Interplay of Strategic Decision Making and Spread of Epidemics

Prof. Ashish R. Hota

Indian Institute of Technology Kharagpur

Research Seminar, Decision, Games and Evolution Program, ICTS

March 9, 2025

Introduction

- Infectious diseases spread through society by exploiting interaction among individuals.
- In the absence of effective treatments, epidemic containment relies on limiting social interaction.
- Centralized approaches for epidemic containment are often not scalable and adequate.
- Individuals take decisions such as adopting protection or vaccination, in a selfish and decentralized manner, which shapes epidemic evolution and vice versa.
- Our research has examined the **coupled evolution of epidemics and human decision-making** in the frameworks of game theory and control theory.

Part 1: Coupled Evolution of Epidemic and Protection Adoption Behavior

- Wearing a mask reduces infection rate, but may cause personal discomfort
- As disease spread decreases, individuals may choose not to wear mask
- How does such behavior influence the spread of disease?
- What is the role of the timescales of behavioral and epidemic dynamics?

IDEAS

The Texas Mask Mystery

When the governor lifted the state's mandate, liberals predicted disaster. But it never came. Why?

By Derek Thompson

Nearly a dozen states move to end masking mandates as COVID-19 infection rates fall

Since Monday, 11 states have announced changes to their statewide mask policies.

By **Arielle Mitropoulos**
February 10, 2022, 6:48 PM • 14 min read

Reference

A. Satapathi, N. K. Dhar, A. R. Hota, V. Srivastava, "Epidemic Propagation under Evolutionary Behavioral Dynamics: Stability and Bifurcation Analysis." American Control Conference (ACC), 2022.

SIS Epidemic Propagation Model

- In the Susceptible-Infected-Susceptible (SIS) epidemic model, there are two epidemic states. An individual is either healthy/susceptible or sick/infected.
- $y(t) =$ fraction of infected individuals, $1 - y(t) =$ fraction of susceptible individuals.
- Let β denote the probability of a susceptible individual becoming infected upon interacting with an infected individual.
- Let γ be the probability with which an infected individual recovers and becomes susceptible.

SIS Epidemic Propagation Model

- In the Susceptible-Infected-Susceptible (SIS) epidemic model, there are two epidemic states. An individual is either healthy/susceptible or sick/infected.
- $y(t)$ = fraction of infected individuals, $1 - y(t)$ = fraction of susceptible individuals.
- Let β denote the probability of a susceptible individual becoming infected upon interacting with an infected individual.
- Let γ be the probability with which an infected individual recovers and becomes susceptible.

SIS Infection Dynamics

$$\dot{y}(t) = \Big[(1 - y(t))\beta - \gamma\Big] y(t).$$

There are two equilibrium points: $y^*_{\text{IFE}} = 0$ and $y^*_{\text{EE}} = 1 - \frac{\gamma}{\beta}$; the latter exists when $\beta > \gamma$ and is stable when it exists.

SIS Epidemic with Behavioral Heterophily

- Each individual decides whether to adopt protection or not.
- Let $z_I(t)/z_S(t)$ denote the fraction infected/susceptible individuals who remain unprotected.

SIS Epidemic with Behavioral Heterophily

- Each individual decides whether to adopt protection or not.
- Let $z_I(t)/z_S(t)$ denote the fraction infected/susceptible individuals who remain unprotected.
- Let β_u/β_p denote the probabilities of an infected individual causing a new infection if it is unprotected/protected.
- Let protected susceptible individual be $\alpha \in (0, 1)$ times (less) likely to become infected.

SIS Epidemic with Behavioral Heterophily

- Each individual decides whether to adopt protection or not.
- Let $z_I(t)/z_S(t)$ denote the fraction infected/susceptible individuals who remain unprotected.
- Let β_u/β_p denote the probabilities of an infected individual causing a new infection if it is unprotected/protected.
- Let protected susceptible individual be $\alpha \in (0, 1)$ times (less) likely to become infected.

SIS infection dynamics with behavioral heterophily

$$\dot{y}(t) = \Big((1 - y(t))(z_S(t) + \alpha(1 - z_S(t))) \cdot (\beta_u z_I(t) + \beta_p(1 - z_I(t))) - \gamma\Big) y(t).$$

The effective infection rate $\beta_{\texttt{eff}}$ now depends on the strategies adopted by the individuals $(z_S(t), z_I(t))$.

Payoffs

- The payoff for an individual depends on its individual infection state and chosen action, as well as the social state $(y(t), z_S(t), z_I(t))$.
- For an infected individual, we define the payoffs to be

$$F_{\mathtt{IU}}(y, z_S, z_I) = -c_{\mathtt{IU}}, \qquad F_{\mathtt{IP}}(y, z_S, z_I) = -c_{\mathtt{IP}}$$

  if it remains unprotected and adopts protection, respectively. We assume $c_{\mathtt{IU}} > c_{\mathtt{IP}} \geq 0$

