Overcoming Limitations of Game-Theoretic Distributed Control

Jason R. Marden
California Institute of Technology

(joint work with Adam Wierman)

Southern California Network Economics and Game Theory Symposium
October 1, 2009

Trend: Transition from centralized to local decision making

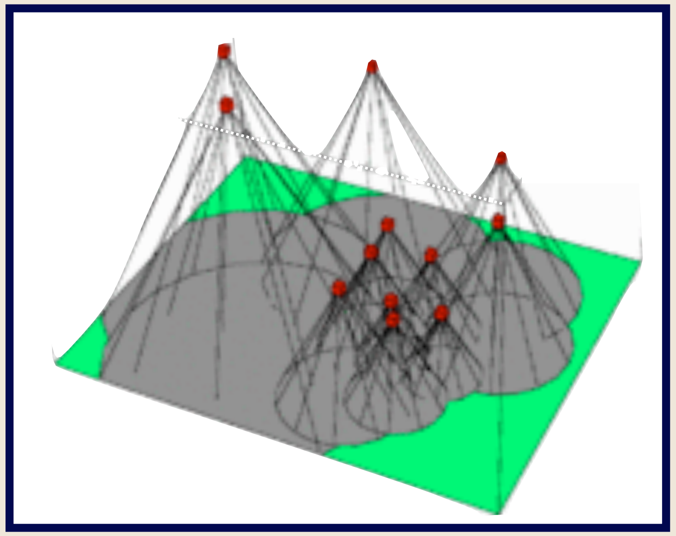

Sensor coverage

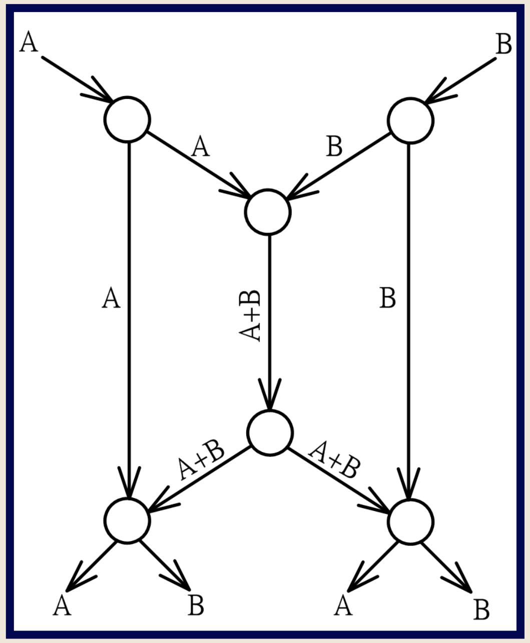

Network Coding

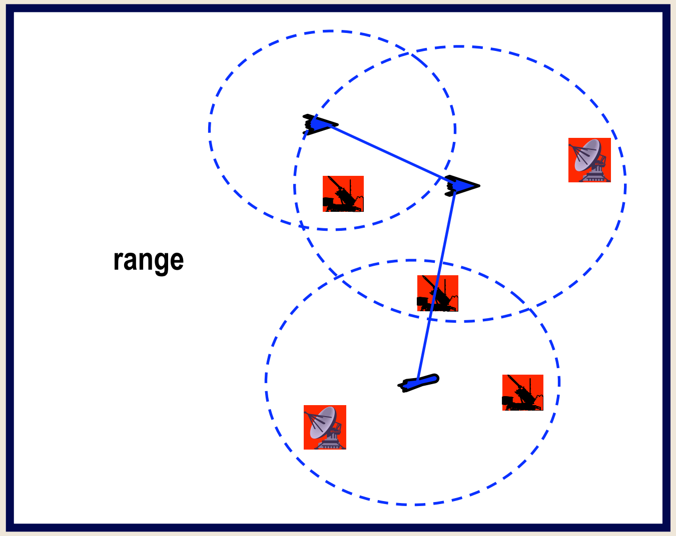

Vehicle Target Assignment

Appeal

Local processing (manageable)

Reduces communication

Robustness

Challenges

Characterization

Coordination

Efficiency

How should we design distributed engineering systems?

Trend: Transition from centralized to local decision making

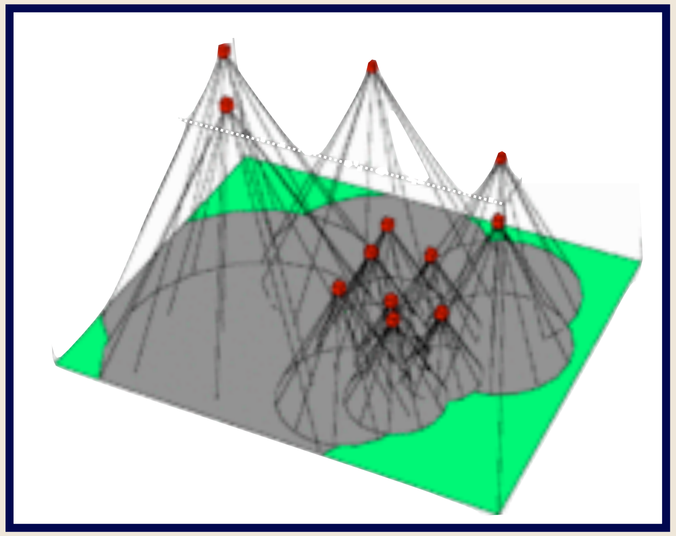

Sensor coverage

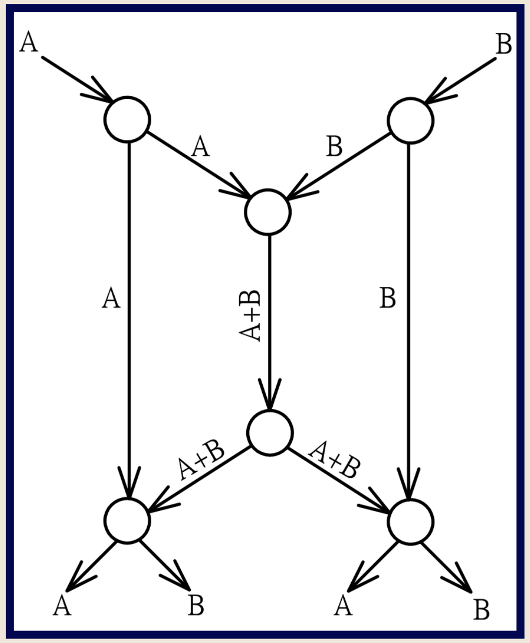

Network Coding

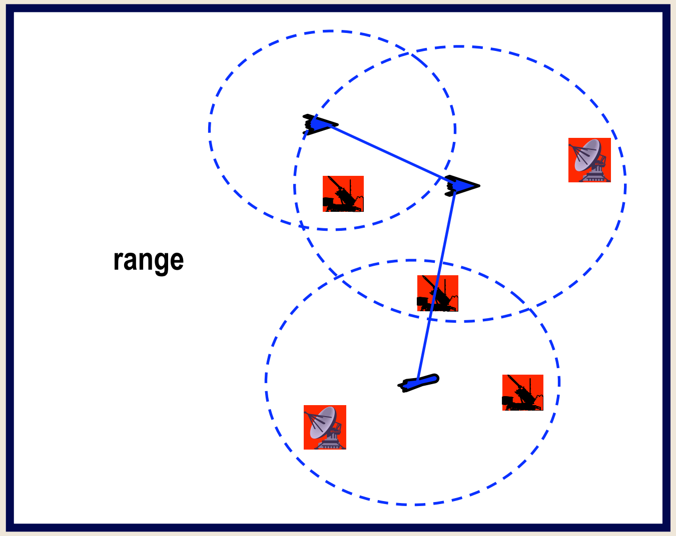

Vehicle Target Assignment

Features of distributed design:

Local decisions

Local information

Global behavior depends on local decisions

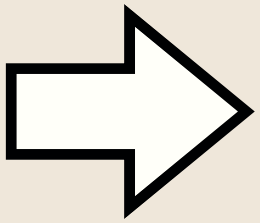

Game Theory

social system

Descriptive Agenda:

Modeling

model as "game"

game theory

Decision Makers → ***Global Behavior***

Metrics:

Reasonable description of sociocultural phenomena?

Matches available experimental/observational data?

social system

Prescriptive Agenda:

Distributed robust optimization

engineering system

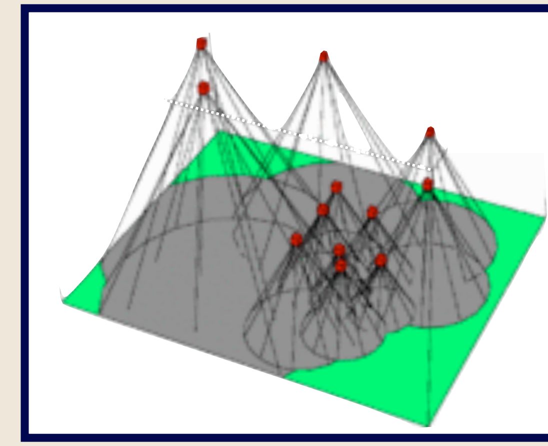

model as "game"

game theory

desired global behavior

Decision Makers

Global Behavior

distributed control

Metrics:

- Asymptotic global behavior?
- Communication/Information requirement?
- Computation requirement?
- Convergence rates?

Design parameters:

- Decision makers
- Objective/Utility functions
- Decision/Learning rule

Game theory for distributed robust optimization

Part #1: model interactions as game

decision makers / players
possible choices
local objective functions

Part #2: local agent decision rules

informational dependencies
processing requirements

Goal: Emergent global behavior is desirable

Appeal:

available distributed learning algorithms
robustness to uncertainties
self-interested users?

Challenges:

convergence rates?

Game theory for distributed robust optimization

Part #1:
model interactions as game

decision makers / players
possible choices
local objective functions

Part #2:
local agent decision rules

informational dependencies
processing requirements

Goal: Emergent global behavior is desirable

Appeal:

available distributed learning algorithms
robustness to uncertainties
self-interested users?

Challenges:

convergence rates?

Outline

Goal: Establish methodology for designing ***desirable*** utility functions

Existence of (pure) NE

Efficiency of NE

Locality of information

Tractability

Budget balance

Outline:

- Propose framework to study utility design: ***Distributed welfare games***
- Identify methodologies that guarantees desirable properties
- Identify ***fundamental limitations***
- Propose new framework to overcome limitations

Non-cooperative game:

- Players: $N = \{1, 2, ..., n\}$
- Actions: $a_i \in \mathcal{A}_i$
- Joint actions: $\mathcal{A} = \mathcal{A}_1 \times ... \times \mathcal{A}_n$
- Utilities: *(preferences)* $U_i : \mathcal{A} \rightarrow R$

  $U_i(a) = U_i(a_i, a_{-i})$

(Pure) Nash equilibrium:

$$U_i(a_i^*, a_{-i}^*) = \max_{a_i \in \mathcal{A}_i} U_i(a_i, a_{-i}^*)$$

Setup:

- *Resources:* $\mathcal{R}$
- *Players:* N
- *Actions:* $\mathcal{A}_i \subseteq 2^{\mathcal{R}}$
- *Welfare* $W^r : 2^N \rightarrow R^+$
- *Global Welfare:* $W(a) = \sum_r W^r(a^r)$

↑ **player set that chose resource r**

Game design = Utility design

Framework is common to many application domains

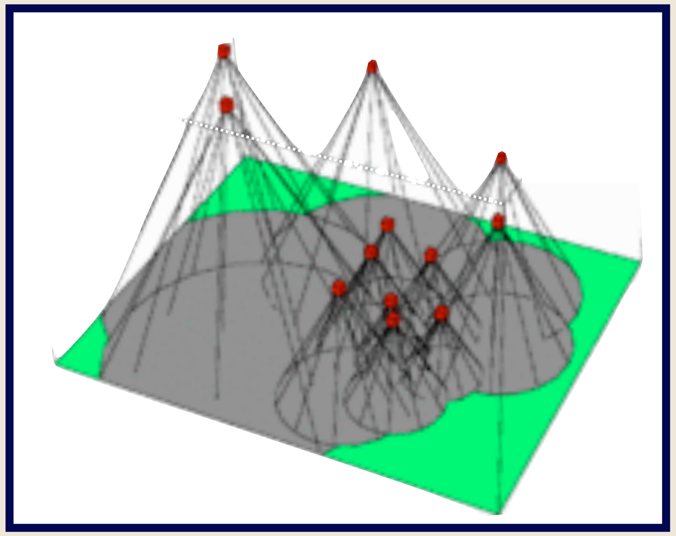

Sensor coverage

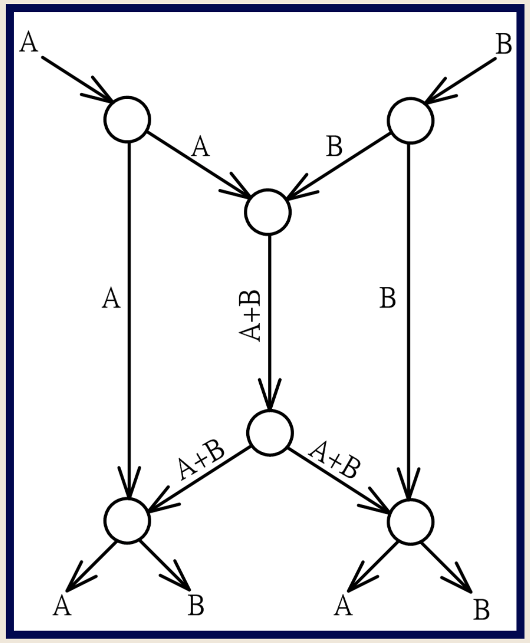

Network Coding

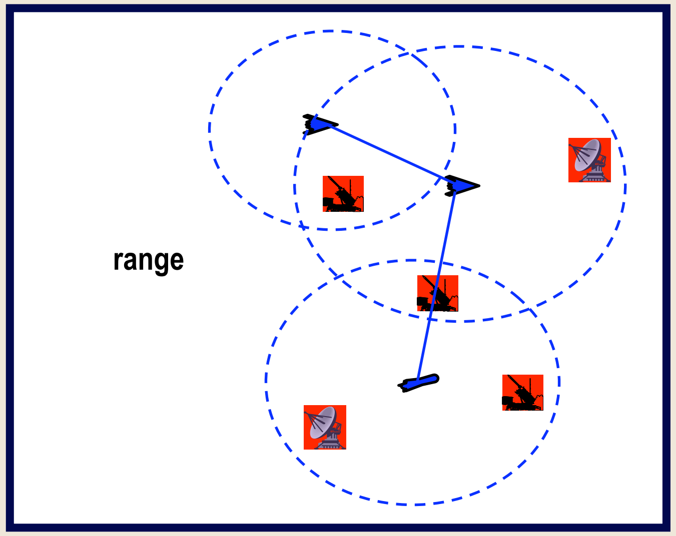

Vehicle Target Assignment

Akella et al., 2002. (Congestion control)
Goemans et al., 2004 (Content distribution)
Kesselman et al., 2005. (Switching/congestion control)
Komali and MacKenzie, 2007. (Topology control in ad-hoc networks)
Campos-Nanez et al., 2008. (Power management in sensor networks)

