Social Laws for Multi-Agent Systems: Logic and Games

Lecture 4: Coordinating Self-Interested Agents

Thomas Ågotnes[1]

[1] Department of Information Science and Media Studies
University of Bergen, Norway

NII Tokyo 10 January 2012

Contents

Incentive Compatibility

- Key idea: design the social law so that *compliance* is in everybody's interests
- For this, we need a model of everybody's *preferences*
- Here we model preferences as a a *prioritised list* of goal formulae.

Motivation

- A new norm is suggested
- Each agent can deside whether or not to commit to it, i.e., to always comply with it in the future
- Decision (commit/not commit): made design time, before the system "starts"
- Each agent has its own goals about the future of the system
- Is it rational to commit? Depends on what the other agents will do.
- $\Rightarrow$ Game theoretic setting

Contents

Introduction Goals The Value of Social Laws Compliance Games Decision Problems References

Goals

- We model agent's goals as a *prioritised list* of CTL formulae

$$\gamma = \langle \varphi_0, \ldots, \varphi_k \rangle$$

- Goals further up (higher index) are more desired
- Kripke structure K satisfies a goal x in goal hierarchy γ iff

$$K \models \gamma[x]$$

- Assume: if a goal is satisfied, unconcerned about goals further down

An Example

- A system with a single non-sharable resource, which is desired by two agents.
- We have two states, *s* and *t*, and two corresponding Boolean variables p_1 and p_2, which are mutually exclusive.

$$p_i \text{ means “agent } i \text{ has control”}$$

- Each agent has two possible actions, when in possession of the resource: either give it away, or keep it.

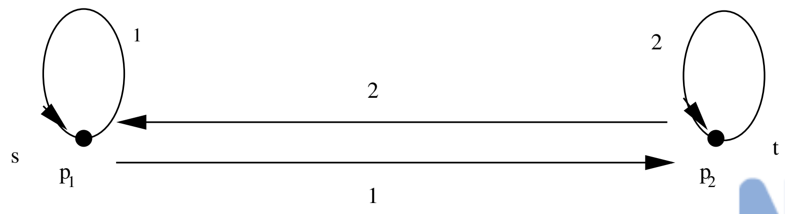

Example

Egent *i* wants to keep the source as often and long as possible for himself:

$$\gamma_i = \langle \varphi_0^i, \ldots, \varphi_8^i \rangle$$

- $\varphi_8^i = \mathsf{A}\Box p_i$
- $\varphi_7^i = \mathsf{A}\Diamond\mathsf{E}\Box p_i$
- $\varphi_6^i = \mathsf{A}\Box\mathsf{A}\Diamond p_i$
- $\varphi_5^i = \mathsf{E}\Diamond\mathsf{A}\Box p_i$
- $\varphi_4^i = \mathsf{A}\Box\mathsf{E}\Diamond p_i$
- $\varphi_3^i = \mathsf{E}\Diamond\mathsf{E}\Box p_i$
- $\varphi_2^i = \mathsf{E}\Box\mathsf{E}\Diamond p_i$
- $\varphi_1^i = \mathsf{E}\Diamond p_i$
- $\varphi_0^i = \top$

Example

Egent i wants to keep the source as often and long as possible for himself:

$$\gamma_i = \langle \varphi^i_0, \ldots, \varphi^i_8 \rangle$$

- $\varphi^i_8 = \mathsf{A}\Box p_i$
- $\varphi^i_7 = \mathsf{A}\Diamond\mathsf{E}\Box p_i$
- $\varphi^i_6 = \mathsf{A}\Box\mathsf{A}\Diamond p_i$
- $\varphi^i_5 = \mathsf{E}\Diamond\mathsf{A}\Box p_i$
- $\varphi^i_4 = \mathsf{A}\Box\mathsf{E}\Diamond p_i$
- $\varphi^i_3 = \mathsf{E}\Diamond\mathsf{E}\Box p_i$
- $\varphi^i_2 = \mathsf{E}\Box\mathsf{E}\Diamond p_i$
- $\varphi^i_1 = \mathsf{E}\Diamond p_i$
- $\varphi^i_0 = \top$

