Negotiation and Bargaining

S.S. Fatima and I. Rahwan

Outline

- Introduction
- Aspects of negotiation
- Game-theoretic models of single-issue negotiation
- Game-theoretic models of multi-issue negotiation
- Heuristic approaches
- Negotiation with humans
- Argumentation-based negotiation

Negotiation

A form of interaction in which a group of agents with conflicting interests try to come to a mutually acceptable agreement over some outcome

The outcome is typically represented in terms of the allocation of resources such as

- Commodities
- Services
- Time
- Money
- CPU cycles

Why negotiate?

- The agents' preferences for the possible outcomes are conflicting

- None of the agents has the power to decide an outcome on its own

Aspects of negotiation

Aspects of negotiation

- The agents conducting negotiation
- The set of issues under negotiation
- The set of possible outcomes and an agent's preferences for them
- The protocol according to which agents search for a specific agreement, and
- The individual strategies that determine the agent's behavior

Utility functions – single issue

- An agent's preferences over the set of possible outcomes *O* is defined as a utility function

- An agent's (say *i*'s) utility function assigns a real number to each possible outcome. It is a mapping of the form

$$U^i : O \longrightarrow R$$

Utility functions – multiple issues

- In case of multiple issues, the utility function is a is a mapping of the form

$$U^i : A_1 \times \ldots \times A_m \rightarrow R$$

where A_k is set of possible values for the issue k

- Typically, the cumulative utility is the weighted sum of the utilities from the individual issues and is of the form

$$U^i = \sum_{k=1}^{m} w_k^i u_k^i(a_k) \quad where \quad u_k^i : A_k \rightarrow R$$

Negotiation: Key approaches

- Game theoretic
  - Cooperative (Axiomatic)
  - Non-cooperative
- Heuristic

Game theoretic approaches for single-issue negotiation

Axiomatic approach

- The idea is to describe the players and utility functions, decide on some characteristics an equilibrium should have (ex., fairness/ efficiency), mathematize the characteristics and show what outcome will result.
- There is a set of possible or feasible outcomes some of which are acceptable/ reasonable outcomes. The problem then is to find a bargaining function that maps the set of possible outcomes to the set of acceptable ones
- In his seminal work, Nash analyzed the bargaining problem and defined a solution for it using this approach

Nash's axiomatic model (1)

- Assumption: perfect rationality
- The players *a* and *b* want to come to an agreement over the alternatives in an arbitrary set *A*
- Failure to reach an agreement, i.e., disagreement is represented by a designated outcome denoted *D*
- For agent $i \in \{a, b\}$, the utility is

$$U^i : \{ A \cup \{D\}\} \rightarrow R$$

Nash's axiomatic model (2)

The set of all utility pairs that result from an agreement is the bargaining set S:

$$S = \{(U^a(z), U^b(z)) \subset R^2 : z \in A\}$$

$d^i = U^i(D)$ and $d \in R^2$ is called the *disagreement/ threat* point

The bargaining problem

- A bargaining problem is defined as a pair *(S, d)*.
- A bargaining solution is a function *f* that maps every bargaining problem *(S, d)* to an outcome in *S*, i.e.,

$$f: \quad (S,d) \rightarrow S$$

Nash's bargaining axioms (1)

- Individual rationality: The bargaining solution should give neither player less than what he/she would get from disagreement
- Symmetry: If the players' utility functions and their disagreement utilities are the same, they receive equal shares
- Efficiency: The bargaining solution should be feasible and Pareto optimal

Nash's bargaining axioms (2)

- Invariance: The solution should not change as a result of linear changes to the utility of either player
- Independence of irrelevant alternatives: Eliminating feasible alternatives, other than the disagreement point, that would not have been chosen should not affect the solution

The bargaining solution

The solution that satisfies the above axioms is given by

$$f(S,d) \in \arg\max_{x \in S, x \geq d} (x^a - d^a)(x^b - d^b)$$

and the solution is unique

Non-cooperative model of single-issue negotiation

- Unlike the axiomatic model, the non-cooperative model specifies a bargaining protocol and analyses the strategic behavior in terms of it

- Perhaps the most influential non-cooperative model is that of Rubinstein

Rubinstein's model

- Two players *a* and *b* and a unit of good to split
- The good is divisible
- If a gets a share of x^a, *b* will get $x^b = 1 - x^a$
- A strategic form game is played over a series of discrete time periods $t = 1, 2, \ldots$
- The players take turns in making offers (alternating offers protocol)
- A player's utility gets discounted in every time period

