Probabilistic Model Checking for Games of imperfect information

P. Ballarini, M. Fisher, M. Wooldridge

University of Liverpool

What is this work about

Uncertainty is relevant for a specific class of Games (game of imperfect information)

Q: Can we apply probabilistic model checking for analysing games in which players' behaviour is characterised by uncertainty ?

Motivation

Model checking for Multi-Agent Systems (MAS)

LTL model checking of BDI MAS (Bordini)

AgentSpeak -> Promela (SPIN)

AgentSpeak -> Java (Java-PathFinder)

can we extend it to probabilistic model checking so that uncertain behaviour can be accounted for?

We need a new language for uncertain MAS (Probmela)

Can we use any existing Probabilistic Modelling Framework (PRISM?) to reason about uncertain MAS ?

Outline

- Games, strategies, equilibria
  - strategic games, equilibria
  - extensive games (perfect/imperfect information)
    - Alternating offers negotiation game
- Markovian model of the Alternating offers game
- Analysis through Model Checking
- Conclusion

Strategic Games

the outcome of the game is achieved in one-shot

- set of players: N={1,..,n}
- players actions: Ai={a1, a2,....ak}
- players preferences: a relation over outcome utilities $\succeq_i$

$$G = \langle N, (A_i), (\succeq_i) \rangle$$

an action profile is combination of actions: $a = (a^1, a^2, \ldots, a^n)$

the outcome of an action profile is denoted: $\mathcal{O}(a^1, a^2, \ldots, a^n)$

Example- Battle of Sexes

- two people wish to go out together to a concert of music by either the "Red Hot Chili Peppers" or "Bach"
- their main concern is to go out together but one prefers the "Peppers" and the other one "Bach"
- individual's preferences are represented by payoff functions

		Jane: Peppers	Jane: Bach
Stephen	Peppers	(2,1)	(0,0)
Stephen	Bach	(0,0)	(1,2)

Example- Battle of Sexes

		Jane	
		Peppers	Bach
Stephen	Peppers	(2,1)	(0,0)
	Bach	(0,0)	(1,2)

Stephen's preferences

$$(\text{Peppers},\text{Peppers}) \succeq_S (\text{Bach},\text{Bach}) \succeq_S (\text{Peppers},\text{Bach}) \sim_S (\text{Bach},\text{Peppers})$$

Jane's preferences

$$(\text{Bach},\text{Bach}) \succeq_J (\text{Peppers},\text{Peppers}) \succeq_J (\text{Peppers},\text{Bach}) \sim_J (\text{Bach},\text{Peppers})$$

Nash Equilibria

- a profile of actions is a Nash Equilibria iff no player has interest in adopting another strategy assuming the other player sticks to his one

Jane

		Peppers	Bach
Stephen	Peppers	(2,1)	(0,0)
	Bach	(0,0)	(1,2)

the Battle of Sexes has 2 Equilibria: (Peppers,Peppers), (Bach,Bach)

i.e. : togetherness rules

Extensive Games

They are sequential strategic games

(the decision problem is iterated over time)

- set of players N={1,..,n}
- set of histories H
- preferences over histories (rather than over action profiles)
- a player function: P(h) is the player who takes an action of history h

$$G = \langle N, H, P, (\succeq_i) \rangle$$

Extensive Games as Trees

Ext. Game example: two people propose different allocations for 2 indivisible items