Example

Egent i wants to keep the source as often and long as possible for himself:

$$\gamma_i = \langle \varphi^i_0, \ldots, \varphi^i_8 \rangle$$

- $\varphi^i_8 = \mathsf{A}\Box p_i$
- $\varphi^i_7 = \mathsf{A}\Diamond\mathsf{E}\Box p_i$
- $\varphi^i_6 = \mathsf{A}\Box\mathsf{A}\Diamond p_i$
- $\varphi^i_5 = \mathsf{E}\Diamond\mathsf{A}\Box p_i$
- $\varphi^i_4 = \mathsf{A}\Box\mathsf{E}\Diamond p_i$
- $\varphi^i_3 = \mathsf{E}\Diamond\mathsf{E}\Box p_i$
- $\varphi^i_2 = \mathsf{E}\Box\mathsf{E}\Diamond p_i$
- $\varphi^i_1 = \mathsf{E}\Diamond p_i$
- $\varphi^i_0 = \top$

Example

Egent *i* wants to keep the source as often and long as possible for himself:

$$\gamma_i = \langle \varphi_0^i, \ldots, \varphi_8^i \rangle$$

- $\varphi_8^i = A\Box p_i$
- $\varphi_7^i = A\Diamond E\Box p_i$
- $\varphi_6^i = A\Box A\Diamond p_i$
- $\varphi_5^i = E\Diamond A\Box p_i$
- $\varphi_4^i = A\Box E\Diamond p_i$
- $\varphi_3^i = E\Diamond E\Box p_i$
- $\varphi_2^i = E\Box E\Diamond p_i$
- $\varphi_1^i = E\Diamond p_i$
- $\varphi_0^i = \top$

Introduction | Goals | The Value of Social Laws | Compliance Games | Decision Problems | References

Example

Egent *i* wants to keep the source as often and long as possible for himself:

$$\gamma_i = \langle \varphi_0^i, \ldots, \varphi_8^i \rangle$$

- $\varphi_8^i = A\Box p_i$
- $\varphi_7^i = A\Diamond E\Box p_i$
- $\varphi_6^i = A\Box A\Diamond p_i$
- $\varphi_5^i = E\Diamond A\Box p_i$
- $\varphi_4^i = A\Box E\Diamond p_i$
- $\varphi_3^i = E\Diamond E\Box p_i$
- $\varphi_2^i = E\Box E\Diamond p_i$
- $\varphi_1^i = E\Diamond p_i$
- $\varphi_0^i = \top$

Example

Egent i wants to keep the source as often and long as possible for himself:

$$\gamma_i = \langle \varphi_0^i, \ldots, \varphi_8^i \rangle$$

- $\varphi_8^i = \mathsf{A}\Box p_i$
- $\varphi_7^i = \mathsf{A}\Diamond\mathsf{E}\Box p_i$
- $\varphi_6^i = \mathsf{A}\Box\mathsf{A}\Diamond p_i$
- $\varphi_5^i = \mathsf{E}\Diamond\mathsf{A}\Box p_i$
- $\varphi_4^i = \mathsf{A}\Box\mathsf{E}\Diamond p_i$
- $\varphi_3^i = \mathsf{E}\Diamond\mathsf{E}\Box p_i$
- $\varphi_2^i = \mathsf{E}\Box\mathsf{E}\Diamond p_i$
- $\varphi_1^i = \mathsf{E}\Diamond p_i$
- $\varphi_0^i = \top$

Example

Egent i wants to keep the source as often and long as possible for himself:

$$\gamma_i = \langle \varphi_0^i, \ldots, \varphi_8^i \rangle$$

- $\varphi_8^i = \mathsf{A}\Box p_i$
- $\varphi_7^i = \mathsf{A}\Diamond\mathsf{E}\Box p_i$
- $\varphi_6^i = \mathsf{A}\Box\mathsf{A}\Diamond p_i$
- $\varphi_5^i = \mathsf{E}\Diamond\mathsf{A}\Box p_i$
- $\varphi_4^i = \mathsf{A}\Box\mathsf{E}\Diamond p_i$
- $\varphi_3^i = \mathsf{E}\Diamond\mathsf{E}\Box p_i$
- $\varphi_2^i = \mathsf{E}\Box\mathsf{E}\Diamond p_i$
- $\varphi_1^i = \mathsf{E}\Diamond p_i$
- $\varphi_0^i = \top$

