Science of Security and Game Theory

Jeff S Shamma
Georgia Institute of Technology

JASON 2010 Summer Study
Science of Cyber Security
June 29, 2010

- Introduction to game theory:
  - Modeling formalisms
  - Intuition
- Illustrative examples:
  - Traditional
  - Cyber security (simplistic)
- References:
  - Game theory texts & monographs (many!)
  - Alpcan & Başar, *Network Security: A Decision and Game Theory Approach*, online
  - Roy et al., “A survey of game theory as applied to network security”, 2010

What is game theory?

- Myerson, *Game Theory: Analysis of Conflict*, 1997:

  “the study of mathematical models of conflict and cooperation between intelligent rational decision makers”

- Popular perception:



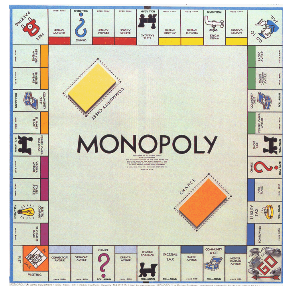

- Broader view: Auctions & markets, conventions, social networks, traffic,...



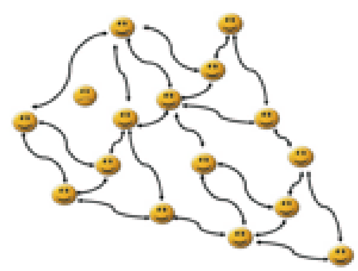


- Players (actors, agents):

$$\mathcal{P} = \{1, 2, ..., p\}$$

- Strategies (choices):
  - Individual:

$$s_i \in \mathcal{S}_i$$

  - Collective:

$$(s_1, ..., s_p) \in \mathcal{S} = \mathcal{S}_1 \times ... \times \mathcal{S}_p$$

- Preferences, expressed as utility function:

$$u_i : \mathcal{S} \to \mathbf{R}$$

$$s \succeq_i s' \quad \Leftrightarrow \quad u_i(s) \geq u_i(s')$$

- Essential feature: Preferences over **collective** strategies:

$$\max_{s_i \in \mathcal{S}_i} u_i(s_i) \quad \text{VS} \quad \max_{s_i \in \mathcal{S}_i} u_i(s_i, s_{-i})$$

- Modeling formalisms:
  - Static games w/ Perfect information
  - Static games w/ Imperfect information
  - Dynamic games w/ Perfect information
  - Dynamic games w/ Imperfect information
- Full rationality vs bounded rationality
- Throughout:
  - Players
  - Strategies
  - Preferences
- *Omission: Cooperative game theory*

Example: Proportional allocation (static w/ perfect info)

- Setup:
  - Players bid b_i for shared resource
  - Resource allocated to player i is:

$$\frac{b_i}{b_1 + ... + b_p}$$

  - Player utility is:

$$u_i(b) = \phi_i\Big(\frac{b_i}{b_1 + ... + b_p}\Big) - b_i$$

  for specified $\phi_i(\cdot)$.

- Proportional allocation is one (of several) *mechanisms* for resource allocation.

Example: Network monitoring (static w/ perfect info)

- Players & strategies:
  - Administrator: $\{$Monitor, Not Monitor$\}$
  - Attacker: $\{$Attack, Not Attack$\}$
- Preferences/utility function:

	M	NM
A	$-c_f - c_a, w - c_m$	$w - c_a, 0$
NA	$0, w - c_m$	$0, w$

where

- $w =$ value of asset
- $c_f =$ cost of failed attack
- $c_a =$ cost to execute attack
- $c_m =$ cost to monitor

Example: Network monitoring (dynamic w/ perfect info)[1]

