Game Theoretic Learning and Social Influence

Jeff S. Shamma

Georgia Institute of Technology

&

جامعة الملك عبدالله
للعلوم والتقنية
King Abdullah University of
Science and Technology

IPAM Workshop on Decision Support for Traffic
November 16–20, 2015

- *Traffic decision makers:*
  - Users
  - Fleet managers
  - Infrastructure planners

cf., Ritter: "Traffic decision support", Thursday AM

- *Traffic decision makers:*
  - Users
  - Fleet managers
  - Infrastructure planners

cf., Ritter: "Traffic decision support", Thursday AM

Satisfaction with decision depends on decisions of others.

- Commuting:
  - Decision = Path
  - Satisfaction depends on paths of others
- Planning:
  - Decision = Infrastructure allocation
  - Satisfaction depends on utilization

"... the study of mathematical models of conflict and cooperation between intelligent rational decision-makers"

Myerson (1991), *Game Theory*.

Viewpoint: Game theory

"... the study of mathematical models of conflict and cooperation between intelligent rational decision-makers"

Myerson (1991), *Game Theory*.

- *Extensive literature...*

- *Game elements:*
  - Players/Agents/Actors
  - Actions/Strategies/Choices
  - Preferences over *joint* choices

Outline

- *Game models*
- *Influence models*

- *Game elements:*
  - Players/Agents/Actors
  - Actions/Strategies/Choices
  - Preferences over *joint* choices

- *Game elements:*
  - Players/Agents/Actors
  - Actions/Strategies/Choices
  - Preferences over *joint* choices

- *Solution concept:* What to expect?

Making decisions vs modeling decisions

- *Single decision maker:*
  - Choice set: $x \in \mathcal{X}$
  - Utility function: $U(x)$

$$x = \arg\max_{x' \in \mathcal{X}} U(x')$$

Linear/Convex/Semidefinite/Integer/Dynamic...Programming

- *Single decision maker:*
  - Choice set: $x \in \mathcal{X}$
  - Utility function: $U(x)$

$$x = \arg\max_{x' \in \mathcal{X}} U(x')$$

Linear/Convex/Semidefinite/Integer/Dynamic...Programming

- Good *model*?

- *Single decision maker:*
  - Choice set: $x \in \mathcal{X}$
  - Utility function: $U(x)$

$$x = \arg\max_{x' \in \mathcal{X}} U(x')$$

Linear/Convex/Semidefinite/Integer/Dynamic...Programming

- Good *model*?

- *Issues:*
  - Complexity, randomness, incompleteness, framing...
  - Furthermore...are preferences even consistent with utility function?

- *Elements:* Players, choices, and preferences over *joint* choices:

$$U_i(x) = U_i(x_i, x_{-i})$$

- *Solution concept:* What to expect?

- *Elements:* Players, choices, and preferences over *joint* choices:

$$U_i(x) = U_i(x_i, x_{-i})$$

- *Solution concept:* What to expect?

Nash Equilibrium

Everyone's choice is optimal *given* the choices of others.

- *Elements:* Players, choices, and preferences over *joint* choices:

$$U_i(x) = U_i(x_i, x_{-i})$$

- *Solution concept:* What to expect?

Nash Equilibrium

Everyone's choice is optimal *given* the choices of others.

- *Alternatives:*
  - Bounded rationality models
  - Hannan consistency, correlated equilibrium, ...

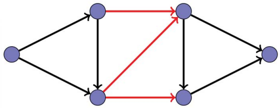

- *Setup:*
  - Players: $1, 2, ..., n$
  - Set of resources: $\mathcal{R} = \{r_1, ..., r_m\}$ (roads)
  - Action sets: $\mathcal{A}_i \subset 2^{\mathcal{R}}$ (paths)
  - Joint action: $(a_1, a_2, ..., a_n)$

- *Setup:*
  - Players: $1, 2, \ldots, n$
  - Set of resources: $\mathcal{R} = \{r_1, \ldots, r_m\}$ (roads)
  - Action sets: $\mathcal{A}_i \subset 2^{\mathcal{R}}$ (paths)
  - Joint action: $(a_1, a_2, \ldots, a_n)$