Example: Vehicle target assignment

***Resources**: Targets*

***Players**: Vehicles / Weapons*

***Actions**: Possible engagements*

***Welfare**: worth, expected damage and loss.*

Welfare
$W^r(1)$
$W^r(2)$
$W^r(3)$
$W^r(1,2)$
$W^r(1,3)$
$W^r(2,3)$
$W^r(1,2,3)$

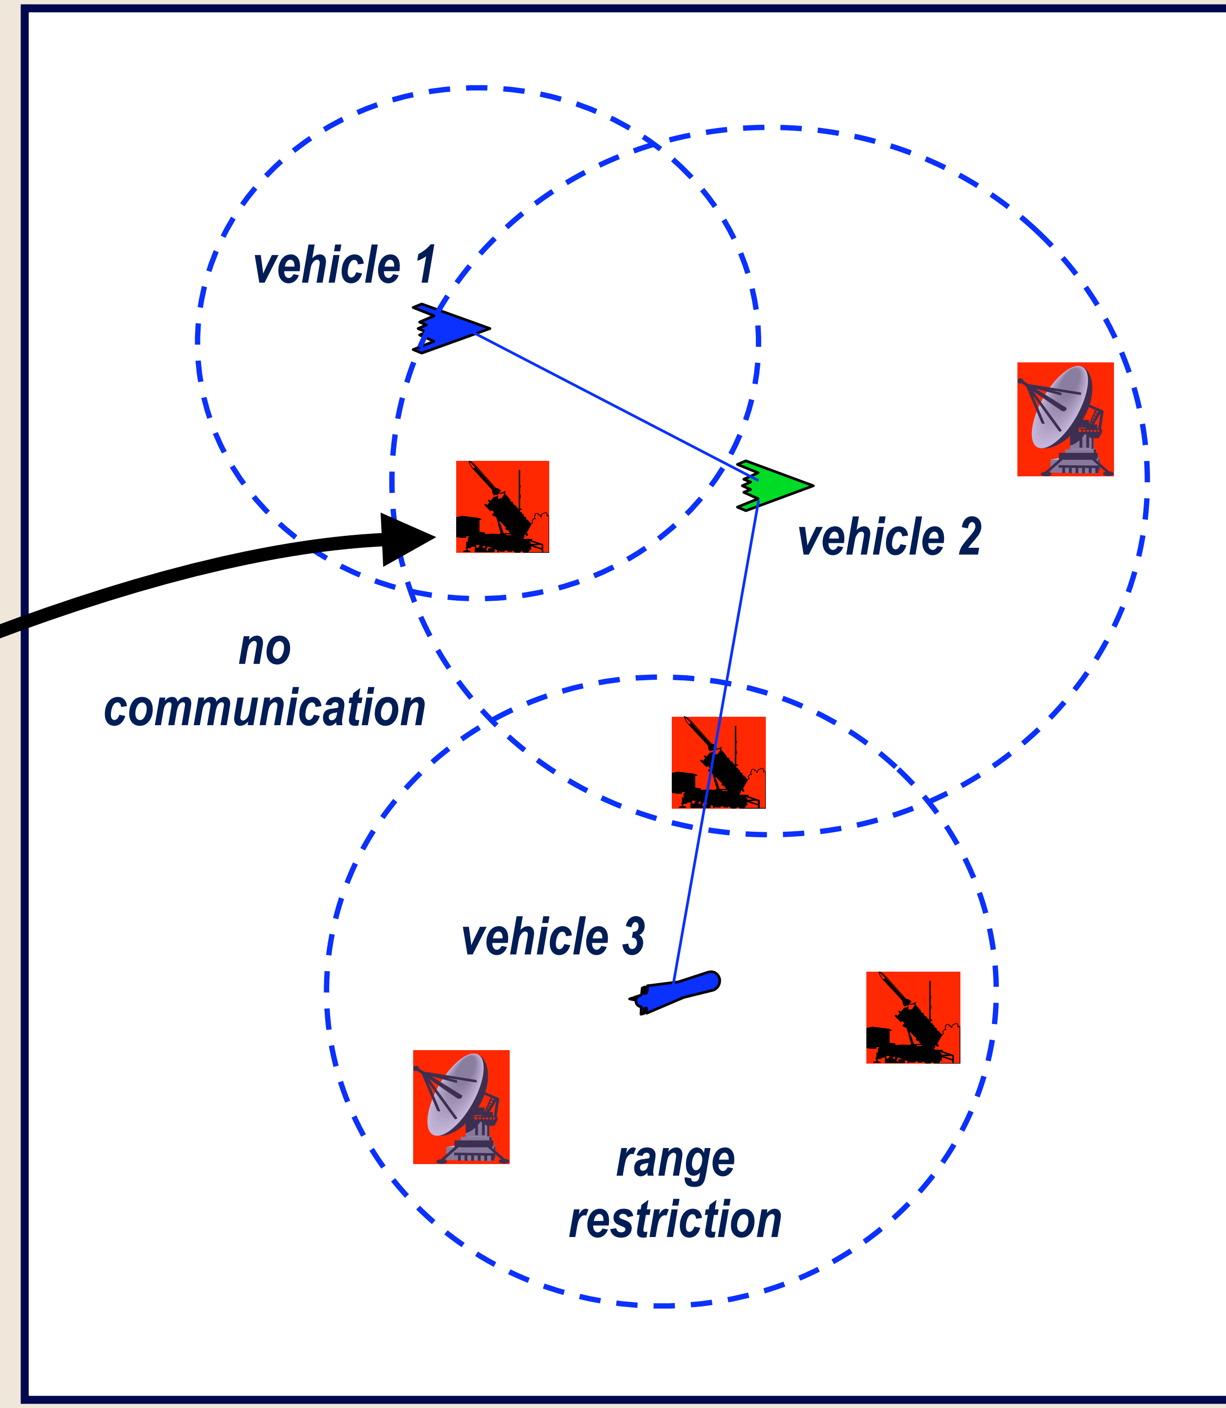

G. Arslan et al., "Autonomous vehicle-target assignment: a game theoretical formulation," 2007.

Example: Vehicle target assignment

Resources: *Targets*

Players: *Vehicles / Weapons*

Actions: *Possible engagements*

Welfare: *worth, expected damage and loss.*

Welfare
$W^r(1)$
$W^r(2)$
$W^r(3)$
$W^r(1,2)$
$W^r(1,3)$
$W^r(2,3)$
$W^r(1,2,3)$

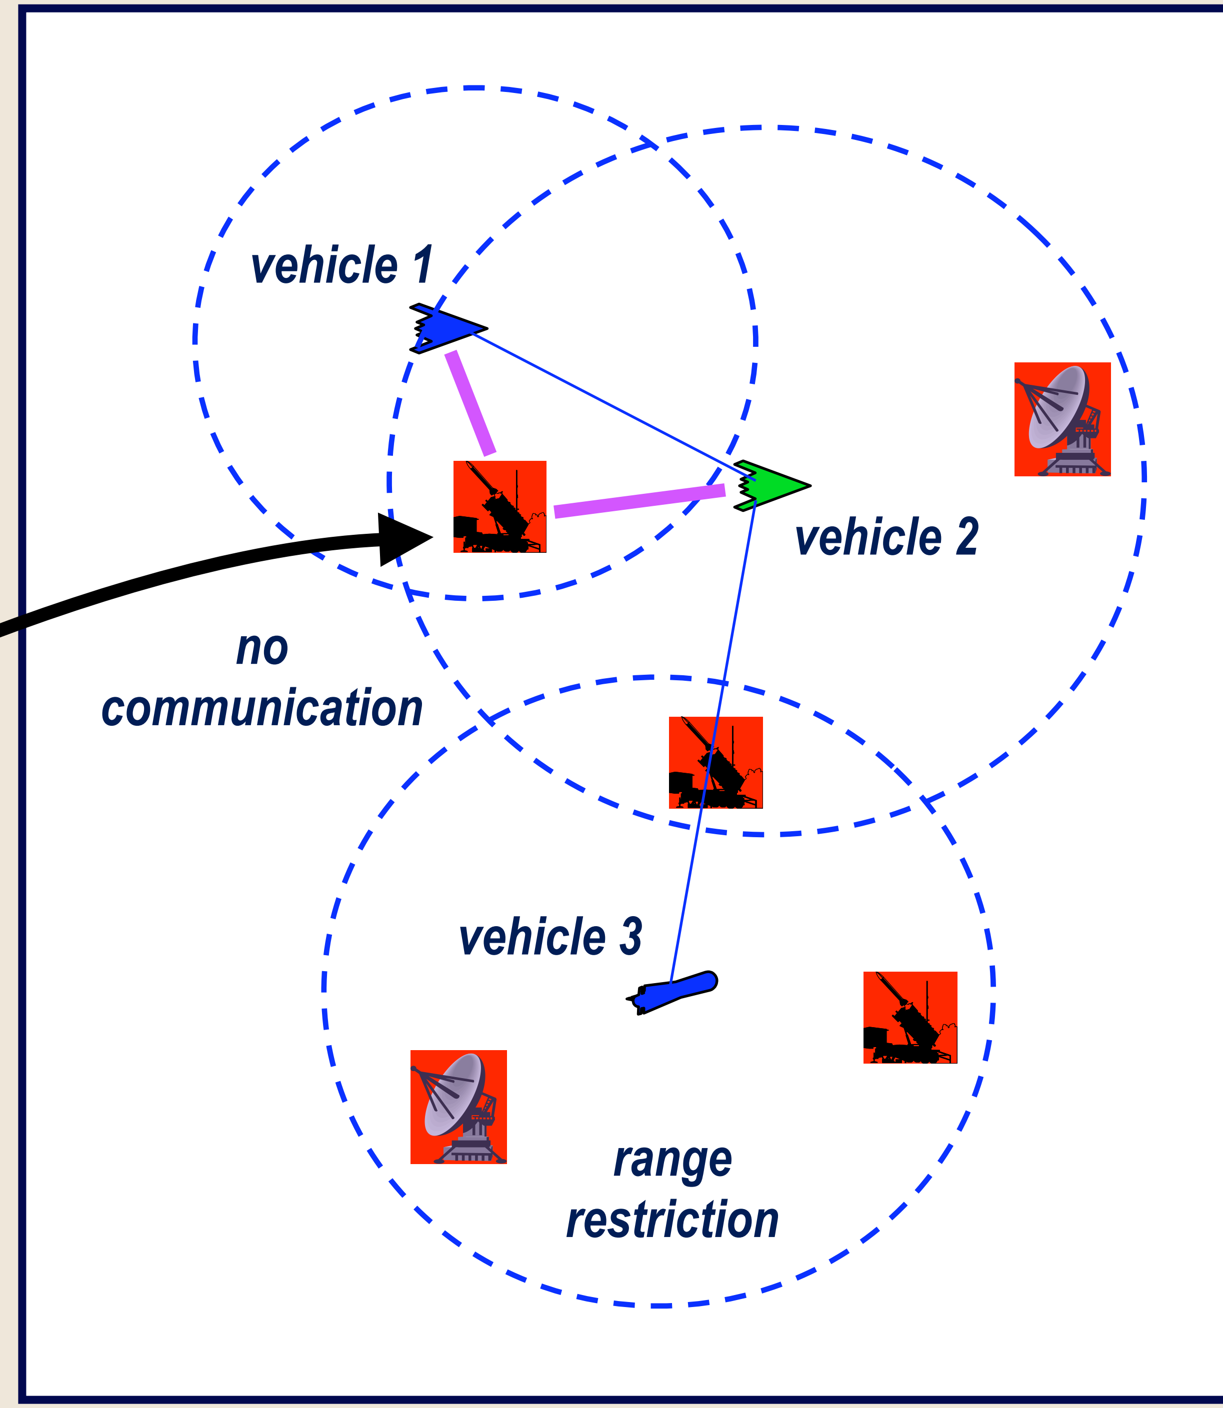

G. Arslan et al., "Autonomous vehicle-target assignment: a game theoretical formulation," 2007.

Example: Vehicle target assignment

Resources: *Targets*

Players: *Vehicles / Weapons*

Actions: *Possible engagements*

Welfare: *worth, expected damage and loss.*