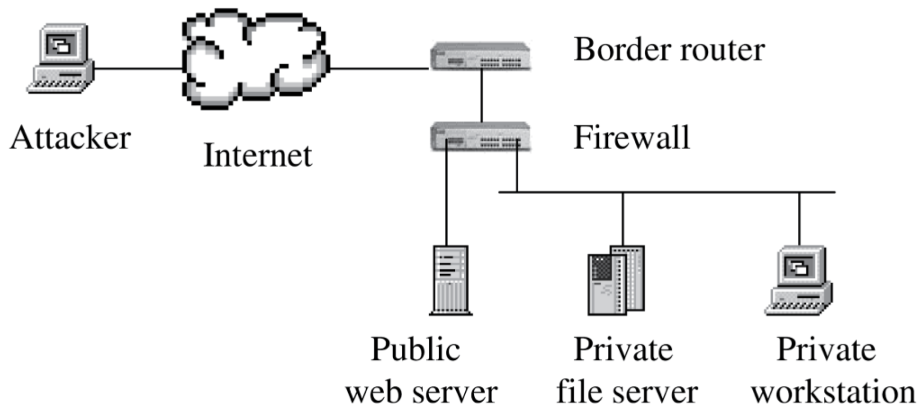

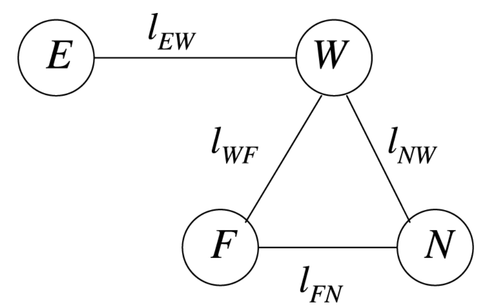

- Setup: External world (E), Web server (W), File server (F), Workstation (N)
- States:
  - Software: `ftpd, httpd, nfsd, process, sniffer, virus`
  - Flags: User account compromised & data compromised
  - 4 Traffic levels per edge
  - Number of states $\approx$ 4 billion

[1]Source: Lye & Wing, "Game strategies in network security", *Int J Inf Secur*, 2005.

- Actions (per state):

$$A^{\text{Attacker}} = \{\text{Attack\_httpd}, \text{Attack\_ftpd}, \text{Continue\_attacking}, \text{Deface\_website\_leave}, \text{Install\_sniffer}, \text{Run\_DOS\_virus}, \text{Crack\_file\_server\_root\_password}, \text{Crack\_workstation\_root\_password}, \text{Capture\_data}, \text{Shutdown\_network}, \varphi\}$$

$$A^{\text{Administrator}} = \{\text{Remove\_compromised\_account\_restart\_httpd}, \text{Restore\_website\_remove\_compromised\_account}, \text{Remove\_virus\_and\_compromised\_account}, \text{Install\_sniffer\_detector}, \text{Remove\_sniffer\_detector}, \text{Remove\_compromised\_account\_restart\_ftpd}, \text{Remove\_compromised\_account\_sniffer}, \varphi\}$$

- *Note:* "Action" $\neq$ "Strategy"

- Dynamics:
  - State/action dependent transition probabilities
  - Transition dependent rewards/costs
- Stochastic Markov game:
  - Stategy = state dependent action rules
  - Preferences = Expected future discounted rewards/costs
- Compare:

	M	NM
A	$-c_f - c_a, w - c_m$	$w - c_a, 0$
NA	$0, w - c_m$	$0, w$

(blurred distinction)