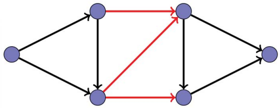

- *Cost:* (vs Utility)
  - Resource level:

$$c_r(a) = \phi_r(\underbrace{\sigma_r(a)}_{\#users})$$

  - User level:

$$C_i(a_i, a_{-i}) = \sum_{r \in a_i} c_r(a)$$

- *Setup:*
  - Players: $1, 2, ..., n$
  - Set of resources: $\mathcal{R} = \{r_1, ..., r_m\}$ (roads)
  - Action sets: $\mathcal{A}_i \subset 2^{\mathcal{R}}$ (paths)
  - Joint action: $(a_1, a_2, ..., a_n)$

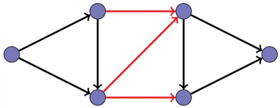

- *Cost:* (vs Utility)
  - Resource level:

$$c_r(a) = \phi_r(\underbrace{\sigma_r(a)}_{\#users})$$

  - User level:

$$C_i(a_i, a_{-i}) = \sum_{r \in a_i} c_r(a)$$

- *Nash equilibrium:* $a^* = (a_1^*, ..., a_n^*)$

$$C_i(a_i^*, a_{-i}^*) \leq C_i(a_i', a_{-i}^*)$$

- Claim: For such congestion games, NE minimizes

$$P(a) = \sum_{r} \sum_{k=0}^{\sigma_r(a)} \phi_r(k)$$

Social influence: Equilibrium shaping

- Claim: For such congestion games, NE minimizes

$$P(a) = \sum_{r} \sum_{k=0}^{\sigma_r(a)} \phi_r(k)$$

- *Overall congestion:*

$$G(a) = \sum_{r} \underbrace{\phi_r(\sigma_r(a))}_{\text{cost}} \cdot \underbrace{\sigma_r(a)}_{\#\ \text{users}}$$

- Claim: For such congestion games, NE minimizes

$$P(a) = \sum_r \sum_{k=0}^{\sigma_r(a)} \phi_r(k)$$

- *Overall congestion:*

$$G(a) = \sum_r \underbrace{\phi_r(\sigma_r(a))}_{\text{cost}} \cdot \underbrace{\sigma_r(a)}_{\#\text{ users}}$$

- Claim: Modified resource cost

$$\phi_r(\sigma_r(a)) + \underbrace{(\sigma_r(a) - 1) \cdot \Big(\phi_r(\sigma_r(a)) - \phi_r(\sigma_r(a) - 1)\Big)}_{\text{imposition toll}}$$

results in NE that minimizes overall congestion.

- *Discussion:*
  - Model presumes NE as outcome
  - How does NE compare to social planner optimal measured by $G(\cdot)$?

- *Discussion:*
  - Model presumes NE as outcome
  - How does NE compare to social planner optimal measured by $G(\cdot)$?

- *Price-of-Anarchy (PoA):*

$$\mathrm{PoA} = \frac{\max_{a \in \mathrm{NE}} G(a)}{\min_a G(a)}$$

pessimistic ratio of performance at NE vs optimal performance

- *Discussion:*
  - Model presumes NE as outcome
  - How does NE compare to social planner optimal measured by $G(\cdot)$?

- *Price-of-Anarchy (PoA):*

$$\mathrm{PoA} = \frac{\max_{a \in \mathrm{NE}} G(a)}{\min_a G(a)}$$

pessimistic ratio of performance at NE vs optimal performance

- *Price-of-Stability (PoS):*

$$\mathrm{PoS} = \frac{\min_{a \in \mathrm{NE}} G(a)}{\min_a G(a)}$$

optimistic ratio of performacne at NE vs optimal performance

- *Setup:*
  - Equally shared cost of resource:

