Uncertainty in Coalition Formation

"Sequential Decision Making in Repeated Coalition Formation under Uncertainty" by Chalkiadakis and Boutilier "Bayesian Coalitional Games" by Ieong and Shoham

Costin Ionita
Matthias Schuurmans

May 17, 2010

Outline

Introduction

- Given a coalition, how do we distribute the payoff according to some stability concepts ?
- Simplifying assumption: the generated value of the coalition and the potential of each member are common knowledge
- Issue: this is rarely the case in real-life scenarios
- How do we deal with uncertainty regarding other agents' capabilities ?
  - A reinforcement learning model, where agents update information about the others through repeated interaction
  - Generalize classical Coalitional Games Theory to settings with uncertainties

Outline

Markov Decision Processes (MDP)

- The environment can be modelled as an MDP (S,A,R,P):
  - *S* - the set of states
  - *A* - the set of actions
  - *R* - a reward function (R: S x A $\rightarrow \mathbb{R}$)
  - *P* - a stochastic transition model (P: S x A $\rightarrow$ S)
- After each action, the agent "interprets" the response from the environment and changes accordingly his beliefs and his next possible moves.
- We define the **discounted future reward** as $\sum_{t=0}^{\infty} \gamma^t \cdot R(s, a)$ where $\gamma \in [0, 1]$ is a discount factor
- Our aim is to find a **policy** π: S $\rightarrow$ A, such that the discounted future reward is maximised
- A PO-MDP is a generalization of an MDP: the underlying state cannot be directly observed, so the agent maintains a probability distribution over the set of possible states

Exploration vs exploitation

- Each agent has two complementary objectives:
  - form efficient, profitable coalitions
  - gather as much reward as possible
- This means that agents should not seek to reduce uncertainty for *its own sake*, by employing crude exploration policies
- A better strategy is to explore only promising partners, while for the others some type uncertainty will still remain in the end
- It can be shown that Bayesian exploration finds an optimal balance between exploration and exploitation

Bayesian Coalition Formation Model

- Given
  - N - the set of agents
  - T_i - a finite set of possible types for agent i
  - $B_i(t_{-i})$ - the beliefs of agent i over the types of the other agents in the coalition
  - A_C - the set of coalitional actions
  - O - the set of outcomes
- The value of the coalition can be expressed as:
  - $V(C|t_C) = \max_{\alpha_C \in A_C} \sum_o Pr(o|\alpha, t_C) \cdot R(o) = max_{\alpha_C \in A_C} Q(C, \alpha|t_C)$
- However, each agent is uncertain about the types of its partners:
  - $V_i(C|t_C) = \max_{\alpha_C \in A_C} \sum_{t_C in T_C} B_i(t_C) \cdot Q_i(C, \alpha|t_C) = max_{\alpha_C \in A_C} Q_i(C, \alpha)$

Example

- Assume we have the following game:
  - 3 major types: *interface designer (ID)*, *programmer (P)*, and *systems engineer (SE)*
  - 3 quality types: *bad*, *average*, and *expert*
  - 3 actions: bid for *large/medium/small* project
  - 3 outcomes: make *large/medium/small* profit
- The more members a coalition has, the more likely it is to be successful in getting higher profits if it tries to bid for large projects.
- However, it is not only about the size of the coalition, but also about how competent the team members are
- Given 5 agents:
  - a1 = bad *P*, a2 = average *P*, a3 = expert *P*, a4 = bad *ID*, and a5 = bad *SE*
- Then, the best coalition structure is:
  - ((a1), (a2, a3), (a4),(a5))

Optimal repeated coalition formation

- Bellman-like equations:
  - $Q_i(C, \alpha, d_C, B_i) = \sum_o Pr(o|C, \alpha, B_i) \cdot [r_i R(o) + \gamma V_i(B_i^{o,\alpha})] = \sum_{t_C} B_i(t_C) \sum_o Pr(o|\alpha, t_C) \cdot [r_i R(o) + \gamma V_i(B_i^{o,\alpha})]$
  - $V_i(B_i) = \sum_{C|i \in C, d_C} Pr(C, \alpha, d_C|B_i) \cdot Q_i(C, \alpha, d_C, B_i)$
- where $r_i = \frac{d_i}{\sum_{j \in C} d_j}$ is the relative demand of agent i
- After an action is taken, the agent observes the resulting state and updates its beliefs concerning its partners' types:
  - $B_i^{o,\alpha} = B_i^{t+1}(t_C) = zPr(o|\alpha, t_C)B_i^t(t_C)$
    (z is a normalization constant)