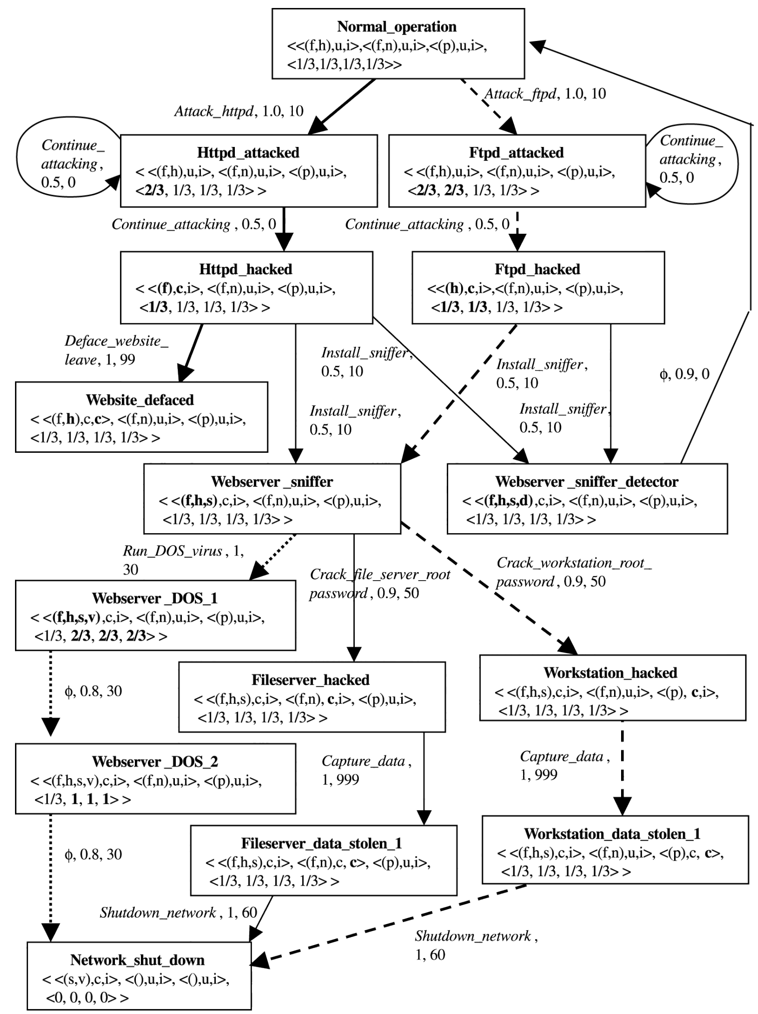

Descriptive agenda, solution concepts, & Nash equilibrium

- Single decision maker:
  - Strategy: $\mathcal{S}$
  - Preferences: $u(s)$
  - Model of rational agent:

$$s^* = \arg\max_{s' \in \mathcal{S}} u(s')$$

- Multiple decision makers:
  - Model of collective = "Solution concept"
  - Prevalent solution concept: **Nash equilibrium**
  - Others: No regret set, correlated equilibrium, cognitive hierarchy
- The action profile a^* is a *Nash equilibrium* if for every player i,

$$u_i(s^*) = u_i(s_i^*, s_{-i}^*) \geq u_i(s_i, s_{-i}^*)$$

  for every $s_i \in \mathcal{S}_i$.
- No player has a *unilateral* incentive to change action

Nash equilibrium (NE) discussion

- Existence (Nash theorem)
- Multiple equilibria:

	S	H
S	$3, 3$	$0, 2$
H	$2, 0$	$2, 2$

Stag hunt

  - NE (S, S) is "payoff dominant"
  - NE (H, H) is "risk dominant"
- Descriptive value, e.g. "beauty contest":
  - Players select number between 0 & 100
  - Player closest to 2/3 of average wins
- Computational complexity in large games

	M	NM
A	$-c_f - c_a, w - c_m$	$w - c_a, 0$
NA	$0, w - c_m$	$0, w$

- No NE for "pure" strategies
- Introduce "mixed" strategies
  - $\mathbf{Pr}\,[\text{A}] = p \quad \& \quad \mathbf{Pr}\,[\text{NA}] = 1 - p$
  - $\mathbf{Pr}\,[\text{M}] = q \quad \& \quad \mathbf{Pr}\,[\text{NM}] = 1 - q$
  - Restate preferences as expected utility
- NE: Solve (p, q)

$$w - c_m = (1 - p) \cdot w$$

$$q \cdot (-c_f - c_a) + (1 - q) \cdot (w - c_a) = 0$$

- Implications:
  - At NE, both players are *indifferent*
  - Specific probabilities depend on *opponent's* utility

- Case I: Dominant strategy
  - s_i^* is a (weakly) **dominant strategy** if for *all* s_{-i}:

    $$u_i(s_i^*, s_{-i}) \geq u_i(s_i', s_{-i})$$

    i.e., s_i^* is always optimal
  - Example: 2nd price sealed bid auction
    - Players have private valuations, v_i
    - Players bid b_i
    - High bid wins and pays second highest bid
    - Fact: $b_i = v_i$ is a dominant strategy
- Case II: Security strategy (hedge against worst case)

  $$s_i^{\text{sec}} = \arg\max_{s_i} \min_{s_{-i}} u_i(s_i, s_{-i})$$

  - Idea: Select s_i^{sec} to maximize *guaranteed* utility
  - Special cases: Security strategies define NE
  - Example: Zero-sum games with mixed strategies (minimax theorem)

- Modeling formulations:
  - Static games w/ Perfect information
  - **Static games w/ Imperfect information**
  - Dynamic games w/ Perfect information
  - Dynamic games w/ Imperfect information
- Full rationality vs bounded rationality
- Throughout:
  - Players
  - Strategies
  - Preferences