$$\frac{\alpha}{\#\text{ users}}$$

    (vs increasing congestion)
  - High road: $n - \epsilon$
  - Low road: 1

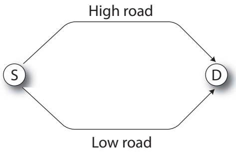

- *Setup:*
  - Equally shared cost of resource:

$$\frac{\alpha}{\#\ \text{users}}$$

    (vs increasing congestion)
  - High road: $n - \epsilon$
  - Low road: 1

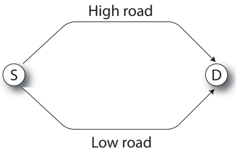

- *NE:*
  - All use High road at individual cost $\frac{n-\epsilon}{n} < 1$
  - All use Low road at individual cost $\frac{1}{n}$

- *Setup:*
  - Equally shared cost of resource:

$$\frac{\alpha}{\#\ \text{users}}$$

    (vs increasing congestion)
  - High road: $n - \epsilon$
  - Low road: 1

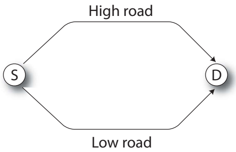

- *NE:*
  - All use High road at individual cost $\frac{n-\epsilon}{n} < 1$
  - All use Low road at individual cost $\frac{1}{n}$

- *PoX:* $G(a)$ is sum of individual costs

$$\text{PoA} \approx n \ \ \& \ \ \text{PoS} = 1$$

- *Setup:*
  - Equally shared cost as before
  - User specific starting points

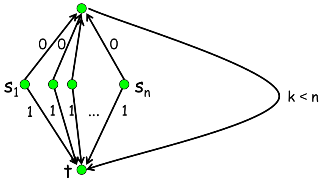

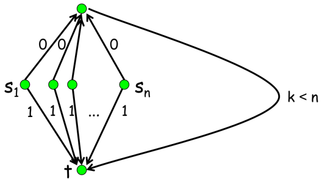

- *Setup:*
  - Equally shared cost as before
  - User specific starting points

- *NE:*
  - All use private resource at individual cost 1
  - All use shared resource at individual cost $\frac{k}{n}$

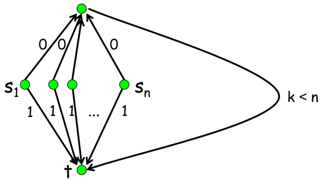

- *Setup:*
  - Equally shared cost as before
  - User specific starting points
- *NE:*
  - All use private resource at individual cost 1
  - All use shared resource at individual cost $\frac{k}{n}$
- *PoX:* $G(a)$ is sum of individual costs

$$\mathrm{PoA} = n/k \ \ \& \ \ \mathrm{PoS} = 1$$

Balcan, Blum, & Mansour (2013), "Circumventing the Price of Anarchy: Leading Dynamics to Good Behavior"

- *Extensions:* Broader solution concepts, various families of games, price-of-uncertainty, price-of-byzantine, price-of- ...

- *Extensions:* Broader solution concepts, various families of games, price-of-uncertainty, price-of-byzantine, price-of- ...

- *(λ, μ)-smoothness:* For any two action profiles, a^* & a'

$$\sum_i C_i(a_i^*, a'_{-i}) \leq \lambda \sum_i C_i(a^*) + \mu \sum_i C(a')$$

- *Extensions:* Broader solution concepts, various families of games, price-of-uncertainty, price-of-byzantine, price-of- ...

- *(λ, μ)-smoothness:* For any two action profiles, a^* & a'

$$\sum_i C_i(a_i^*, a'_{-i}) \leq \lambda \sum_i C_i(a^*) + \mu \sum_i C(a')$$

- *Theorem:* Under (λ, μ)-smoothness,

$$\mathrm{PoA} \leq \frac{\lambda}{1-\mu}$$

Think of a' as NE and a^* as central optimum

Roughgarden (2009), "Intrinsic robustness of the price of anarchy"

