Asymmetric Information Security Games

Jeff S. Shamma
with Lichun Li & Malachi Jones

Georgia Institute of Technology

&

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

IPAM Graduate Summer School
Games and Contracts for Cyber-Physical Security
7–23 July 2015

- *Motivation:* One player has superior information
  - Attacker knowns own skill set
  - Defender knows resource characteristics

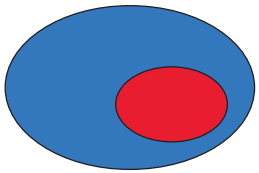

- *Motivation:* One player has superior information
  - Attacker knowns own skill set
  - Defender knows resource characteristics

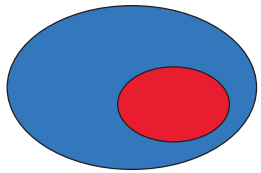

- *Tension:* Exploiting information also reveals information

- *Motivation:* One player has superior information
  - Attacker knowns own skill set
  - Defender knows resource characteristics

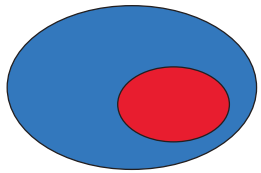

- *Tension:* Exploiting information also reveals information

- *Solution:*
  - Randomized deception
  - *Deliberately do not utilize best resources*

Illustration: Network interdiction

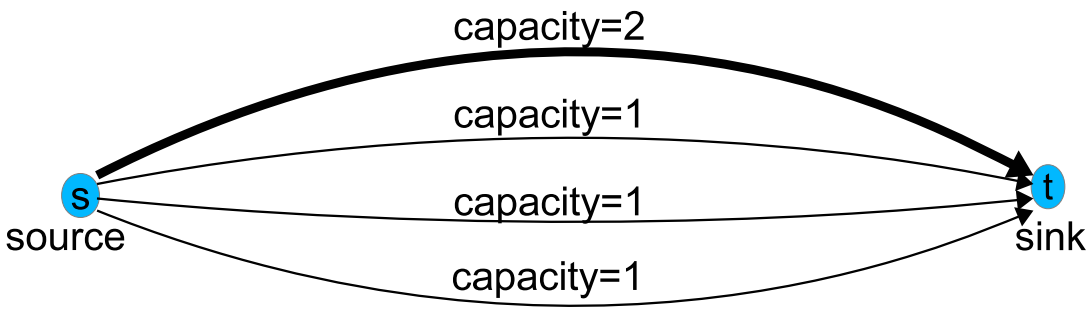

- *System:* One high capacity resource and several low capacity (unknown to attacker)
- *Attacker:*
  - Observes usage (binary) during Phase I
  - Disables selected resource for Phase II
- *Tension:* Initial vs future usage

- *Players:*
  - Maximizer ROW with actions

$$a \in A = \{1, 2, ..., |A|\}$$

  - Minimizer COL with actions

$$b \in B = \{1, 2, ..., |B|\}$$

- *Players:*
  - Maximizer ROW with actions

$$a \in A = \{1, 2, ..., |A|\}$$

  - Minimizer COL with actions

$$b \in B = \{1, 2, ..., |B|\}$$

- *Stages:* $t = 0, 1, 2, ...$

- *Players:*
  - Maximizer ROW with actions

$$a \in A = \{1, 2, ..., |A|\}$$

  - Minimizer COL with actions

$$b \in B = \{1, 2, ..., |B|\}$$

- *Stages:* $t = 0, 1, 2, ...$

- *History:*
  - $h_t = \{(a_0, b_0), (a_1, b_1), ..., (a_{t-1}, b_{t-1})\} \in \mathcal{H}_t$
  - Set of finite histories: $\mathcal{H}_*$

- *Stage payoff:* Matrix $M = [m_{ab}]$
  - m_{ab} = Payoff to ROW under action pair (a, b)
  - m_{ab} = Penalty to COL under action pair (a, b)

- *Stage payoff:* Matrix $M = [m_{ab}]$
  - m_{ab} = Payoff to ROW under action pair (a, b)
  - m_{ab} = Penalty to COL under action pair (a, b)

- *State:*
  - $k \in K = \{1, 2, ..., |K|\}$
  - Index of stage game: $\left\{M^1, M^2, ..., M^{|K|}\right\}$

- *Stage payoff:* Matrix $M = [m_{ab}]$
  - m_{ab} = Payoff to ROW under action pair (a, b)
  - m_{ab} = Penalty to COL under action pair (a, b)

- *State:*
  - $k \in K = \{1, 2, ..., |K|\}$
  - Index of stage game: $\left\{M^1, M^2, ..., M^{|K|}\right\}$

- *Asymmetric information:*
  - At $t = 0$, nature selects $M \in \left\{M^1, M^2, ..., M^{|K|}\right\}$
  - Prior probabilities $\left\{p^1, p^2, ..., p^{|K|}\right\}$
  - ROW informed of selected game

- *Behavioral strategies:*

$$\sigma : \mathcal{H}_* \times K \to \Delta(A) \quad (\textsc{Row})$$
$$\tau : \mathcal{H}_* \to \Delta(B) \quad (\textsc{Col})$$

- *Behavioral strategies:*

$$\sigma : \mathcal{H}_* \times K \to \Delta(A) \quad (\textsc{Row})$$
$$\tau : \mathcal{H}_* \to \Delta(B) \quad (\textsc{Col})$$

Note: *Do not measure payoffs!*

- *Behavioral strategies:*

$$\sigma : \mathcal{H}_* \times K \to \Delta(A) \quad (\textsc{Row})$$
$$\tau : \mathcal{H}_* \to \Delta(B) \quad (\textsc{Col})$$

Note: *Do not measure payoffs!*

- *T-stage game, $\Gamma_T(p)$:*

$$\gamma_t(\sigma, \tau) = \mathbf{E}\left[\frac{1}{T}\sum_{t=0}^{T-1} M^k(a_t, b_t)\right]$$

($\Gamma_\infty(p)$ later...)

Extensive literature

- *Foundations:*
  - Aumann & Maschler (1967), "Repeated games with incomplete information: A survey of resent results", *Report to US Arms Control and Disarmament Agency*.
- *Surveys:*
  - *Zamir (1992), "Repeated games of incomplete information: Zero-sum", Handbook of Game Theory, v. I.*
  - Laraki & Sorin (2014), "Advances in zero-sum dynamic games", *Handbook of Game Theory, v. IV*.
- *Monographs:*
  - Aumann & Maschler (1967/1995), *Repeated Games with Incomplete Information*.
  - Mertens, Sorin & Zamir (1994/2015), *Repeated Games*.
  - Sorin (2002), *A First Course on Zero-Sum Repeated Games*.
- *Variations:* Evolving state, signal monitoring, two-sided incomplete information, ...