Payoffs

- The payoff for an individual depends on its individual infection state and chosen action, as well as the social state $(y(t), z_S(t), z_I(t))$.
- For an infected individual, we define the payoffs to be

$$F_{\mathrm{IU}}(y, z_S, z_I) = -c_{\mathrm{IU}}, \qquad F_{\mathrm{IP}}(y, z_S, z_I) = -c_{\mathrm{IP}}$$

  if it remains unprotected and adopts protection, respectively. We assume $c_{\mathrm{IU}} > c_{\mathrm{IP}} \geq 0$
- For a susceptible individual, let
  - $c_P > 0$ be cost of adopting protection
  - $L > 0$ be the cost of getting infected
  - $(\beta_u z_I + \beta_p(1 - z_I))y$ being the instantaneous probability of becoming infected for a susceptible unprotected node.
- The payoffs for a susceptible individual are now defined as

$$F_{\mathrm{SU}}(y, z_S, z_I) = -L(\beta_u z_I + \beta_p(1 - z_I))y$$
$$F_{\mathrm{SP}}(y, z_S, z_I) = -c_P - L\alpha(\beta_u z_I + \beta_p(1 - z_I))y.$$

- Recall that $\begin{bmatrix} F_{\mathrm{SU}}(y,z_S,z_I) \\ F_{\mathrm{SP}}(y,z_S,z_I) \\ F_{\mathrm{IU}}(y,z_S,z_I) \\ F_{\mathrm{IP}}(y,z_S,z_I) \end{bmatrix} = \begin{bmatrix} -L(\beta_u z_I + \beta_p(1-z_I))y \\ -c_P - L\alpha(\beta_u z_I + \beta_p(1-z_I))y \\ -c_{\mathrm{IU}} \\ -c_{\mathrm{IP}} \end{bmatrix}$

Replicator dynamics

The fraction of unprotected individuals follow the following dynamics

$$\begin{aligned} \dot{z}_S(t) &= z_S(t)\Big[F_{\mathrm{SU}} - (z_S F_{\mathrm{SU}} + (1 - z_S)F_{\mathrm{SP}})\Big] \\ &= z_S(t)(1 - z_S(t))\Big(c_P - L(1-\alpha)(\beta_u z_I(t) + \beta_p(1 - z_I(t)))y(t)\Big) \\ \dot{z}_I(t) &= z_I(t)\Big[F_{\mathrm{IU}} - (z_I F_{\mathrm{IU}} + (1 - z_I)F_{\mathrm{IP}})\Big] \\ &= z_I(t)(1 - z_I(t))(c_{\mathrm{IP}} - c_{\mathrm{IU}}) \end{aligned}$$

Coupled SIS Epidemic and Evolutionary Behavioral Dynamics

Coupled SIS epidemic and evolutionary behavioral dynamics

$$\begin{aligned}
\dot{y}(t) &= \Big((1-y(t))(z_S(t)+\alpha(1-z_S(t)))\cdot(\beta_u z_I(t)+\beta_p(1-z_I(t)))-\gamma\Big)y(t)\\
\dot{z}_S(t) &= z_S(t)(1-z_S(t))\Big(c_P - L(1-\alpha)(\beta_u z_I(t)+\beta_p(1-z_I(t)))y(t)\Big)\\
\dot{z}_I(t) &= z_I(t)(1-z_I(t))(c_{\mathrm{IP}}-c_{\mathrm{IU}}).
\end{aligned}$$

Coupled SIS Epidemic and Evolutionary Behavioral Dynamics

Coupled SIS epidemic and evolutionary behavioral dynamics

$$\dot{y}(t) = \Big((1 - y(t))(z_S(t) + \alpha(1 - z_S(t))) \cdot (\beta_u z_I(t) + \beta_p(1 - z_I(t))) - \gamma\Big) y(t)$$

$$\dot{z}_S(t) = z_S(t)(1 - z_S(t))\Big(c_P - L(1 - \alpha)(\beta_u z_I(t) + \beta_p(1 - z_I(t)))y(t)\Big)$$

$$\dot{z}_I(t) = z_I(t)(1 - z_I(t))(c_{\mathrm{IP}} - c_{\mathrm{IU}}).$$

Equilibrium Points

The equilibria (y^*, z_S^*, z_I^*) corresponding to $z_I^* = 0$ are:

$$\mathbf{E0} = (0, 0, 0), \quad \mathbf{E1} = (0, 1, 0), \quad \mathbf{E2} = (y_u^*, 1, 0), \mathbf{E3} = (y_{\mathrm{int}}^*, z_{S,\mathrm{int}}^*, 0), \text{ and } \quad \mathbf{E4} = (y_p^*, 0, 0).$$

Here $y_u^* := 1 - \frac{\gamma}{\beta_p}$, $\quad y_{\mathrm{int}}^* := \frac{c_P}{L(1-\alpha)\beta_p}$, $\quad y_p^* := 1 - \frac{\gamma}{\alpha\beta_p}$, and $z_{S,\mathrm{int}}^* := \frac{1}{1-\alpha}\left(\frac{\gamma}{\beta_p(1-y_{\mathrm{int}}^*)} - \alpha\right)$

Equilibria of Coupled Dynamics

Equilibrium Points

The equilibria (y^*, z_S^*, z_I^*) corresponding to $z_I^* = 0$ are:

$$\mathbf{E0} = (0,0,0), \quad \mathbf{E1} = (0,1,0), \quad \mathbf{E2} = (y_u^*, 1, 0), \mathbf{E3} = (y_{\text{int}}^*, z_{S,\text{int}}^*, 0), \text{ and } \quad \mathbf{E4} = (y_p^*, 0, 0).$$

- **E0** is infection-free and everyone adopts protection.
- **E1** is infection-free and susceptible individuals do not adopt protection.
- **E2** is an endemic equilibrium where susceptible individuals continue to remain unprotected.
- **E3** is an endemic equilibrium where a fraction of susceptible individuals adopt protection.
- **E4** is an endemic equilibrium where all susceptible individuals adopt protection.