Example

Egent i wants to keep the source as often and long as possible for himself:

$$\gamma_i = \langle \varphi_0^i, \ldots, \varphi_8^i \rangle$$

- $\varphi_8^i = A\Box p_i$
- $\varphi_7^i = A\Diamond E\Box p_i$
- $\varphi_6^i = A\Box A\Diamond p_i$
- $\varphi_5^i = E\Diamond A\Box p_i$
- $\varphi_4^i = A\Box E\Diamond p_i$
- $\varphi_3^i = E\Diamond E\Box p_i$
- $\varphi_2^i = E\Box E\Diamond p_i$
- $\varphi_1^i = E\Diamond p_i$
- $\varphi_0^i = \top$

Introduction Goals The Value of Social Laws Compliance Games Decision Problems References

Example

Egent i wants to keep the source as often and long as possible for himself:

$$\gamma_i = \langle \varphi_0^i, \ldots, \varphi_8^i \rangle$$

- $\varphi_8^i = A\Box p_i$
- $\varphi_7^i = A\Diamond E\Box p_i$
- $\varphi_6^i = A\Box A\Diamond p_i$
- $\varphi_5^i = E\Diamond A\Box p_i$
- $\varphi_4^i = A\Box E\Diamond p_i$
- $\varphi_3^i = E\Diamond E\Box p_i$
- $\varphi_2^i = E\Box E\Diamond p_i$
- $\varphi_1^i = E\Diamond p_i$
- $\varphi_0^i = \top$

Contents

Ordinal Utilities

- The *utility of a Kripke structure* for i is the highest index of any goal that is guaranteed for i in the Kripke structure.

$$u_i(K) = \max\{j : 0 \leq j \leq |\gamma_i| \; \& \; K \models \gamma_i[j]\}$$

- These are ordinal values: you can't compare utility between agents.

Ordinal Utilities

- The *utility of a Kripke structure* for i is the highest index of any goal that is guaranteed for i in the Kripke structure.

$$u_i(K) = \max\{j : 0 \leq j \leq |\gamma_i| \ \& \ K \models \gamma_i[j]\}$$

- These are ordinal values: you can't compare utility between agents.

Example

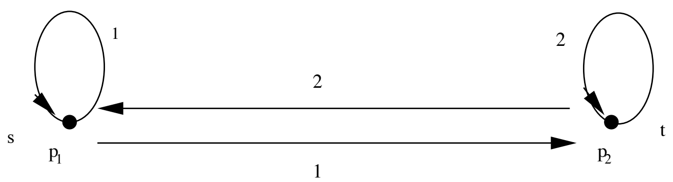

$$\gamma_i = (\quad \varphi_0^i = \top, \quad \varphi_1^i = \mathsf{E}\Diamond p_i,$$
$$\varphi_2^i = \mathsf{E}\Box\mathsf{E}\Diamond p_i, \quad \varphi_3^i = \mathsf{E}\Diamond\mathsf{E}\Box p_i,$$
$$\varphi_4^i = \mathsf{A}\Box\mathsf{E}\Diamond p_i, \quad \varphi_5^i = \mathsf{E}\Diamond\mathsf{A}\Box p_i$$
$$\varphi_6^i = \mathsf{A}\Box\mathsf{A}\Diamond p_i, \quad \varphi_7^i = \mathsf{A}\Diamond\mathsf{E}\Box p_i,$$
$$\varphi_8^i = \mathsf{A}\Box p_i \quad)$$

$$u_1(K) = u_2(K) = 4$$

The Utility of a Social Law

- We can now ask *how good* a social law is for an agent
- The value

$$u_i(K \dagger \eta)$$

gives us agent *i*'s utility of implementing the social law

Introduction Goals The Value of Social Laws Compliance Games Decision Problems References