	L	R
T	4	3
B	2	1

Case I

	L	R
T	4	2
B	1	3

Case II

- *Case I:* T is a *dominant* strategy, i.e., oblivious best response
- *Case II:* No dominant strategy, i.e., contingent best response

	L	R
T	1	0
B	0	0

M^1

	L	R
T	0	0
B	0	1

M^2

Should ROW *use (weakly) dominant strategy?*

	L	R
T	1	0
B	0	0

M^1

	L	R
T	0	0
B	0	1

M^2

Should ROW *use (weakly) dominant strategy?*

- Dominant strategy payoff stream: 1-or-0, 0, 0, ...

	L	R
T	1	0
B	0	0

M^1

	L	R
T	0	0
B	0	1

M^2

Should ROW *use (weakly) dominant strategy?*

- Dominant strategy payoff stream: 1-or-0, 0, 0, ...
- Compare to perpetual 50-50 strategy:
  - COL faced with averaged game

$$\begin{pmatrix} 1/2 & 0 \\ 0 & 1/2 \end{pmatrix}$$

  - Expected payoff of $1/4$

	L	R
T	1	0
B	0	0

M^1

	L	R
T	0	0
B	0	1

M^2

Should ROW *use (weakly) dominant strategy?*

- Dominant strategy payoff stream: 1-or-0, 0, 0, ...
- Compare to perpetual 50-50 strategy:
  - COL faced with averaged game

$$\begin{pmatrix} 1/2 & 0 \\ 0 & 1/2 \end{pmatrix}$$

  - Expected payoff of $1/4$
- Conclusion: ROW better off *ignoring* information

Example II: Full revelation

	L	R
T	-1	0
B	0	0

M^1

	L	R
T	0	0
B	0	-1

M^2

- Dominant strategy (fully revealing) payoff: 0
- Non-revealing game:

$$\begin{pmatrix} -1/2 & 0 \\ 0 & -1/2 \end{pmatrix}$$

  has expected payoff $-1/4$
- Conclusion: $\textsc{Row}$ better off *fully revealing*

	L	C	R
T	4	0	2
B	4	0	-2

M^1

	L	C	R
T	0	4	-2
B	0	4	2

M^2

- Dominant strategy (long run) payoff: 0
- Non-revealing game:

$$\begin{pmatrix} 2 & 2 & 0 \\ 2 & 2 & 0 \end{pmatrix}$$

  has expected payoff 0
- Can Row do better?

Example III: Partial revelation, cont

	L	C	R
T	4	0	2
B	4	0	-2

M^1

	L	C	R
T	0	4	-2
B	0	4	2

M^2

- State dependent lottery $(\mathtt{h}, \mathtt{t})$:
  - $k = 1$ probabilities: $(3/4, 1/4)$
  - $k = 2$ probabilities: $(1/4, 3/4)$
- Outcome dependent strategy:
  - $\mathtt{h}$: Play T forever
  - $\mathtt{t}$: Play B forever

	L	C	R
T	4	0	2
B	4	0	-2

M^1

	L	C	R
T	0	4	-2
B	0	4	2

M^2

- After stage 0, COL knows outcome
  - **Pr** $\left[k = 1 \;\middle|\; \mathtt{h}\right] = 3/4$ leading to average game

$$\begin{pmatrix} 3 & 1 & 1 \\ 3 & 1 & -1 \end{pmatrix}$$

    in which ROW plays T.
  - **Pr** $\left[k = 1 \;\middle|\; \mathtt{t}\right] = 1/4$ leading to average game

$$\begin{pmatrix} 3 & 1 & -1 \\ 3 & 1 & 1 \end{pmatrix}$$

    in which ROW plays B.
- Long run payoff: 1

- Selected basic results
- Computational approaches

- *Payoff matrix:* $M = [m_{ab}]$

- *Payoff matrix:* $M = [m_{ab}]$
- *Mixed strategies:* $x, y \in \Delta$
  - $x(a) = \mathbf{Pr}[a]$ & $y(b) = \mathbf{Pr}[b]$
  - $x^{\mathrm{T}} M y =$ Expected payoff to $\textsc{Row}$ under strategies (x, y)
  - $x^{\mathrm{T}} M y =$ Expected penalty to $\textsc{Col}$ under strategies (x, y)

- *Payoff matrix:* $M = [m_{ab}]$
- *Mixed strategies:* $x, y \in \Delta$
  - $x(a) = \mathbf{Pr}[a]$ & $y(b) = \mathbf{Pr}[b]$
  - $x^{\mathrm{T}} M y =$ Expected payoff to $\textsc{Row}$ under strategies (x, y)
  - $x^{\mathrm{T}} M y =$ Expected penalty to $\textsc{Col}$ under strategies (x, y)
- *Security levels and strategies:*

$$\begin{aligned} \underline{v} &= \max_{x} \min_{y} x^{\mathrm{T}} M y \\ &= \max_{x,\ell} \ell \\ x^{\mathrm{T}} M &\geq \ell \cdot \mathbf{1}^{\mathrm{T}} \\ x &\in \Delta \end{aligned} \qquad \begin{aligned} \overline{v} &= \min_{y} \max_{x} x^{\mathrm{T}} M y \\ &= \min_{y,\ell} \ell \\ M y &\leq \ell \cdot \mathbf{1} \\ y &\in \Delta \end{aligned}$$

- *Payoff matrix:* $M = [m_{ab}]$
- *Mixed strategies:* $x, y \in \Delta$
  - $x(a) = \mathbf{Pr}[a]$ & $y(b) = \mathbf{Pr}[b]$
  - $x^{\mathrm{T}} M y =$ Expected payoff to $\textsc{Row}$ under strategies (x, y)
  - $x^{\mathrm{T}} M y =$ Expected penalty to $\textsc{Col}$ under strategies (x, y)
- *Security levels and strategies:*

$$\begin{aligned} \underline{v} &= \max_{x} \min_{y} x^{\mathrm{T}} M y \\ &= \max_{x,\ell} \ell \\ x^{\mathrm{T}} M &\geq \ell \cdot \mathbf{1}^{\mathrm{T}} \\ x &\in \Delta \end{aligned} \qquad \begin{aligned} \overline{v} &= \min_{y} \max_{x} x^{\mathrm{T}} M y \\ &= \min_{y,\ell} \ell \\ M y &\leq \ell \cdot \mathbf{1} \\ y &\in \Delta \end{aligned}$$

- *Value:* $\mathtt{val}[M] = \underline{v} = \overline{v}$

- *Pure strategies:*

$$\sigma : \mathcal{H}_* \times K \to A \quad (\textsc{Row})$$
$$\tau : \mathcal{H}_* \to B \quad (\textsc{Col})$$