Welfare
$W^r(1)$
$W^r(2)$
$W^r(3)$
$W^r(1,2)$
$W^r(1,3)$
$W^r(2,3)$
$W^r(1,2,3)$

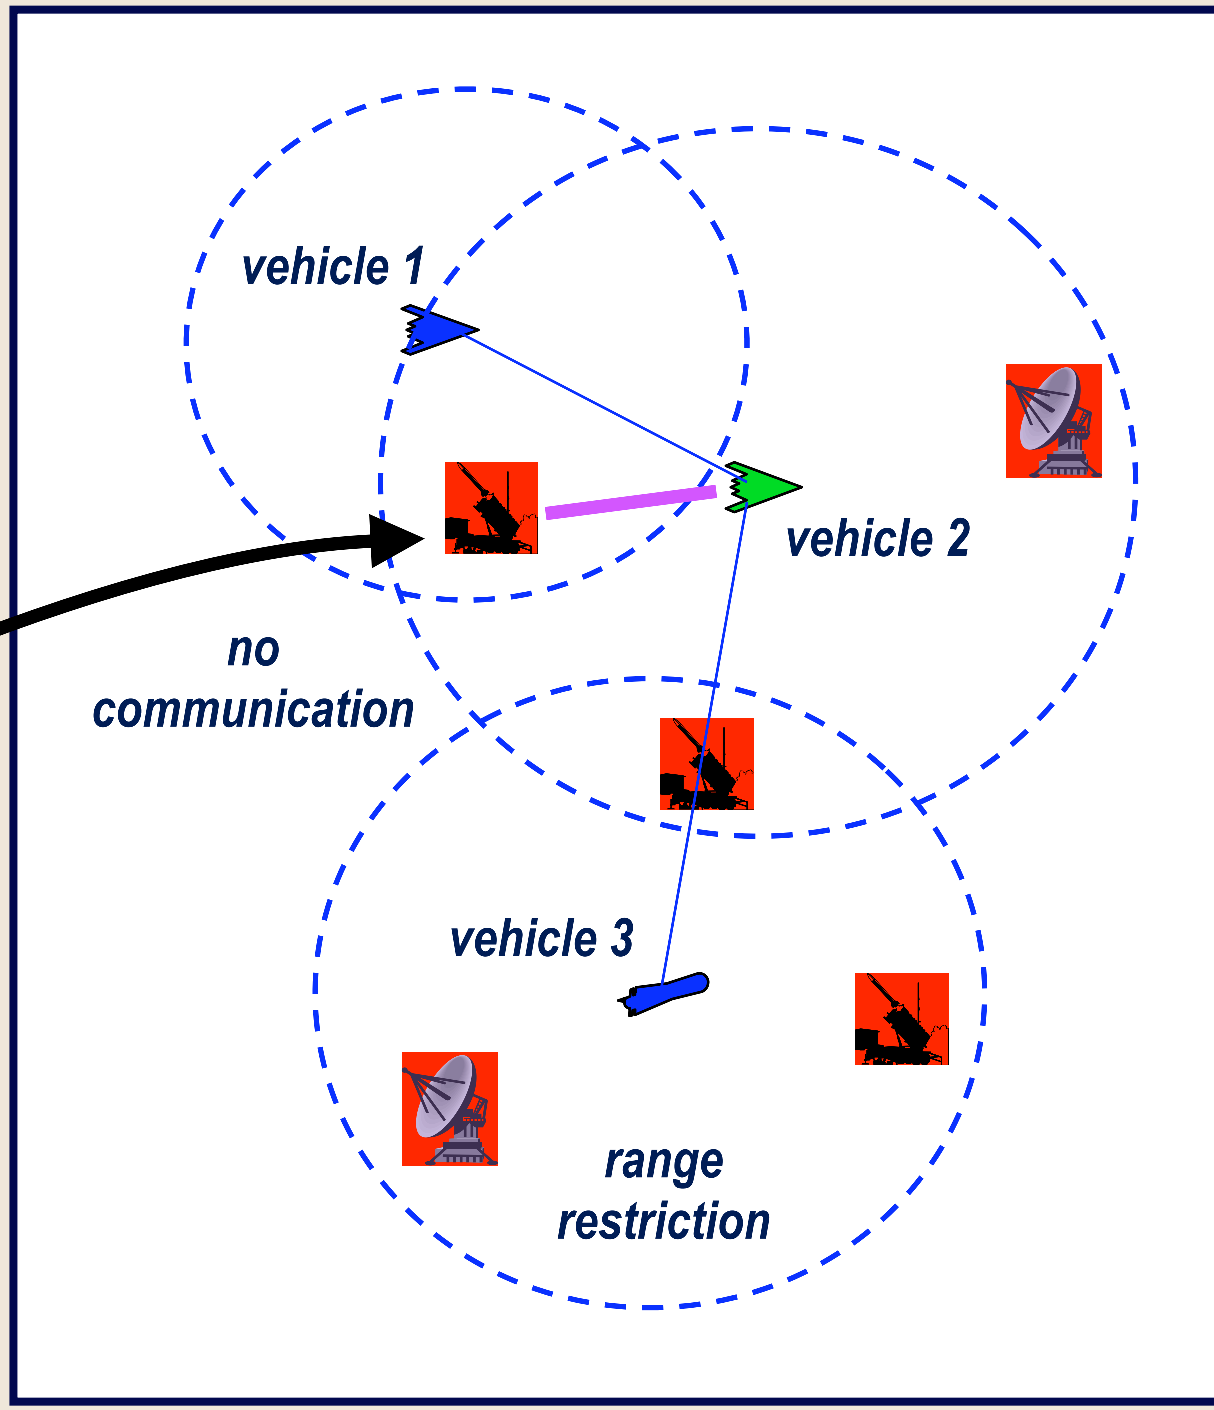

G. Arslan et al., "Autonomous vehicle-target assignment: a game theoretical formulation," 2007.

Example: Vehicle target assignment

***Resources**: Targets*

***Players**: Vehicles / Weapons*

***Actions**: Possible engagements*

***Welfare**: worth, expected damage and loss.*

Welfare
$W^r(1)$
$W^r(2)$
$W^r(3)$
$W^r(1,2)$
$W^r(1,3)$
$W^r(2,3)$
$W^r(1,2,3)$

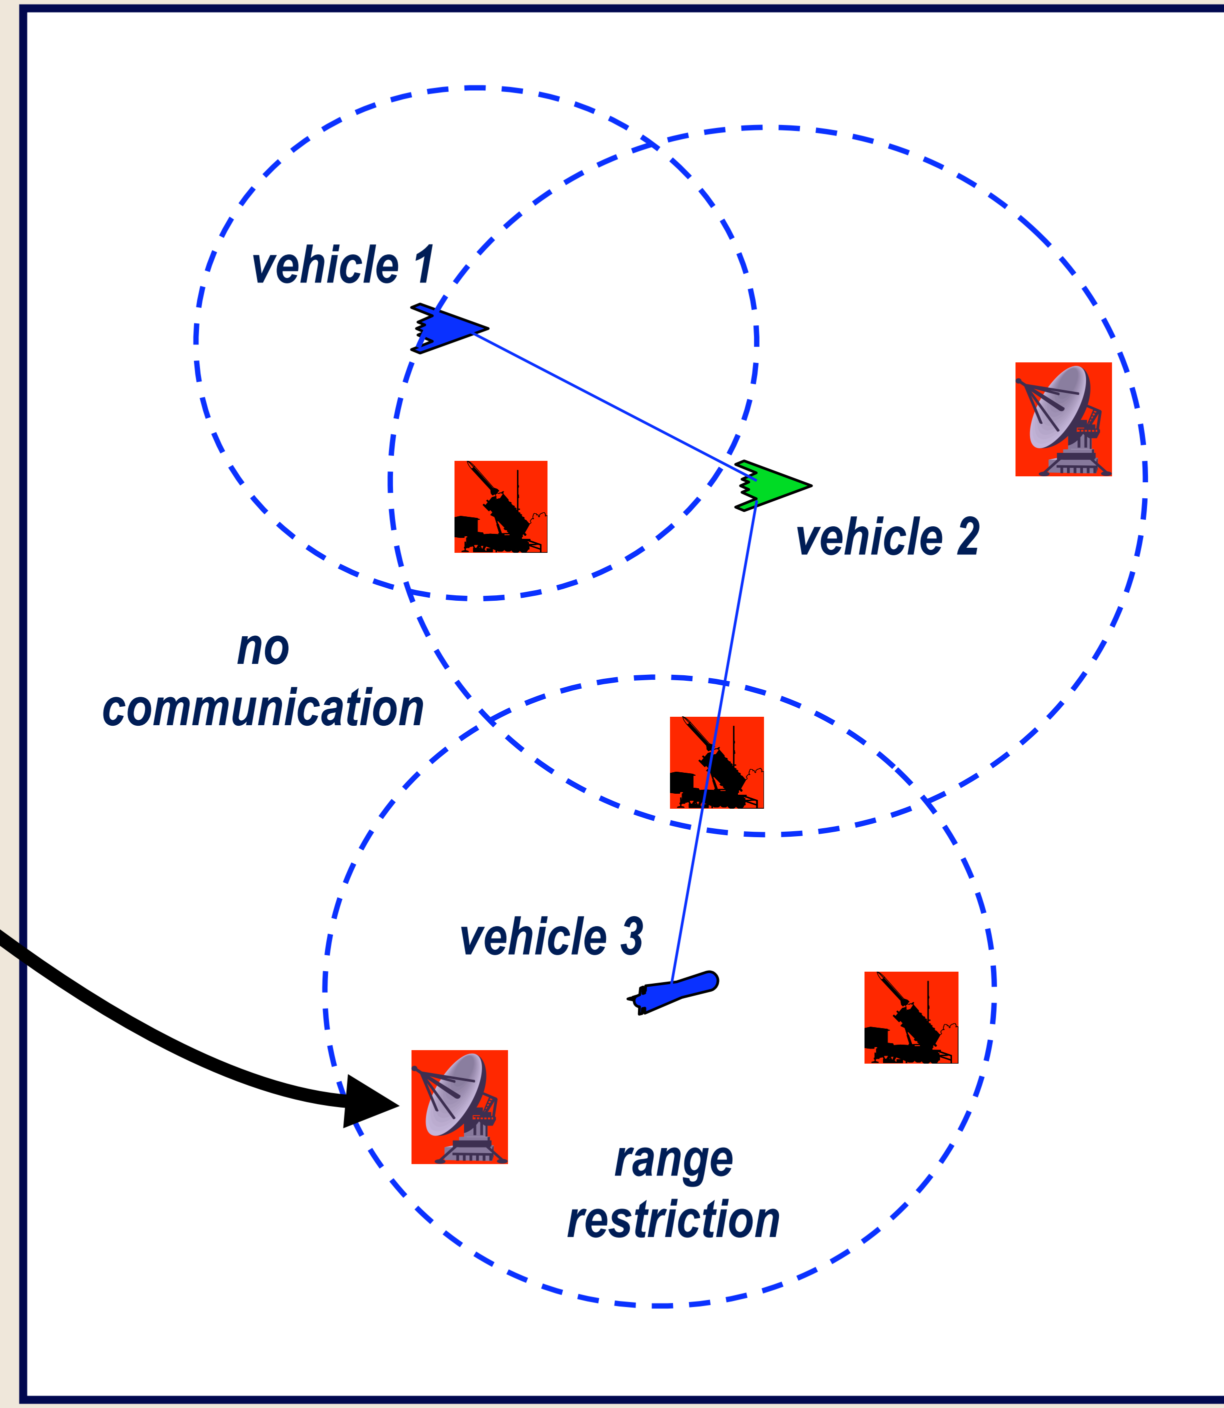

G. Arslan et al., "Autonomous vehicle-target assignment: a game theoretical formulation," 2007.

Example: Vehicle target assignment

Resources: *Targets*

Players: *Vehicles / Weapons*

Actions: *Possible engagements*

Welfare: *worth, expected damage and loss.*

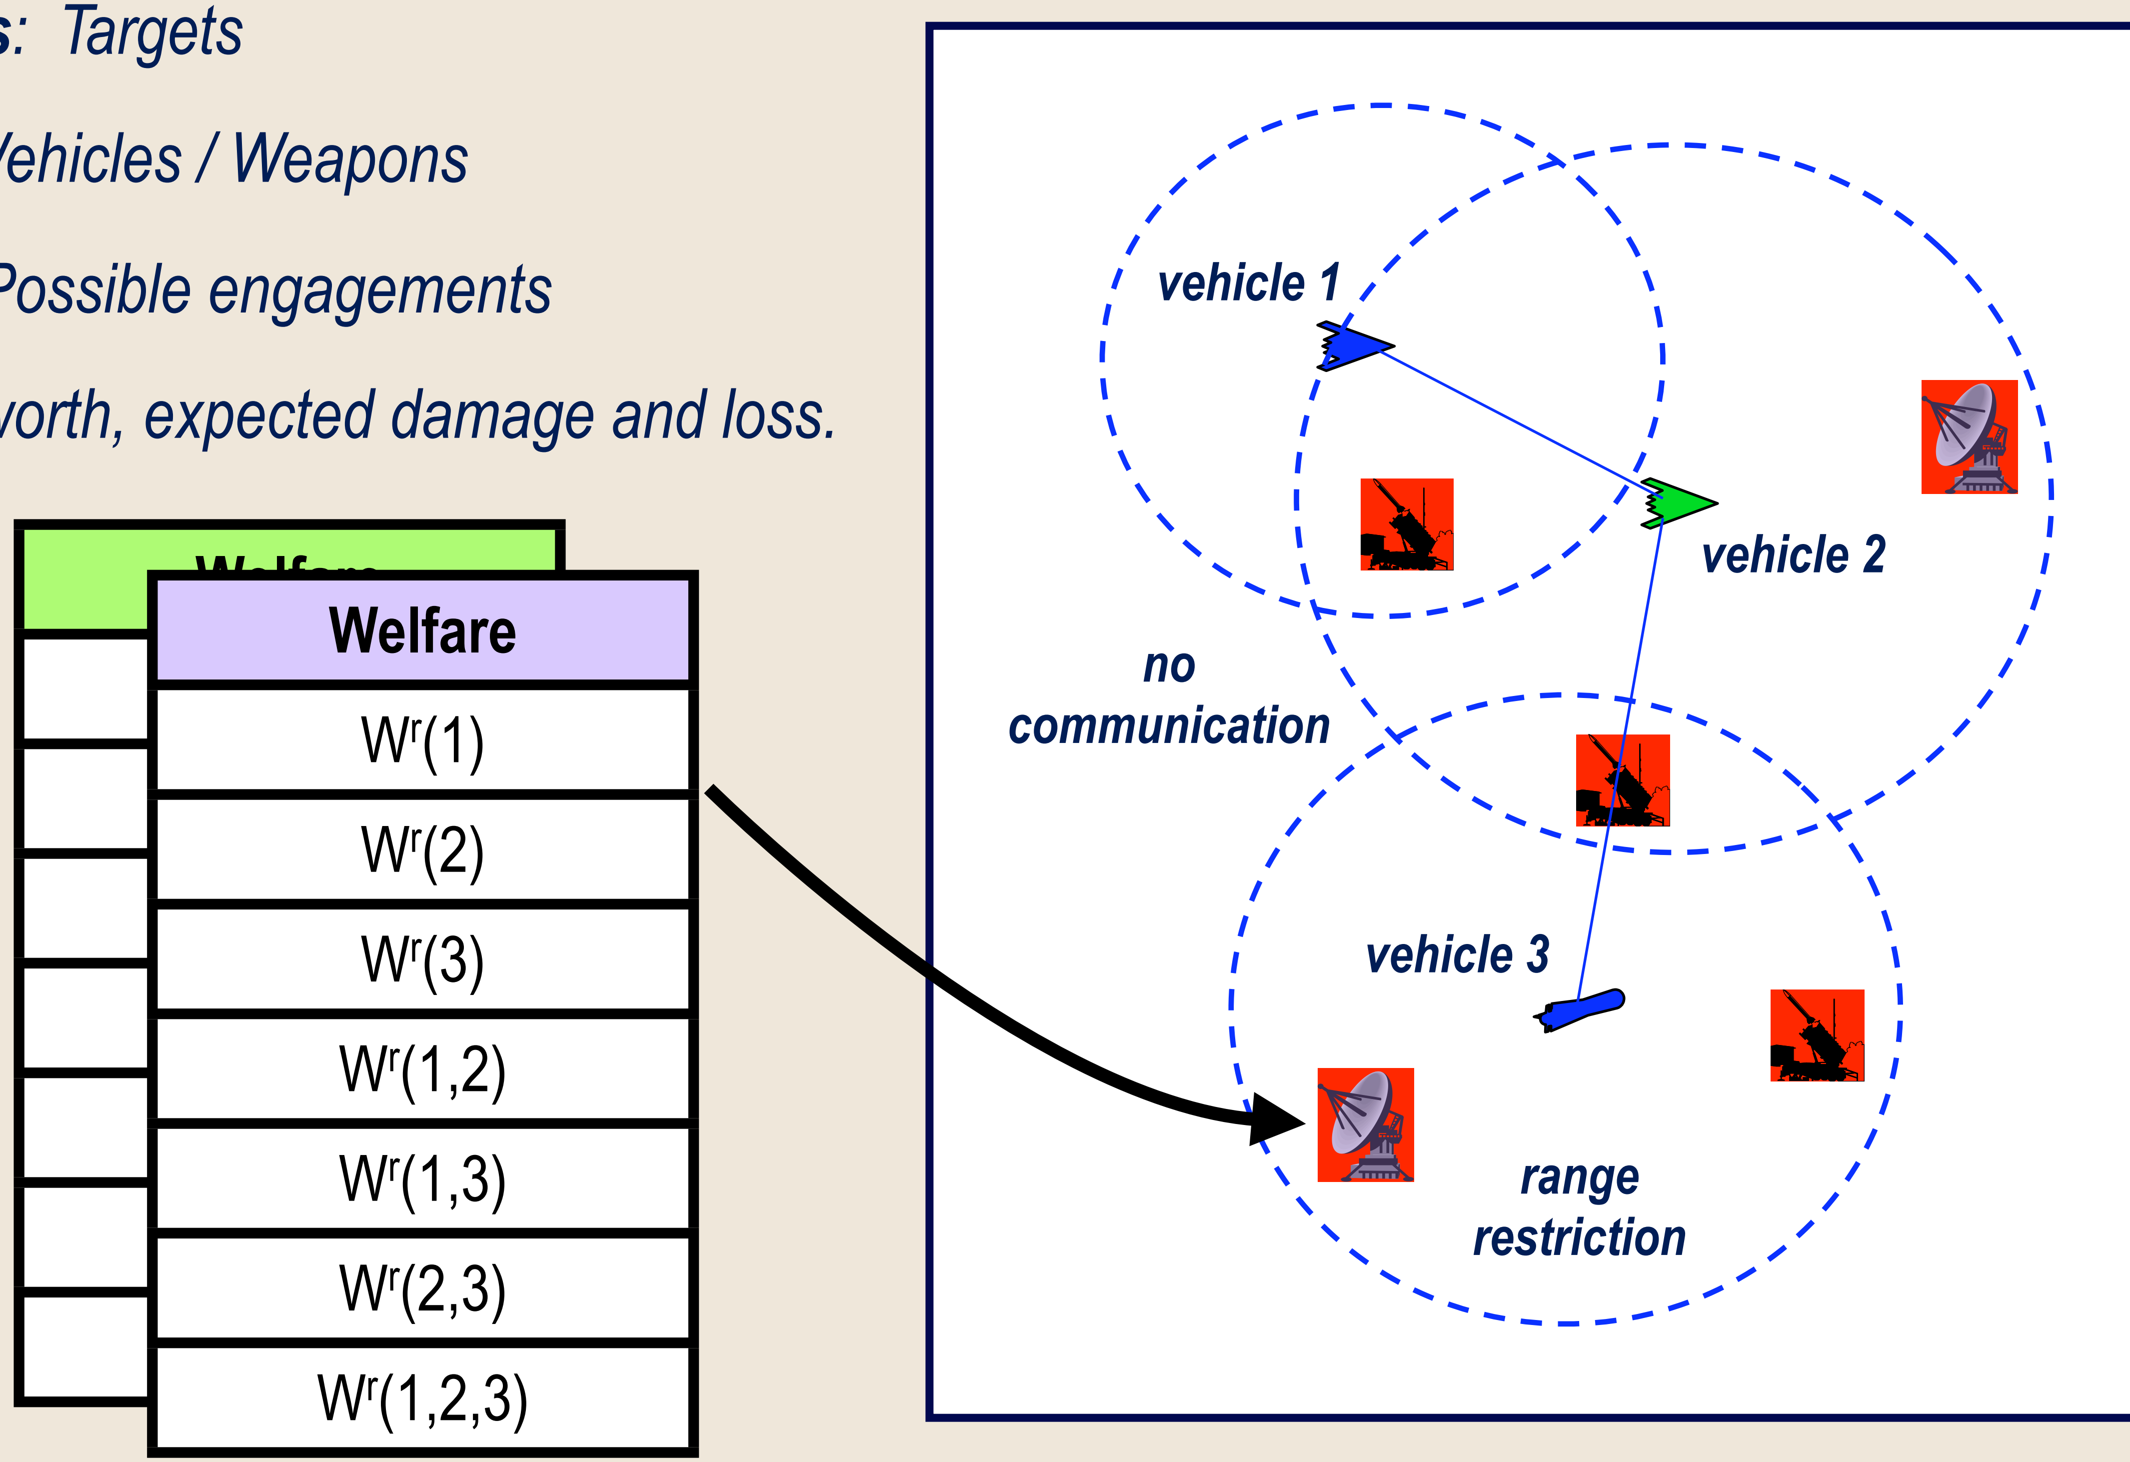

Global objective: *Maximize sum of welfare (centralized assignment not feasible)*