Utilities

If *a* gets a share of x^a and *b* gets x^b, their utilities will be:

$$U^a = x^a \delta_a^{t-1} \quad and \qquad U^b = x^b \delta_b^{t-1}$$

where δ_a is *a*'s discount factor and δ_b that of *b*

Infinite horizon game

If the discounted game is played infinitely over time, then the unique perfect equilibrium outcome will be

$$x^a = \frac{1-\delta_b}{1-\delta_a\delta_b} \qquad and \qquad x^b = \frac{\delta_b-\delta_a\delta_b}{1-\delta_a\delta_b}$$

and will result in the first time period

Bargaining with deadlines

- It is stipulated that negotiation will end after a fixed number of rounds (n)

- Equilibrium is obtained using backward induction

Equilibrium strategies (1)

Agent *a*'s equilibrium strategy for the last time period (*n*) will be:

if *a*'s turn to offer: OFFER (δ^{n-1}, 0)

if *b*'s turn to offer: ACCEPT

Equilibrium strategies (2)

- For the periods $t < n$, it will be:
  if a's turn to offer:
  OFFER $(\delta^{t-1} - X^b(t+1), X^b(t+1))$
  if b receives the offer (x^a, x^b):
  If $(U^a(x^a, t) \geq UA(t+1))$
  ACCEPT
  else REJECT

  where $X^b(t)$ denotes *b*'s equilibrium share for time *t* and $UA(t)$ *a*'s equilibrium utility for time *t*

- An agreement results in the first time period

Game theoretic approaches for multi-issue negotiation

Multi-issue procedures (1)

Multiple issues can be bargained using one of several different procedures:

- Global bargaining procedure (also called package deal procedure): Here, all the issues are addressed at once
- Independent/separate bargaining procedure: Negotiations over the individual issues separate and independent with each having no effect on the other

Multi-issue procedures (2)

- Sequential bargaining with independent implementation: The issues considered sequentially, one at a time, and an agreement on an issue goes into effect immediately (i.e., before negotiation begins on the next issue)

- Sequential bargaining with simultaneous implementation: The issues considered sequentially, one at a time, and an agreement on an issue does not take effect until an agreement is reached on all the subsequent issues

Cooperative and non-cooperative models of multi-issue negotiation

- Cooperative models: Give axioms that relate the outcomes of different procedures and characterize solutions satisfying those axioms

- Non-cooperative models: Analyze the strategic behavior of agents in terms of a negotiation protocol

Strategic multi-issue negotiation

- For the simultaneous procedure, the equilibrium strategies for the individual issues remain the same as that for single issue negotiation
- For the sequential procedure with independent implementation, the equilibrium strategies for the individual issues remain the same as that for single issue negotiation
- For the package deal procedure, the equilibrium strategies are obtained by solving a tradeoff problem

Divisible issues: making tradeoffs

For making tradeoffs at time *t*, the following optimization problem must be solved:

$$Max \sum_{k=1}^{n} w_k^a x_k^a$$

$$s.t. \quad \sum_{k=1}^{n} w_k^b (1 - x_k^a) \geq UB(t+1) \qquad 0 \leq x_k^a \leq 1$$

(Agent *a* is the offering agent for time *t* and UB(t) denotes b's equilibrium utility for time *t*)

This is the real knapsack problem

Indivisible issues: making tradeoffs

For making tradeoffs at time *t*, the following optimization problem must be solved:

$$Max \quad \sum_{k=1}^{n} w_k^a x_k^a$$

$$s.t. \quad \sum_{k=1}^{n} w_k^b (1 - x_k^a) \geq UB(t+1) \qquad x_k^a \in \{0,1\}$$

This is the integer knapsack problem (NP hard)

Heuristic Approaches

Heuristic approaches for multi-issue negotiation

- This approach is particularly useful when there are multiple issues to negotiate, and finding an equilibrium offer is computationally hard

- Heuristics can be used for
  - Generating counter-offers
  - Predicting opponent preferences/strategies
  - Generating negotiation agendas
  - Reasoning about deliberation cost

Heuristics for counter-offer generation

- Counter-offers are typically generated using negotiation decision functions (NDFs) as heuristics

- An NDF is a mapping from some parameters of negotiation (such as the players' reserve prices, the current time, and the deadline) to a counter offer