Malicious

	M	NM
A	$-c_f - c_a, w - c_m$	$w - c_a, 0$
NA	$0, w - c_m$	$0, w$

Normal

	M	NM
NA	$0, w - c_m$	$0, w$
NA	$0, w - c_m$	$0, w$

- Example[2]:
  - System user knows own "type"
  - Administrator receives signals (e.g,. $\{G, Y, R\}$) and forms "beliefs"
    - $G \Rightarrow \mathbf{Pr}$ [Malicious $= 0.05$]
    - $Y \Rightarrow \mathbf{Pr}$ [Malicious $= 0.25$]
    - $R \Rightarrow \mathbf{Pr}$ [Malicious $= 0.8$]
- Can introduce uncertainty to either or both players (e.g., "honey pot or not")
- Standard example: Auctions

[2] Source: Liu et al., "A Bayesian game approach for intrusion detection in wireless ad hoc networks", *GameNets*, 2006.

- Strategy: Mapping from signal to action probabilities
- Note distinction between "strategy" and "action"
- **Bayesian NE:** Mutually optimal strategies
- Common knowledge, e.g.[3],

$$\begin{array}{c|c|c|}
 & L & R \\ \hline
L & 2,2 & 0,0 \\ \hline
R & 3,0 & 1,1 \\ \hline
\multicolumn{3}{c}{\alpha}
\end{array}
\qquad
\begin{array}{c|c|c|}
 & L & R \\ \hline
L & 2,2 & 0,0 \\ \hline
R & 0,0 & 1,1 \\ \hline
\multicolumn{3}{c}{\beta}
\end{array}
\qquad
\begin{array}{c|c|c|}
 & L & R \\ \hline
L & 2,2 & 0,0 \\ \hline
R & 0,0 & 1,1 \\ \hline
\multicolumn{3}{c}{\gamma}
\end{array}$$

  - Beliefs:
    - Player 1: $\mathbf{Pr}\left[\omega|\overline{\alpha}\right] = \{1, 0, 0\} \quad \& \quad \mathbf{Pr}\left[\omega|\overline{\beta\gamma}\right] = \{0, 3/4, 1/4\}$
    - Player 2: $\mathbf{Pr}\left[\omega|\overline{\alpha\beta}\right] = \{3/4, 1/4/0\} \quad \& \quad \mathbf{Pr}\left[\omega|\overline{\gamma}\right] = \{0, 0, 1\}$
  - Examine "knowledge" in state γ
- Value of information: More accurate signals can lead to lower utility.
- Sensitivity: NE depend on belief probabilities and signal structure of opponents.

[3] Source: Osborne, *An Introduction to Game Theory*, 2003.

- Setup:

$$\text{Private info} \quad \overset{\mathcal{D}}{\Longrightarrow} \quad \text{Social decision}$$

vs

$$\text{Private info} \quad \overset{\mathcal{S}}{\Longrightarrow} \quad \text{Messages} \quad \overset{\mathcal{M}}{\Longrightarrow} \quad \text{Social decision}$$

  - A "mechanism" $\mathcal{M}$ is a rule from reports to decisions.
  - Basis:
    - $*$ Solution concept $\mathcal{S}$ for induced game
    - $*$ Probabilistic model of agent views of environment
  - $\mathcal{D} = \mathcal{M} \circ \mathcal{S}$?
- Standard example: 2nd price auction

- Modeling formulations:
  - Static games w/ Perfect information
  - Static games w/ Imperfect information
  - **Dynamic games w/ Perfect information**
  - Dynamic games w/ Imperfect information
- Full rationality vs bounded rationality
- Throughout:
  - Players
  - Strategies
  - Preferences

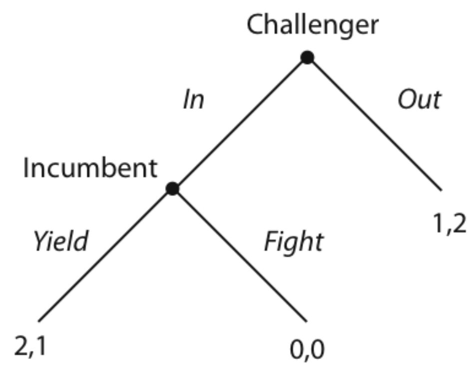

- Entry game:
  - Challenger (Player 1) determines whether or not to compete
  - Incumbent (Player 2) determines whether or not to oppose challenger
  - Payoffs to (player 1, player 2)
- Strategy = Player's action at *every* node
- Strategic form representation:

	Yield	Fight
In	$2, 1$	$0, 0$
Out	$1, 2$	$1, 2$

- NE of strategic form representation: (In,Yield) & (Out,Fight)
- Issue: Non-credible threats!