Equilibrium Characterization: Existence and Stability

Equilibrium Points

The equilibria (y^*, z_S^*, z_I^*) corresponding to $z_I^* = 0$ are:

$$\mathbf{E0} = (0,0,0), \quad \mathbf{E1} = (0,1,0), \quad \mathbf{E2} = (y_u^*, 1, 0), \mathbf{E3} = (y_{\text{int}}^*, z_{S,\text{int}}^*, 0), \text{ and } \quad \mathbf{E4} = (y_p^*, 0, 0).$$

Here $y_u^* := 1 - \frac{\gamma}{\beta_p}$, $\quad y_{\text{int}}^* := \frac{c_P}{L(1-\alpha)\beta_p}$, $\quad y_p^* := 1 - \frac{\gamma}{\alpha\beta_p}$, and $z_{S,\text{int}}^* := \frac{1}{1-\alpha}\left(\frac{\gamma}{\beta_p(1-y_{\text{int}}^*)} - \alpha\right)$

Theorem

Epidemic Parameters	*Endemic Infection Level*	*Equilibria*			
		$\mathbf{E1}: (0,1,0)$	$\mathbf{E2}: (y_u^*, 1, 0)$	$\mathbf{E3}: (y_{\text{int}}^*, z_{S,int}^*, 0)$	$\mathbf{E4}: (y_p^*, 0, 0)$
$\gamma > \beta_p$	–	✓, *stable*	–	–	–
$\alpha\beta_p < \gamma < \beta_p$	$y_u^* < y_{\text{int}}^*$	✓, *unstable*	✓, *stable*	–	–
	$y_{\text{int}}^* < y_u^*$	✓, *unstable*	✓, *unstable*	✓, *stable*	–
$\gamma < \alpha\beta_p$	$y_u^* < y_{\text{int}}^*$	✓, *unstable*	✓, *stable*	–	✓, *unstable*
	$y_p^* < y_{\text{int}}^* < y_u^*$	✓, *unstable*	✓, *unstable*	✓, *stable*	✓, *unstable*
	$y_{\text{int}}^* < y_p^*$	✓, *unstable*	✓, *unstable*	–	✓, *stable*

Illustration I

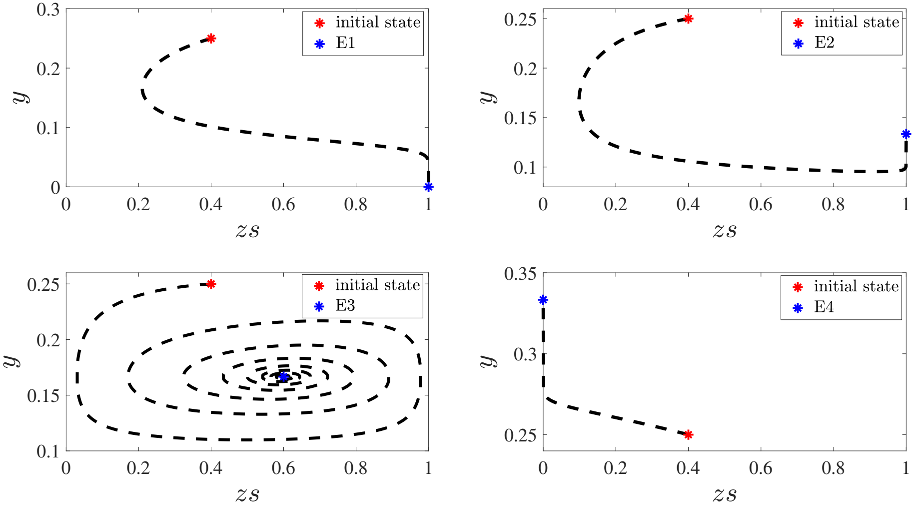

Four possible type of system trajectories

($\gamma = 0.2, 0.13$ in the top row, $\gamma = 0.1, 0.05$ in the bottom row).

When Epidemic Dynamics are Slower than Behavior Dynamics

When epidemic dynamics are slower, then using timescale separation ideas they reduce to

$$\dot{y}(t) = \begin{cases} \big((1 - y(t))\beta_p - \gamma\big) y(t), & \text{if } y < y^*_{\text{int}}, \\ \big((1 - y(t))\alpha\beta_p - \gamma\big) y(t), & \text{if } y > y^*_{\text{int}}. \end{cases}$$

When $y < y^*_{\text{int}}$, it is optimal for susceptible individuals to remain unprotected, and vice versa.

When Epidemic Dynamics are Slower than Behavior Dynamics

When epidemic dynamics are slower, then using timescale separation ideas they reduce to

$$\dot{y}(t) = \begin{cases} \big((1 - y(t))\beta_p - \gamma\big)y(t), & \text{if } y < y^*_{\text{int}}, \\ \big((1 - y(t))\alpha\beta_p - \gamma\big)y(t), & \text{if } y > y^*_{\text{int}}. \end{cases}$$

When $y < y^*_{\text{int}}$, it is optimal for susceptible individuals to remain unprotected, and vice versa.