G. Arslan et al., "Autonomous vehicle-target assignment: a game theoretical formulation," 2007.

Goal: *Assign each agent a utility such that the resulting game is desirable*

- Existence of NE
- Efficiency of NE
- Locality of information
- Tractability
- Budget balance

Approach: *View like a cost sharing problem*

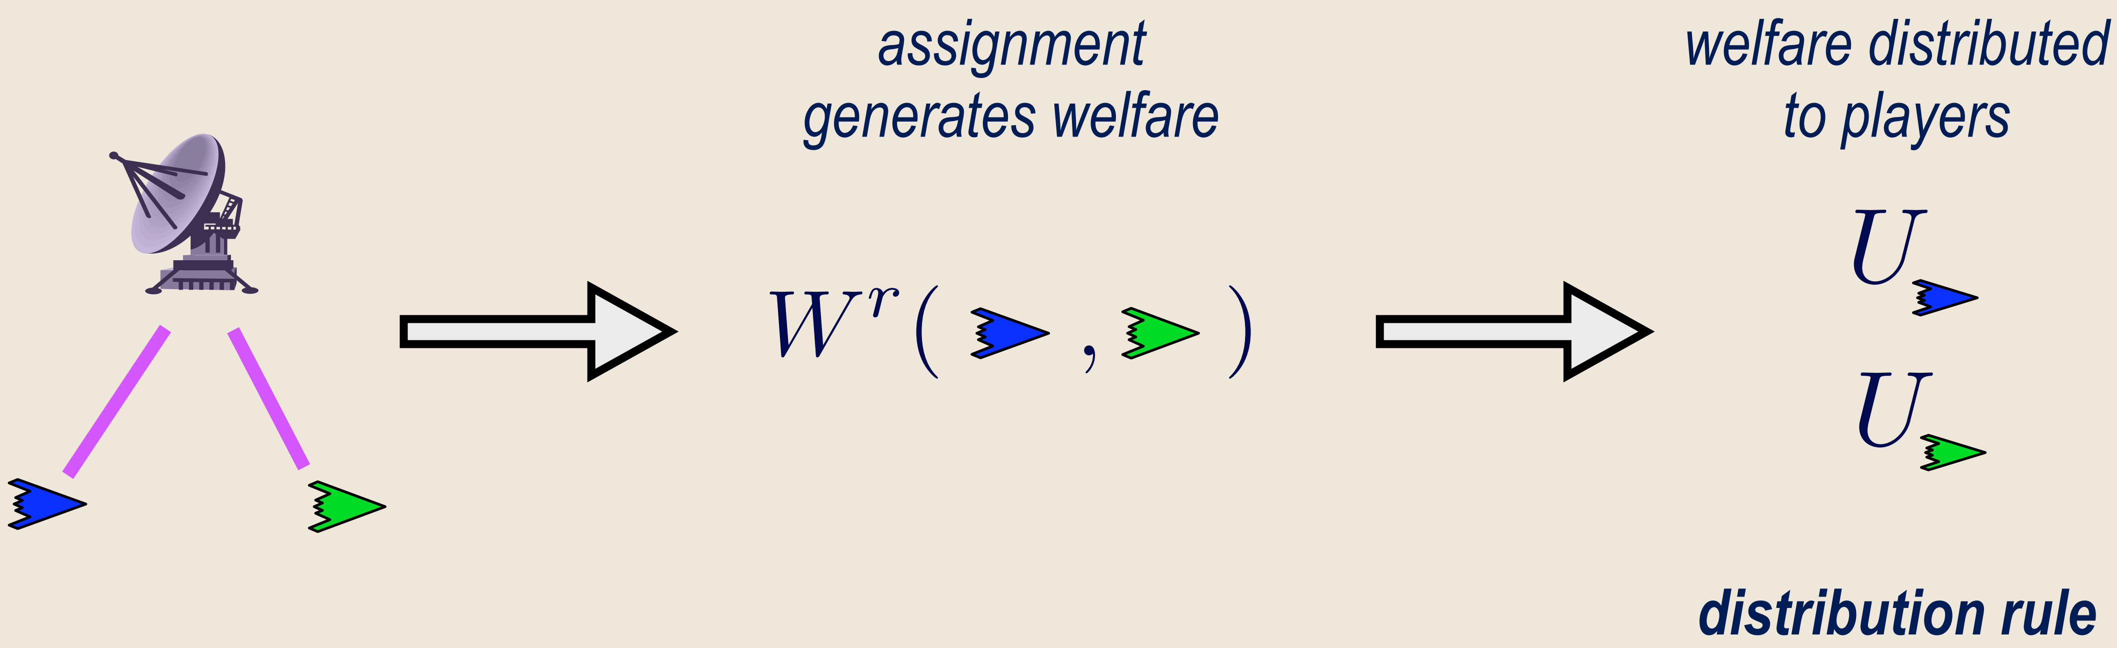

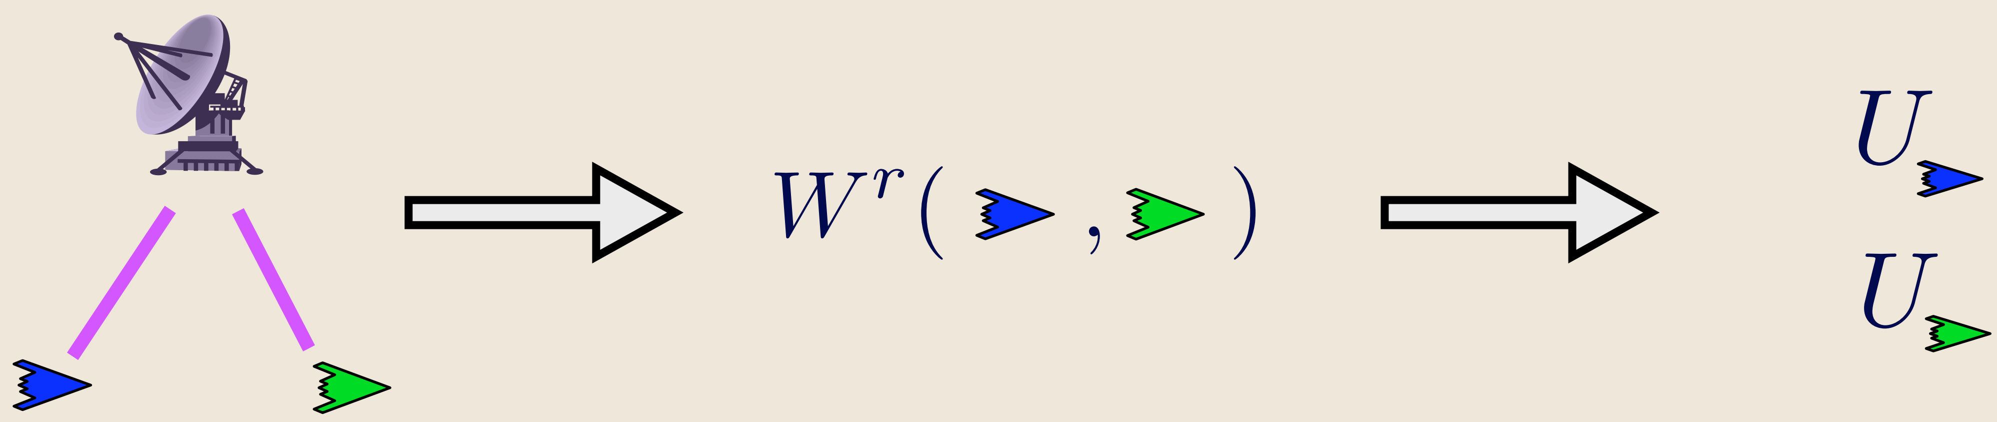

distribution rule

Utility structure:

$$U_i(a) = \sum_{r \in a_i} f^r(i, a^r) W^r(a^r)$$

depends only on local information

Properties of distribution rule:

1. $f^r(i, a^r) \geq 0$
2. $r \notin a_i \Rightarrow f^r(i, a^r) = 0$
3. $\sum_i f^r(i, a^r) \leq 1$

Budget Balanced:

$$W(a) = \sum U_i(a)$$

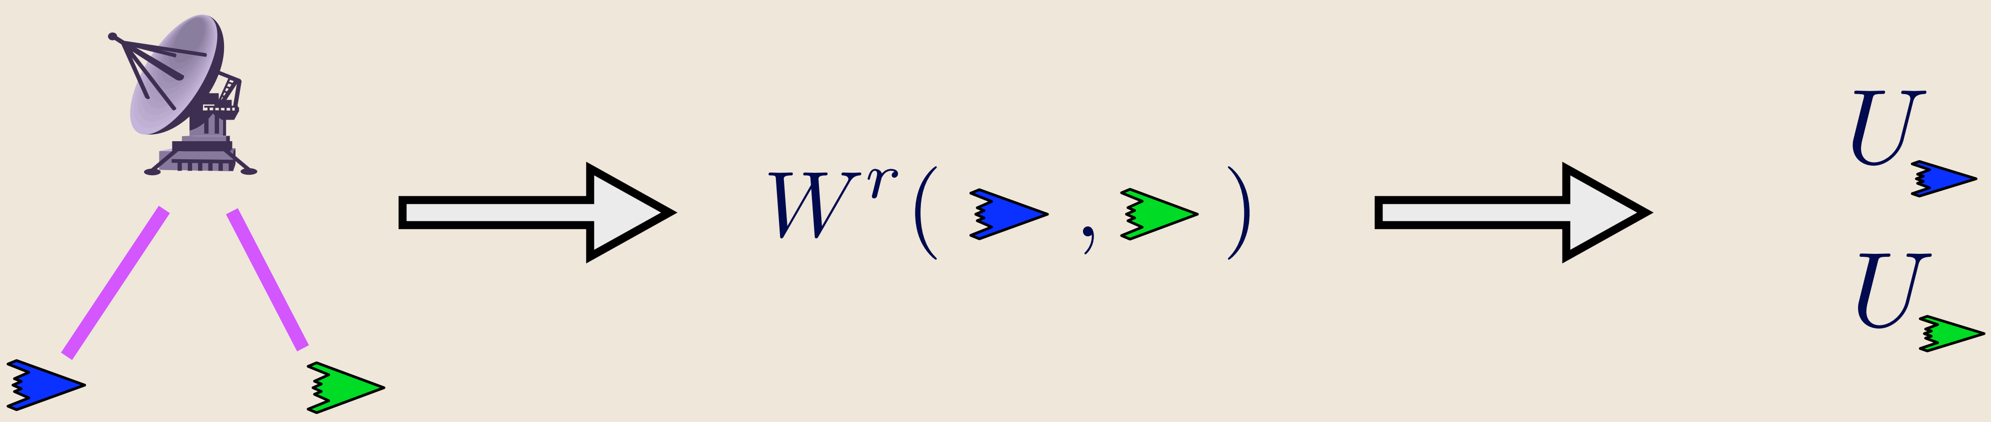

distribution rule

Utility structure:

$$U_i(a) = \sum_{r \in a_i} f^r(i, a^r) W^r(a^r)$$

Properties of distribution rule:

1. $f^r(i, a^r) \geq 0$
2. $r \notin a_i \Rightarrow f^r(i, a^r) = 0$
3. $\sum_i f^r(i, a^r) \leq 1$

Are cost sharing methodologies useful in designing utilities?

$$U_i(a_i, a_{-i}) = \sum_{r \in a_i} \frac{1}{|a^r|} W^r(a^r)$$

	NE exists	Budget Balanced	Complexity
Equal share		✓	**low**

***** If welfare function is anonymous, then NE exists.*** *(Monderer and Shapley, 1996)*

$$W^r(a^r) = W^r(|a^r|)$$

Marginal contribution

$$U_i(a_i, a_{-i}) = \sum_{r \in a_i} W^r(a^r) - W^r(a^r \setminus i)$$

	NE exists	Budget Balanced	Complexity
Equal share		✓	**low**
Marginal contribution	✓		**medium**

(Wolpert and Tumor, 1999)

$$U_i(a_i, a_{-i}) = \sum_{r \in a_i} Sh^r(i, a^r)$$

	NE exists	Budget Balanced	Complexity
Equal share		✓	**low**
Marginal contribution	✓		**medium**
Shapley value	✓	✓	**high**

(builds upon Hart and Mas-Collell, 1989)

$$U_i(a_i, a_{-i}) = \sum_{r \in a_i} Sh^r(i, a^r)$$

$$Sh^r(i, N) = \sum_{S \subseteq N : i \in S} \omega_S \left(W^r(S) - W^r(S \setminus i)\right)$$

summation over
all player subset

marginal contribution
to player subset

intractable for large N

Summary

	NE exists	Budget Balanced	Complexity
Equal share		✓	***low***
Marginal contribution	✓		***medium***
Shapley value	✓	✓	***high***

Tradeoff: Properties vs. Complexity

Is there anything else?

No, (weighted) SV only rule that guarantees NE + BB in all games.

[Chen, Roughgarden & Valiant, 2008]: Network formation games (uniform)

Yes if we restrict attention to special classes of games

[JRM & Wierman, 2008]: Not restricted to SV in some settings

Efficiency

Can we provide efficiency guarantees for general welfare functions?

Price of Anarchy

$$POA = \inf_{G} \min_{a^{\mathrm{ne}} \in G} \frac{W(a^{\mathrm{ne}})}{W(a^{\mathrm{opt}})}$$

worst case performance of any NE

Price of Stability

$$POS = \inf_{G} \max_{a^{\mathrm{ne}} \in G} \frac{W(a^{\mathrm{ne}})}{W(a^{\mathrm{opt}})}$$

worst case performance of best NE

(independent of number of game specifics)

No. In general a NE can be arbitrarily bad.