Optimal repeated coalition formation (2)

- In typical Bellman equations:
  $V_i(s') = Q_i(s', a*)$, where $a* = argmax_a \sum_{i=1}^{n} Q_i(s', a)$
- However, this idea is not applicable in our case because the choices that dictate reward, namely, the coalition that is formed, are not in complete control of agent i
- Instead, i must predict, based on his beliefs, probability $Pr(C, \alpha, d_C | B_i)$ with which a specific coalitional agreement (C, α, d_C) will arise as a result of the negotiation

Challenges

- The state space and action space grow exponentially with the number of agents
- An exact solution to the repeated coalition formation is generally infeasible
  - Computational approximations are needed

Outline

Computational bottlenecks

- Computing $Pr(C, \alpha, d_C | B_i)$
  - Agents can only observe the outcome of their own coalition action only, and thus it is not possible for them to monitor how the beliefs of others are changing
- Solving the optimal exploration PO-MDP
  - The number of future states that need to be considered is too large

One-step Lookahead Algorithm (OSLA)

- Deals with only immediate successor belief states
- Approximates $Pr(C, \alpha, d_C | B_i)$ by viewing it as the probability of reaching an agreement after one negotiation step, rather than after a whole negotiation process
- Advantages:
  - Allows more flexibility in investigating the space of coalition structures, without forcing the agents to reach a stable structure at each stage before acting
  - Applies best when "real-time" performance is required
  - It can be shown that one-step methods converge to the Bayesian Core of the game (if that exists)

VPI Exploration method

- The main idea is to consider the gain achieved by learning the true value of some coalitional agreement $\sigma = (C, \alpha, d_C)$
- Assume that by adopting σ we obtain exact evidence regarding the types of the agents in C (t^*_C)
- Based on t^*_C we define the true value of σ as:
  - $q^*_\sigma = Q_i(C, \alpha, d_C | t^*_C) = r_i \cdot \sum_o Pr(o|t^*_C) \cdot R(o)$

VPI Exploration method (2)

- Let σ_1 be the current best coalitional agreement, and σ_2 the second-best option
- Then, the gain achieved by learning the true value of σ is:

$$gain_\sigma(q^*_\sigma, t^*_C) = \begin{cases} q^*_\sigma - q_{\sigma_1} & \text{if } \sigma \neq \sigma_1 \text{ and } q^*_\sigma > q_{\sigma_1} \\ q_{\sigma_2} - q^*_\sigma & \text{if } \sigma = \sigma_1 \text{ and } q^*_\sigma < q_{\sigma_2} \\ 0 & \text{otherwise} \end{cases}$$

- which is also known as the *Value of Perfect Information*

EVPI

- However, the agent does not know in advance what types will be revealed for σ
- What it actually considers is the *Expected Value of Perfect Information*:
  - $EVPI(\sigma|B_i) = \sum_{t_C^*} gain(q_\sigma^*|t_C^*) \cdot B_i(t_C^*)$
- EVPI gives the value of exploring σ, therefore agent should have a preference for maximizing:
  - $QV_i(\sigma|B_i) = Q_i(\sigma|B_i) + EVPI(\sigma|B_i)$

OSLA vs VPI

- Both compute the value of the agreements in a myopic manner
- VPI exploits the value of perfect information regarding the *types*, which is in contrast to OSLA which estimates the value of specific coalitional actions
  - VPI does not have to explicitly incorporate the prior hypothesis (B_i) in the calculation of the Q-values
  - VPI does not need to account for the probability of agreement when transitioning to future belief states
- None of them is tied to specific coalition formation processes, therefore can be used with any bargaining processes

Outline

Bayesian Coalitional Games concept

- BCG: different way of dealing with uncertainty in same situation

Bayesian Coalitional Games concept

- BCG: different way of dealing with uncertainty in same situation
- Extension of Coalitional Game Theory (CGT) to include uncertainty about other agents