Trajectories under slow epidemic dynamics

- if $y^*_u < 0$, then $y(t)$ monotonically decreases and converges to the origin;
- if $0 < y^*_u < y^*_{\text{int}}$, then $y(t)$ monotonically converges to y^*_u;
- if $y^*_p < y^*_{\text{int}} < y^*_u$, then $y(t)$ converges to y^*_{int}, and the convergence may not be monotonic;
- if $y^*_{\text{int}} < y^*_p$, then $y(t)$ monotonically converges to y^*_p.

Illustration II: Trajectories under slow epidemic dynamics

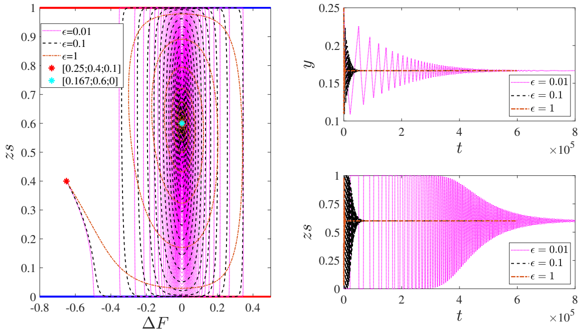

- Smaller ϵ signifies faster behavioral dynamics.
- Here $\Delta F = c_P - L(1-\alpha)(\beta_u z_I(t) + \beta_p(1 - z_I(t)))y(t)$.
- When $\Delta F < 0$, susceptible individuals prefer to adopt protection ($z_S = 0$). Consequently, y decreases, which makes $\Delta F > 0$.
- Similarly, individuals stop adopting protection when $\Delta F > 0$, which makes $z_S = 1$, increases y, and reduces ΔF. This leads to the oscillatory behavior shown above.

Summary

- We proposed a population game model that captures the interaction of epidemic propagation and human behavioral dynamics.
- We derived a complete characterization of the equilibria, their stability, and bifurcations.
- The stability of these equilibria is a function of infection and recovery rate as well as the utility functions of the individuals.
- We investigated the transient behavior of the model under timescale separation.

Reference

A. Satapathi, N. K. Dhar, A. R. Hota, V. Srivastava, “Epidemic Propagation under Evolutionary Behavioral Dynamics: Stability and Bifurcation Analysis.” American Control Conference (ACC), 2022.

A. Satapathi, N. K. Dhar, A. R. Hota and V. Srivastava, “Coupled Evolutionary Behavioral and Disease Dynamics under Reinfection Risk,” IEEE Transactions on Control of Network Systems, vol. 11, no. 2, pages: 795 - 807, 2024.

Part 2: Dynamic Population Game Model for Forward-Looking Agents

In the setting considered in Part 1,

- individuals are homogeneous, i.e., belong to a single well-mixed population
- individuals are myopic
- the disease model does not impart any immunity against future infection and infected individuals are aware of being infected.

Part 2: Dynamic Population Game Model for Forward-Looking Agents

In the setting considered in Part 1,

- individuals are homogeneous, i.e., belong to a single well-mixed population
- individuals are myopic
- the disease model does not impart any immunity against future infection and infected individuals are aware of being infected.

In the remainder of this talk, we propose a **dynamic population game** model that generalizes the above features.

Reference

E. Elokda, S. Bolognani and A. R. Hota, "A Dynamic Population Model of Strategic Interaction and Migration under Epidemic Risk," IEEE Conference on Decision and Control (CDC), 2021.

Augmented SAIR Epidemic

We consider an augmented Susceptible-Asymptomatic-Infected-Recovered (SAIR) epidemic.

There are five epidemic states.

- S: susceptible
- X: asymptomatically infected, infectious
- Y: symptomatically infected, infectious
- U: recovered without exhibiting symptoms, not aware of its recovery
- R: recovered after exhibiting symptoms.

β_x, β_y denote probability of becoming infected from an asymptomatic/symptomatic neighbor.

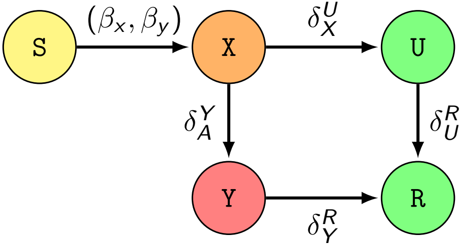

Figure: Augmented SAIR epidemic and transition probabilities

Population State, Actions, Policies

- The state of an agent is denoted $(z, l) \in \mathtt{Z} \times \mathtt{L}$ where $\mathtt{Z} = \{\mathtt{S}, \mathtt{X}, \mathtt{Y}, \mathtt{R}, \mathtt{U}\}$ denotes infection state and $\mathtt{L} = \{1, 2, \ldots, L\}$ denotes L locations.
- State distribution $d \in \mathcal{D} := \Delta(\mathtt{Z} \times \mathtt{L})$ where $d(z, l)$ denotes the proportion of nodes in state z in location l.

