.

Acknowledgements

- JOBS! Penn state is hiring junior and senior researchers in infectious disease dynamics. I'm looking for a postdoc and graduate students for theoretical AND APPLIED research.

- Collaborators
  - Darla Lindberg (PSU)
  - Rachel Smith (PSU)

- Funding
  - Bill and Melinda Gates Foundation
  - National Institutes of Health

References

John D. Sterman. Learning from evidence in a complex world. *American Journal of Public Health*, 96:505–514, 2006.

Ronald A. Howard. *Dynamic Programming and Markov Processes*. MIT Press, Cambridge, MA, 1960.

H. W. Hethcote and I. A. Yorke. *Gonorrhea Transmission and Control*, volume 56 of *Lecture Notes in Biomathematics*. Springer-Verlag, Heidelberg, 1984.

E. Beretta and V. Capasso. On the general structure of epidemic systems - global asymptotic stability. *Computers and Mathematics with Applications-part A*, 12(6):677–694, 1986.