Bayesian Coalitional Games concept

- BCG: different way of dealing with uncertainty in same situation
- Extension of Coalitional Game Theory (CGT) to include uncertainty about other agents
  - Payoff vector vs *contract*

Bayesian Coalitional Games concept

- BCG: different way of dealing with uncertainty in same situation
- Extension of Coalitional Game Theory (CGT) to include uncertainty about other agents
  - Payoff vector vs *contract*
  - Same solution concepts

Bayesian Coalitional Games concept

- BCG: different way of dealing with uncertainty in same situation
- Extension of Coalitional Game Theory (CGT) to include uncertainty about other agents
  - Payoff vector vs *contract*
  - Same solution concepts
  - Lots of notions needed to solve a BCG

Bayesian Coalitional Games concept

- BCG: different way of dealing with uncertainty in same situation
- Extension of Coalitional Game Theory (CGT) to include uncertainty about other agents
  - Payoff vector vs *contract*
  - Same solution concepts
  - Lots of notions needed to solve a BCG
- Meant to be a combination of Bayesian Theory and CGT to bring real life and CGT closer

BCG definition

$(N, \Omega, \mathbb{P}, (\mathcal{I}_j), (\succeq_j))$ with

- $N = \{1, 2, \ldots, n\}$ set of agents participating in the game

BCG definition

$(N, \Omega, \mathbb{P}, (\mathcal{I}_j), (\succeq_j))$ with

- $N = \{1, 2, \ldots, n\}$ set of agents participating in the game
- $\Omega = \{\omega_1, \omega_2, \ldots, \omega_m\}$ set of possible worlds

BCG definition

$(N, \Omega, \mathbb{P}, (\mathcal{I}_j), (\succeq_j))$ with

- $N = \{1, 2, \ldots, n\}$ set of agents participating in the game
- $\Omega = \{\omega_1, \omega_2, \ldots, \omega_m\}$ set of possible worlds
- $\mathbb{P} = \{p_1, p_2, \ldots, p_m\}$ prior over the worlds in Ω (*common knowledge*)

BCG definition

$(N, \Omega, \mathbb{P}, (\mathcal{I}_j), (\succeq_j))$ with

- $N = \{1, 2, \ldots, n\}$ set of agents participating in the game
- $\Omega = \{\omega_1, \omega_2, \ldots, \omega_m\}$ set of possible worlds
- $\mathbb{P} = \{p_1, p_2, \ldots, p_m\}$ prior over the worlds in Ω (*common knowledge*)
- $\mathcal{I}_j$ agent j's information partition of the worlds Ω (*common knowledge*)

BCG definition

$(N, \Omega, \mathbb{P}, (\mathcal{I}_j), (\succeq_j))$ with

- $N = \{1, 2, \ldots, n\}$ set of agents participating in the game
- $\Omega = \{\omega_1, \omega_2, \ldots, \omega_m\}$ set of possible worlds
- $\mathbb{P} = \{p_1, p_2, \ldots, p_m\}$ prior over the worlds in Ω (*common knowledge*)
- $\mathcal{I}_j$ agent j's information partition of the worlds Ω (*common knowledge*)
- $\succeq_j$ agent j's preference over distributions of payoffs

BCG definition interpretation

$(N, \Omega, \mathbb{P}, (\mathcal{I}_j), (\succeq_j))$ with

- Ω represents uncertainty: multiple possible worlds for possible agent types

BCG definition interpretation

$(N, \Omega, \mathbb{P}, (\mathcal{I}_j), (\succeq_j))$ with

- Ω represents uncertainty: multiple possible worlds for possible agent types
- $\mathbb{P}$ represents educated guesses on agent types

BCG definition interpretation

$(N, \Omega, \mathbb{P}, (\mathcal{I}_j), (\succeq_j))$ with

- Ω represents uncertainty: multiple possible worlds for possible agent types
- $\mathbb{P}$ represents educated guesses on agent types
- $\mathcal{I}_j$ is similar to Coalition Structure: partition over possible worlds

BCG definition interpretation

$(N, \Omega, \mathbb{P}, (\mathcal{I}_j), (\succeq_j))$ with

- Ω represents uncertainty: multiple possible worlds for possible agent types
- $\mathbb{P}$ represents educated guesses on agent types
- $\mathcal{I}_j$ is similar to Coalition Structure: partition over possible worlds
  - $\{\{\omega_1, \omega_2\}, \{\omega_3\}, \{\omega_4, \omega_5\}\}$