Yes if welfare is ***submodular*** (decreasing marginal welfare)

- Submodularity (decreasing marginal welfare)

$$W(S+s) - W(S) \geq W(S'+s) - W(S') \qquad S \subset S' \subset N$$

- Submodularity can be exploited to improve efficiency

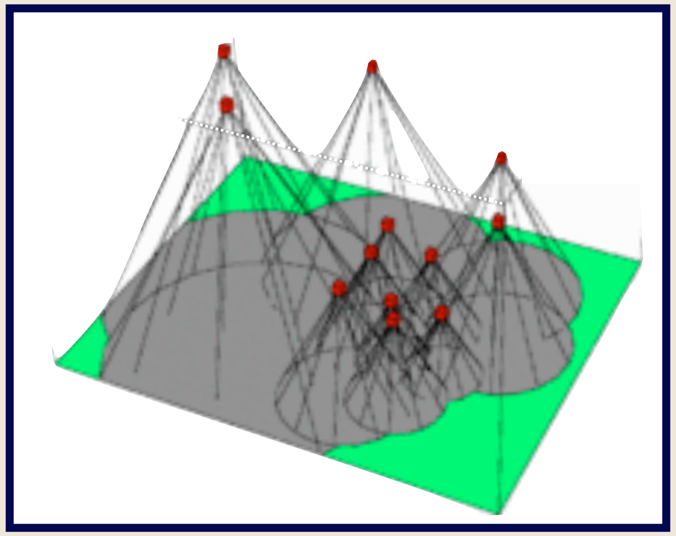

Sensor coverage

Andreas Krause (Caltech)

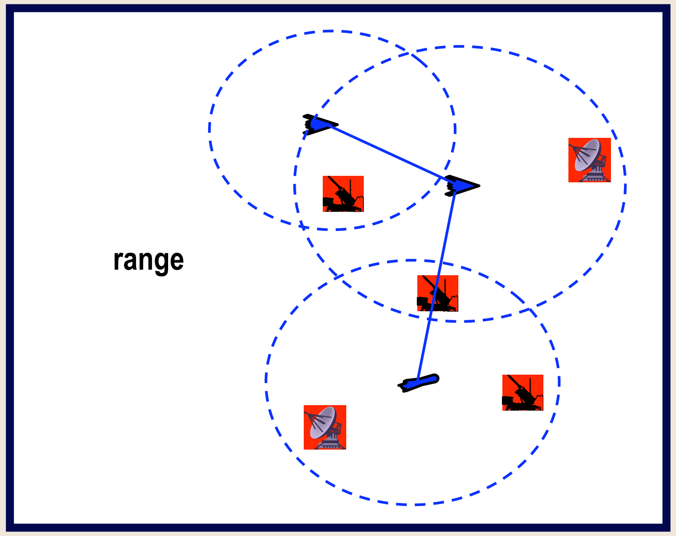

Vehicle Target Assignment

- Submodularity (decreasing marginal welfare)

$$W(S+s) - W(S) \geq W(S'+s) - W(S') \qquad S \subset S' \subset N$$

- Submodularity can be exploited to improve efficiency

Theorem: *For* ***any*** *distributed welfare game where*

[JRM & Wierman, 2008]
[Vetta, 2002]

(i) Resource specific welfare functions are submodular

(ii) Utilities are greater than or equal to marginal contribution

$$U_i(a_i, a_{-i}) \geq W(a_i, a_{-i}) - W(\emptyset, a_{-i})$$

then if a NE exists, the price of anarchy is $\geq$ *1/2, i.e.,*

$$\frac{W(a^{\text{ne}})}{W(a^{\text{opt}})} \geq \frac{1}{2}$$

Efficiency

	NE exists	*Budget Balanced*	*Complexity*	***POS***	***POA***
Marginal contribution	✓		***medium***		**1/2**
Shapley value	✓	✓	***high***		**1/2**

Theorem: *For **any** distributed welfare game where*

[JRM & Wierman, 2008]
[Vetta, 2002]

(i) Resource specific welfare functions are submodular

(ii) Utilities are greater than or equal to marginal contribution

$$U_i(a_i, a_{-i}) \geq W(a_i, a_{-i}) - W(\emptyset, a_{-i})$$

then if a NE exists, the price of anarchy is $\geq$ *1/2, i.e.,*

$$\frac{W(a^{\text{ne}})}{W(a^{\text{opt}})} \geq \frac{1}{2}$$

Efficiency

	NE exists	*Budget Balanced*	*Complexity*	***POS***	***POA***
Marginal contribution	✓		***medium***	**1**	**1/2**
Shapley value	✓	✓	***high***	**?**	**1/2**

Best known ***centralized*** approximation algorithms: (1-1/e) = 0.63

What about price of stability?

	NE exists	*Budget Balanced*	*Complexity*	***POS***	***POA***
Marginal contribution	✓		***medium***	**1**	**1/2**
Shapley value	✓	✓	***high***	**?**	**1/2**

Best known ***centralized*** approximation algorithms: (1-1/e) = 0.63

What about price of stability?

Fundamental Limitation: [JRM & Wierman, 2009]

Existence of NE
Budget balance
→
POS < 1
POS = 1/2 *(submodular)*

Direction: distribution rule → game (POS=1)

Submodular welfare functions of the form $W^r(a^r) = c$ for all $a^r \neq \emptyset$

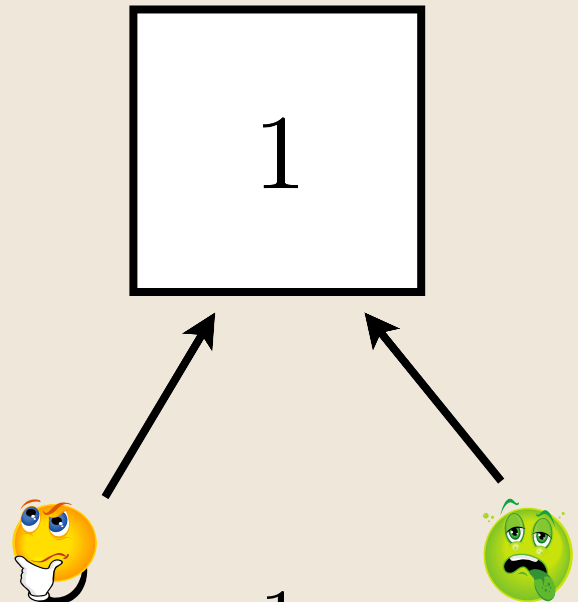

$$x \geq \frac{1}{2} \geq y$$

Direction: distribution rule → game (POS=1)

Submodular welfare functions of the form $W^r(a^r) = c$ for all $a^r \neq \emptyset$

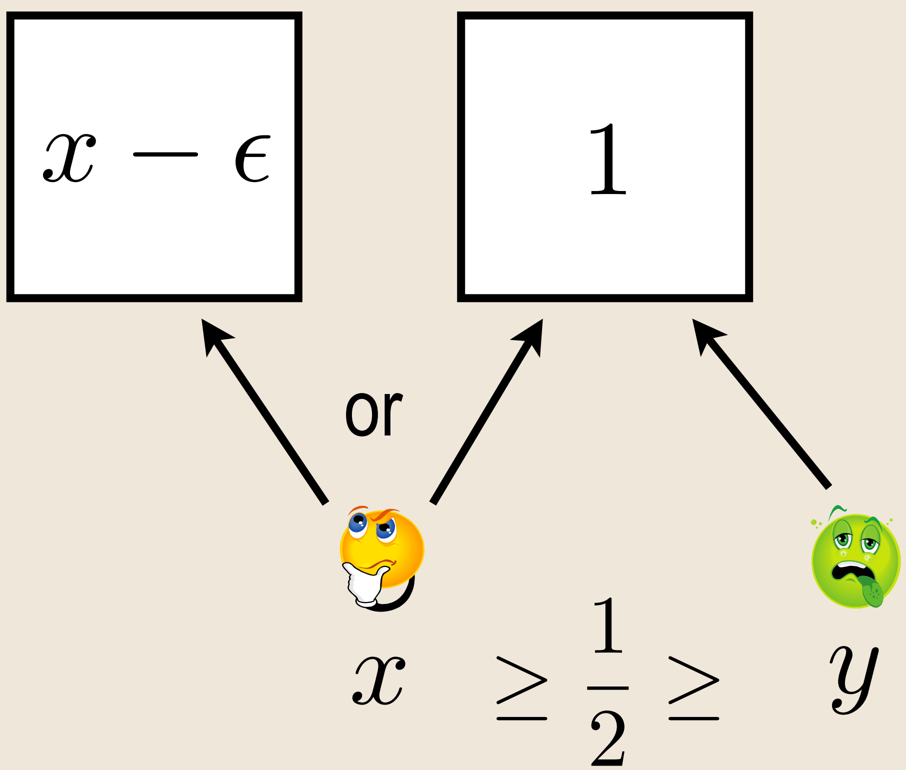

Direction: distribution rule game (POS=1)

Submodular welfare functions of the form $W^r(a^r) = c$ for all $a^r \neq \emptyset$

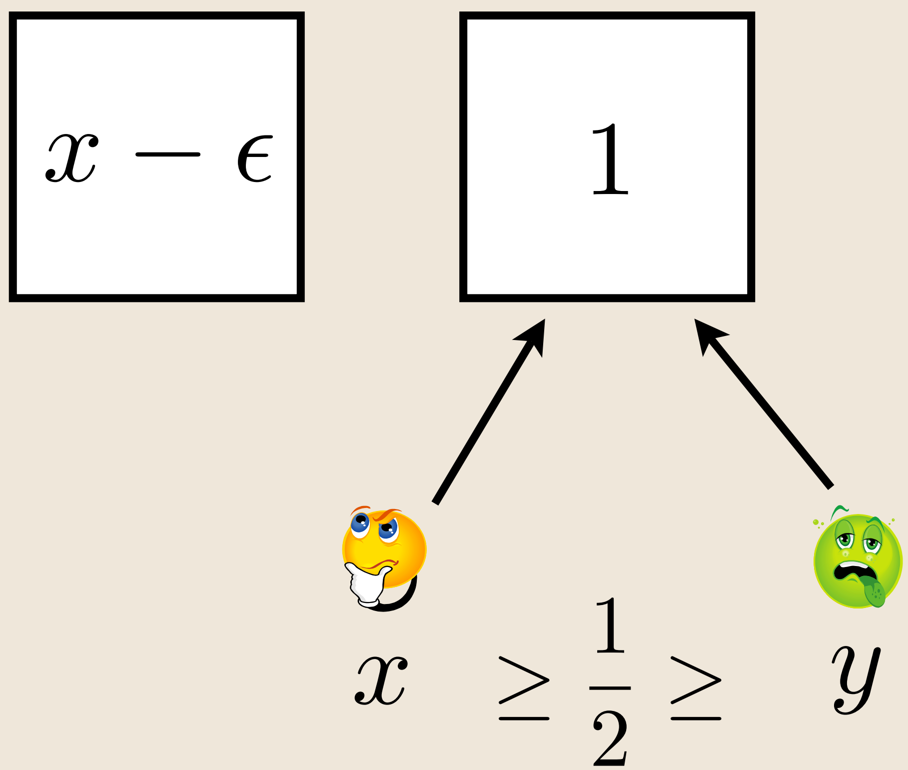

Unique NE

W=1

Proof

Direction: distribution rule → game (POS=1)

Submodular welfare functions of the form $W^r(a^r) = c$ for all $a^r \neq \emptyset$

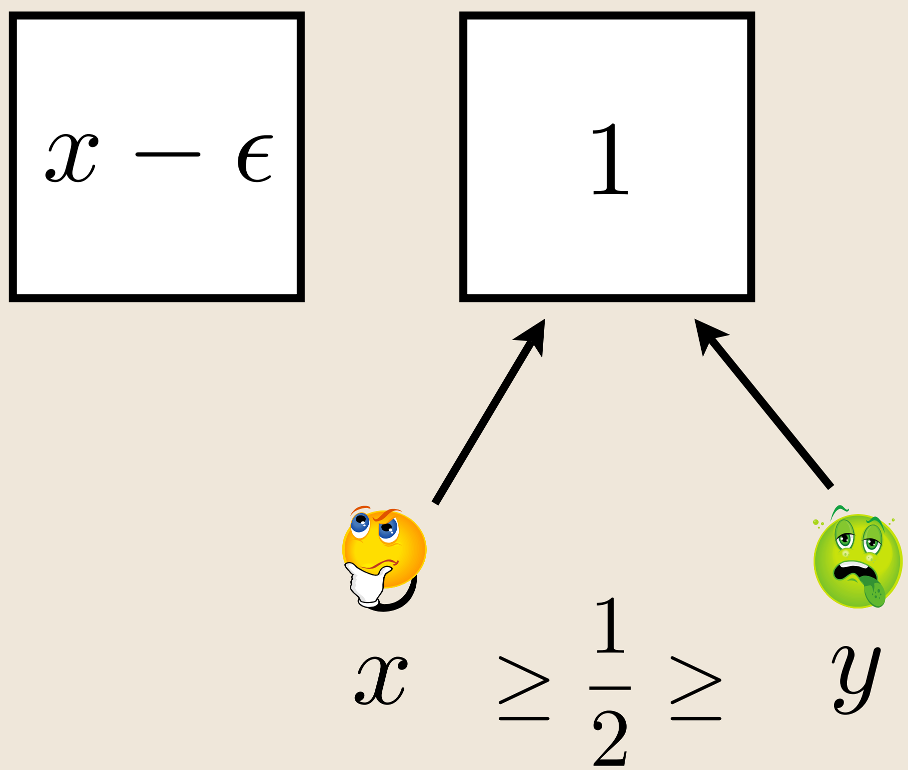

Unique NE

W=1

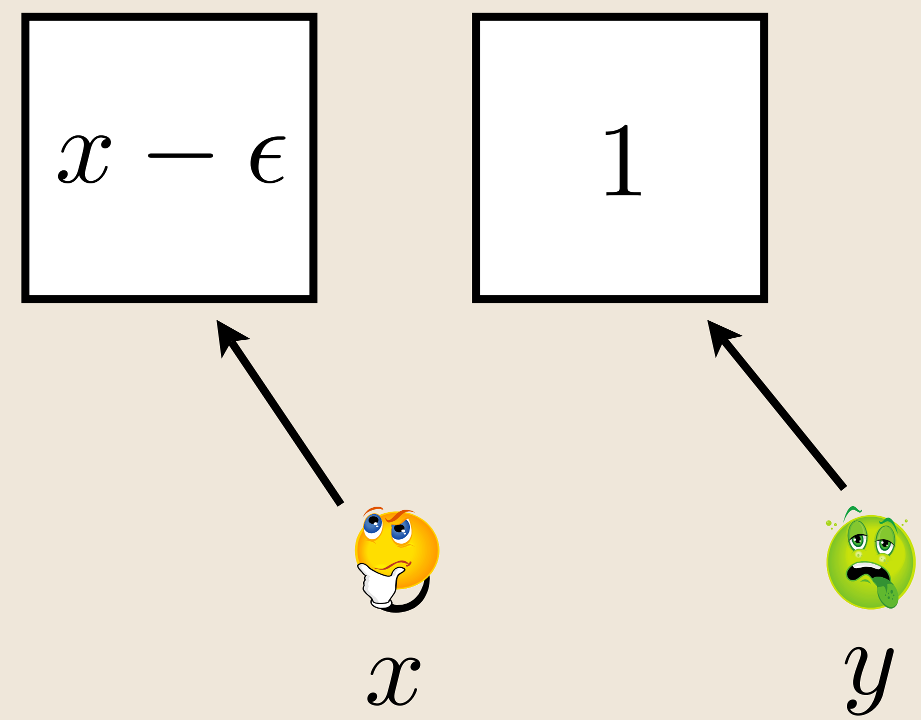

Optimal

W=1+x

Direction: distribution rule → game (POS=1)

Submodular welfare functions of the form $W^r(a^r) = c$ for all $a^r \neq \emptyset$