Negotiation decision functions

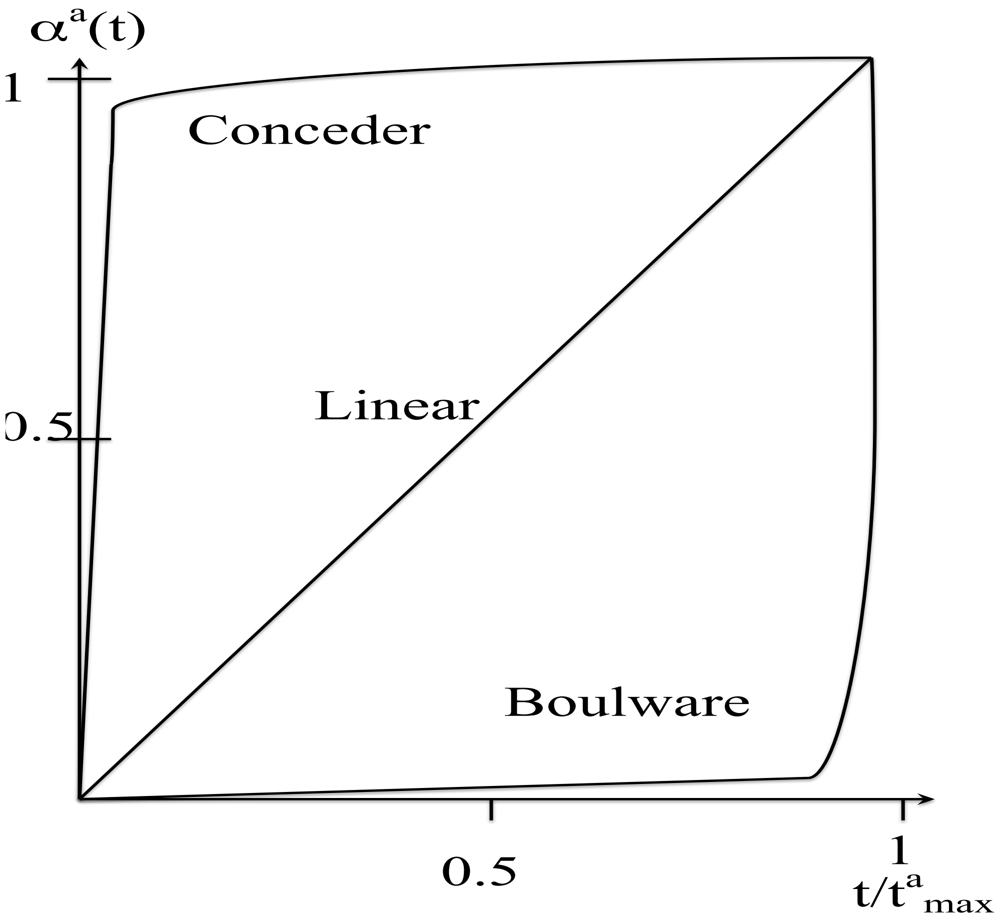

$\alpha^a(t)$: Price offered by agent *a* at time *t*

$t^a{}_{max}$: *a*'s deadline

Heuristics for predicting opponent preferences

- Use past offers to figure out what the opponent really wants, then generate similar offers that you like
  - -Use hill climbing on space of possible offers
  - -Generate similar offers e.g. using fuzzy measure

- Use mediator:
  - - Parties send offer to (trusted) mediator
  - - Mediator uses hill-climbing or simulated annealing to search for Pareto improving deals

Heuristics for generating an optimal agenda

- The set of issues included in a negotiation is called the negotiation agenda
- In competitive negotiations, different agents may prefer different agendas
- An agenda that maximizes an agent's utility is called its optimal agenda
- Finding an agent's optimal agenda is a computationally hard problem
- The use of evolutionary methods such as genetic algorithms for determining an agent's optimal agenda

Heuristics for reasoning about deliberation cost

- Agents negotiating over resources
- To evaluate an offer, an agent may have to solve an NP-complete problem E.g. with the resources I get, what is the utility of the best vehicle routing schedule I can get?
- Possible to reason explicitly about the cost of deliberation
  - Control an anytime algorithm based on the degree to which it is expected to improve the solution
  - New concept of deliberation equilibrium

Negotiating with Humans

Why would agents negotiate with people?

- Provide negotiation training environment for people

- Negotiate in Web-based markets that include both humans and bots

Why is it different?