- *Pure strategies:*

$$\sigma : \mathcal{H}_* \times K \to A \quad (\textsc{Row})$$
$$\tau : \mathcal{H}_* \to B \quad (\textsc{Col})$$

- *Strategic form conversion:*
  - Enumerate all pure strategies
  - Define $\mathcal{M}(p)$ as associated (large) matrix game for $p \in \Delta(K)$

- *Pure strategies:*

$$\sigma : \mathcal{H}_* \times K \to A \quad (\textsc{Row})$$
$$\tau : \mathcal{H}_* \to B \quad (\textsc{Col})$$

- *Strategic form conversion:*
  - Enumerate all pure strategies
  - Define $\mathcal{M}(p)$ as associated (large) matrix game for $p \in \Delta(K)$

- *Consequences:*
  - Value
$$v_t(p) = \texttt{val}[\mathcal{M}(p)]$$
    exists along with associated (mixed) security strategies
  - Equivalence to behavioral strategies (Kuhn's theorem)

- *Average game:*

$$D(p) = \sum_k p^k M^k \quad \& \quad u(p) = \texttt{val}[D(p)]$$

- *Average game:*

$$D(p) = \sum_k p^k M^k \quad \& \quad u(p) = \texttt{val}[D(p)]$$

- *Claim:* $v_t(p) \geq u(p)$
  Proof: ROW plays security strategy for $D(p)$

- *Average game:*

$$D(p) = \sum_k p^k M^k \quad \& \quad u(p) = \texttt{val}[D(p)]$$

- *Claim:* $v_t(p) \geq u(p)$
  Proof: ROW plays security strategy for $D(p)$

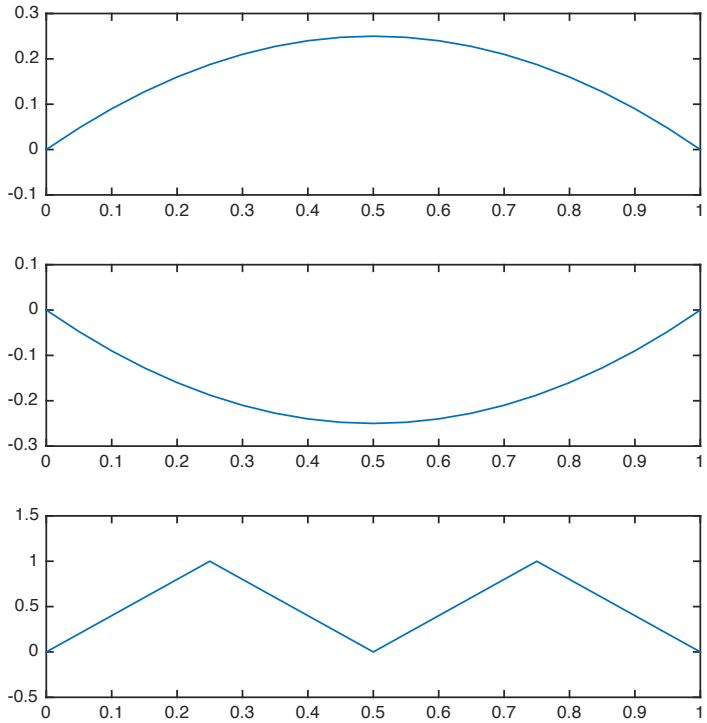

$$u(p) = \max_{x \in \Delta} \min_{y \in \Delta} x^{\mathrm{T}} \left(\sum_k p^k M^k \right) y$$

$$\Leftrightarrow$$

$$\max_{x, \ell} \ell$$

$$x^{\mathrm{T}} \left(\sum_k p^k M^k \right) \geq \ell \cdot \mathbf{1}^{\mathrm{T}}$$

$$x \in \Delta$$

- *Claim:* Suppose

$$p = \sum_{\ell=1}^{L} \lambda_\ell p_\ell$$

There exists a ROW strategy such that

$$v_t(p) \geq \sum_{\ell=1}^{L} \lambda_\ell u(p_\ell)$$

- *Claim:* Suppose

$$p = \sum_{\ell=1}^{L} \lambda_\ell p_\ell$$

There exists a $\textsc{Row}$ strategy such that

$$v_t(p) \geq \sum_{\ell=1}^{L} \lambda_\ell u(p_\ell)$$

- *Implication:* By optimally selecting mixtures

$$v_T(p) \geq \texttt{Cav}u(p)$$

where $\texttt{Cav}u(p) \geq u(p)$ is pointwise smallest concave function

Cav$u(p)$, *cont*

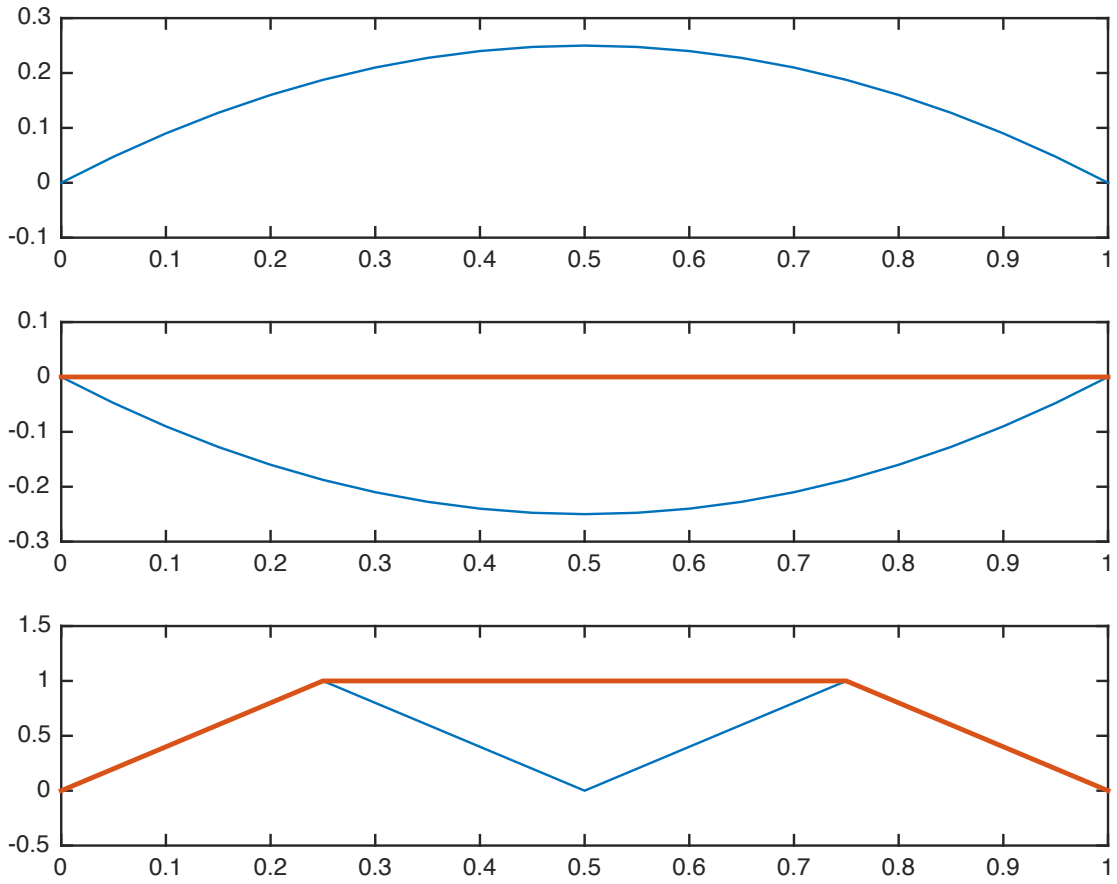
0.3
0.2
0.1
0
-0.1
0 0.1 0.2 0.3 0.4 0.5 0.6 0.7 0.8 0.9 1
0.1
0
-0.1
-0.2
-0.3
0 0.1 0.2 0.3 0.4 0.5 0.6 0.7 0.8 0.9 1
1.5
1
0.5
0
-0.5
0 0.1 0.2 0.3 0.4 0.5 0.6 0.7 0.8 0.9 1