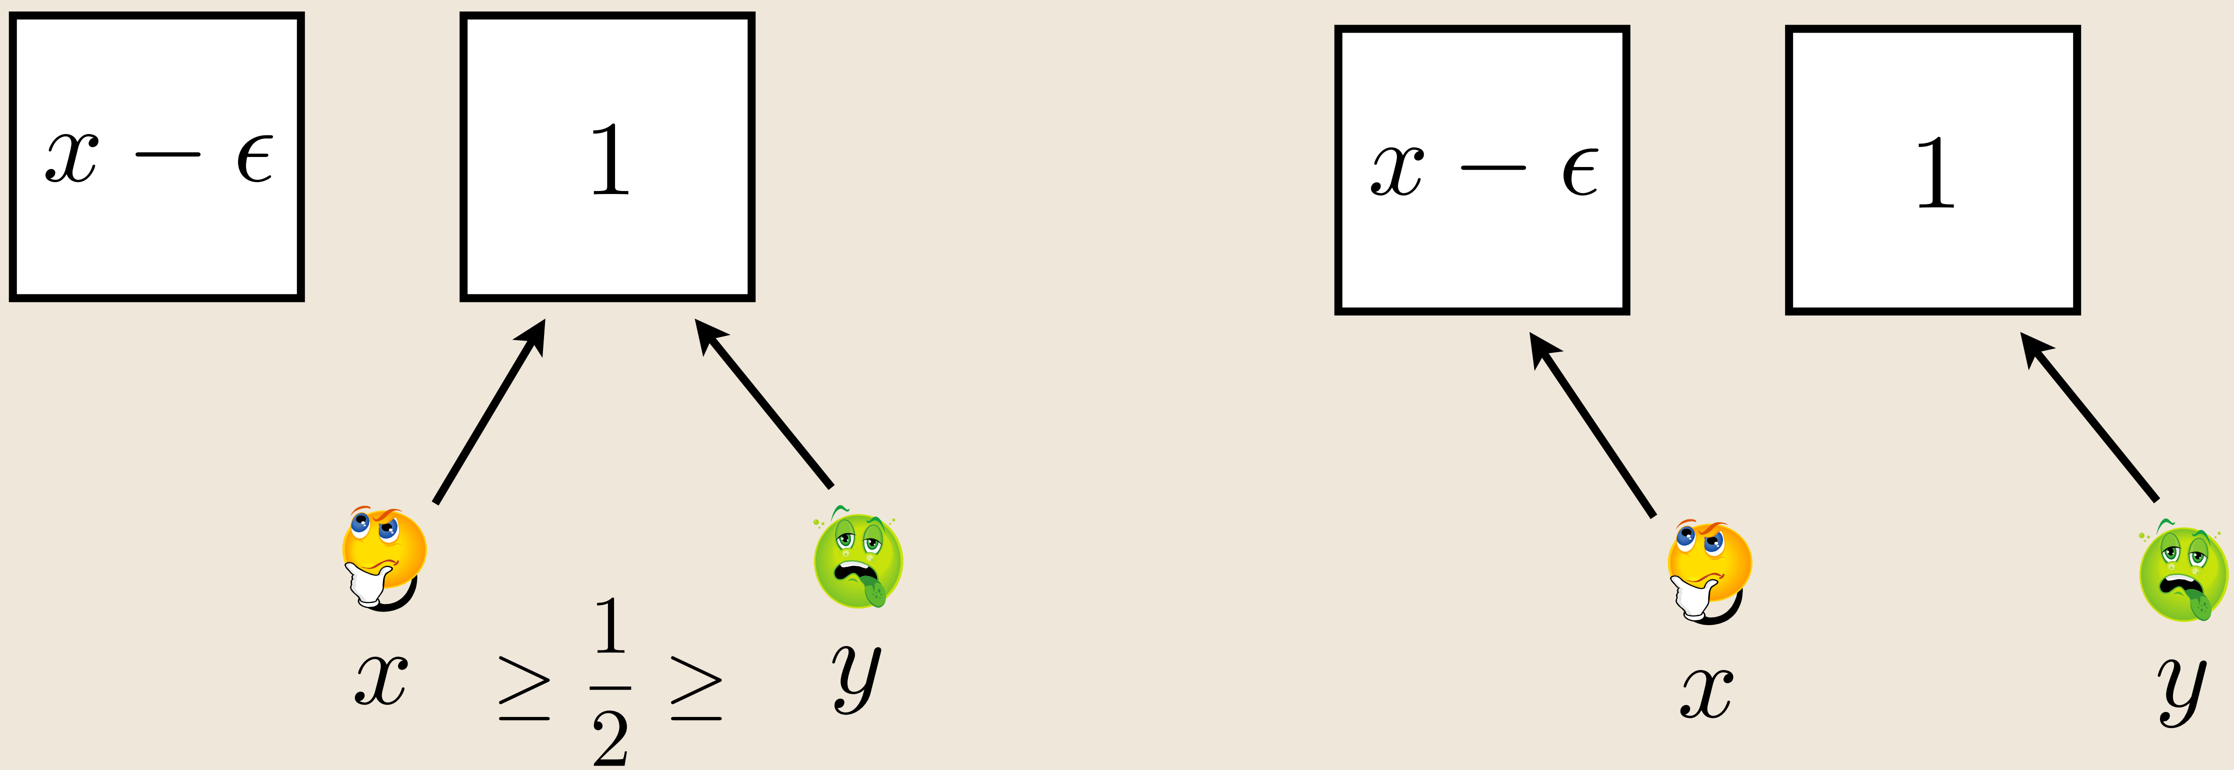

Unique NE
W=1

$$POS \leq \frac{2}{3}$$

Optimal
W=1+x

By increasing the number of players we can drive POS to 1/2

Efficiency

	NE exists	Budget Balanced	Complexity	POS	POA
Marginal contribution	✓		*medium*	1	1/2
Shapley value	✓	✓	*high*	1/2	1/2

conflict between budget balanced and efficiency

Is it possible to overcome limitations by conditioning utilities on more information?

Recap

Proved: distribution rule → game POS<1

Possible? game → distribution rule POS=1

ordered protocols

Ordered Protocol

Ordered Protocol

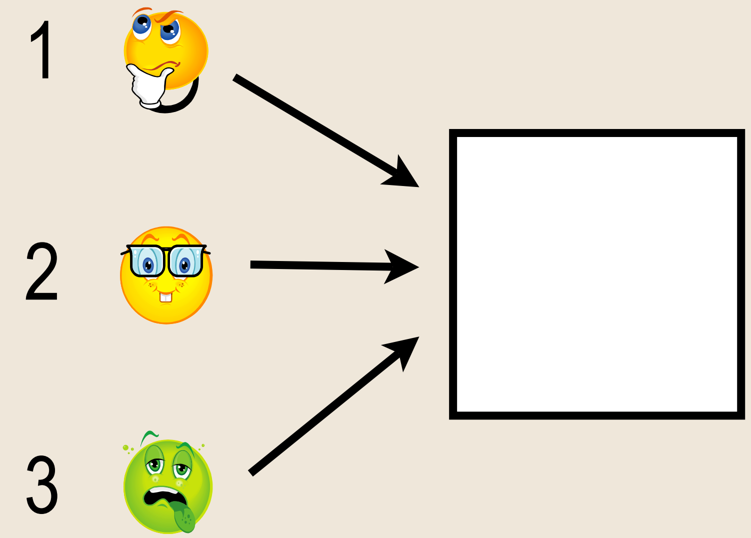

Payoffs

(1st) $W^r(1)$

(2nd) $W^r(1,2) - W^r(1)$

(3rd) $W^r(1,2,3) - W^r(1,2)$

Welfare
$W^r(1)$
$W^r(2)$
$W^r(3)$
$W^r(1,2)$
$W^r(1,3)$
$W^r(2,3)$
$W^r(1,2,3)$

Ordered Protocol

Properties

Budget Balanced

Payoffs

 (1st) $W^r(1)$

 (2nd) $W^r(1,2) - W^r(1)$

 (3rd) $W^r(1,2,3) - W^r(1,2)$

= $W^r(1,2,3)$

Ordered Protocol

Properties

Budget Balanced

U ≥ Marginal Contribution

Payoffs

(1st) $W^r(1)$ $\geq W^r(1,2,3) - W^r(2,3)$

(2nd) $W^r(1,2) - W^r(1)$ $\geq W^r(1,2,3) - W^r(1,3)$

(3rd) $W^r(1,2,3) - W^r(1,2)$ $= W^r(1,2,3) - W^r(1,2)$

Last player's utility **equal** to marginal contribution

Can we use ordered protocols to guarantee POS = 1 for a given game?

Direction: game → distribution rule (POS = 1)

Direction: game → distribution rule (POS = 1)

Create ordered protocol

(1) Consider OPT

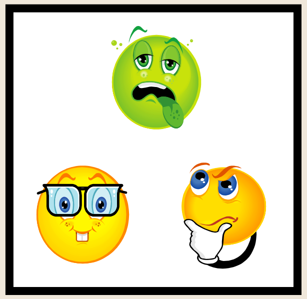

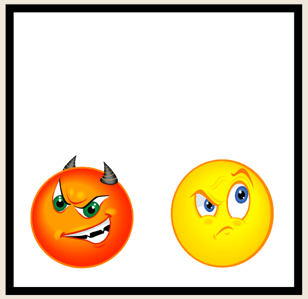

Direction: game ➡ distribution rule (POS = 1)

Create ordered protocol

(1) Consider OPT

(2) Specify any order

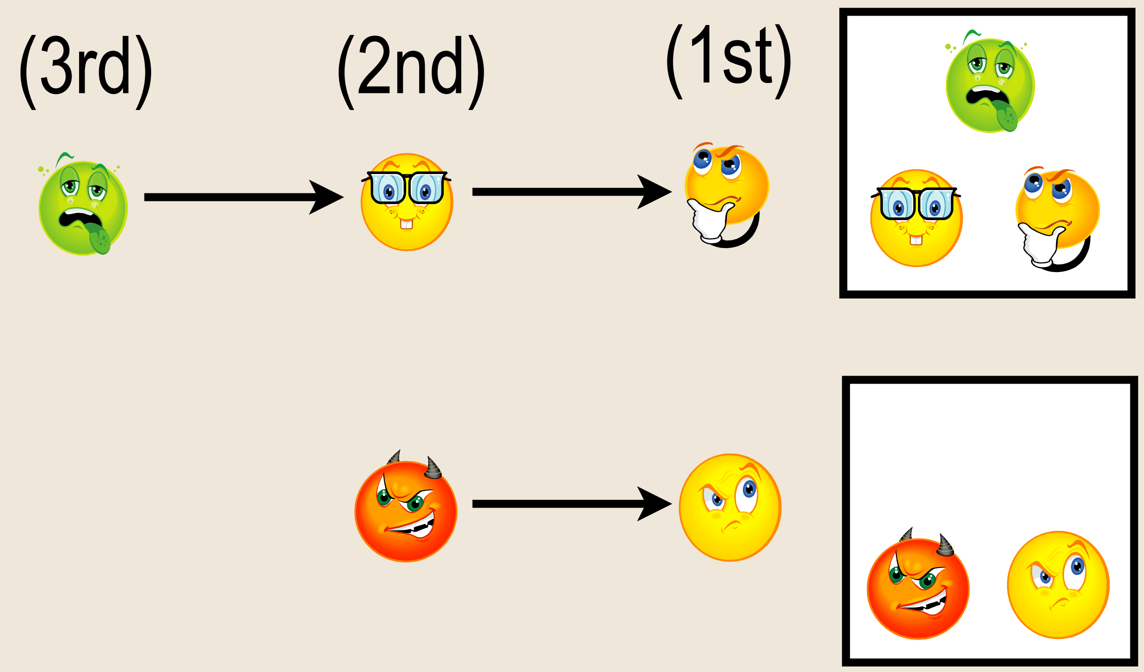

Direction: game → distribution rule (POS = 1)

Create ordered protocol

(1) Consider OPT

(2) Specify any order

(3) Extend order by alternatives last

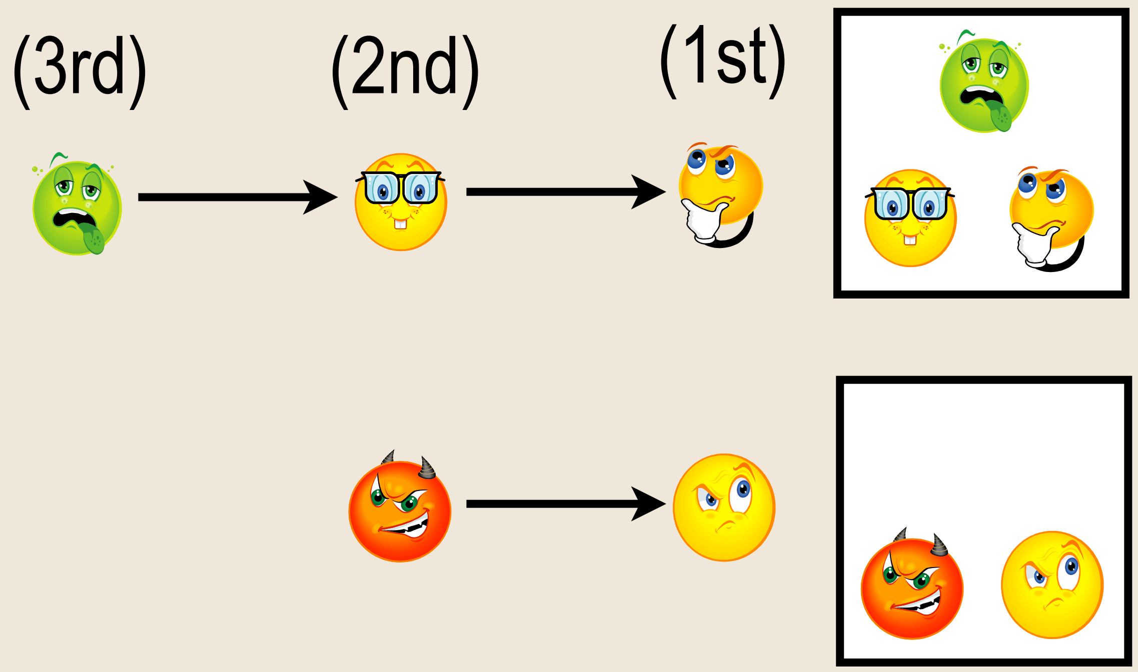

Efficiency

Direction: game → distribution rule (POS = 1)

Create ordered protocol

(1) Consider OPT

(2) Specify any order

(3) Extend order by alternatives last

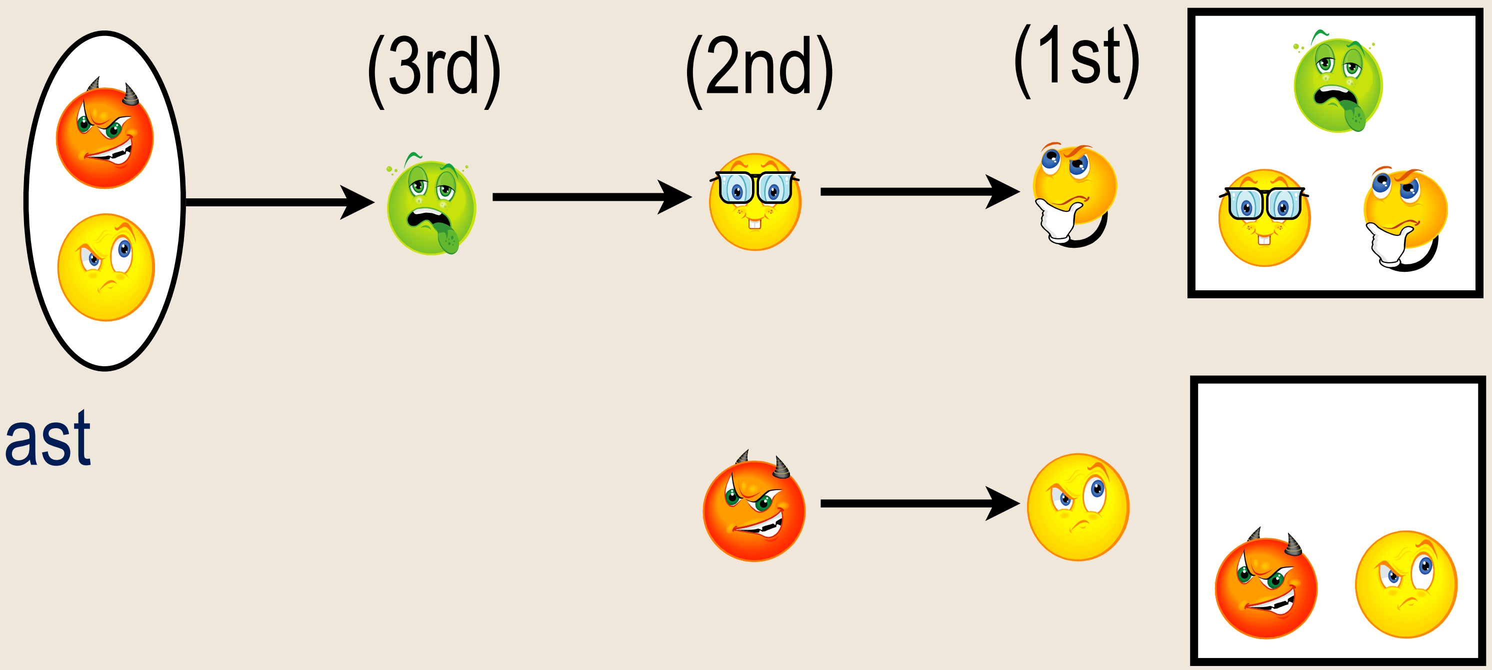

Direction: game → distribution rule (POS = 1)