BCG definition interpretation

$(N, \Omega, \mathbb{P}, (\mathcal{I}_j), (\succeq_j))$ with

- Ω represents uncertainty: multiple possible worlds for possible agent types
- $\mathbb{P}$ represents educated guesses on agent types
- $\mathcal{I}_j$ is similar to Coalition Structure: partition over possible worlds
  - $\{\{\omega_1, \omega_2\}, \{\omega_3\}, \{\omega_4, \omega_5\}\}$
  - each particular subset is called *information set*

BCG definition interpretation

$(N, \Omega, \mathbb{P}, (\mathcal{I}_j), (\succeq_j))$ with

- Ω represents uncertainty: multiple possible worlds for possible agent types
- $\mathbb{P}$ represents educated guesses on agent types
- $\mathcal{I}_j$ is similar to Coalition Structure: partition over possible worlds
    - $\{\{\omega_1, \omega_2\}, \{\omega_3\}, \{\omega_4, \omega_5\}\}$
    - each particular subset is called *information set*
- $\succeq_j$ represents preference over distributions of payoffs, (more or less) i.e contracts

Contracts

- A contract maps possible worlds to payoff vectors: $\mathbf{c}^S : \Omega \mapsto \mathbb{R}^S$

Contracts

- A contract maps possible worlds to payoff vectors: $\mathbf{c}^S : \Omega \mapsto \mathbb{R}^S$
  - c^S is called S-contract: contract over coalition S

Contracts

- A contract maps possible worlds to payoff vectors: $\mathbf{c}^S : \Omega \mapsto \mathbb{R}^S$
  - c^S is called S-contract: contract over coalition S
  - $c_j^S(\omega)$ denotes the payoff to agent $j \in S$ in world ω

Contracts

- A contract maps possible worlds to payoff vectors: $\mathbf{c}^S : \Omega \mapsto \mathbb{R}^S$
  - c^S is called S-contract: contract over coalition S
  - $c_j^S(\omega)$ denotes the payoff to agent $j \in S$ in world ω
-

Coalition	ω_1	ω_2
{1, 2}	(2.5, 3.0)	(6.0, 2.5)
{3}	0.5	7.0

Contracts

- A contract maps possible worlds to payoff vectors: $\mathbf{c}^S : \Omega \mapsto \mathbb{R}^S$
  - c^S is called S-contract: contract over coalition S
  - $c_j^S(\omega)$ denotes the payoff to agent $j \in S$ in world ω
-

Coalition	ω_1	ω_2
{1, 2}	(2.5, 3.0)	(6.0, 2.5)
{3}	0.5	7.0

- "A contract is a collection of payoff distributions for each possible world for the agents engaged in that contract"

Contracts

- A contract maps possible worlds to payoff vectors: $\mathbf{c}^S : \Omega \mapsto \mathbb{R}^S$
  - c^S is called S-contract: contract over coalition S
  - $c_j^S(\omega)$ denotes the payoff to agent $j \in S$ in world ω
-

Coalition	ω_1	ω_2
{1, 2}	(2.5, 3.0)	(6.0, 2.5)
{3}	0.5	7.0

- "A contract is a collection of payoff distributions for each possible world for the agents engaged in that contract"
- Captures the uncertainty in numbers

Contracts

- A contract maps possible worlds to payoff vectors: $\mathbf{c}^S : \Omega \mapsto \mathbb{R}^S$
  - c^S is called S-contract: contract over coalition S
  - $c_j^S(\omega)$ denotes the payoff to agent $j \in S$ in world ω
-

Coalition	ω_1	ω_2
{1, 2}	(2.5, 3.0)	(6.0, 2.5)
{3}	0.5	7.0

- "A contract is a collection of payoff distributions for each possible world for the agents engaged in that contract"
- Captures the uncertainty in numbers
- Let us focus on grand contracts, c^N

Timing of evaluation of contract

- Same principle, much more difficult to execute: solution concept *The Core*

Timing of evaluation of contract

- Same principle, much more difficult to execute: solution concept *The Core*
- Preference $\succeq_{j \in S}$: coalition prefers one contract over another if all agents in it get more payoff than they would in the other contract

Timing of evaluation of contract

- Same principle, much more difficult to execute: solution concept *The Core*
- Preference $\succeq_{j \in S}$: coalition prefers one contract over another if all agents in it get more payoff than they would in the other contract
- Imagine a contract with high payoff in lots of worlds except one, and that world has a high prior probability

Timing of evaluation of contract

- Same principle, much more difficult to execute: solution concept *The Core*
- Preference $\succeq_{j \in S}$: coalition prefers one contract over another if all agents in it get more payoff than they would in the other contract
- Imagine a contract with high payoff in lots of worlds except one, and that world has a high prior probability
- What do you know about the true state of the world ω^*?