- *Game models*
  - Setup & equilibrium
  - Price-of-X

- *Influence models*
  - Equilibrium shaping

Lingering issues:
Uncertain landscapes
Equilibrium analysis

- *Marginal contribution utility:*
  - Assume global objetive, $G(\cdot)$
  - Define "null" action $\emptyset$
  - Set
$$U(a_i, a_{-i}) = G(a_i, a_{-i}) - G(\emptyset, a_{-i})$$
  - Claim: PoS $= 1$

- *Marginal contribution utility:*
  - Assume global objetive, $G(\cdot)$
  - Define "null" action $\emptyset$
  - Set

$$U(a_i, a_{-i}) = G(a_i, a_{-i}) - G(\emptyset, a_{-i})$$

  - Claim: PoS $= 1$

- *Recall:*

$$\phi_r(\sigma_r(a)) + \overbrace{\beta}^{\text{new term}} \underbrace{(\sigma_r(a) - 1) \cdot \Big(\phi_r(\sigma_r(a)) - \phi_r(\sigma_r(a) - 1)\Big)}_{\text{imposition toll:}\tau(a)}$$

- *Marginal contribution utility:*
  - Assume global objetive, $G(\cdot)$
  - Define "null" action $\emptyset$
  - Set
$$U(a_i, a_{-i}) = G(a_i, a_{-i}) - G(\emptyset, a_{-i})$$
  - Claim: PoS $= 1$

- *Recall:*

$$\phi_r(\sigma_r(a)) + \overbrace{\beta}^{\text{new term}} \underbrace{(\sigma_r(a) - 1) \cdot \Big(\phi_r(\sigma_r(a)) - \phi_r(\sigma_r(a) - 1)\Big)}_{\text{imposition toll:}\tau(a)}$$

- *Uncertain landscape:* What if users have different β's?
  many more sources of uncertainty...

$$\phi_r(\sigma_r(a)) + \overbrace{\beta}^{\text{new term}} \underbrace{(\sigma_r(a) - 1) \cdot \Big(\phi_r(\sigma_r(a)) - \phi_r(\sigma_r(a) - 1)\Big)}_{\text{imposition toll:}\tau(a)}$$

- Theorem: As $\kappa \to \infty$

$$\kappa \cdot (\phi_r(\sigma_r(a)) + \tau(a))$$

leads to $\mathrm{PoA} \to 1$

$$\phi_r(\sigma_r(a)) + \overbrace{\beta}^{\text{new term}} \underbrace{(\sigma_r(a) - 1) \cdot \Big(\phi_r(\sigma_r(a)) - \phi_r(\sigma_r(a) - 1)\Big)}_{\text{imposition toll:}\tau(a)}$$

- Theorem: As $\kappa \to \infty$

$$\kappa \cdot (\phi_r(\sigma_r(a)) + \tau(a))$$

  leads to $\mathrm{PoA} \to 1$

- *Extensions:* Optimal bounded tolls in special case of parallel links & affine costs.

Brown & Marden (2015), "Optimal mechanisms for robust coordination in congestion games"

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

vs

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

- A "mechanism" $\mathcal{M}$ is a rule from reports to decisions.
- Mechanism $\mathcal{M}$ induces a game in reporting strategies.
- Seek to implement $\mathcal{D}$ as solution of game, i.e.,

$$\mathcal{D} = \mathcal{M} \circ \mathcal{S}?$$

Standard illustration: Sealed bid/second price auctions

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

vs

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

- *Planner objective:* Assign item to highest private valuation
- *Agent objective:* Item value minus payment
- *Messages:* Bids
- *Social decision:*
  - Item to high bidder
  - Payment from high bidder = Second highest bid
- *Claim:* Truthful bidding is a NE.
- *Special case of broad discussion...*

Illustration: Sequential resource allocation

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

vs

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

- Private info is revealed *sequentially*
- Agents do not know own valuations in advance
- Decisions based on sequential messages

- Two users: $\{1, 2\}$
- User's need for resource is low or high: $\theta_i \in \{L, H\}$

- Two users: $\{1, 2\}$
- User's need for resource is low or high: $\theta_i \in \{L, H\}$
- Planner allocates resource $a \in \{1, 2\}$