Create ordered protocol

(1) Consider OPT

(2) Specify any order

(3) Extend order by alternatives last

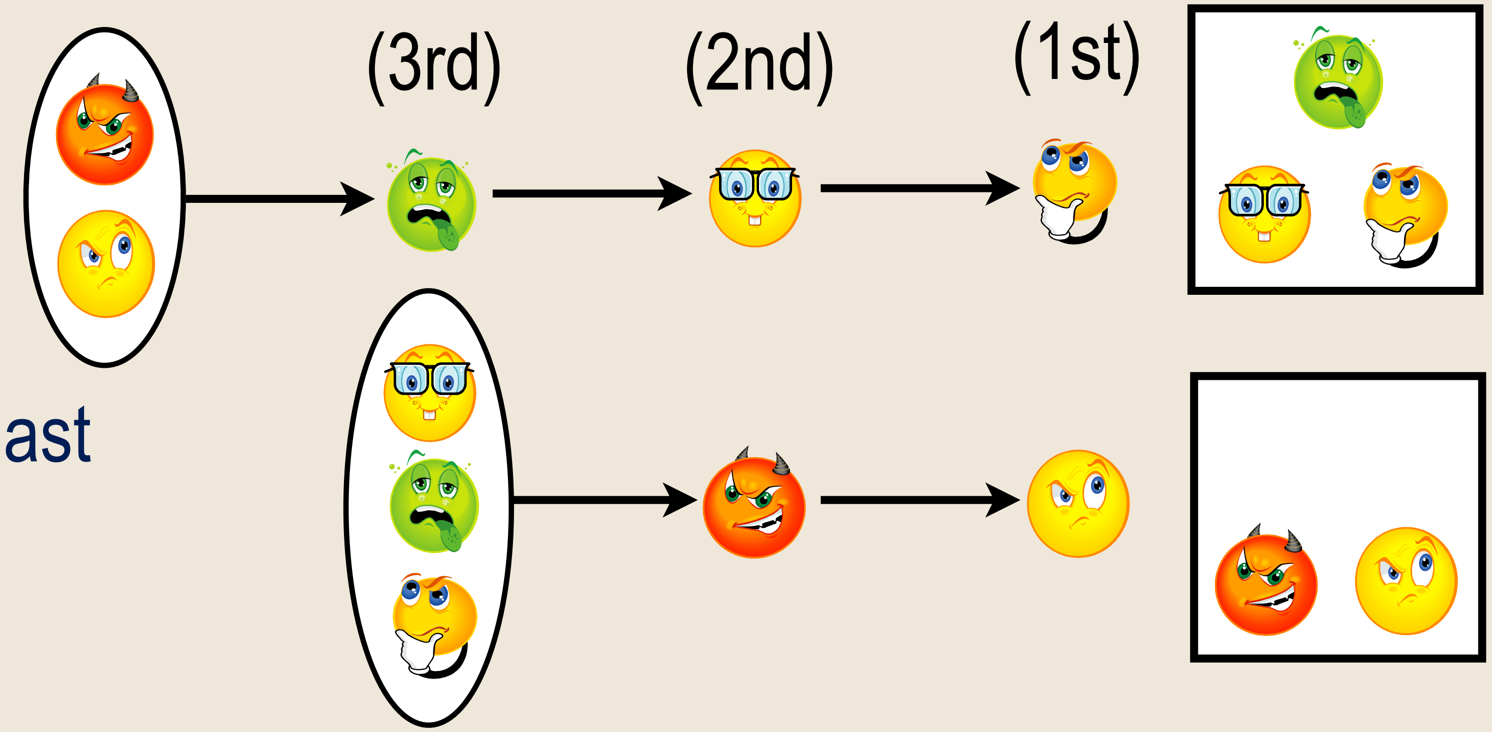

Direction: game ➡ distribution rule (POS = 1)

Create ordered protocol

(1) Consider OPT

(2) Specify any order

(3) Extend order by alternatives last

(4) Remaining order anything

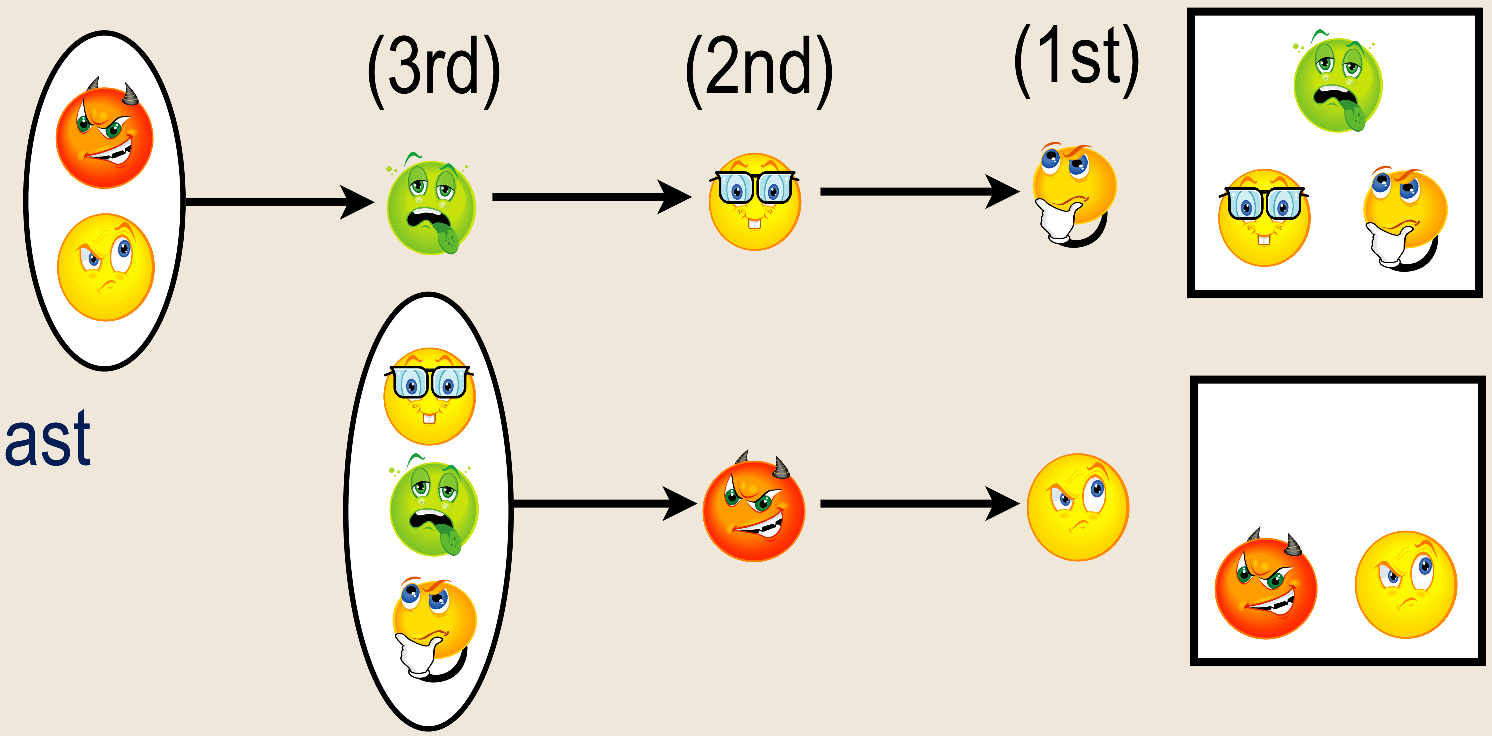

Efficiency

Direction: game → distribution rule (POS = 1)

Create ordered protocol

(1) Consider OPT

(2) Specify any order

(3) Extend order by alternatives last

(4) Remaining order anything

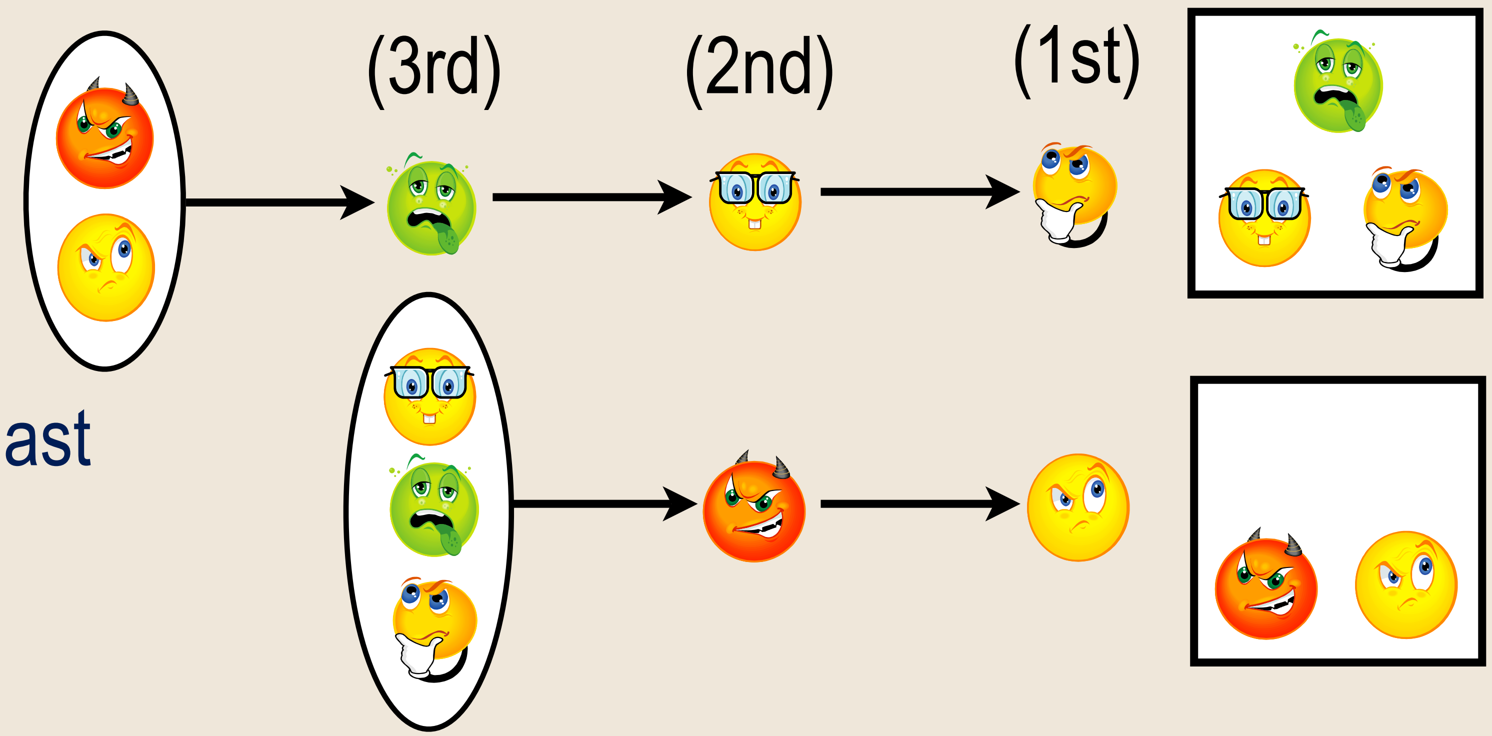

Utility at OPT satisfies

$$U_i(a^{\mathrm{opt}}) \geq W(a^{\mathrm{opt}}) - W(\emptyset, a_{-i}^{\mathrm{opt}})$$

$$U_i(a_i', a_{-i}^{\mathrm{opt}}) = W(a_i', a_{-i}^{\mathrm{opt}}) - W(\emptyset, a_{-i}^{\mathrm{opt}})$$

Efficiency

Direction: game ➡ distribution rule (POS = 1)

Create ordered protocol

(1) Consider OPT

(2) Specify any order

(3) Extend order by alternatives last

(4) Remaining order anything

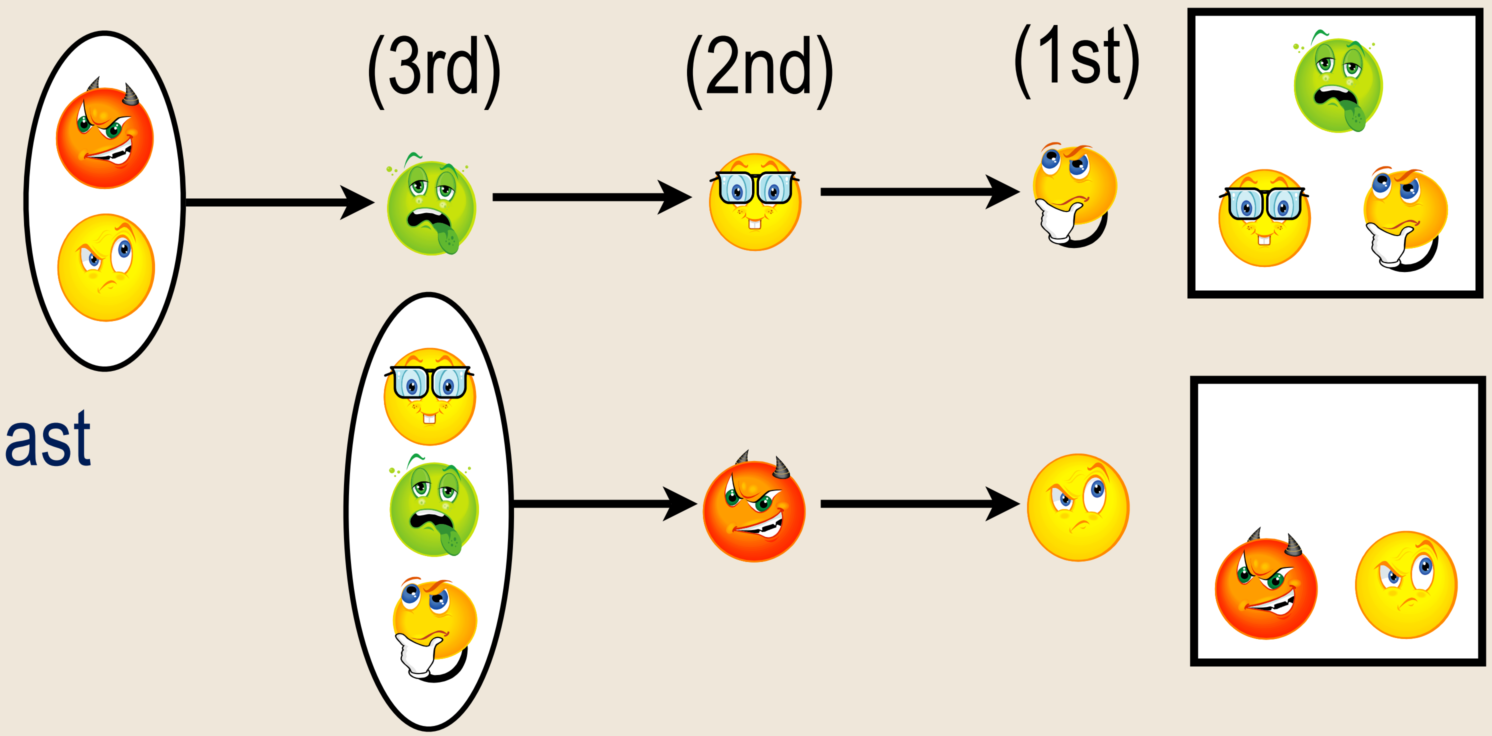

Utility at OPT satisfies

$$U_i(a^{\text{opt}}) \geq W(a^{\text{opt}}) - W(\emptyset, a_{-i}^{\text{opt}})$$

$$U_i(a_i', a_{-i}^{\text{opt}}) = W(a_i', a_{-i}^{\text{opt}}) - W(\emptyset, a_{-i}^{\text{opt}})$$

$$U_i(a_i', a_{-i}^{\text{opt}}) > U_i(a^{\text{opt}}) \;\Rightarrow W(a_i', a_{-i}^{\text{opt}}) > W(a^{\text{opt}})$$