Example

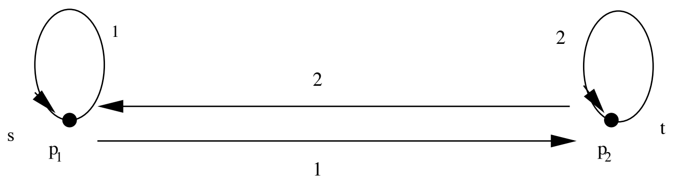

$\eta_0 = \emptyset$ $\quad \eta_1 = \{(s,s)\}$ $\quad \eta_2 = \{(t,t)\}$

$\eta_3 = \{(s,s),(s,t)\}$ $\quad \eta_4 = \{(s,t)\}$ $\quad \eta_5 = \{(t,s)\}$

$\eta_6 = \{(s,s),(t,s)\}$ $\quad \eta_7 = \{(t,t),(s,t)\}$ $\quad \eta_8 = \{(s,t),(t,s)\}$

	η_0	η_1	η_2	η_3	η_4	η_5	η_6	η_7	η_8
$u_1(K \dagger \eta)$	4	4	7	6	5	0	0	7	0
$u_2(K \dagger \eta)$	4	7	4	6	0	5	7	0	0

The Benefit of a Social Law

- Given K, K', i, measure the the *difference* in utility when moving from K to K':

$$\delta_i(K, K') = u_i(K') - u_i(K)$$

- The benefit for agent i of implementing η in K is then

$$\delta_i(K, K \dagger \eta)$$

Introduction Goals The Value of Social Laws Compliance Games Decision Problems References

The Benefit of a Social Law

- Given K, K', i, measure the the *difference* in utility when moving from K to K':

$$\delta_i(K, K') = u_i(K') - u_i(K)$$

- The benefit for agent i of implementing η in K is then

$$\delta_i(K, K \dagger \eta)$$

Example

	η_0	η_1	η_2	η_3	η_4	η_5	η_6	η_7	η_8
$u_1(K \dagger \eta)$	4	4	7	6	5	0	0	7	0
$u_2(K \dagger \eta)$	4	7	4	6	0	5	7	0	0

η	$\delta_1(K, K \dagger \eta)$	$\delta_2(K, K \dagger \eta)$
$\eta_\emptyset$	0	0
η_1	0	3
η_2	3	0
η_3	2	2

Example

	η_0	η_1	η_2	η_3	η_4	η_5	η_6	η_7	η_8
$u_1(K \dagger \eta)$	4	4	7	6	5	0	0	7	0
$u_2(K \dagger \eta)$	4	7	4	6	0	5	7	0	0

η	$\delta_1(K, K \dagger \eta)$	$\delta_2(K, K \dagger \eta)$
$\eta_\emptyset$	0	0
η_1	0	3
η_2	3	0
η_3	2	2

Universal and Existential Goals

Universal and existential fragment of CTL, respectively:

$$\mu ::= \top \mid p \mid \neg p \mid \mu \vee \mu \mid \mathsf{A}\bigcirc\mu \mid \mathsf{A}\Box\mu \mid \mathsf{A}(\mu\,\mathcal{U}\,\mu)$$
$$\varepsilon ::= \top \mid p \mid \neg p \mid \varepsilon \vee \varepsilon \mid \mathsf{E}\bigcirc\varepsilon \mid \mathsf{E}\Box\varepsilon \mid \mathsf{E}(\varepsilon\,\mathcal{U}\,\varepsilon)$$

- If $u_i(K) = n$ and $\gamma_i[n]$ is universal, then $\delta_i(K, K \dagger \eta) \geq 0$ for any social law η
- If $u_i(K \dagger \eta) = n$ for some social law η and $\gamma_i[n]$ is existential, then $\delta_i(K \dagger \eta, K) \geq 0$.