- Two users: $\{1, 2\}$
- User's need for resource is low or high: $\theta_i \in \{L, H\}$
- Planner allocates resource $a \in \{1, 2\}$
- Planner objective: Fair allocation with User 2 priority

- Two users: $\{1, 2\}$
- User's need for resource is low or high: $\theta_i \in \{L, H\}$
- Planner allocates resource $a \in \{1, 2\}$
- Planner objective: Fair allocation with User 2 priority
- Dynamics: Coupled state transitions according to 4×4 matrices over set

$$\{(L, L), (H, L), (L, H), (H, H)\}$$

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

vs

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

- *Planner objective:* Induce truthful reporting
- *Agent objective:* Future discounted resource access minus payments

$$\sum_{t=0}^{\infty} \delta^t \Big(v_i(\pi(r_1^t, ..., r_n^t), \theta_i^t) - q_i^t(r_1^t, ..., r_n^t) \Big)$$

- *Messages:* High/Low resource need
- *Theorem:* LP computations so that truthful reporting is a NE.

Kotsalis & Shamma (2013): "Dynamic mechanism design in correlated environments"

- Optimal efficient policy: Favor Agent 2

$$\pi(L,L) = 1/2, \quad \pi(H,L) = 1, \quad \pi(L,H) = 2, \quad \pi(H,H) = 2$$

- Agent 2 can monopolize by misreporting

- Payment rule:

$$q_1(\cdot,\cdot) = 0$$

$$q_2(L,L) < 0, \quad q_2(L,H) < 0, \quad q_2(H,L) < 0$$

$$q_2(H,H) > 0$$

- Ex ante payment from agent $2 = 0$

Lingering issues:
Uncertain landscapes
Equilibrium analysis

- Tuned for behavior only at specific Nash equilibrium
- Presumes agents solve coordinated (dynamic) optimization
- Presumes knowledge of full system dynamics available to all
- Neglects model mismatch

- *Game models*
  - Setup & equilibrium
  - Price-of-X

- *Influence models*
  - Equilibrium shaping
  - Mechanism design

Lingering issues:
Uncertain landscapes
Equilibrium analysis

Shift of focus:

- Away from equilibrium—Nash equilibrium
- Towards how players might arrive to solution—i.e., *dynamics*

Shift of focus:

- Away from equilibrium—Nash equilibrium
- Towards how players might arrive to solution—i.e., *dynamics*

> "The attainment of equilibrium requires a disequilibrium process."

Arrow, 1987.

> "The explanatory significance of the equilibrium concept depends on the underlying dynamics."

Skyrms, 1992.

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.

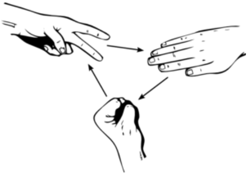

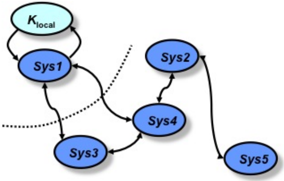

Caution! Single agent learning $\neq$ Multiagent learning

Sato, Akiyama, & Farmer, "Chaos in a simple two-person game", PNAS, 2002.
Piliouras & JSS, "Optimization despite chaos: Convex relaxations to complete limit sets via Poincare recurrence", SODA, 2014.

Best reply dynamics (with inertia)

$$a_i(t) \in \begin{cases} B_i(a_{-i}(t-1)) & \text{w.p. } 0 < \rho < 1 \\ a_i(t-1) & \text{w.p. } 1-\rho \end{cases}$$

$$a_i(t) \in \begin{cases} B_i(a_{-i}(t-1)) & \text{w.p. } 0 < \rho < 1 \\ a_i(t-1) & \text{w.p. } 1 - \rho \end{cases}$$

Features:

- Pure NE is a stationary point
- Based on greedy response to myopic forecast:

$$a_{-i}^{\text{guess}}(t) = a_{-i}(t-1)?$$

- Need *not* converge to NE

$$a_i(t) \in \begin{cases} B_i(a_{-i}(t-1)) & \text{w.p. } 0 < \rho < 1 \\ a_i(t-1) & \text{w.p. } 1-\rho \end{cases}$$

Features:

- Pure NE is a stationary point
- Based on greedy response to myopic forecast:

$$a_{-i}^{\text{guess}}(t) = a_{-i}(t-1)?$$

- Need *not* converge to NE

Theorem

For *finite-improvement-property games* under *best reply with inertia*, player strategies *converge to NE*.