- Humans make systematic deviations from (theoretically) optimal behavior
  - Preferences depend on the *framing* of offers
  - Aversion in unfair offers
  - Willing to punish greed even at cost to themselves
- Therefore, AI agents need to:
  - Predict human peculiar negotiation behavior
  - React appropriately given human perception

Colored trails experimental platform

- Computer lab-based experiments
- Human-human, or human vs agents
- Task-based environment:
  - *n x m* board of colored tiles
  - Each player possesses colored chips
  - Chips used to move on same-colored tiles
  - Negotiation to exchange chips
  - Control information availability: can see other's location, other's goal, other's chips?

“me” wants to get to goal “G”

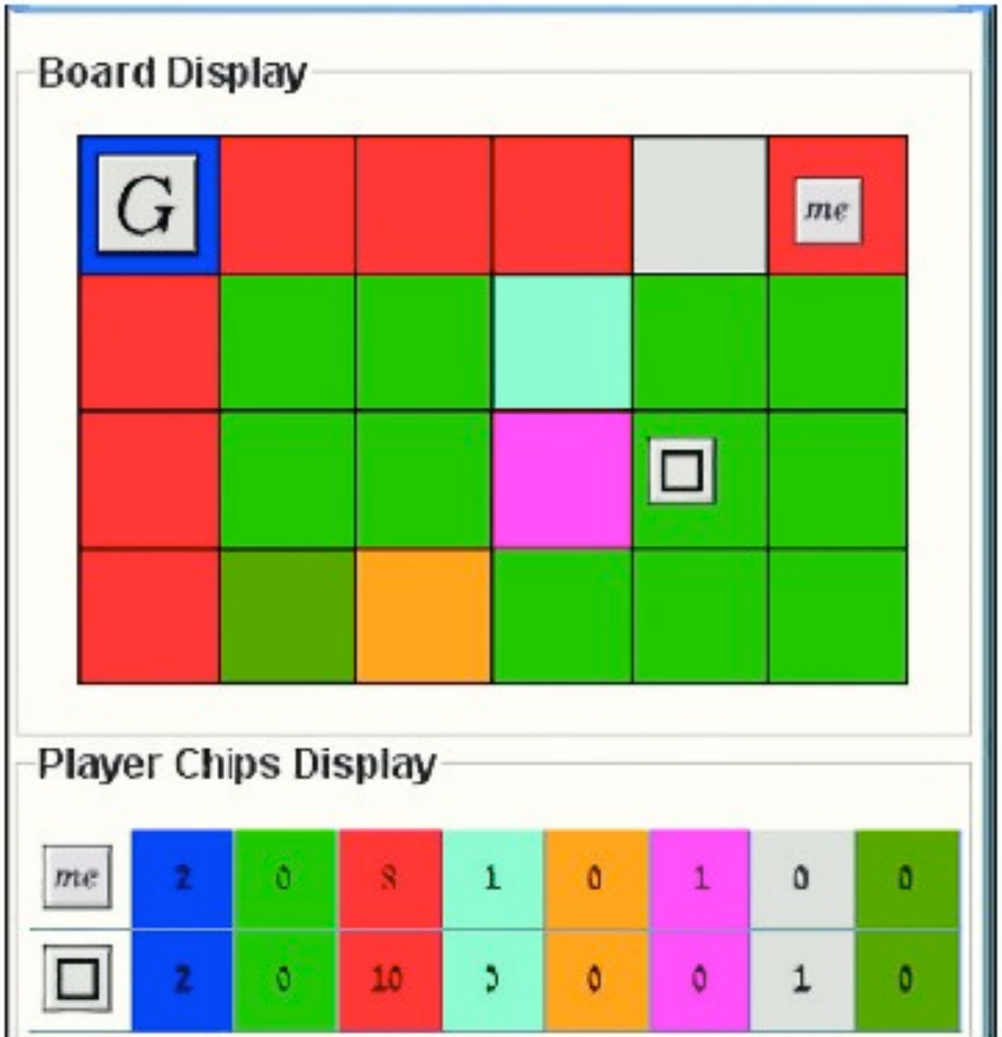

Proposal to exchange chips

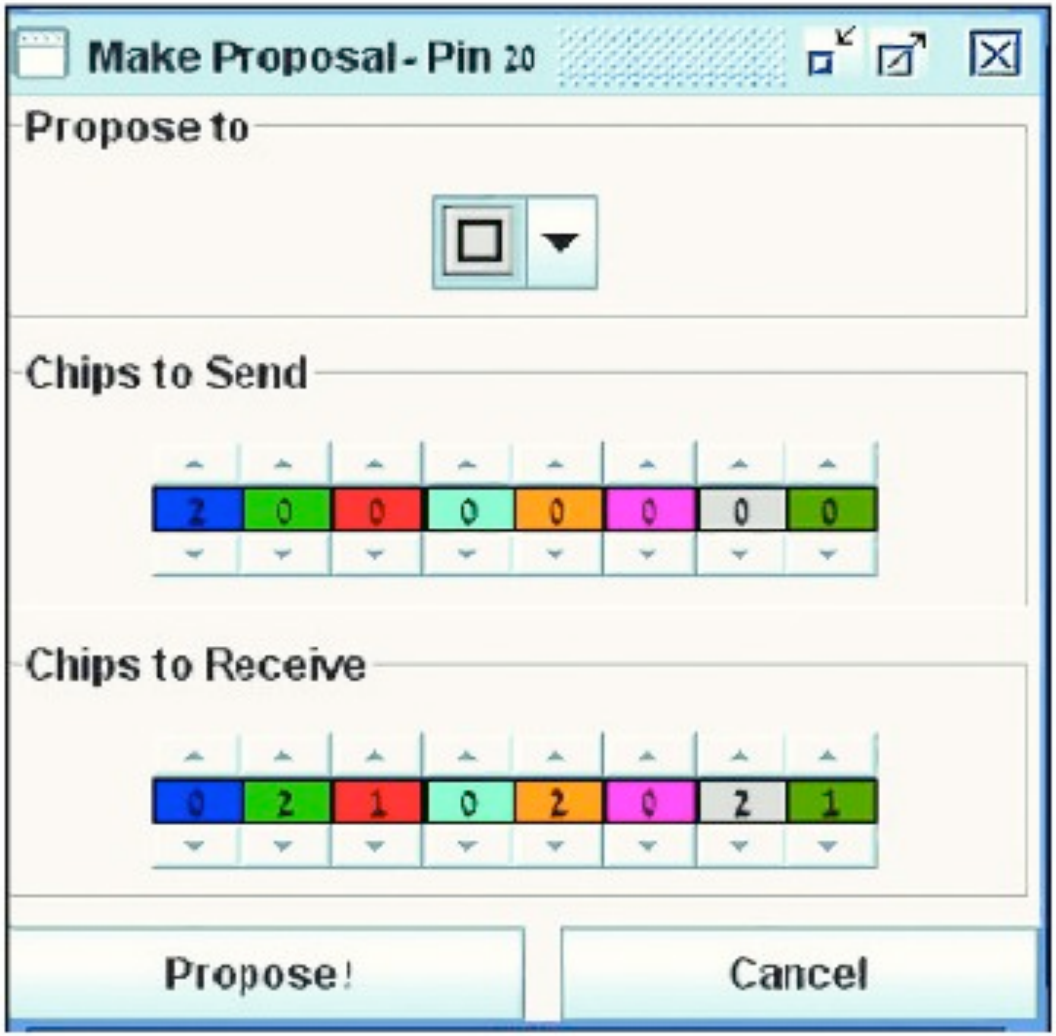

Some colored trails studies

- Predictive models of human reciprocity
- Predictive models of human's willingness to reveal their goals
- Agents that adapt to gender-specific and culture-specific negotiation behavior
- Study how humans *design* strategies

Argumentation-based Negotiation

Proposal vs. argument-based negotiation

- In *proposal*-based negotiation:
  -Exchange offers (potential deals)
- In *argument*-based negotiation:
  - Exchange meta-information (about goals, beliefs, plans etc.)
  - Reasons why you want to exchange resources in the first place
- Argument can shape/expand the space of possible agreement

Hanging pictures [Parsons et al]