Introduction | Goals | The Value of Social Laws | Compliance Games | Decision Problems | References

Universal and Existential Goals

Universal and existential fragment of CTL, respectively:

$$\mu ::= \top \mid p \mid \neg p \mid \mu \vee \mu \mid \mathsf{A}\bigcirc\mu \mid \mathsf{A}\Box\mu \mid \mathsf{A}(\mu\,\mathcal{U}\,\mu)$$
$$\varepsilon ::= \top \mid p \mid \neg p \mid \varepsilon \vee \varepsilon \mid \mathsf{E}\bigcirc\varepsilon \mid \mathsf{E}\Box\varepsilon \mid \mathsf{E}(\varepsilon\,\mathcal{U}\,\varepsilon)$$

- If $u_i(K) = n$ and $\gamma_i[n]$ is universal, then $\delta_i(K, K \dagger \eta) \geq 0$ for any social law η
- If $u_i(K \dagger \eta) = n$ for some social law η and $\gamma_i[n]$ is existential, then $\delta_i(K \dagger \eta, K) \geq 0$.

Universal and Existential Goals

Universal and existential fragment of CTL, respectively:

$$\mu ::= \top \mid p \mid \neg p \mid \mu \vee \mu \mid \mathsf{A}\mu \mid \mathsf{A}\Box\mu \mid \mathsf{A}(\mu\,\mathcal{U}\,\mu)$$
$$\varepsilon ::= \top \mid p \mid \neg p \mid \varepsilon \vee \varepsilon \mid \mathsf{E}\varepsilon \mid \mathsf{E}\Box\varepsilon \mid \mathsf{E}(\varepsilon\,\mathcal{U}\,\varepsilon)$$

- If $u_i(K) = n$ and $\gamma_i[n]$ is universal, then $\delta_i(K, K \dagger \eta) \geq 0$ for any social law η
- If $u_i(K \dagger \eta) = n$ for some social law η and $\gamma_i[n]$ is existential, then $\delta_i(K \dagger \eta, K) \geq 0$.

Contents

Setting

- Given:
  - Kripke structure K
  - Goals $\gamma_i = \langle \varphi_0^i, \ldots, \varphi_{k_i}^i \rangle$
  - Social law η over K
- It is proposed η should be imposed
- Each agent must decide: should it commit to η or not
- Before the system "starts"

Restrictions on Social Laws

Define operators on social laws which correspond to groups of agents "defecting" from the social law.

$$\eta \upharpoonright C$$

is the social law that is the same as η except that it only contains the arcs of η that correspond to the actions of agents in C.

$$\eta \upharpoonleft C$$

denotes the social law that is the same as η except that it only contains the arcs of η that *do not* correspond to actions of agents in C.

Example

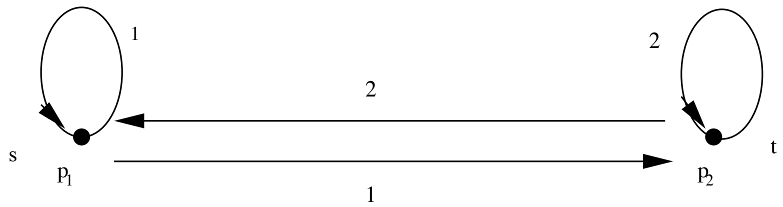

$$\eta_0 = \emptyset$$
$$\eta_1 = \{(s, s)\}$$
$$\eta_2 = \{(t, t)\}$$
$$\eta_3 = \{(s, s), (t, t)\}$$

- $\eta_3 \upharpoonright \{1\} = \eta_3 \upharpoonleft \{2\} = \eta_1$
- $\eta_3 \upharpoonright \{2\} = \eta_3 \upharpoonleft \{1\} = \eta_2$

Strategic Form Games

A game in strategic form :

$$\mathcal{G} = \langle \mathcal{AG}, \mathcal{S}_1, \ldots, \mathcal{S}_n, \mathcal{U}_1, \ldots, \mathcal{U}_n \rangle \text{ where:}$$

$\mathcal{AG} = \{1, \ldots, n\}$ is a set of players
$\mathcal{S}_i$ is the set of strategies for each agent $i \in \mathcal{AG}$
$\mathcal{U}_i : (\mathcal{S}_1 \times \cdots \times \mathcal{S}_n) \rightarrow \mathbb{R}$ is the utility function for agent $i \in \mathcal{AG}$