(Includes anonymous congestion games...)

- *Each player:*
  - Maintain empirical frequencies (histograms) of opposing actions
  - Forecasts (incorrectly) that others play independently according to observed empirical frequencies
  - Selects an action that maximizes expected payoff

- *Each player:*
  - Maintain empirical frequencies (histograms) of opposing actions
  - Forecasts (incorrectly) that others play independently according to observed empirical frequencies
  - Selects an action that maximizes expected payoff

- *Compare to best reply:*

$$a_{-i}^{\text{guess}}(t) = a_{-i}(t-1)?$$

vs

$$a_{-i}^{\text{guess}}(t) \sim q_{-i}(t) \in \Delta(\mathcal{A}_{-i})$$

Theorem

For *zero-sum games (1951), 2×2 games (1961), potential games (1996), and $2 \times N$ games (2003)* under *fictitious play*, player empirical frequencies *converge to NE*.

Fictitious play: Convergence results

Theorem

For *zero-sum games (1951), 2×2 games (1961), potential games (1996), and $2 \times N$ games (2003)* under *fictitious play*, player empirical frequencies *converge to NE*.

Theorem

For *Shapley "fashion game" (1964), Jordan anti-coordination game (1993), Foster & Young merry-go-round game (1998)* under *fictitious play*, player empirical frequencies *DO NOT converge to NE*.

Detail: Discussion extended to mixed/randomized NE

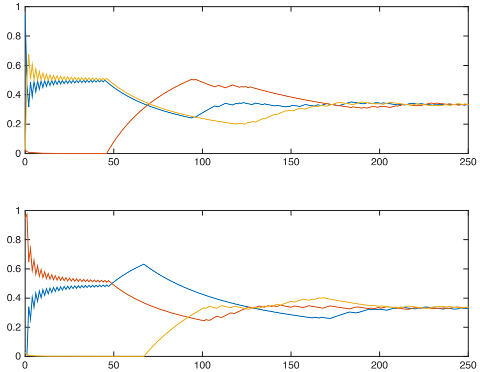

Rock-Paper-Scissors

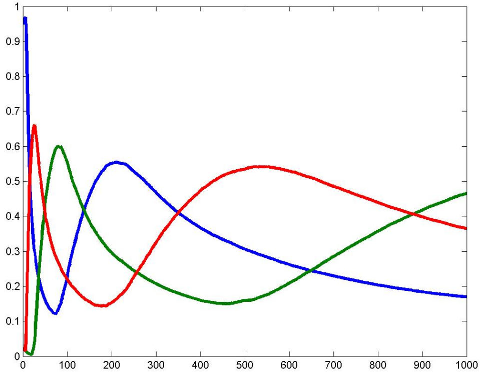

Shapley "Fashion Game"

	R	G	B
R	0, 1	0, 0	1, 0
G	1, 0	0, 1	0, 0
B	0, 0	1, 0	0, 1

FP bookkeeping

- Observe actions of all players
- Construct probability distribution of all possible opponent configurations

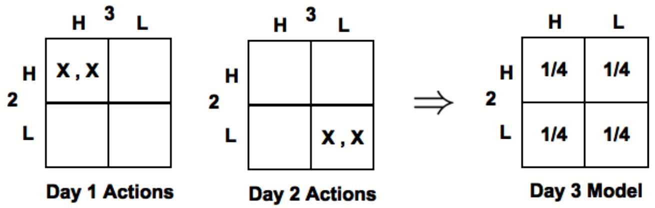

FP bookkeeping

- Observe actions of all players
- Construct probability distribution of all possible opponent configurations