Efficiency

Direction: game ⇒ distribution rule (POS = 1)

Create ordered protocol

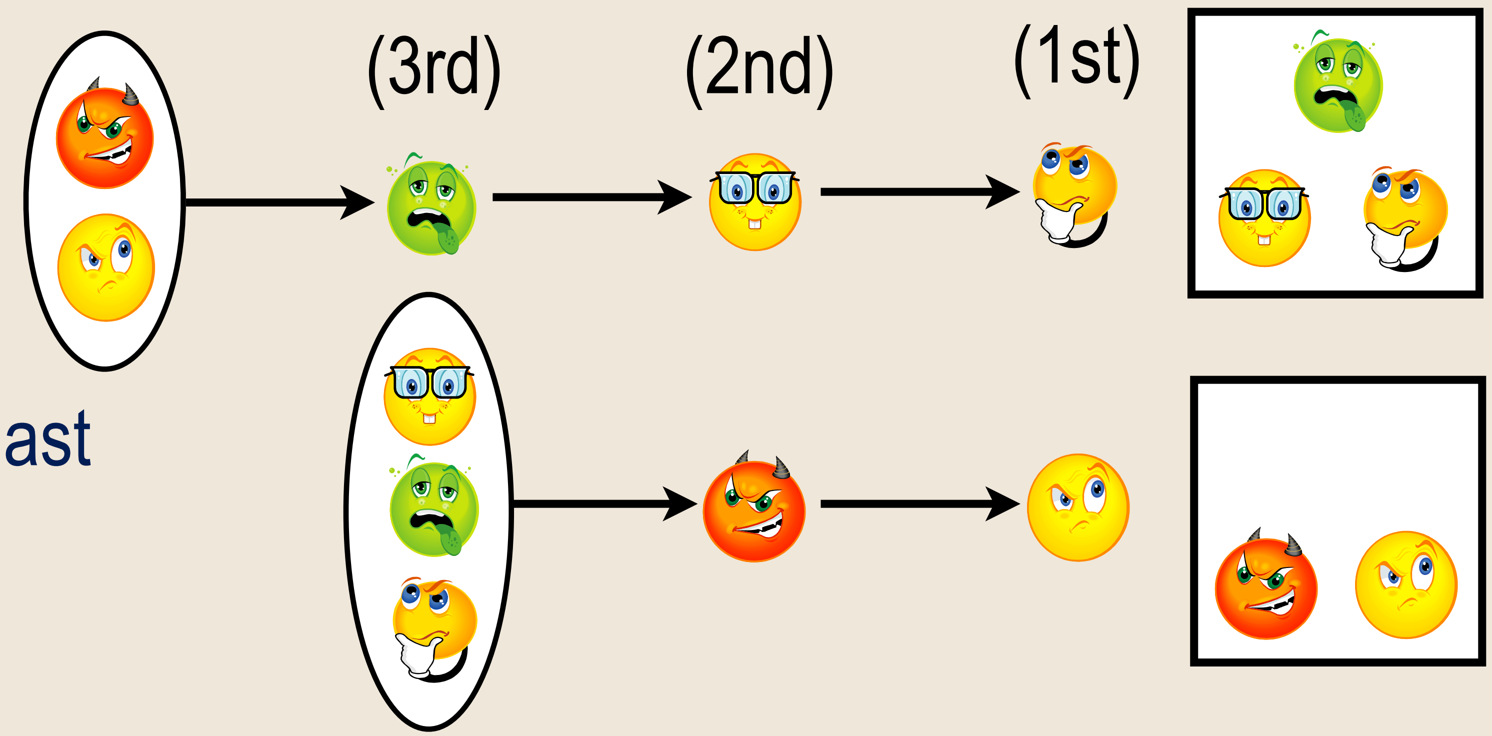

(1) Consider OPT

(2) Specify any order

(3) Extend order by alternatives last

(4) Remaining order anything

Utility at OPT satisfies

$$U_i(a^{\mathrm{opt}}) \geq W(a^{\mathrm{opt}}) - W(\emptyset, a_{-i}^{\mathrm{opt}})$$

$$U_i(a_i', a_{-i}^{\mathrm{opt}}) = W(a_i', a_{-i}^{\mathrm{opt}}) - W(\emptyset, a_{-i}^{\mathrm{opt}})$$

$$U_i(a_i', a_{-i}^{\mathrm{opt}}) > U_i(a^{\mathrm{opt}}) \;\Rightarrow W(a_i', a_{-i}^{\mathrm{opt}}) > W(a^{\mathrm{opt}})$$

(OPT = NE)

Recap

Proved: distribution rule → game POS<1

Possible? game → distribution rule POS=1

ordered protocols

Do we need to condition the distribution rule on the game?

No. Simple adaptive dynamics can find desired distribution rule.

Priority Based Distribution Rule

(1) Define an auxiliary state for each resource that specifies the order

Priority Based Distribution Rule

(1) Define an auxiliary state for each resource that specifies the order
(2) If user leaves resource, all player behind him move up one spot in the queue

Priority Based Distribution Rule

(1) Define an auxiliary state for each resource that specifies the order
(2) If user leaves resource, all player behind him move up one spot in the queue
(3) If user joins resource, user enter last spot in queue

Priority Based Distribution Rule

(1) Define an auxiliary state for each resource that specifies the order
(2) If user leaves resource, all player behind him move up one spot in the queue
(3) If user joins resource, user enter last spot in queue

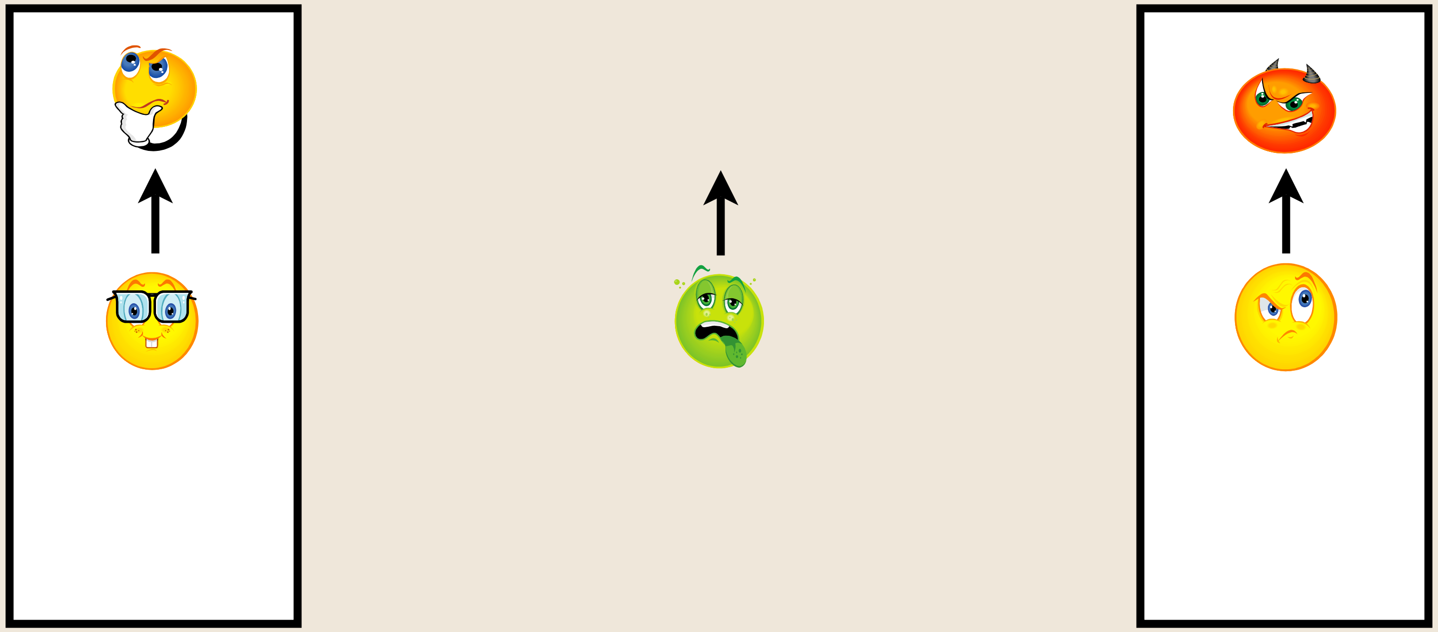

Priority Based Distribution Rule

(1) Define an auxiliary state for each resource that specifies the order
(2) If user leaves resource, all player behind him move up one spot in the queue
(3) If user joins resource, user enter last spot in queue

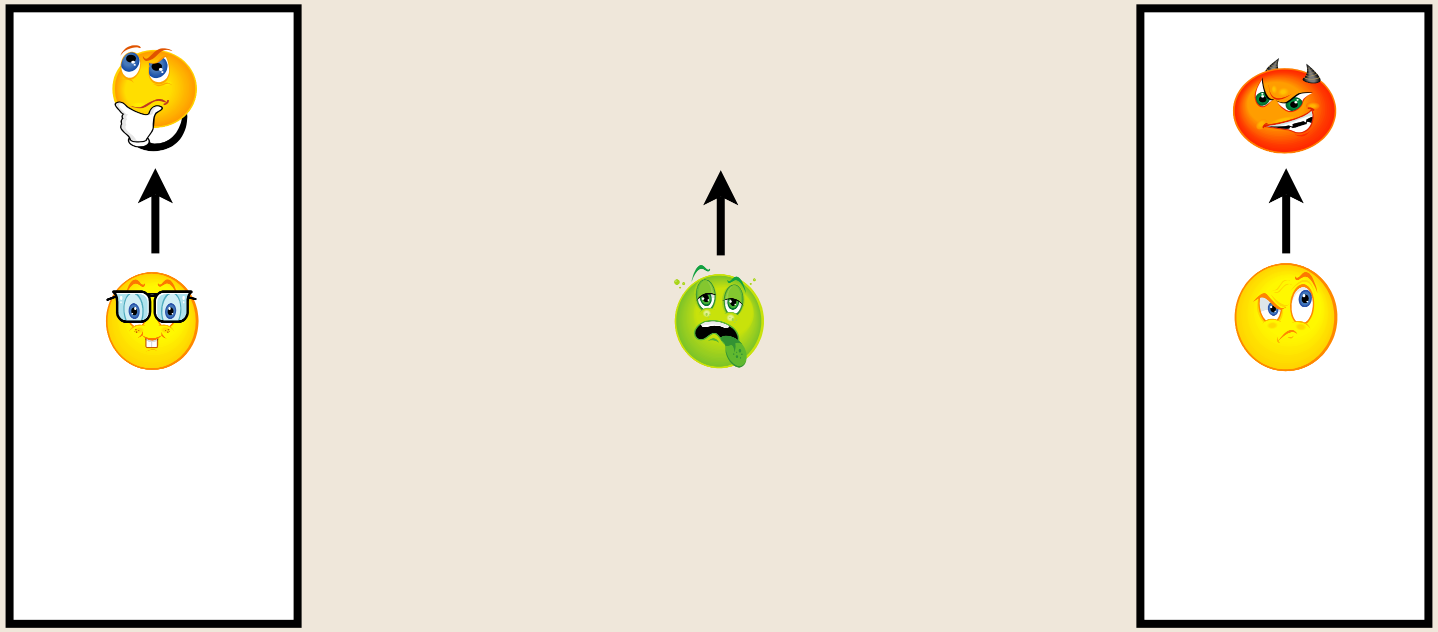

Priority Based Distribution Rule

(1) Define an auxiliary state for each resource that specifies the order
(2) If user leaves resource, all player behind him move up one spot in the queue
(3) If user joins resource, user enter last spot in queue

Priority Based Distribution Rule

(1) Define an auxiliary state for each resource that specifies the order
(2) If user leaves resource, all player behind him move up one spot in the queue
(3) If user joins resource, user enter last spot in queue

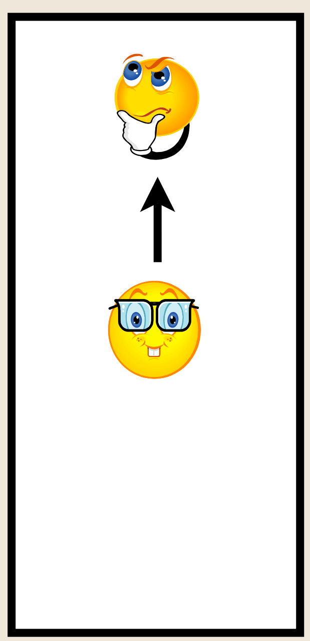

If OPT is played then it is a NE

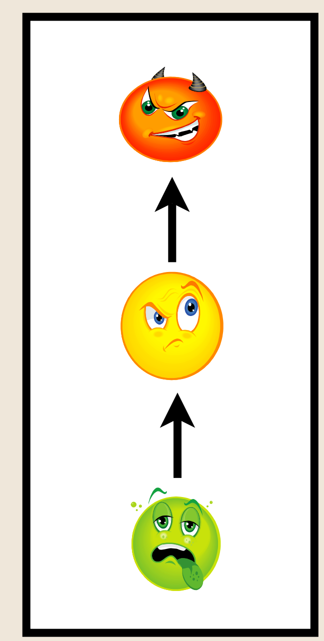

Summary

	NE exists	Budget Balanced	Complexity	POS	POA
Marginal contribution	✓		*medium*	**1**	**1/2**
Shapley value	✓	✓	*high*	**1/2**	**1/2**
Priority based	✓	✓	*medium*	**1**	**1/2**

Take Away Points:

(1) Noncooperative game theory has **inherent limitation** with respect to distributed control
(2) Utilizing noncooperative game theory for distributed control is a **design choice,** not a requirement
(3) Many of the limitations can be overcome by moving beyond noncooperative games (introducing auxiliary state variable)

	Noncooperative Game	State Based Game
Players	$\{1, ..., n\}$	$\{1, ..., n\}$
Actions	$\mathcal{A}_i$	$\mathcal{A}_i$
Utilities	$U_i : \mathcal{A} \rightarrow R$	$U_i : \mathcal{A} \times X \rightarrow R$
States		X
State Transition		$P : X \times \mathcal{A} \rightarrow \Delta(X)$

Extra flexibility in design can be utilized to improve performance

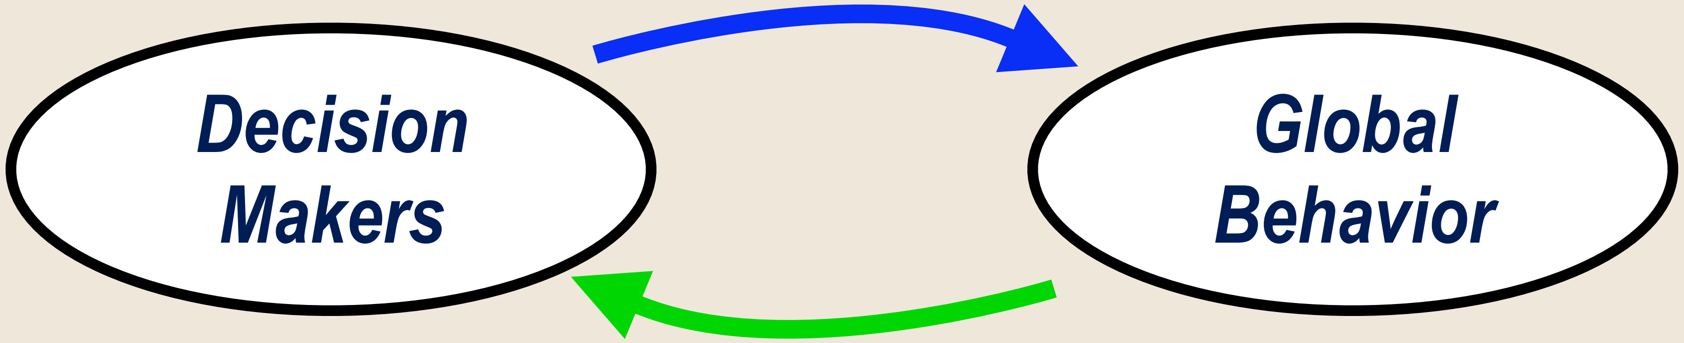

Thank You!