Timing of evaluation of contract

- Same principle, much more difficult to execute: solution concept *The Core*
- Preference $\succeq_{j \in S}$: coalition prefers one contract over another if all agents in it get more payoff than they would in the other contract
- Imagine a contract with high payoff in lots of worlds except one, and that world has a high prior probability
- What do you know about the true state of the world ω^*?
  - *ex ante:* nothing: world is not yet drawn: c_j^S

Timing of evaluation of contract

- Same principle, much more difficult to execute: solution concept *The Core*
- Preference $\succeq_{j \in S}$: coalition prefers one contract over another if all agents in it get more payoff than they would in the other contract
- Imagine a contract with high payoff in lots of worlds except one, and that world has a high prior probability
- What do you know about the true state of the world ω^*?
  - *ex ante:* nothing: world is not yet drawn: c_j^S
  - *ex interim:* what *information set* the world belongs to: world is drawn, but not made common knowledge: $c_j^S(\omega)$ with probability $\mathbb{P}(\omega|\mathcal{I}_j(\omega^*))$, denoted $c_j^S|\mathcal{I}_j(\omega^*)$

Timing of evaluation of contract

- Same principle, much more difficult to execute: solution concept *The Core*
- Preference $\succeq_{j \in S}$: coalition prefers one contract over another if all agents in it get more payoff than they would in the other contract
- Imagine a contract with high payoff in lots of worlds except one, and that world has a high prior probability
- What do you know about the true state of the world ω^*?
  - *ex ante:* nothing: world is not yet drawn: c_j^S
  - *ex interim:* what *information set* the world belongs to: world is drawn, but not made common knowledge: $c_j^S(\omega)$ with probability $\mathbb{P}(\omega|\mathcal{I}_j(\omega^*))$, denoted $c_j^S|\mathcal{I}_j(\omega^*)$
  - *ex post:* everything: world is drawn and made common knowledge: $c_j^S|\omega^*$

Timing of evaluation of contract

- Same principle, much more difficult to execute: solution concept *The Core*
- Preference $\succeq_{j \in S}$: coalition prefers one contract over another if all agents in it get more payoff than they would in the other contract
- Imagine a contract with high payoff in lots of worlds except one, and that world has a high prior probability
- What do you know about the true state of the world ω^*?
  - *ex ante:* nothing: world is not yet drawn: c_j^S
  - *ex interim:* what *information set* the world belongs to: world is drawn, but not made common knowledge: $c_j^S(\omega)$ with probability $\mathbb{P}(\omega|\mathcal{I}_j(\omega^*))$, denoted $c_j^S|\mathcal{I}_j(\omega^*)$
  - *ex post:* everything: world is drawn and made common knowledge: $c_j^S|\omega^*$
- Possibly different solution for all three timings

Blocking a grand contract and the core

- Three timings, three forms of blocking, three forms of the core

Blocking a grand contract and the core

- Three timings, three forms of blocking, three forms of the core
- "Given a BCG $(N, \Omega, \mathbb{P}, (\mathcal{I}_j), (\succeq_j))$ and a grand contract c^N, a coalition S *ex-ante blocks* c^N if there exists an S-contract c^S such that $c^S_S \succ_S c^N_S$."

Blocking a grand contract and the core

- Three timings, three forms of blocking, three forms of the core
- "Given a BCG $(N, \Omega, \mathbb{P}, (\mathcal{I}_j), (\succeq_j))$ and a grand contract c^N, a coalition S *ex-ante blocks* c^N if there exists an S-contract c^S such that $c^S_S \succ_S c^N_S$."
- "It *ex-post blocks* c^N if there exists a world $\omega^* \in \Omega$ and an S-contract c^S such that $c^S_S|\omega^* \succ_S c^N_S|\omega^*$."