Population State, Actions, Policies

- The state of an agent is denoted $(z, l) \in \mathtt{Z} \times \mathtt{L}$ where $\mathtt{Z} = \{\mathtt{S}, \mathtt{X}, \mathtt{Y}, \mathtt{R}, \mathtt{U}\}$ denotes infection state and $\mathtt{L} = \{1, 2, \ldots, L\}$ denotes L locations.
- State distribution $d \in \mathcal{D} := \Delta(\mathtt{Z} \times \mathtt{L})$ where $d(z, l)$ denotes the proportion of nodes in state z in location l.
- Each node chooses:
  - its activation degree $a \in \mathcal{A} = \{0, 1, \ldots, a_{\max}\}$ ($a = 0$ means no activation)
  - next zone $\tilde{l} \in \mathtt{L}$ to move to.

Population State, Actions, Policies

- The state of an agent is denoted $(z, l) \in \mathtt{Z} \times \mathtt{L}$ where $\mathtt{Z} = \{\mathtt{S}, \mathtt{X}, \mathtt{Y}, \mathtt{R}, \mathtt{U}\}$ denotes infection state and $\mathtt{L} = \{1, 2, \ldots, L\}$ denotes L locations.
- State distribution $d \in \mathcal{D} := \Delta(\mathtt{Z} \times \mathtt{L})$ where $d(z, l)$ denotes the proportion of nodes in state z in location l.
- Each node chooses:
  - its activation degree $a \in \mathcal{A} = \{0, 1, \ldots, a_{\max}\}$ ($a = 0$ means no activation)
  - next zone $\tilde{l} \in \mathtt{L}$ to move to.
- A Markovian policy is denoted by $\pi : \mathtt{Z} \times \mathtt{L} \to \Delta(\mathcal{A} \times \mathtt{L})$ where $\pi[a, \tilde{l}|z, l]$ denotes the probability of an agent in state (z, l) choosing action $a, \tilde{l}$.
- We assume $\pi[\cdot|\mathtt{S}, l] = \pi[\cdot|\mathtt{A}, l] = \pi[\cdot|\mathtt{U}, l]$ as agents in these infection states *believe* that they can become infected.
- (π, d) jointly determine the social state.

State Transitions

- The state of an agent evolves probabilistically and is dependent on the social state (π, d) as

$$P[z^+, l^+ \mid z, l](\pi, d) = \sum_{a, \tilde{l}} \pi[a, \tilde{l} \mid z, l] \cdot p[z^+, l^+ \mid z, l, a, \tilde{l}](\pi, d), \tag{1}$$

where

$$p[z^+, l^+ \mid z, l, a, \tilde{l}](\pi, d) = p[z^+ \mid z, l, a](\pi, d) \cdot p[l^+ \mid \tilde{l}].$$

- The first term denotes infection state evolution according to the augmented SAIR model.
- The second term denotes deterministic location evolution with $p[l^+ \mid \tilde{l}] = 1$ if $l^+ = \tilde{l}$.

Infection State Transition

- The infection state evolution probabilities are exogenous constants except for the transition between susceptible to asymptomatic, e.g.,

$$p[\mathtt{R} \mid \mathtt{Y}, l, a](\pi, d) = \delta_Y^R, \qquad p[\mathtt{U} \mid \mathtt{X}, l, a](\pi, d) = \delta_X^U, \quad \text{and so on.}$$

- We define

$$p[\mathtt{X} \mid \mathtt{S}, l, a](\pi, d) = 1 - p[\mathtt{S} \mid \mathtt{S}, l, a](\pi, d) = 1 - \Big[1 - \beta_x \phi_{\mathtt{X},l}(\pi, d) - \beta_y \phi_{\mathtt{Y},l}(\pi, d)\Big]^a$$

where $\phi_{\mathtt{X},l}(\pi, d)$: probability of a randomly chosen neighbor being asymptomatic in location l, given by

$$\phi_{\mathtt{X},l}(\pi, d) = \frac{d(\mathtt{X}, l) \sum_{a,\tilde{l}} a\, \pi(a, \tilde{l} \mid \mathtt{X}, l)}{\sum_z d(z, l) \sum_{a,\tilde{l}} a\, \pi(a, \tilde{l} \mid z, l)}.$$

- Thus, an agent is more likely to connect with an agent with a high activation degree.
- Expression for $\phi_{\mathtt{Y}}(\pi, d)$ is analogous.

Stage Rewards

The stage reward of an agent in state (z, l) choosing action $(a, \tilde{l})$ is:

$$r[z, l, a, \tilde{l}] := r_a[z, l, a] + r_m[l, \tilde{l}] + r_d[z]. \tag{2}$$

- Activation: $r_a[z, l, a] := o[a] - c[z, l, a]$
  - $o[a] \in \mathbb{R}_+$: benefit of social interaction, nondecreasing in degree a
  - $c[z, l, a] \in \mathbb{R}_+$: lockdown cost/restriction in location l. We assume $c[\mathtt{Y}, l, a] \geq c[\mathtt{S}, l, a] = c[\mathtt{X}, l, a] = c[\mathtt{U}, l, a] \geq c[\mathtt{R}, l, a]$.
- Migration: $r_m[l, \tilde{l}] = -c_m$ if $\tilde{l} \neq l$, and 0 otherwise.
- Sickness: $r_d[z] := -c_d$ if $z = \mathtt{Y}$ and 0 otherwise.