- For $p, p_1, ..., p_N \in \Delta$, suppose

$$p = \sum_{\ell=1}^{L} \lambda_\ell p_\ell$$

- Define joint distribution over $\{1, 2, ..., L\} \times \{1, 2, ..., |K|\}$

$$Q(\ell, k) = \lambda_\ell \cdot p_\ell^k$$

- For $p, p_1, ..., p_N \in \Delta$, suppose

$$p = \sum_{\ell=1}^{L} \lambda_\ell p_\ell$$

- Define joint distribution over $\{1, 2, ..., L\} \times \{1, 2, ..., |K|\}$

$$Q(\ell, k) = \lambda_\ell \cdot p_\ell^k$$

- *Properties:*
  - $\mathbf{Pr}[k] = p^k$
  - $\mathbf{Pr}[\ell] = \lambda_\ell$
  - $\mathbf{Pr}\left[k \mid \ell\right] = p_\ell^k$
  - $\mathbf{Pr}\left[\ell \mid k\right] \propto \lambda_\ell p_\ell^k$

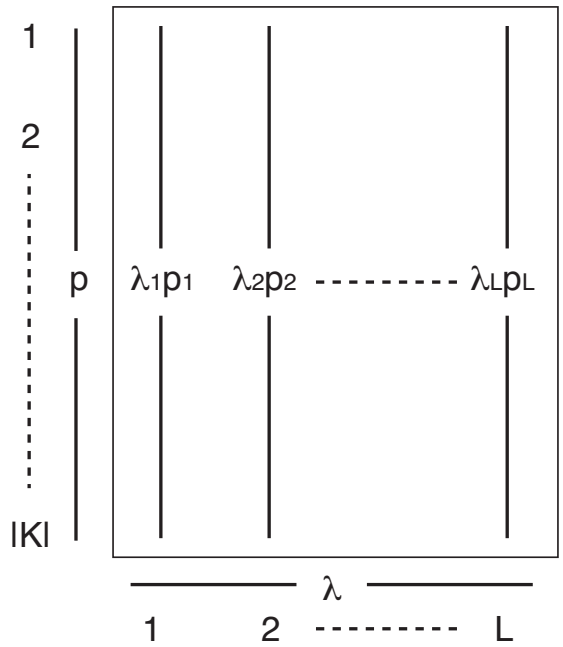

- Starting point: $p = \sum_{\ell=1}^{L} \lambda_\ell p_\ell$
- Row strategy:
  1. Let x_ℓ be Row optimal strategy for $u(p_\ell)$
  2. Select $\ell \propto \lambda_\ell \cdot p_\ell^k$ (k-dependent lottery)
  3. Play selected x_ℓ

- Starting point: $p = \sum_{\ell=1}^{L} \lambda_\ell p_\ell$
- Row strategy:
  1. Let x_ℓ be Row optimal strategy for $u(p_\ell)$
  2. Select $\ell \propto \lambda_\ell \cdot p_\ell^k$ (k-dependent lottery)
  3. Play selected x_ℓ
- Assess Col response *as if* observed ℓ

$$\begin{aligned}
&\min_y \sum_k p^k \left(\sum_\ell \mathbf{Pr}\left[\ell \mid k\right] x_\ell \right)^{\mathrm{T}} M^k y \\
&\quad \geq \min_{y_\ell, \ell=1,2,\ldots,L} \sum_k p^k \left(\mathbf{Pr}\left[\ell \mid k\right] x_\ell \right)^{\mathrm{T}} M^k y_\ell \\
&\quad = \sum_\ell \lambda_\ell u(p_\ell)
\end{aligned}$$

- *Belief splitting:* Posterior belief is p_ℓ with probability λ_ℓ

- *Belief splitting:* Posterior belief is p_ℓ with probability λ_ℓ
- *General setup:*
  - Prior beliefs p over K
  - k-dependent strategy for ROW, $\mathbf{X} = (x^1, ..., x^{|K|}) \in \Delta(A)^{|K|}$

$$x_a^k = \mathbf{Pr}\left[a \;\middle|\; k\right]$$

- *Belief splitting:* Posterior belief is p_ℓ with probability λ_ℓ
- *General setup:*
  - Prior beliefs p over K
  - k-dependent strategy for ROW, $\mathbf{X} = (x^1, ..., x^{|K|}) \in \Delta(A)^{|K|}$

$$x_a^k = \mathbf{Pr}\left[a \mid k\right]$$

- *Computations:*
  - Probability ROW plays a:

$$\pi(a; \mathbf{X}, p) = \sum_k \mathbf{Pr}\left[a \mid k\right] \mathbf{Pr}\left[k\right] = \sum_k x_a^k p^k$$

  - Posterior belief after ROW plays a:

$$\mathcal{B}^k(a; \mathbf{X}, p) = \frac{\mathbf{Pr}\left[a \mid k\right] \mathbf{Pr}\left[k\right]}{\mathbf{Pr}\left[a\right]} = \frac{x_a^k p^k}{\pi(a; \mathbf{X}, p)}$$

- *Naive defense:*
  - Define $p_t^k = \mathbf{Pr}\left[k \;\middle|\; h_t\right]$ (requires ROW strategy σ)
  - Set $\tau_t(h_t)$ to be security strategy for $D(p_t)$