Social Law Games

Given $\Sigma = \langle K, \gamma_1, \ldots, \gamma_n, \eta \rangle$, the social law game $\mathcal{G}_\Sigma$ is:

- The agents $\mathcal{AG}$ in $\mathcal{G}_\Sigma$ are as in Σ.
- Each agent i has just two strategies available to it:
  - C – *comply* with the norm system;
  - D – *do not comply*.
- $\mathcal{U}_i(S) = \delta_i(K, K \dagger (\eta \restriction \text{agents that play C in } S))$.

Example

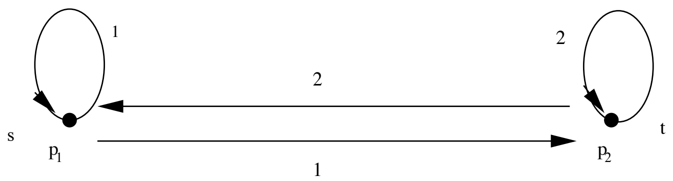

$\eta_0 = \emptyset$
$\eta_1 = \{(s, s)\}$
$\eta_2 = \{(t, t)\}$
$\eta_3 = \{(s, s), (t, t)\}$

η	$\delta_1(K, \eta)$	$\delta_2(K, \eta)$
$\eta_\emptyset$	0	0
η_1	0	3
η_2	3	0
η_3	2	2

Payoff matrix for η_3 :

	C	D
C	$(2, 2)$	$(0, 3)$
D	$(3, 0)$	$(0, 0)$

Contents

Individually Rational Social Laws

- A social law is individually rational if every agent would fare better if the social law was imposed than otherwise.

Example:

	C	D
C	$(2,2)$	$(0,3)$
D	$(3,0)$	$(0,0)$

Individually Rational Social Laws

- A social law is individually rational if every agent would fare better if the social law was imposed than otherwise.

Example:

	C	D
C	$(2,2)$	$(0,3)$
D	$(3,0)$	$(0,0)$

Individually Rational Social Laws

- A social law is individually rational if every agent would fare better if the social law was imposed than otherwise.

INDIVIDUALLY RATIONAL SOCIAL LAW (IRNS):
Given: $K, \gamma_1, \ldots, \gamma_n$.
Question: *Does there exist an individually rational social law?*

Theorem

IRNS *is* NP*-complete, even in one-agent systems.*

Individually Rational Social Laws

- A social law is individually rational if every agent would fare better if the social law was imposed than otherwise.

  INDIVIDUALLY RATIONAL SOCIAL LAW (IRNS):
  Given: $K, \gamma_1, \ldots, \gamma_n$.
  Question: *Does there exist an individually rational social law?*

Theorem

IRNS *is* NP*-complete, even in one-agent systems.*

Pareto Efficient Social Laws

- A benign leader asks: *is it possible to make some agents better off* without *making anybody else worse off?*
- A system is *Pareto efficient* if no such *Pareto improvements* are possible.
- A social law is pareto efficient if there is no other social law under which every agent is better off.

	η_0	η_1	η_2	η_3	η_4	η_5	η_6	η_7	η_8
$u_1(K \dagger \eta)$	4	4	7	6	5	0	0	7	0
$u_2(K \dagger \eta)$	4	7	4	6	0	5	7	0	0

Pareto Efficient Social Laws

- A benign leader asks: *is it possible to make some agents better off* without *making anybody else worse off?*
- A system is *Pareto efficient* if no such *Pareto improvements* are possible.
- A social law is pareto efficient if there is no other social law under which every agent is better off.

	η_0	η_1	η_2	η_3	η_4	η_5	η_6	η_7	η_8
$u_1(K \dagger \eta)$	4	4	7	6	5	0	0	7	0
$u_2(K \dagger \eta)$	4	7	4	6	0	5	7	0	0

Introduction Goals The Value of Social Laws Compliance Games Decision Problems References

Pareto Efficient Social Laws

- A benign leader asks: *is it possible to make some agents better off* without *making anybody else worse off?*
- A system is *Pareto efficient* if no such *Pareto improvements* are possible.
- A social law is pareto efficient if there is no other social law under which every agent is better off.