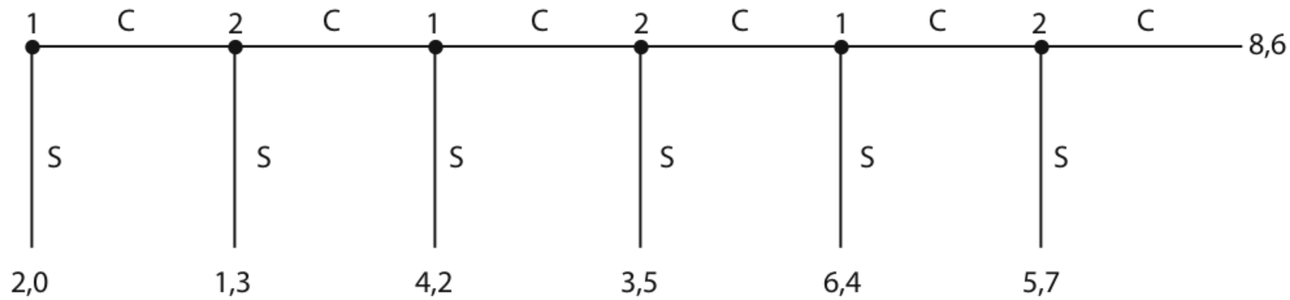

- Backwards induction (i.e., dynamic programming) leads to
  - Construction of Nash equilibrium
  - Exclusion of non-credible threats

  Terminology: **subgame perfect equilibrium**
- Fact: For centipede game, subgame perfect equilibrium is to Stop at any opportunity for both players
- Criticism: Imagine very long centipede game.
  - What should Player 2 do according to subgame perfect equilibrium at interim stage?
  - What should Player 2 do intuitively?

- Players engage in repeated engagements of same game
- *Assumption:* Players observe actions of opponents
- Strategy: Mapping from history to (probabilities of) actions

$$\sigma_i : \mathcal{H} \rightarrow \mathcal{A}_i$$

- Note distinction between "strategy" & "action"
- Network monitoring:

$$\{(NA, NM), (NA, NM), (A, NM)\} \longrightarrow ???$$

- Utilities:
  - Sum of stage payoffs (finite)
  - Discounted future sum of stage payoffs (infinite)

"Infinite" repetition and new equilibria

- Standard example: Long run vs long run Prisoner's dilemma

	C	D
C	$3,3$	$0,4$
D	$4,0$	$1,1$

  - One shot or finitely repeated NE: Play D (dominant strategy)
  - Repeated NE: Play C until observe D, then punish

- Entry game: Long run vs short run players

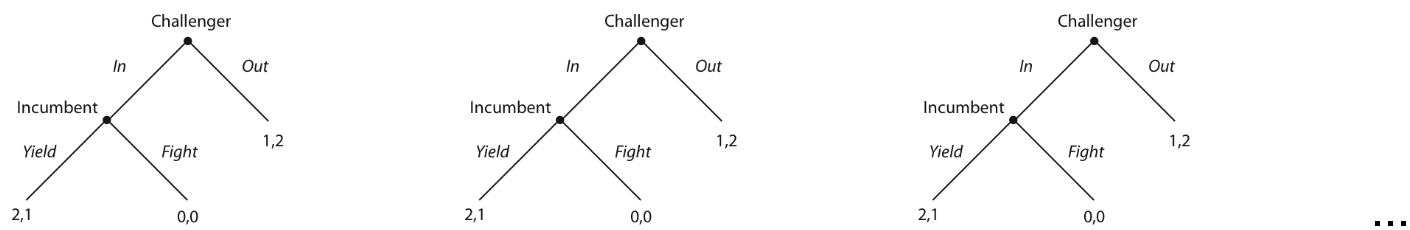

  - One shot or finitely repeated NE: Fight is not credible
  - Repeated NE: Fight is credible

- cf., Repeated game "folk theorems"
- Note: "infinite repetition" equivalent to probabilistic termination

- Modeling formulations:
  - Static games w/ Perfect information
  - Static games w/ Imperfect information
  - Dynamic games w/ Perfect information
  - **Dynamic games w/ Imperfect information**
- Full rationality vs bounded rationality
- Throughout:
  - Players
  - Strategies
  - Preferences