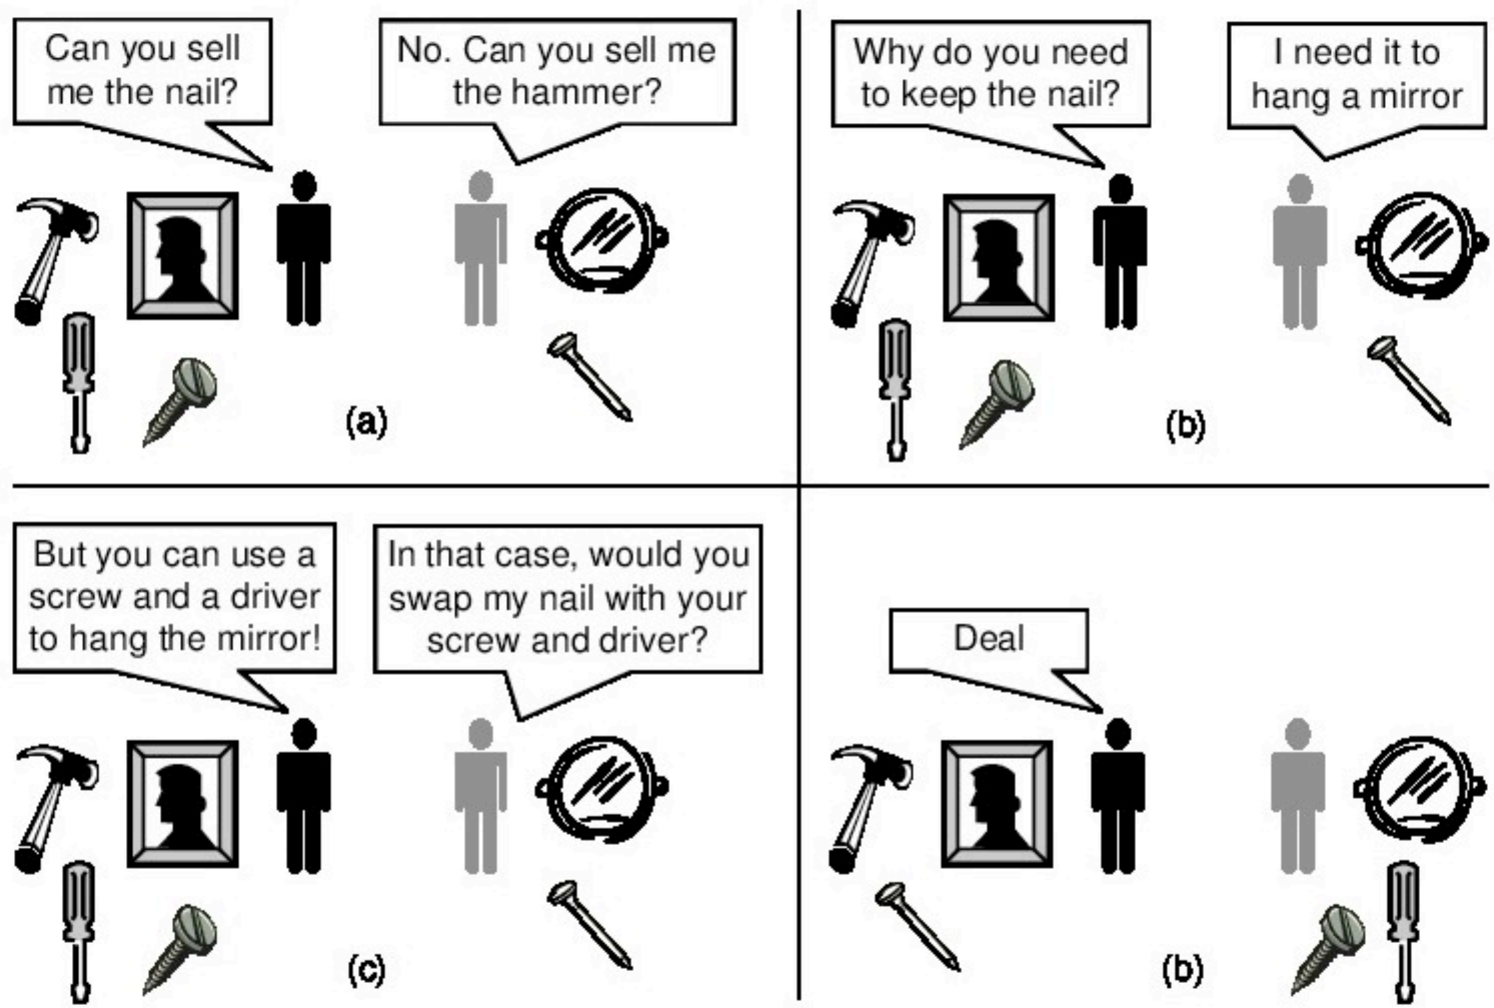

Challenges

- Finding bargaining strategies and solutions that are not only individual rational and Pareto optimal but also computationally feasible
- Designing expressive, tractable, and welfare improving argumentation protocols
- Designing agents that can not only negotiate with other agents but also with humans
- Still much work to be done