Blocking a grand contract and the core

- Three timings, three forms of blocking, three forms of the core
- "Given a BCG $(N, \Omega, \mathbb{P}, (\mathcal{I}_j), (\succeq_j))$ and a grand contract c^N, a coalition S *ex-ante blocks* c^N if there exists an S-contract c^S such that $c^S_S \succ_S c^N_S$."
- "It *ex-post blocks* c^N if there exists a world $\omega^* \in \Omega$ and an S-contract c^S such that $c^S_S|\omega^* \succ_S c^N_S|\omega^*$."
- Ex-interim blocking is more difficult, explained later

Blocking a grand contract and the core

- Three timings, three forms of blocking, three forms of the core
- "Given a BCG $(N, \Omega, \mathbb{P}, (\mathcal{I}_j), (\succeq_j))$ and a grand contract c^N, a coalition S *ex-ante blocks* c^N if there exists an S-contract c^S such that $c_S^S \succ_S c_S^N$."
- "It *ex-post blocks* c^N if there exists a world $\omega^* \in \Omega$ and an S-contract c^S such that $c_S^S|\omega^* \succ_S c_S^N|\omega^*$."
- Ex-interim blocking is more difficult, explained later
- "A grand contract c^N is in the *(ex-ante, ex-interim, ex-post)* core of a BCG if no coalition $S \subseteq N$ *(ex-ante, ex-interim, ex-post)* blocks c^N."

Ex-interim blocking

- Why is ex-interim blocking more difficult? Example:

Ex-interim blocking

- Why is ex-interim blocking more difficult? Example:
  - Two agents $\{1, 2\}$ and two worlds $\{\omega_1, \omega_2\}$, true world ω_1

Ex-interim blocking

- Why is ex-interim blocking more difficult? Example:
  - Two agents {1, 2} and two worlds $\{\omega_1, \omega_2\}$, true world ω_1
  - Information partitions $\mathcal{I}_1 = \{\{\omega_1, \omega_2\}\}$ and $\mathcal{I}_2 = \{\{\omega_1\}, \{\omega_2\}\}$, and grand contract c^N

Ex-interim blocking

- Why is ex-interim blocking more difficult? Example:
  - Two agents $\{1, 2\}$ and two worlds $\{\omega_1, \omega_2\}$, true world ω_1
  - Information partitions $\mathcal{I}_1 = \{\{\omega_1, \omega_2\}\}$ and $\mathcal{I}_2 = \{\{\omega_1\}, \{\omega_2\}\}$, and grand contract c^N
  - Preferences $c_2^S|\mathcal{I}_2(\omega_1) \succ_2 c_2^N|\mathcal{I}_2(\omega_1)$ and $c_2^N|\mathcal{I}_2(\omega_2) \succeq_2 c_2^S|\mathcal{I}_2(\omega_2)$

Ex-interim blocking

- Why is ex-interim blocking more difficult? Example:
  - Two agents $\{1, 2\}$ and two worlds $\{\omega_1, \omega_2\}$, true world ω_1
  - Information partitions $\mathcal{I}_1 = \{\{\omega_1, \omega_2\}\}$ and $\mathcal{I}_2 = \{\{\omega_1\}, \{\omega_2\}\}$, and grand contract c^N
  - Preferences $c_2^S|\mathcal{I}_2(\omega_1) \succ_2 c_2^N|\mathcal{I}_2(\omega_1)$ and $c_2^N|\mathcal{I}_2(\omega_2) \succeq_2 c_2^S|\mathcal{I}_2(\omega_2)$
  - Agent 2 *knows* that ω_1 is the true world! And will prefer c^S

Ex-interim blocking

- Why is ex-interim blocking more difficult? Example:
  - Two agents $\{1, 2\}$ and two worlds $\{\omega_1, \omega_2\}$, true world ω_1
  - Information partitions $\mathcal{I}_1 = \{\{\omega_1, \omega_2\}\}$ and $\mathcal{I}_2 = \{\{\omega_1\}, \{\omega_2\}\}$, and grand contract c^N
  - Preferences $c_2^S|\mathcal{I}_2(\omega_1) \succ_2 c_2^N|\mathcal{I}_2(\omega_1)$ and $c_2^N|\mathcal{I}_2(\omega_2) \succeq_2 c_2^S|\mathcal{I}_2(\omega_2)$
  - Agent 2 *knows* that ω_1 is the true world! And will prefer c^S
  - Agent 1 sees him preferring that, and can deduct that ω_1 is the true world too