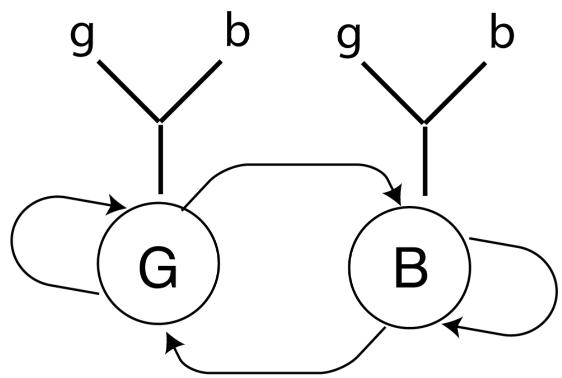

- Setup:
  - Two states & two players
  - Action dependent state transition probabilities
  - Each player has correlated observations about state
  - Strategy: Mapping from *private* history to actions
- Obstruction:
  - Beliefs (of beliefs...) on opponent observations
  - Non-standard information patterns
  - In brief: Intractable
- Positive results for special cases:
  - Repeated games with public monitoring
  - Belief-free equilibria

Simple example: Repeated zero-sum game

	A	B
A	$0, 0$	$1, -1$
B	$0, 0$	$-1, 1$

α

	A	B
A	$-1, 1$	$0, 0$
B	$1, -1$	$0, 0$

β

- Setup:
  - Administrator (row) knows state (allowed behavior)
  - Attacker has probabilistic beliefs
  - Players monitor actions of opponent
  - Two-stages
- NE (depending on specifics...)
  - Administrator does *not* use dominant strategy
  - Rather, use probabilities based on true state (deception?)

- Modeling formulations:
  - Static games w/ Perfect information
  - Static games w/ Imperfect information
  - Dynamic games w/ Perfect information
  - Dynamic games w/ Imperfect information
- **Full rationality vs bounded rationality**
- Throughout:
  - Players
  - Strategies
  - Preferences

- How could agents converge to NE? If so, which NE?

> *Arrow: "The attainment of equilibrium requires a disequilibrium process."*

- Monographs:
  - Weibull, *Evolutionary Game Theory*, 1997.
  - Young, *Individual Strategy and Social Structure*, 1998.
  - Fudenberg & Levine, *The Theory of Learning in Games*, 1998.
  - Samuelson, *Evolutionary Games and Equilibrium Selection*, 1998.
  - Young, *Strategic Learning and Its Limits*, 2004.
  - Sandholm, *Population Dynamics and Evolutionary Games*, 2010.
- Surveys:
  - Hart, "Adaptive heuristics", *Econometrica*, 2005.
  - Fudenberg & Levine, "Learning and equilibrium", *Annual Review of Economics*, 2009.
- Relevance: Online distributed self-configuration

- Single agent adaptation:
  - Stationary environment
  - Asymptotic guarantees
- Multiagent adaptation:

Environment

$=$

Other learning agents

$\Rightarrow$

Non-stationary

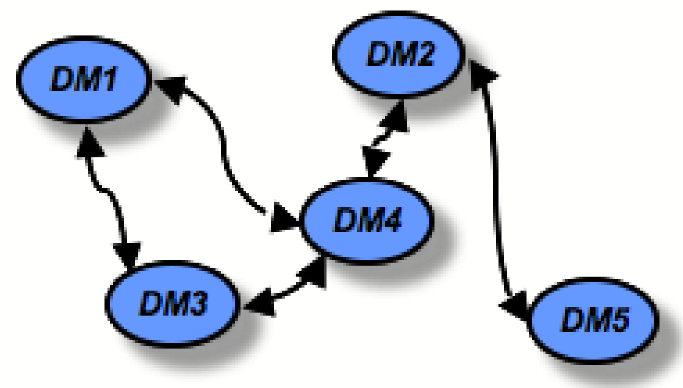

- A is learning about B, whose behavior depends on A, whose behavior depends on B...
- Resulting “feedback loop” has major implications on achievable outcomes.