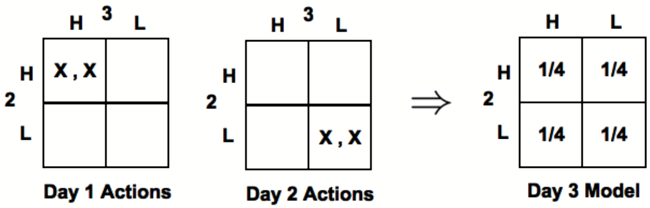

- *Prohibitive* for large games

Modification:

- Maintain empirical frequencies (histograms) of opposing actions
- Forecasts (incorrectly) that others play ~~independently~~ according to observed empirical frequencies
- Selects an action that maximizes expected payoff

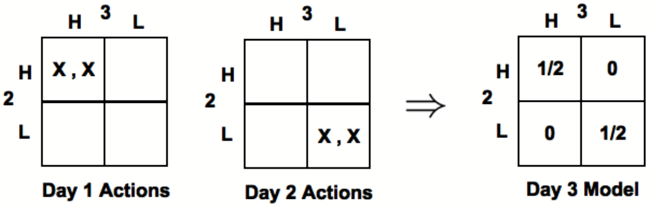

Virtual payoff vector: The payoffs that could have been obtained

$$U_i(t) = \begin{pmatrix} u_i(1, a_{-i}(t)) \\ u_i(2, a_{-i}(t)) \\ \vdots \\ u_i(m, a_{-i}(t)) \end{pmatrix}$$

Time averaged virtual payoff:

$$V_i(t+1) = (1-\rho)V_i(t) + \rho U_i(t)$$

Stepsize ρ is either *constant* (fading) or *diminishing* (averaging)

JSFP bookkeeping

Virtual payoff vector: The payoffs that could have been obtained

$$U_i(t) = \begin{pmatrix} u_i(1, a_{-i}(t)) \\ u_i(2, a_{-i}(t)) \\ \vdots \\ u_i(m, a_{-i}(t)) \end{pmatrix}$$

Time averaged virtual payoff:

$$V_i(t+1) = (1-\rho)V_i(t) + \rho U_i(t)$$

Stepsize ρ is either *constant* (fading) or *diminishing* (averaging)

Equivalent JSFP: At each stage, select best virtual payoff action

Viewpoint: Bookkeeping is oracle based (cf., traffic reports)

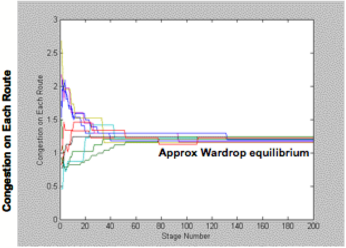

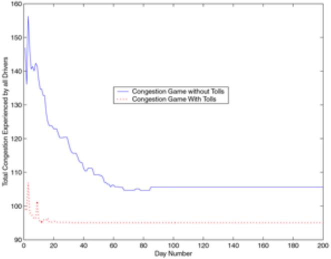

Anonymous congestion

- *Game models*
  - Setup & equilibrium
  - Price-of-X
  - Learning/evolutionary games
- *Influence models*
  - Equilibrium shaping
  - Mechanism design

Lingering issues:
Uncertain landscapes
Equilibrium analysis

- Public service advertising
  - Phase I:
    - Receptive agents: Follow planner's advice (e.g., with probability α)
    - Non-receptive agents: Unilateral best-response dynamics
  - Phase II: Receptive agents may revert to best-response dynamics

- Public service advertising
  - Phase I:
    - Receptive agents: Follow planner's advice (e.g., with probability α)
    - Non-receptive agents: Unilateral best-response dynamics
  - Phase II: Receptive agents may revert to best-response dynamics

- Main results: Desirable bounds on resulting PoA for various settings (anonymous congestion, shared cost, set coverage...)