Stage Rewards

The stage reward of an agent in state (z, l) choosing action $(a, \tilde{l})$ is:

$$r[z, l, a, \tilde{l}] := r_a[z, l, a] + r_m[l, \tilde{l}] + r_d[z]. \tag{2}$$

- Activation: $r_a[z, l, a] := o[a] - c[z, l, a]$
  - $o[a] \in \mathbb{R}_+$: benefit of social interaction, nondecreasing in degree a
  - $c[z, l, a] \in \mathbb{R}_+$: lockdown cost/restriction in location l. We assume $c[\mathtt{Y}, l, a] \geq c[\mathtt{S}, l, a] = c[\mathtt{X}, l, a] = c[\mathtt{U}, l, a] \geq c[\mathtt{R}, l, a]$.
- Migration: $r_m[l, \tilde{l}] = -c_m$ if $\tilde{l} \neq l$, and 0 otherwise.
- Sickness: $r_d[z] := -c_d$ if $z = \mathtt{Y}$ and 0 otherwise.

For an agent in state (z, l), expected stage reward under policy π is:

$$R[z, l](\pi) = \sum_{a, \tilde{l}} \pi[a, \tilde{l} \mid z, l]\, r[z, l, a, \tilde{l}].$$

Value Function

The expected discounted infinite horizon reward of an agent in state (z, I) is:

$$V[z, I](\pi, d) = R[z, I](\pi) + \alpha \sum_{z^+, I^+} P[z^+, I^+ \mid z, I](\pi, d)\; V[z^+, I^+](\pi, d),$$

assuming social state (π, d) remains unchanged with $\alpha \in [0, 1)$.

Value Function

The expected discounted infinite horizon reward of an agent in state (z, I) is:

$$V[z, I](\pi, d) = R[z, I](\pi) + \alpha \sum_{z^+, I^+} P[z^+, I^+ \mid z, I](\pi, d)\ V[z^+, I^+](\pi, d),$$

assuming social state (π, d) remains unchanged with $\alpha \in [0, 1)$.

An agent chooses its action $(a, \tilde{I})$ to maximize a single-stage deviation from the homogeneous policy π [Filar and Vrieze, 2012], i.e., to maximize

$$Q[z, I, a, \tilde{I}](\pi, d) := r[z, I, a, \tilde{I}] + \alpha \sum_{z^+, I^+} p[z^+, I^+ \mid z, I, a, \tilde{I}](\pi, d)\ V[z^+, I^+](\pi, d).$$

Each agent is aware of the effect of its action on the future state; but assumes the (π, d) to be stationary.

Best Response and Stationary Equilibrium

Definition (Best Response)

The best response of an agent in state (z, l) at the social state (π, d) is the set valued correspondence $B_{z,l} : \Pi \times \mathcal{D} \rightrightarrows \Delta(\mathcal{A} \times \mathrm{L})$ given by

$$B_{z,l}(\pi, d) \in \Big\{\sigma \in \Delta(\mathcal{A} \times \mathrm{L}) : \forall \sigma' \in \Delta(\mathcal{A} \times \mathrm{L}) \\ \sum_{a,\tilde{l}} \Big(\sigma[a, \tilde{l}] - \sigma'[a, \tilde{l}]\Big)\ Q[z, l, a, \tilde{l}](\pi, d) \geq 0\Big\}. \quad (3)$$

The best response is a distribution σ over the actions which maximizes its expected single-stage deviation reward Q.

Best Response and Stationary Equilibrium

Definition (Best Response)

The best response of an agent in state (z, l) at the social state (π, d) is the set valued correspondence $B_{z,l} : \Pi \times \mathcal{D} \rightrightarrows \Delta(\mathcal{A} \times \mathrm{L})$ given by

$$B_{z,l}(\pi, d) \in \Big\{\sigma \in \Delta(\mathcal{A} \times \mathrm{L}) : \forall \sigma' \in \Delta(\mathcal{A} \times \mathrm{L}) \\ \sum_{a,\tilde{l}} \Big(\sigma[a, \tilde{l}] - \sigma'[a, \tilde{l}]\Big) \; Q[z, l, a, \tilde{l}](\pi, d) \geq 0\Big\}. \quad (3)$$

The best response is a distribution σ over the actions which maximizes its expected single-stage deviation reward Q.

Definition (Stationary Equilibrium)

A stationary equilibrium is a social state $(\boldsymbol{\pi}, \boldsymbol{d}) \in \Pi \times \mathcal{D}$ which satisfies

$$\boldsymbol{\pi}[\cdot \mid z, l] \in B_{z,l}(\boldsymbol{\pi}, \boldsymbol{d}), \; \forall (z, l) \in \mathrm{Z} \times \mathrm{L}, \quad \text{(SE.1)}$$

$$\boldsymbol{d} = \boldsymbol{d} \, P(\boldsymbol{\pi}, \boldsymbol{d}). \quad \text{(SE.2)}$$

Stationary Equilibrium

Theorem (Elokda et al. [2021])

There exists a stationary equilibrium $(\boldsymbol{\pi}, \boldsymbol{d})$ *in the proposed game.*

Existence proof relies on fixed point arguments.