- *Naive defense:*
  - Define $p_t^k = \mathbf{Pr}\left[k \;\middle|\; h_t\right]$ (requires ROW strategy σ)
  - Set $\tau_t(h_t)$ to be security strategy for $D(p_t)$

- *Claim:*

$$\mathbf{E}\left[\gamma_t(\sigma,\tau)\right] \leq \mathtt{Cav}u(p) + \frac{\|M\|}{\sqrt{t}} \sum_k \sqrt{p^k(1-p^k)}$$

- *Naive defense:*
  - Define $p_t^k = \mathbf{Pr}\left[k \;\middle|\; h_t\right]$ (requires ROW strategy σ)
  - Set $\tau_t(h_t)$ to be security strategy for $D(p_t)$

- *Claim:*

$$\mathbf{E}\left[\gamma_t(\sigma,\tau)\right] \leq \mathtt{Cav}u(p) + \frac{\|M\|}{\sqrt{t}} \sum_k \sqrt{p^k(1-p^k)}$$

- *Implication:*

$$\mathtt{Cav}u(p) \leq v_t(p) \leq \mathtt{Cav}u(p) + \frac{\|M\|}{\sqrt{t}} \sum_k \sqrt{p^k(1-p^k)}$$

- *Defend:* React to σ s.t.

$$\gamma_t(\sigma, \mathsf{BR}(\sigma)) \leq \ell, \quad \forall \sigma$$

e.g., Belief based reaction

Uninformed player: Defend vs guarantee

- *Defend:* React to σ s.t.

$$\gamma_t(\sigma, \mathsf{BR}(\sigma)) \leq \ell, \quad \forall \sigma$$

  e.g., Belief based reaction

- *Guarantee:* Find τ s.t.

$$\gamma_t(\sigma, \tau) \leq \ell, \quad \forall \sigma$$

Uninformed player: Defend vs guarantee

- *Defend:* React to σ s.t.

$$\gamma_t(\sigma, \mathsf{BR}(\sigma)) \leq \ell, \quad \forall \sigma$$

  e.g., Belief based reaction

- *Guarantee:* Find τ s.t.

$$\gamma_t(\sigma, \tau) \leq \ell, \quad \forall \sigma$$

- *Example:*

	L	R
T	1	-1
B	-1	1

- *Defend:* React to σ s.t.

$$\gamma_t(\sigma, \mathrm{BR}(\sigma)) \leq \ell, \quad \forall \sigma$$

  e.g., Belief based reaction

- *Guarantee:* Find τ s.t.

$$\gamma_t(\sigma, \tau) \leq \ell, \quad \forall \sigma$$

- *Example:*

	L	R
T	1	-1
B	-1	1

Defend: BR(T) = R & BR(B) = L & BR(50-50) = L/R ($\ell = 0$)

- *Defend:* React to σ s.t.

$$\gamma_t(\sigma, \mathrm{BR}(\sigma)) \leq \ell, \quad \forall \sigma$$

  e.g., Belief based reaction

- *Guarantee:* Find τ s.t.

$$\gamma_t(\sigma, \tau) \leq \ell, \quad \forall \sigma$$

- *Example:*

	L	R
T	1	-1
B	-1	1

Defend: BR(T) = R & BR(B) = L & BR(50-50) = L/R ($\ell = 0$)

Guarantee: L/R @ 50-50 ($\ell = 0$)

- *Hypothetical payoff vector:* Given observations COL can compute

$$g_t(a_t, b_t) = \begin{pmatrix} M^1(a_t, b_t) & M^2(a_t, b_t) & \dots & M^{|K|}(a_t, b_t) \end{pmatrix}$$

as well as its running average

$$\overline{g}_{t+1} = \overline{g}_t + \frac{1}{t+1}(g_t(a_t, b_t) - \overline{g}_t)$$

- *Hypothetical payoff vector:* Given observations COL can compute

$$g_t(a_t, b_t) = \begin{pmatrix} M^1(a_t, b_t) & M^2(a_t, b_t) & \dots & M^{|K|}(a_t, b_t) \end{pmatrix}$$

as well as its running average

$$\overline{g}_{t+1} = \overline{g}_t + \frac{1}{t+1}(g_t(a_t, b_t) - \overline{g}_t)$$

- *Challenge:* Steer $\overline{g}_t$ so that

$$\limsup_{t \to \infty} p^{\mathrm{T}} \overline{g}_t \leq \mathtt{Cav} u(p)$$

$$\overline{g}_{t+1} = \overline{g}_t + \frac{1}{t+1}(g_t(a_t, b_t) - \overline{g}_t)$$

Approachability: A closed convex set, $\mathcal{C}$, is approachable, i.e.,

$$\mathbf{Pr}\left[\text{dist}(\overline{g}_t, \mathcal{C}) \to 0\right] = 1$$

if and only if for all half-spaces, $\mathcal{H}$, containing $\mathcal{C}$, there exists a $y \in \Delta$ such that

$$\mathbf{Co}\left\{\sum_b y^b g_t(a, b) \;\middle|\; a \in A\right\} \subset \mathcal{H}$$

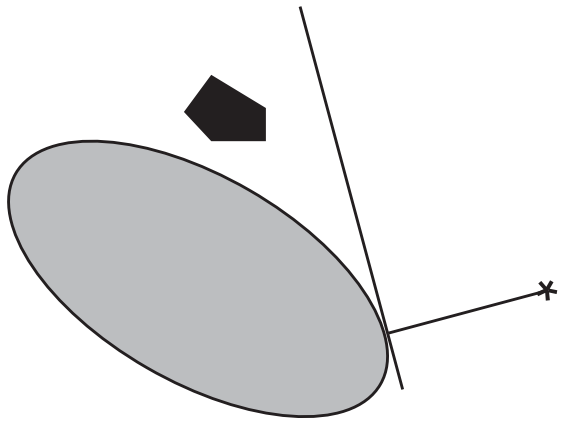

Construction:

1. Find a supporting hyperplane $\mathbf{v} \in \mathbb{R}^{|K|}$ such that

$$\mathbf{v}^{\mathrm{T}} p = \mathtt{Cav} u(p)$$
$$u(q) \leq \mathbf{v}^{\mathrm{T}} q, \forall q \in \Delta$$