Illustration: Marginal foresight & mixed equilibria[4]

- Rock-paper-scissors
- Reinforcement learning/replicator dynamics with & without "marginal foresight"

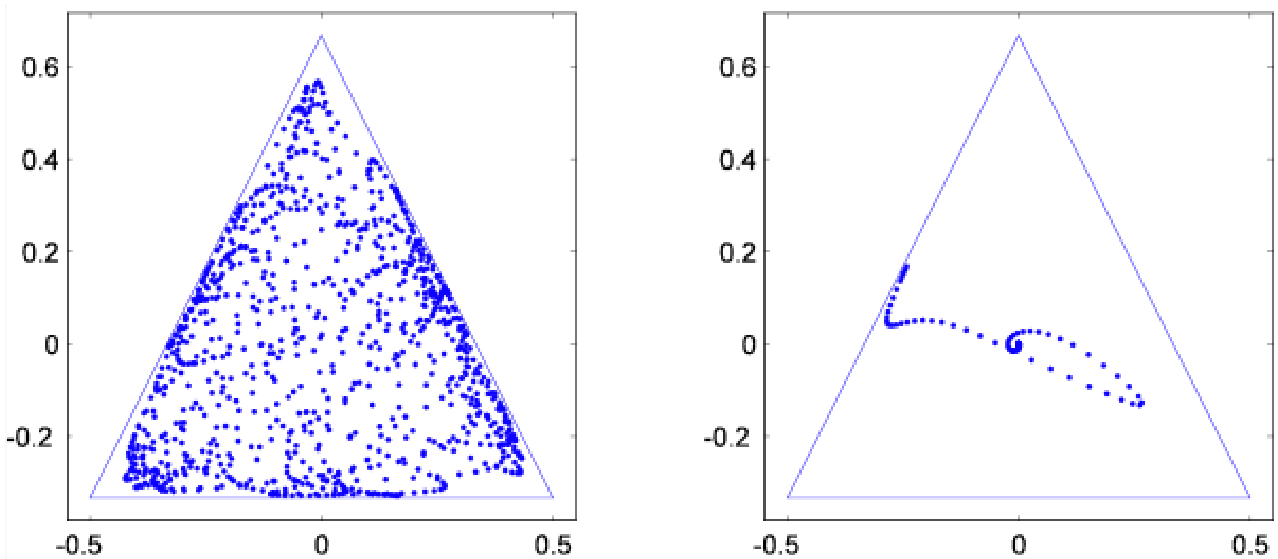

$$\dot{q}_1^j = \left(e_j^\top M_{12}(q_2 + \gamma \dot{r}_2) - q_1^\top M_{12}(q_2 + \gamma \dot{r}_2)\right) q_1^j$$
$$\dot{q}_2^j = \left(e_j^\top M_{21}(q_1 + \gamma \dot{r}_1) - q_2^\top M_{21}(q_1 + \gamma \dot{r}_1)\right) q_2^j$$
$$\dot{r}_1 = \lambda(q_1 - r_1)$$
$$\dot{r}_2 = \lambda(q_2 - r_2)$$

[4] Arslan & Shamma, "Anticipatory learning in general evolutionary games", *IEEE Conference on Decision and Control*, 2006.

- Cyber security and mathematical social sciences:
  - Human decision makers
  - Growing interest in "behavioral game theory" and "neuro-economics"
  - Limitations on repeatable controlled experiments
- Issues:
  - Descriptive vs Prescriptive agenda
  - Computational requirements
  - Full rationality
  - Breaking the symmetry
    - Setup: Repeated game with slightly perturbed payoffs
    - Players monitor opponent actions but do not know opponent perturbation
    - Players play optimal strategies w.r.t. probabilistic forecast models
    - *Theorem*[5]*:* Forecast probabilities are incorrect

[5] Source: Foster & Young, "On the impossibility of predicting the behavior of rational agents," *PNAS*, 2001.

- Cyber security and mathematical social sciences:
  - Human decision makers
  - Growing interest in "behavioral game theory" and "neuro-economics"
  - Limitations on repeatable controlled experiments
- Issues:
  - Descriptive vs Prescriptive agenda
  - Computational requirements
  - Full rationality
  - Breaking the symmetry
    - Setup: Repeated game with slightly perturbed payoffs
    - Players monitor opponent actions but do not know opponent perturbation
    - Players play optimal strategies w.r.t. probabilistic forecast models
    - *Theorem*[5]*:* Forecast probabilities are incorrect

Lou Rawls: "Ain't a horse that can't be rode; ain't a man that can't be throwed."

[5] Source: Foster & Young, "On the impossibility of predicting the behavior of rational agents," *PNAS*, 2001.