Proposition

The stationary distribution $\boldsymbol{d}$ *does not have any asymptomatic or infected agents. If* $\delta_{\mathrm{U}}^{\mathrm{R}} > 0$*, it also does not have any unknowingly recovered agents.*

Societal State Dynamics/Learning Process

- We consider the following coupled state-policy evolution dynamics in discrete-time.
- State evolution: $d^+ = d\,P(\pi, d)$
- Perturbed best response:

$$\tilde{\pi}[a, \tilde{I} \mid z, I](\pi, d) = \frac{\exp\left(\lambda\, Q[z, I, a, \tilde{I}](\pi, d)\right)}{\sum_{a', \tilde{I}'} \exp\left(\lambda\, Q[z, I, a', \tilde{I}'](\pi, d)\right)},$$

  where parameter λ reflects degree of (bounded) rationality.
- Policy evolution:

$$\pi^+[\cdot \mid z, I] = (1 - \eta)\, \pi[\cdot \mid z, I] + \eta\, \tilde{\pi}[\cdot \mid z, I],$$

  where η is an inertia term.

Numerical Results

- We set $\beta_X = \beta_Y = 0.2$, $\delta_A^Y = \delta_A^U = 0.08$, and $\delta_Y^R = 0.04$.
- The agents can activate up to degree $a_{\max} = 6$ with activation reward linearly increasing in the degree.
- Lockdown is implemented by setting $c[z, l, a] = 0$ if $a \leq a_{lock}[z, l]$, and $c[z, l, a] = 3\, o[a_{\max}]$ otherwise.
- We let agents to be highly rational ($\lambda = 10$) and that they update their decisions with an inertia $\eta = 0.2$.

Myopic vs. Farsighted Agents in a Single Zone

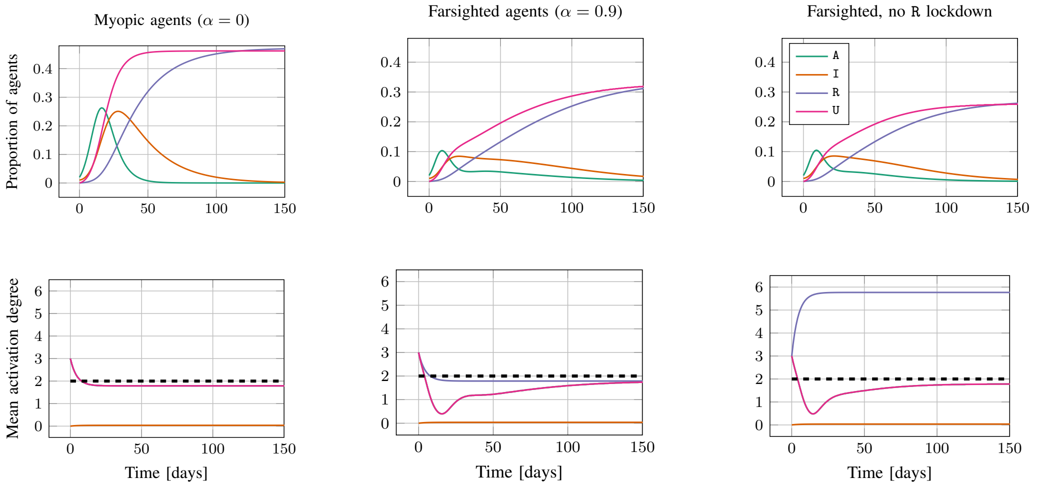

Here maximum allowed activation $a_{lock} = 2$ shown by dashed black line.

Myopic vs. Farsighted Agents

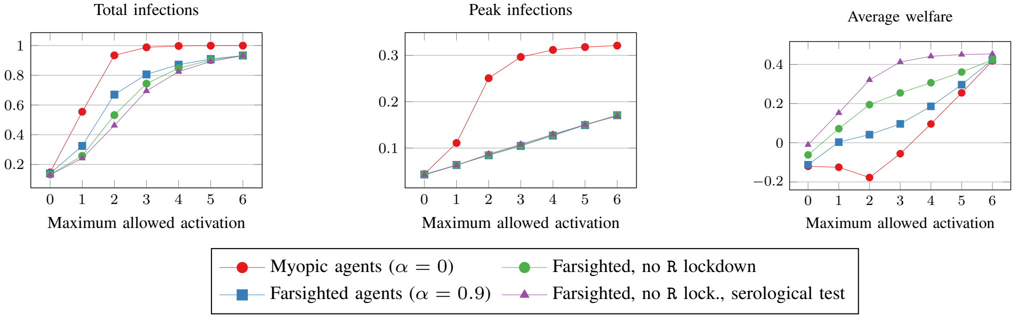

Key Insights:

- Allow recovered agents to activate freely. That would lead to higher activation by susceptible agents without leading to worsening the pandemic.
- Do serological testing and let recovered agents know that they are immune.