2. Define $\mathcal{C} = \left\{ x \in \mathbb{R}^{|K|} \;\middle|\; x \leq \mathbf{v} \right\}$
3. At stage t, if $\overline{g}_t \notin \mathcal{C}$, define

$$q = \frac{1}{Z}\left(\overline{g}_t - \Pi(\overline{g}_t, \mathcal{C})\right)$$

and play optimal strategy for $D(q)$

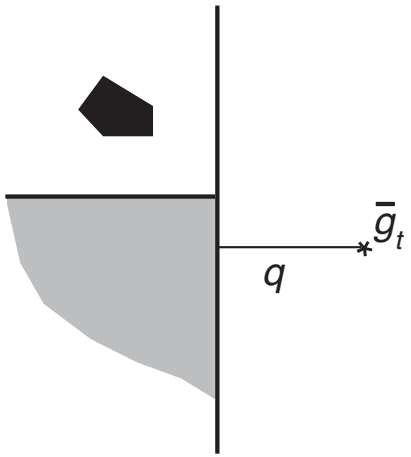

- Row *belief splitting:*

$$\text{Cav}\, u(p) \leq v_t(p)$$

- Row *belief splitting:*

$$\text{Cav}u(p) \leq v_t(p)$$

- Col *naive defense:*

$$v_t(p) \leq \texttt{Cav}u(p) + \frac{\|M\|}{\sqrt{t}} \sum_k \sqrt{p^k(1-p^k)}$$

- Row *belief splitting:*

$$\text{Cav}\, u(p) \leq v_t(p)$$

- Col *naive defense:*

$$v_t(p) \leq \texttt{Cav}\, u(p) + \frac{\|M\|}{\sqrt{t}} \sum_k \sqrt{p^k(1-p^k)}$$

- Col *approachability:*

$$\limsup_{t\to\infty} \mathbf{E}\left[\gamma_t(\sigma,\tau)\right] \leq \text{Cav}\, u(p)$$

- Row *belief splitting:*

$$\textsc{Cav}u(p) \leq v_t(p)$$

- Col *naive defense:*

$$v_t(p) \leq \texttt{Cav}u(p) + \frac{\|M\|}{\sqrt{t}} \sum_k \sqrt{p^k(1-p^k)}$$

- Col *approachability:*

$$\limsup_{t\to\infty} \mathbf{E}\left[\gamma_t(\sigma,\tau)\right] \leq \textsc{Cav}u(p)$$

- *Implication:* Infinite horizon game, $\Gamma_\infty(p)$ has value $\textsc{Cav}u(p)$

- Suppose M is $S \times S$
- Assume "oblivious" $\textsc{Row}$ who ignores $\textsc{Col}$'s actions (optimal)

- Suppose M is $S \times S$
- Assume "oblivious" ROW who ignores COL's actions (optimal)
- Number of strategies:
  - Stage 0: STATE $\rightarrow$ ACTION: S^K
  - Stage 1: (STATE,ACTION) $\rightarrow$ ACTION: $S^{K \cdot S}$
  - $\vdots$
  - Stage T: (STATE, ACTION,..., ACTION) $\rightarrow$ ACTION: $S^{K \cdot S^T}$
  - **Total:**

$$\prod_{t=0}^{T} S^{K \cdot S^t}$$

- "Conversion" to strategic form computationally prohibitive

$$v_{t+1}(p) = \frac{1}{t+1}\left(\max_{\mathbf{X}} \min_{y} \sum_{k} p^k x^k M^k y + t \sum_{a} \pi(a; \mathbf{X}, p) v_t(\mathcal{B}(a; \mathbf{X}, p))\right)$$

- ROW is oblivious to COL
- COL plays myopic defense based on current beliefs

- What is COL's best response to

$$z_{t+1} = Z_t(z_t, a_t)$$
$$a_t \sim \mathbf{X}_t(z_t)$$

- What is COL's best response to

$$z_{t+1} = Z_t(z_t, a_t)$$
$$a_t \sim \mathbf{X}_t(z_t)$$

- Value iteration: COL plays myopic defense

$$V_T(z_T, p_T) = \min_{b_T \sim y} \mathbf{E}\left[M^k(a_T, b_T)\right]$$
$$= \min_y \sum_k p_T^k \mathbf{X}^k(z_T)^{\mathrm{T}} M^k y$$

$$V_t(z_t, p_t) = \min_y \sum_k p_t^k \mathbf{X}^k(z_t)^{\mathrm{T}} M^k y$$
$$+ \mathbf{E}\left[V_{t+1}(Z_t(z_t, a_t), \mathcal{B}(a_t; \mathbf{X}(z_t), p_t))\right]$$

Note: Reversed time indexing and neglected normalization.

$$V_t(z_t, p_t) = \min_y \sum_k p_t^k \mathbf{X}^k(z_t)^{\mathrm{T}} M^k y + \mathbf{E}\left[V_{t+1}(Z_t(z_t, a_t), \mathcal{B}(a_t; \mathbf{X}(z_t), p_t))\right]$$

$$V_t(z_t, p_t) = \min_y \sum_k p_t^k \mathbf{X}^k(z_t)^{\mathrm{T}} M^k y + \mathbf{E}\left[V_{t+1}(Z_t(z_t, a_t), \mathcal{B}(a_t; \mathbf{X}(z_t), p_t))\right]$$

- Row's task as a maximizer

$$v_T(p_T) = \max_{\mathbf{X}} \min_y \sum_k p_T^k x^k M^k y$$

$$v_t(p_t) = \max_{\mathbf{X}} \sum_k \min_y p_t^k x^k M^k y + \mathbf{E}\left[v_{t+1}(\mathcal{B}(a_t; \mathbf{X}, p_t))\right]$$

$$V_t(z_t, p_t) = \min_y \sum_k p_t^k \mathbf{X}^k(z_t)^{\mathrm{T}} M^k y + \mathbf{E}\left[V_{t+1}(Z_t(z_t, a_t), \mathcal{B}(a_t; \mathbf{X}(z_t), p_t))\right]$$

- Row's task as a maximizer

$$v_T(p_T) = \max_{\mathbf{X}} \min_y \sum_k p_T^k x^k M^k y$$

$$v_t(p_t) = \max_{\mathbf{X}} \sum_k \min_y p_t^k x^k M^k y + \mathbf{E}\left[v_{t+1}(\mathcal{B}(a_t; \mathbf{X}, p_t))\right]$$

- Equivalent problem:
  - State space: $p_{t+1} = \mathcal{B}(p_t, a_t)$
  - Action space: $\Delta(A)^{|K|}$
  - Stage reward: $\min_y \sum_k x^k M^k y$