Ex-interim blocking

- Why is ex-interim blocking more difficult? Example:
  - Two agents $\{1, 2\}$ and two worlds $\{\omega_1, \omega_2\}$, true world ω_1
  - Information partitions $\mathcal{I}_1 = \{\{\omega_1, \omega_2\}\}$ and $\mathcal{I}_2 = \{\{\omega_1\}, \{\omega_2\}\}$, and grand contract c^N
  - Preferences $c_2^S|\mathcal{I}_2(\omega_1) \succ_2 c_2^N|\mathcal{I}_2(\omega_1)$ and $c_2^N|\mathcal{I}_2(\omega_2) \succeq_2 c_2^S|\mathcal{I}_2(\omega_2)$
  - Agent 2 *knows* that ω_1 is the true world! And will prefer c^S
  - Agent 1 sees him preferring that, and can deduct that ω_1 is the true world too
- Agents can *observe* and *learn* from other agents

Ex-interim blocking

- Why is ex-interim blocking more difficult? Example:
  - Two agents $\{1, 2\}$ and two worlds $\{\omega_1, \omega_2\}$, true world ω_1
  - Information partitions $\mathcal{I}_1 = \{\{\omega_1, \omega_2\}\}$ and $\mathcal{I}_2 = \{\{\omega_1\}, \{\omega_2\}\}$, and grand contract c^N
  - Preferences $c_2^S|\mathcal{I}_2(\omega_1) \succ_2 c_2^N|\mathcal{I}_2(\omega_1)$ and $c_2^N|\mathcal{I}_2(\omega_2) \succeq_2 c_2^S|\mathcal{I}_2(\omega_2)$
  - Agent 2 *knows* that ω_1 is the true world! And will prefer c^S
  - Agent 1 sees him preferring that, and can deduct that ω_1 is the true world too
- Agents can *observe* and *learn* from other agents
- Hard to formalize, so we won't do that :)

Intuition and definition of ex-interim blocking

- *iterated elimination of dominated information sets*: eliminate worlds that are not 'attractive': worlds that are not preferred by any agents

Intuition and definition of ex-interim blocking

- *iterated elimination of dominated information sets*: eliminate worlds that are not 'attractive': worlds that are not preferred by any agents
- Result: set of worlds $\Omega^*(c^N, c^S)$ that *is* attractive to all agents engaged in contract c^S

Intuition and definition of ex-interim blocking

- *iterated elimination of dominated information sets*: eliminate worlds that are not 'attractive': worlds that are not preferred by any agents
- Result: set of worlds $\Omega^*(c^N, c^S)$ that *is* attractive to all agents engaged in contract c^S
- "Given a BCG $(N, \Omega, \mathbb{P}, (\mathcal{I}_j), (\succeq_j))$ and a grand contract c^N, a coalition S *ex-interim blocks* c^N if there exists an S-contract c^S such that for some $\omega \in \Omega^*(c^N, c^S)$, $c^S_S|\mathcal{I}_j(\omega) \cap \Omega^*(c^N, c^S) \succ_S c^N_S|\mathcal{I}_j(\omega) \cap \Omega^*(c^N, c^S)$."

Intuition and definition of ex-interim blocking

- *iterated elimination of dominated information sets*: eliminate worlds that are not 'attractive': worlds that are not preferred by any agents
- Result: set of worlds $\Omega^*(c^N, c^S)$ that *is* attractive to all agents engaged in contract c^S
- "Given a BCG $(N, \Omega, \mathbb{P}, (\mathcal{I}_j), (\succeq_j))$ and a grand contract c^N, a coalition S *ex-interim blocks* c^N if there exists an S-contract c^S such that for some $\omega \in \Omega^*(c^N, c^S)$, $c^S_S | \mathcal{I}_j(\omega) \cap \Omega^*(c^N, c^S) \succ_S c^N_S | \mathcal{I}_j(\omega) \cap \Omega^*(c^N, c^S)$."
- "Agents only consider contracts for worlds that other agents, given their information partitions, would find attractive"