	η_0	η_1	η_2	η_3	η_4	η_5	η_6	η_7	η_8
$u_1(K \dagger \eta)$	4	4	7	6	5	0	0	7	0
$u_2(K \dagger \eta)$	4	7	4	6	0	5	7	0	0

Pareto Efficient Social Laws

- A social law is pareto efficient if there is no other social law under which every agent is better off.

PARETO EFFICIENT SOCIAL LAW (PENS):
Given: $K, \gamma_1, \ldots, \gamma_n, \eta$.
Question: Is η Pareto efficient?

Theorem

PENS *is co*-NP*-complete, even for one-agent systems.*

Pareto Efficient Social Laws

- A social law is pareto efficient if there is no other social law under which every agent is better off.

PARETO EFFICIENT SOCIAL LAW (PENS):
Given: $K, \gamma_1, \ldots, \gamma_n, \eta$.
Question: Is η Pareto efficient?

Theorem

PENS *is co*-NP*-complete, even for one-agent systems.*

Pareto Efficient Social Laws

- A social law is pareto efficient if there is no other social law under which every agent is better off.

PARETO EFFICIENT SOCIAL LAW (PENS):
Given: $K, \gamma_1, \ldots, \gamma_n, \eta$.
Question: Is η Pareto efficient?

Theorem

PENS *is co-*NP*-complete, even for one-agent systems.*

Nash Implementation

- A social law is a Nash implementation if everyone complying is a Nash equilibrium $\mathcal{G}_\Sigma$.

Example: Pay-off matrix for η_3:

	C	D
C	$(2, 2)$	$(0, 3)$
D	$(3, 0)$	$(0, 0)$

Is η_3 a Nash implementation?

Nash Implementation

- A social law is a Nash implementation if everyone complying is a Nash equilibrium $\mathcal{G}_\Sigma$.

Example: Pay-off matrix for η_3:

	C	D
C	$(2,2)$	$(0,3)$
D	$(3,0)$	$(0,0)$

Is η_3 a Nash implementation?

Nash Implementation

- A social law is a *Nash implementation* if everyone complying is a Nash equilibrium $\mathcal{G}_\Sigma$.

NASH IMPLEMENTATION (NI) :
Given: $K, \gamma_1, \ldots, \gamma_n$.
Question: *Does there exist a Nash implementation?*

Theorem

NI *is* NP*-complete, even for two-agent systems.*

Nash Implementation

- A social law is a *Nash implementation* if everyone complying is a Nash equilibrium $\mathcal{G}_\Sigma$.

NASH IMPLEMENTATION (NI) *:*
Given*:* $K, \gamma_1, \ldots, \gamma_n$.
Question*: Does there exist a Nash implementation?*

Theorem

NI *is* NP*-complete, even for two-agent systems.*

Introduction Goals The Value of Social Laws Compliance Games Decision Problems References

Nash Implementable Mechanism Design

- Suppose a system designer has an *objective*, φ, and wants to know whether this objective is achievable via a Nash implementation.

NASH IMPLEMENTABLE (NASH) *:*
Given*:* $K, \gamma_1, \ldots, \gamma_n, \varphi$.
Question*: Can φ be achieved through a Nash implementation?*

Theorem

NASH *is* NP*-complete.*

Nash Implementable Mechanism Design

- Suppose a system designer has an *objective*, φ, and wants to know whether this objective is achievable via a Nash implementation.

  NASH IMPLEMENTABLE (NASH) :
  Given: $K, \gamma_1, \ldots, \gamma_n, \varphi$.
  Question: *Can φ be achieved through a Nash implementation?*

Theorem

NASH *is* NP*-complete.*

Contents

Some references I

Thomas Ågotnes, Wiebe van der Hoek, and Michael Wooldridge.
Normative system games.
In M. Huhns and O. Shehory, editors, *Proceedings of the Sixth International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2007)*, pages 876–883. May 2007.