$S \cdot K + 1$ *variables &* S *constraints*

$$\max_{\mathbf{x} \in \Delta^{|K|}} \min_{y \in \Delta} \Big(\sum_k p^k (x^k)^{\mathrm{T}} M^k \Big) y$$

$$\Leftrightarrow$$

$$\max_{\mathbf{x}, \ell} \ell$$

$$\sum_k p^k x^k M^k \geq \ell \cdot \mathbf{1}^{\mathrm{T}}$$

$$x^k \in \Delta, \quad \forall k$$

$S \cdot K + 1$ *variables &* S *constraints*

$$\max_{\mathbf{x} \in \Delta^{|K|}} \min_{y \in \Delta} \Big(\sum_k p^k (x^k)^{\mathrm{T}} M^k \Big) y$$

$$\Leftrightarrow$$

$$\max_{\mathbf{x}, \ell} \ell$$

$$\sum_k p^k x^k M^k \geq \ell \cdot \mathbf{1}^{\mathrm{T}}$$

$$x^k \in \Delta, \quad \forall k$$

Extended $v_1(\cdot)$*:*

- Redefine $v_1(q)$ over positive $q \in \mathbb{R}_+^{|K|}$
- *Positive homogeneity:* $c \cdot v_1(q) = v_1(c \cdot q)$

$$v_2(p) = \max_{\mathbf{X}} \min_{y} \theta \sum_k p^k x^k M^k y + (1-\theta) \sum_a \pi(a; \mathbf{X}, p) v_1(\mathcal{B}(a; \mathbf{X}, p))$$

$$= \max_{\mathbf{X}} \min_{y} \theta \sum_k p^k x^k M^k y + (1-\theta) \sum_a v_1(x_a^1 p^1, ..., x_a^{|K|} p^{|K|})$$

$$v_2(p) = \max_{\mathbf{X}} \min_{y} \theta \sum_k p^k x^k M^k y + (1-\theta) \sum_a \pi(a; \mathbf{X}, p) v_1(\mathcal{B}(a; \mathbf{X}, p))$$

$$= \max_{\mathbf{X}} \min_{y} \theta \sum_k p^k x^k M^k y + (1-\theta) \sum_a v_1(x_a^1 p^1, ..., x_a^{|K|} p^{|K|})$$

$$\Leftrightarrow$$

$$v_2(p) = \max_{\mathbf{X}, \ell_0, ..., \ell_{|A|}} \theta \ell_0 + (1-\theta) \sum_a \ell_a$$

$$\sum_k p^k x^k M^k \geq \ell_0 \cdot \mathbf{1}^{\mathrm{T}}$$

$$v_1(x_a^1 p^1, ..., x_a^{|K|} p^{|K|}) \geq \ell_a, \quad \forall a$$

$$x^k \in \Delta, \quad \forall k$$

$$v_2(p) = \max_{\mathbf{X}} \min_{y} \theta \sum_k p^k x^k M^k y + (1-\theta) \sum_a \pi(a; \mathbf{X}, p) v_1(\mathcal{B}(a; \mathbf{X}, p))$$

$$= \max_{\mathbf{X}} \min_{y} \theta \sum_k p^k x^k M^k y + (1-\theta) \sum_a v_1(x_a^1 p^1, ..., x_a^{|K|} p^{|K|})$$

$$\Leftrightarrow$$

$$v_2(p) = \max_{\mathbf{X}, \ell_0, \ldots, \ell_{|A|}} \theta \ell_0 + (1-\theta) \sum_a \ell_a$$

$$\sum_k p^k x^k M^k \geq \ell_0 \cdot \mathbf{1}^{\mathrm{T}}$$

$$v_1(x_a^1 p^1, ..., x_a^{|K|} p^{|K|}) \geq \ell_a, \quad \forall a$$

$$x^k \in \Delta, \quad \forall k$$

Constraints on $v_1(\cdot)$ result in product terms in LP

- LHS Nested LP:

$$\max_{v} c^{\mathrm{T}} v$$

$$Av \leq b$$

$$\sum_{i} v_i F_i w \leq f$$

- LHS Nested LP:

$$\max_{v} c^{\mathrm{T}} v$$
$$Av \leq b$$
$$\sum_{i} v_i F_i w \leq f$$

LP structure lost!

- LHS Nested LP:

$$\max_{v} c^{\mathrm{T}} v$$
$$Av \leq b$$
$$\sum_{i} v_i F_i w \leq f$$

LP structure lost!

- RHS Nested LP:

$$\max_{v} c^{\mathrm{T}} v$$
$$Av \leq b$$
$$Fw \leq \sum_{i} v_i f_i$$

LP structure preserved!

$$v_1(p) = \max_{\mathbf{X},\ell} \ell$$

$$\sum_k p^k x^k M^k \geq \ell \cdot \mathbf{1}^{\mathrm{T}}$$

$$x^k \in \Delta, \quad \forall k$$

$$\Leftrightarrow$$

$$v_1(p) = \max_{\mathbf{Z},\ell} \ell$$

$$\sum_k z^k M^k \geq \ell \cdot \mathbf{1}^{\mathrm{T}}$$

$$\mathbf{1}^{\mathrm{T}} z^k = p^k, \quad \forall k$$

Change of variables:

- $z^k = p^k x^k \geq 0$
- Probabilities now enter in RHS

$$v_2(p) = \max_{\mathbf{x}, \ell_0, \ldots, \ell_{|A|}} \theta \ell_0 + (1-\theta) \sum_a \ell_a$$

$$\sum_k p^k x^k M^k \geq \ell_0 \cdot \mathbf{1}^{\mathrm{T}}$$

$$v_1(x_a^1 p^1, \ldots, x_a^{|K|} p^{|K|}) \geq \ell_a, \quad \forall a$$

$$x^k \in \Delta, \quad \forall k$$

Each constraint on $v_1(\cdot)$ is an LP:

$$\sum_k z^k(a) M^k \geq \ell_a \cdot \mathbf{1}^{\mathrm{T}}$$

$$\mathbf{1}^{\mathrm{T}} z^k(a) = x_a^k p^k$$

$$(S+1) \cdot (S \cdot K \text{ vars } \& \ S \text{ cons })$$

- *Claim:*
  - *Recursive structure:* If $v_t(p)$ has RHS LP-dependence on p, then $v_{t+1}(p)$ has RHS LP-dependence on p.
  - *Growth in LP size:* LP size grows with

$$\text{size}[v_t] \approx S \cdot \textit{size}[v_{t-1}]$$

- *Claim:*
  - *Recursive structure:* If $v_t(p)$ has RHS LP-dependence on p, then $v_{t+1}(p)$ has RHS LP-dependence on p.
  - *Growth in LP size:* LP size grows with

$$\text{size}[v_t] \approx S \cdot \textit{size}[v_{t-1}]$$

- *Polynomial dependence* on size of game (but not length):

$$(S + S^2 + \ldots + S^T) \cdot K \quad \text{vs} \quad \underbrace{\prod_{t=0}^{T} S^{K \cdot S^t}}_{\text{face value}}$$

(cf., "sequence form" of Koller, von Stengel, & Megiddo, 1996)