Outline

Comparison between Repeated Coalition Forming and Bayesian Coalitional Games

- Motivation exactly the same

Comparison between Repeated Coalition Forming and Bayesian Coalitional Games

- Motivation exactly the same
- Approaches vastly different

Comparison between Repeated Coalition Forming and Bayesian Coalitional Games

- Motivation exactly the same
- Approaches vastly different
  - **Process:** sequential vs. single

Comparison between Repeated Coalition Forming and Bayesian Coalitional Games

- Motivation exactly the same
- Approaches vastly different
  - **Process:** sequential vs. single
  - **Idea behind:** learning vs. not learning

Comparison between Repeated Coalition Forming and Bayesian Coalitional Games

- Motivation exactly the same
- Approaches vastly different
  - **Process:** sequential vs. single
  - **Idea behind:** learning vs. not learning
  - **Time focus:** Predicting future vs. rationalizing about now

Comparison between Repeated Coalition Forming and Bayesian Coalitional Games

- Motivation exactly the same
- Approaches vastly different
  - **Process:** sequential vs. single
  - **Idea behind:** learning vs. not learning
  - **Time focus:** Predicting future vs. rationalizing about now
  - **Theory:** Multi-Agent System oriented vs. Cooperative Games oriented

Comparison between Repeated Coalition Forming and Bayesian Coalitional Games

- Motivation exactly the same
- Approaches vastly different
  - **Process:** sequential vs. single
  - **Idea behind:** learning vs. not learning
  - **Time focus:** Predicting future vs. rationalizing about now
  - **Theory:** Multi-Agent System oriented vs. Cooperative Games oriented
  - **Approach:** Determine optimal policy through POMDP vs. Solution concept *the core*

Comparison between Repeated Coalition Forming and Bayesian Coalitional Games

- Motivation exactly the same
- Approaches vastly different
  - **Process:** sequential vs. single
  - **Idea behind:** learning vs. not learning
  - **Time focus:** Predicting future vs. rationalizing about now
  - **Theory:** Multi-Agent System oriented vs. Cooperative Games oriented
  - **Approach:** Determine optimal policy through POMDP vs. Solution concept *the core*
  - **Uncertainty representation:** through beliefs vs. through prior probabilities

Comparison between Repeated Coalition Forming and Bayesian Coalitional Games

- Motivation exactly the same
- Approaches vastly different
  - **Process:** sequential vs. single
  - **Idea behind:** learning vs. not learning
  - **Time focus:** Predicting future vs. rationalizing about now
  - **Theory:** Multi-Agent System oriented vs. Cooperative Games oriented
  - **Approach:** Determine optimal policy through POMDP vs. Solution concept *the core*
  - **Uncertainty representation:** through beliefs vs. through prior probabilities
  - **Solution method:** maximize discounted reward vs. preferences and blocking

Comparison between Repeated Coalition Forming and Bayesian Coalitional Games

- Motivation exactly the same
- Approaches vastly different
  - **Process:** sequential vs. single
  - **Idea behind:** learning vs. not learning
  - **Time focus:** Predicting future vs. rationalizing about now
  - **Theory:** Multi-Agent System oriented vs. Cooperative Games oriented
  - **Approach:** Determine optimal policy through POMDP vs. Solution concept *the core*
  - **Uncertainty representation:** through beliefs vs. through prior probabilities
  - **Solution method:** maximize discounted reward vs. preferences and blocking
  - **Final outcome:** Best/fairest coalition!

Outline

Conclusion

- CGT falls short in real life due to not being able to represent uncertainty

Conclusion

- CGT falls short in real life due to not being able to represent uncertainty
- Two possible solutions: Repeated Coalition Forming and Bayesian Coalitional Games

Conclusion

- CGT falls short in real life due to not being able to represent uncertainty
- Two possible solutions: Repeated Coalition Forming and Bayesian Coalitional Games
- Completely different models and ways of finding the best coalition

Conclusion

- CGT falls short in real life due to not being able to represent uncertainty
- Two possible solutions: Repeated Coalition Forming and Bayesian Coalitional Games
- Completely different models and ways of finding the best coalition
- Many other possible ways of dealing with uncertainty in coalition formation

Thank you!!

Questions?