Impact of Strategic Migration

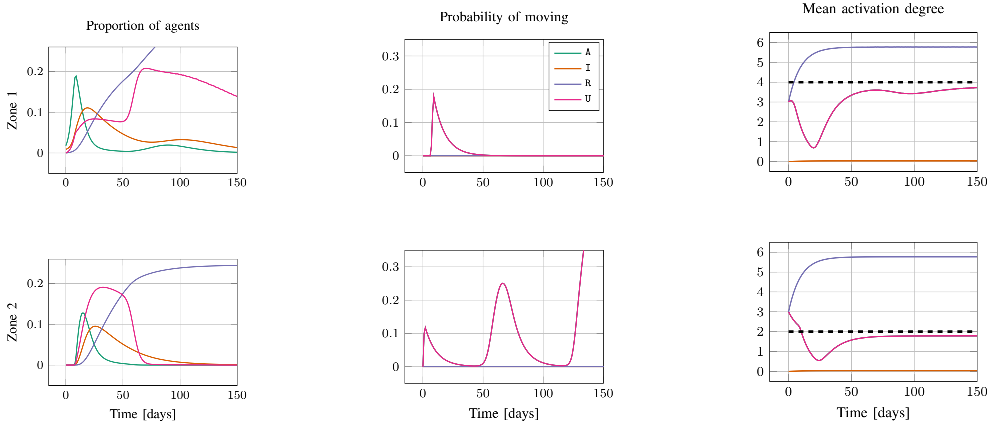

- Zone 1 allows activation degree up to 4 while Zone 2 allows up to 2.
- Zone 2 was initially infection free, yet seems a significant outbreak due to movement from Zone 1.
- Zone 1 sees a minor second wave (around day 50) as agents move there to benefit from relaxed lockdown.

Summary

- We investigated the behavior of forward-looking agents against an epidemic model in a dynamic population game framework.
- Far-sighted agents perform better than myopic ones.
- Permitting recovered agents interact freely does not worsen or subdue the epidemic, but improves welfare.
- Strategic migration (by asymptomatic carriers) can cause multiple outbreaks.

Reference

E. Elokda, S. Bolognani and A. R. Hota, "A Dynamic Population Model of Strategic Interaction and Migration under Epidemic Risk," IEEE Conference on Decision and Control (CDC), 2021.

A. R. Hota, U. Maitra, E. Elokda and S. Bolognani, "Learning to mitigate epidemic risks: A dynamic population game approach," Dynamic Games and Applications, vol. 13, pages: 1106 - 1129, 2023.

Ongoing and Future Work

In addition, we have

- studied the effect of behavioral perception of probabilities on vaccination decisions **[IEEE Trans. Control of Network Systems 2019]**,
- analyzed the behavior of non-myopic or forward-looking agents against the SIS epidemic **[IEEE Control Systems Letters 2023]**,
- examined how the spreading processes can be contained via Bayesian persuasion **[IEEE Control Systems Letters 2024, Under Revision at IEEE Trans. Automatic Control]**, and
- developed algorithms to estimate parameters of spreading processes on networks **[IEEE Tran. Network Science and Engineering 2021, Under Revision at IEEE Trans. Automatic Control]**.

Collaborators

Urmee Maitra

Ezzat Elokda

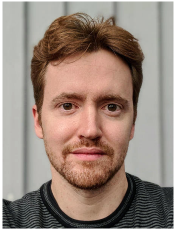

Saverio Bolognani

Abhisek Satapathi

Vaibhav Srivastava

Philip Pare

Many Thanks!

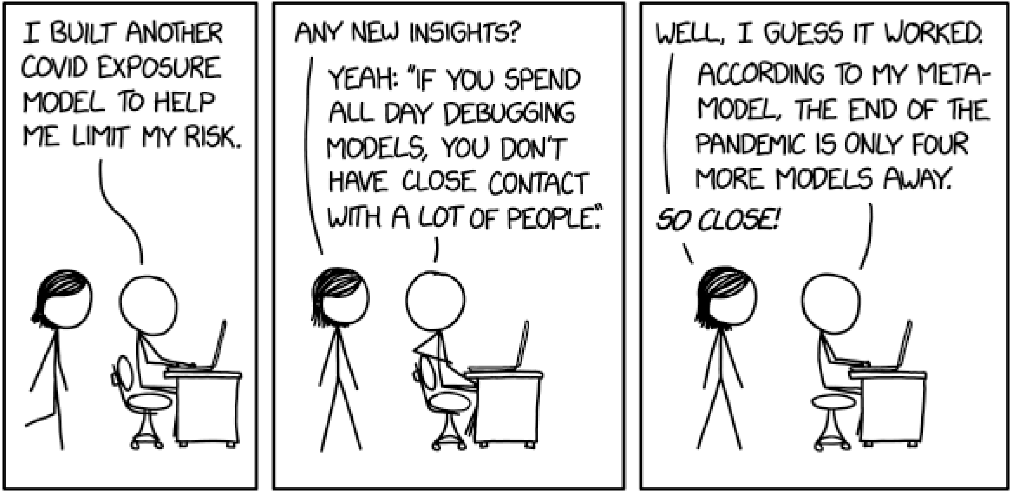

I BUILT ANOTHER COVID EXPOSURE MODEL TO HELP ME LIMIT MY RISK.
ANY NEW INSIGHTS?
YEAH: "IF YOU SPEND ALL DAY DEBUGGING MODELS, YOU DON'T HAVE CLOSE CONTACT WITH A LOT OF PEOPLE."
WELL, I GUESS IT WORKED.
ACCORDING TO MY META-MODEL, THE END OF THE PANDEMIC IS ONLY FOUR MORE MODELS AWAY.
SO CLOSE!

References I

Ezzat Elokda, Andrea Censi, and Saverio Bolognani. Dynamic population games. *arXiv preprint arXiv:2104.14662*, 2021. URL `https://arxiv.org/abs/2104.14662`.

Jerzy Filar and Koos Vrieze. *Competitive Markov decision processes*. Springer Science & Business Media, 2012.