- *Claim:*
  - *Recursive structure:* If $v_t(p)$ has RHS LP-dependence on p, then $v_{t+1}(p)$ has RHS LP-dependence on p.
  - *Growth in LP size:* LP size grows with

$$\text{size}[v_t] \approx S \cdot \textit{size}[v_{t-1}]$$

- *Polynomial dependence* on size of game (but not length):

$$(S + S^2 + \ldots + S^T) \cdot K \quad \text{vs} \quad \underbrace{\prod_{t=0}^{T} S^{K \cdot S^t}}_{\text{face value}}$$

(cf., "sequence form" of Koller, von Stengel, & Megiddo, 1996)

- Recursive computation applicable to *time-varying* repeated games (i.e., changing M-matrices)

	L	R
T	1	0
B	0	0

M^1

	L	R
T	0	0
B	0	1

M^2

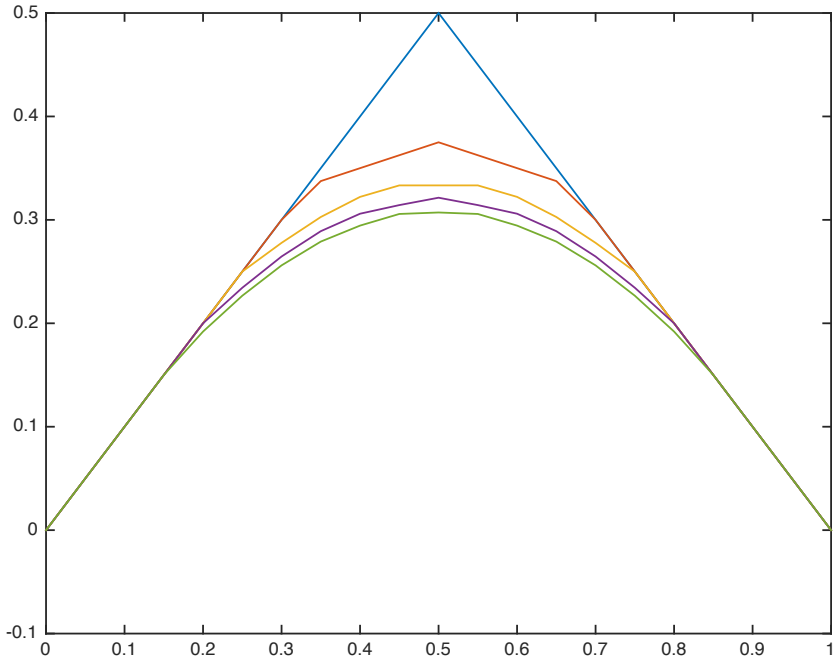

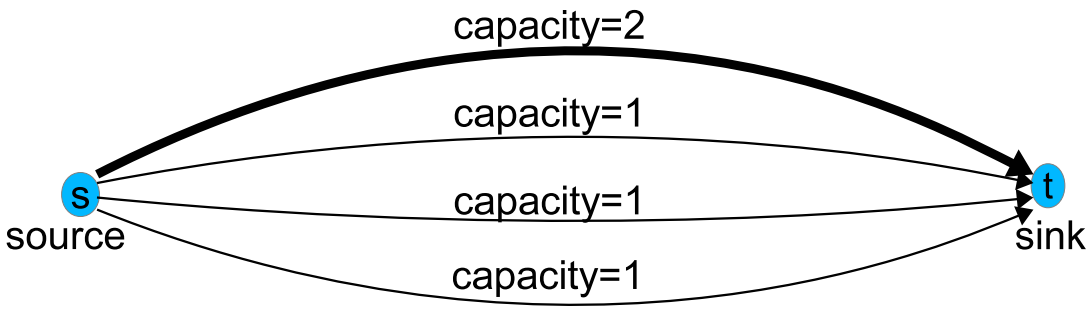

- *System:* One high capacity resource and several low capacity (unknown to attacker)
- *Attacker:*
  - Observes usage (binary) during Phase I
  - Disables selected resource for Phase II
- *Tension:* Initial vs future usage
- *Note:* *Time-varying* M-matrices

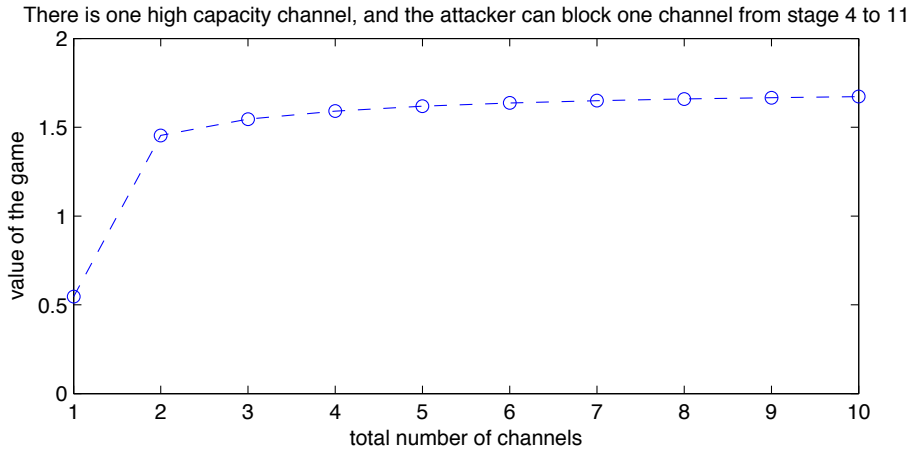

- Initial phase: 3 stages
- Remaining phase: R stages
- Probability of using high capacity resource:

$r \setminus t$	1	2	3	4
1	1	1	1	0.1429
2	1	1	1	0
3	0.5799	0.5799	0.5799	0.4201
4	0.3846	0.3846	0.3846	0.6154
5	0.3529	0.3529	0.3529	0.6471
6	0.2553	0.2553	0.2553	0.9787
7	0.25	0.25	0.25	1
8	0.25	0.25	0.25	1

- *Recap:*
  - Examples
  - Basic results
  - Computational approach

- *Recap:*
  - Examples
  - Basic results
  - Computational approach
- *Extensions:*
  - Discounted problems:

$$\gamma_\lambda(\sigma,\tau) = \mathbf{E}\left[(1-\lambda)\sum_{t=0}^{\infty}\lambda^t M^k(a_t,b_t)\right]$$

  - Markov chains with informed controller:

$$k_{t+1} \sim \phi(k_t, a_t)$$

  - Receding horizon implementation

- *Recap:*
  - Examples
  - Basic results
  - Computational approach
- *Extensions:*
  - Discounted problems:

$$\gamma_\lambda(\sigma,\tau) = \mathbf{E}\left[(1-\lambda)\sum_{t=0}^{\infty}\lambda^t M^k(a_t,b_t)\right]$$

  - Markov chains with informed controller:

$$k_{t+1} \sim \phi(k_t, a_t)$$

  - Receding horizon implementation
- *Lingering issue:* Computational policies for uninformed player