- Public service advertising
  - Phase I:
    - Receptive agents: Follow planner's advice (e.g., with probability α)
    - Non-receptive agents: Unilateral best-response dynamics
  - Phase II: Receptive agents may revert to best-response dynamics

- Main results: Desirable bounds on resulting PoA for various settings (anonymous congestion, shared cost, set coverage...)

- Compare: Nash equilibrium vs learning agents

Balcan, Blum, & Mansour (2013), "Circumventing the Price of Anarchy: Leading Dynamics to Good Behavior"
Balcan, Krehbiel, Piliouras, and Shin (2012), "Minimally invasive mechanism design: Distributed covering with carefully chosen advice"

- *Game models*
  - Setup & equilibrium
  - Price-of-X
  - Learning/evolutionary games

- *Influence models*
  - Equilibrium shaping
  - Mechanism design
  - Dynamic incentives

Lingering issues:
Uncertain landscapes
Equilibrium analysis

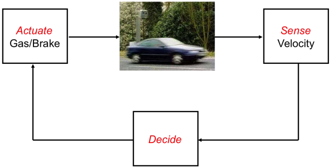
Actuate
Gas/Brake
Sense
Velocity
Decide

vs

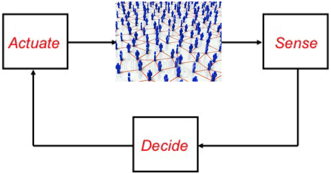
Actuate
Sense
Decide

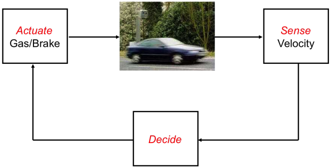

vs

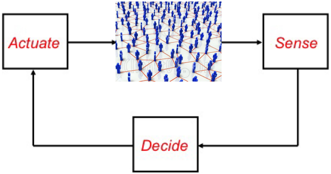

Benefits of feedback (Astrom)

- Reliable behavior from unreliable ~~components~~ *humans*
- Mitigate disturbances
- Shape dynamic behavior

- *Modeling:*
  - Order?
  - Time-scale?
  - Heterogeneity?
  - Non-stationarity?
  - Resolution?
  - Social network effects?

- *Modeling:*
  - Order?
  - Time-scale?
  - Heterogeneity?
  - Non-stationarity?
  - Resolution?
  - Social network effects?

> "Thus unless we know quite a lot about the topology of interaction and the agents' decision-making processes, estimates of the speed of adjustment could be off by many orders of magnitude."

Young, "Social dynamics: Theory and applications", 2001.

- *Modeling:*
  - Order?
  - Time-scale?
  - Heterogeneity?
  - Non-stationarity?
  - Resolution?
  - Social network effects?

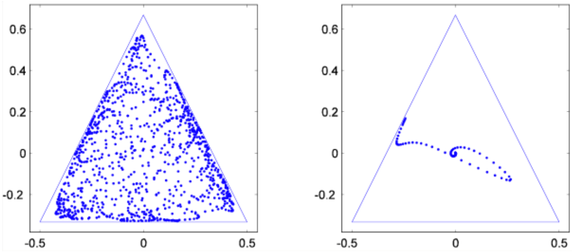

Reinforcement learning vs Trend-based reinforcement learning

- *Modeling:*
  - Order?
  - Time-scale?
  - Heterogeneity?
  - Non-stationarity?
  - Resolution?
  - Social network effects?

- *Measurement & actuation*

- *Game models*
  - Setup & equilibrium
  - Price-of-X
  - Learning/evolutionary games
- *Influence models*
  - Equilibrium shaping
  - Mechanism design
  - Dynamic incentives
- *Challenges*
  - Modeling
  - Measurement & actuation

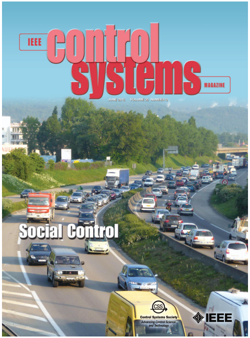

June 2015