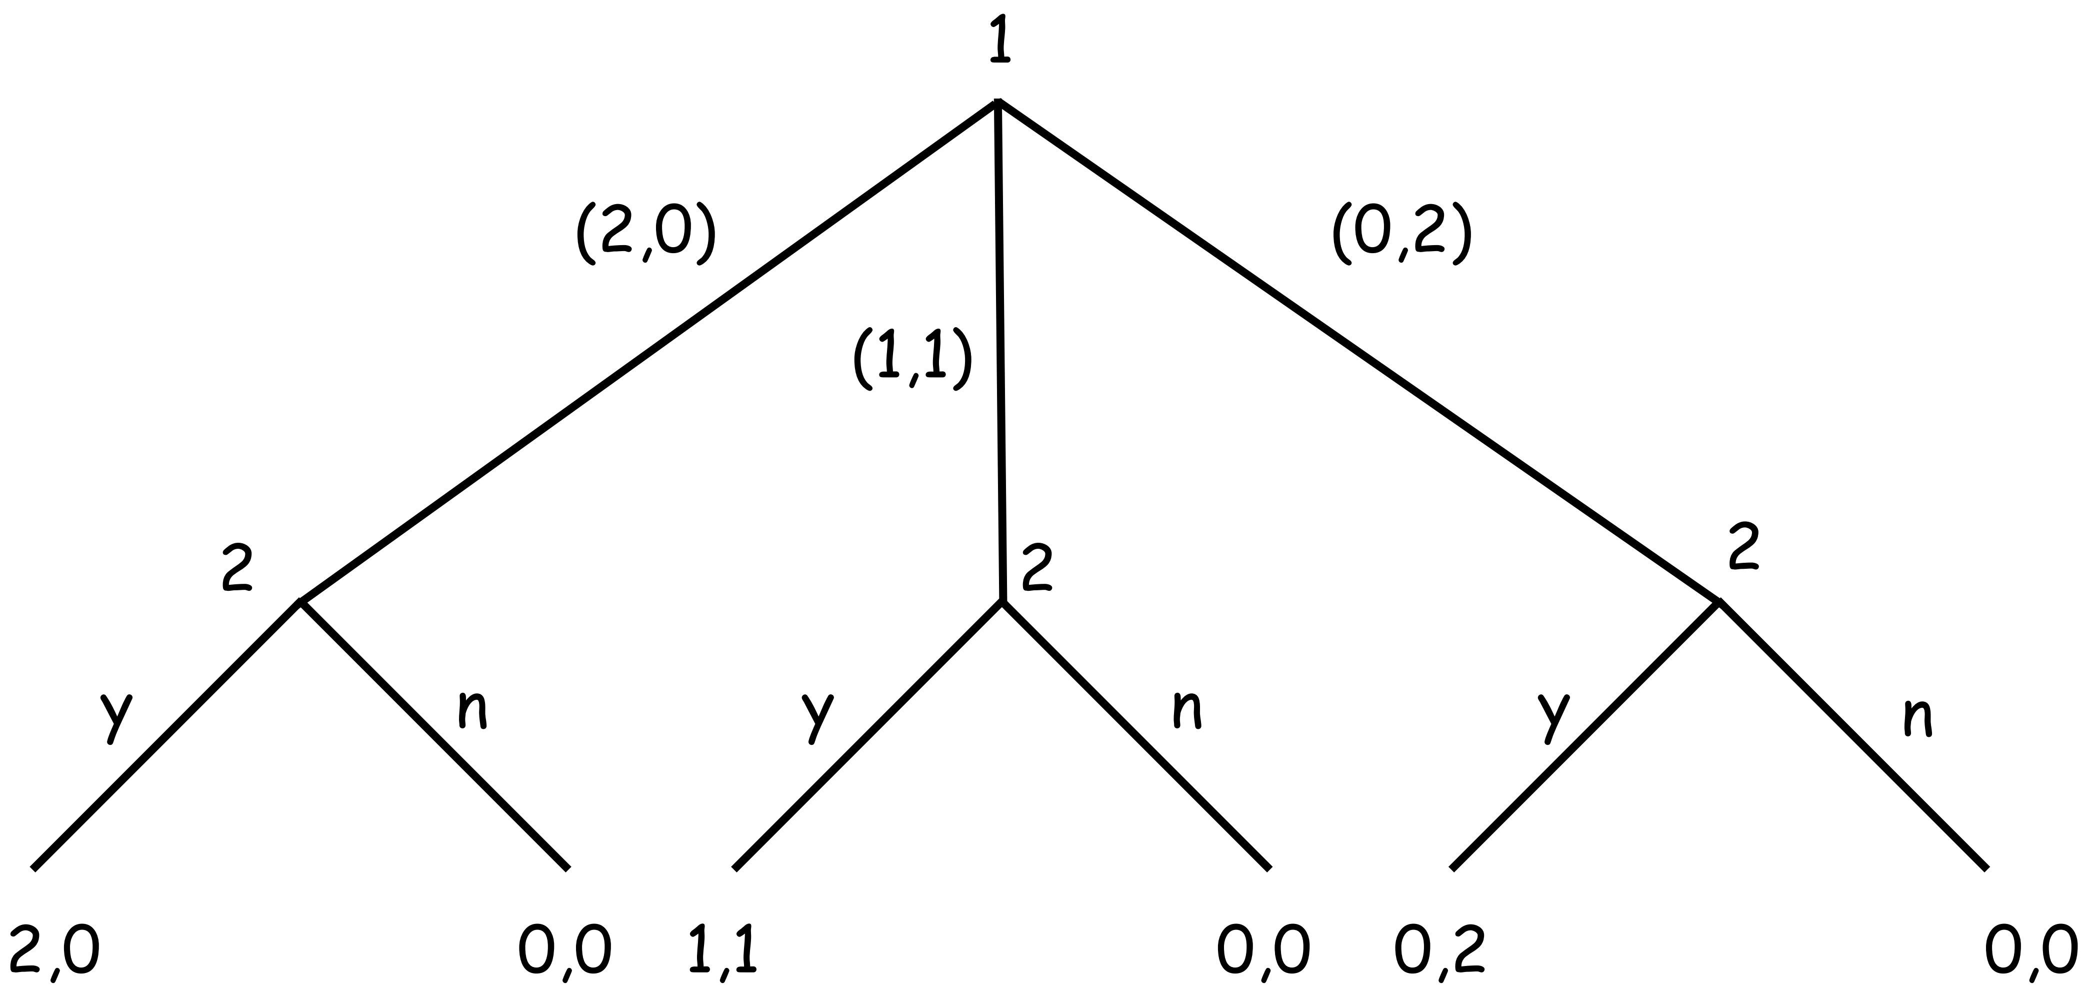

Perfect information: strategies

a strategy in an Ext. Game of perfect information is a function that assign an action to each non-terminal history

(Perf.Inf. assumption: players are completely informed on past actions)

strategies examples

1- $s_1(e)=(2,0)$ $s_2((2,0))=y$ $s_2((1,1))=n$ ..... $\xrightarrow{\text{outcome}}$ (2,0)

2- $s_1(e)=(2,0)$ $s_2((2,0))=n$ $s_2((1,1))=y$ ..... $\xrightarrow{\text{outcome}}$ (?)

perfect information: equilibria

a Nash Equilibria of an Ext. Game of perfect information is a strategy profile s=(s1,s2,..,sn) such that no player would get a better outcome by choosing a different strategy assuming all other players are sticking with their ones

Formally: a profile $s^* = (s_1^*, \ldots, s_n^*)$ is a Nash Equilibria iff

$$\mathcal{O}(s_{-i}^*, s_i^*) \succeq_i \mathcal{O}(s_{-i}^*, s_i) \quad \text{for all strategy } s_i \text{ of player } i$$

$\mathcal{O}(s^*)$: outcome for $s^* = (s_1^*, \ldots, s_n^*)$

Alternating offers game

(Rubinstein)

two players aim to split a pie (or bargain over an item)

players alternatively propose agreements in the set:

$$X = \{(x_1, x_2) | x_i \geq 0 \text{ and } x_1 + x_2 = 1\}$$

D: disagreement

players either accept (Y) or Reject (N) the most recent offer they receive

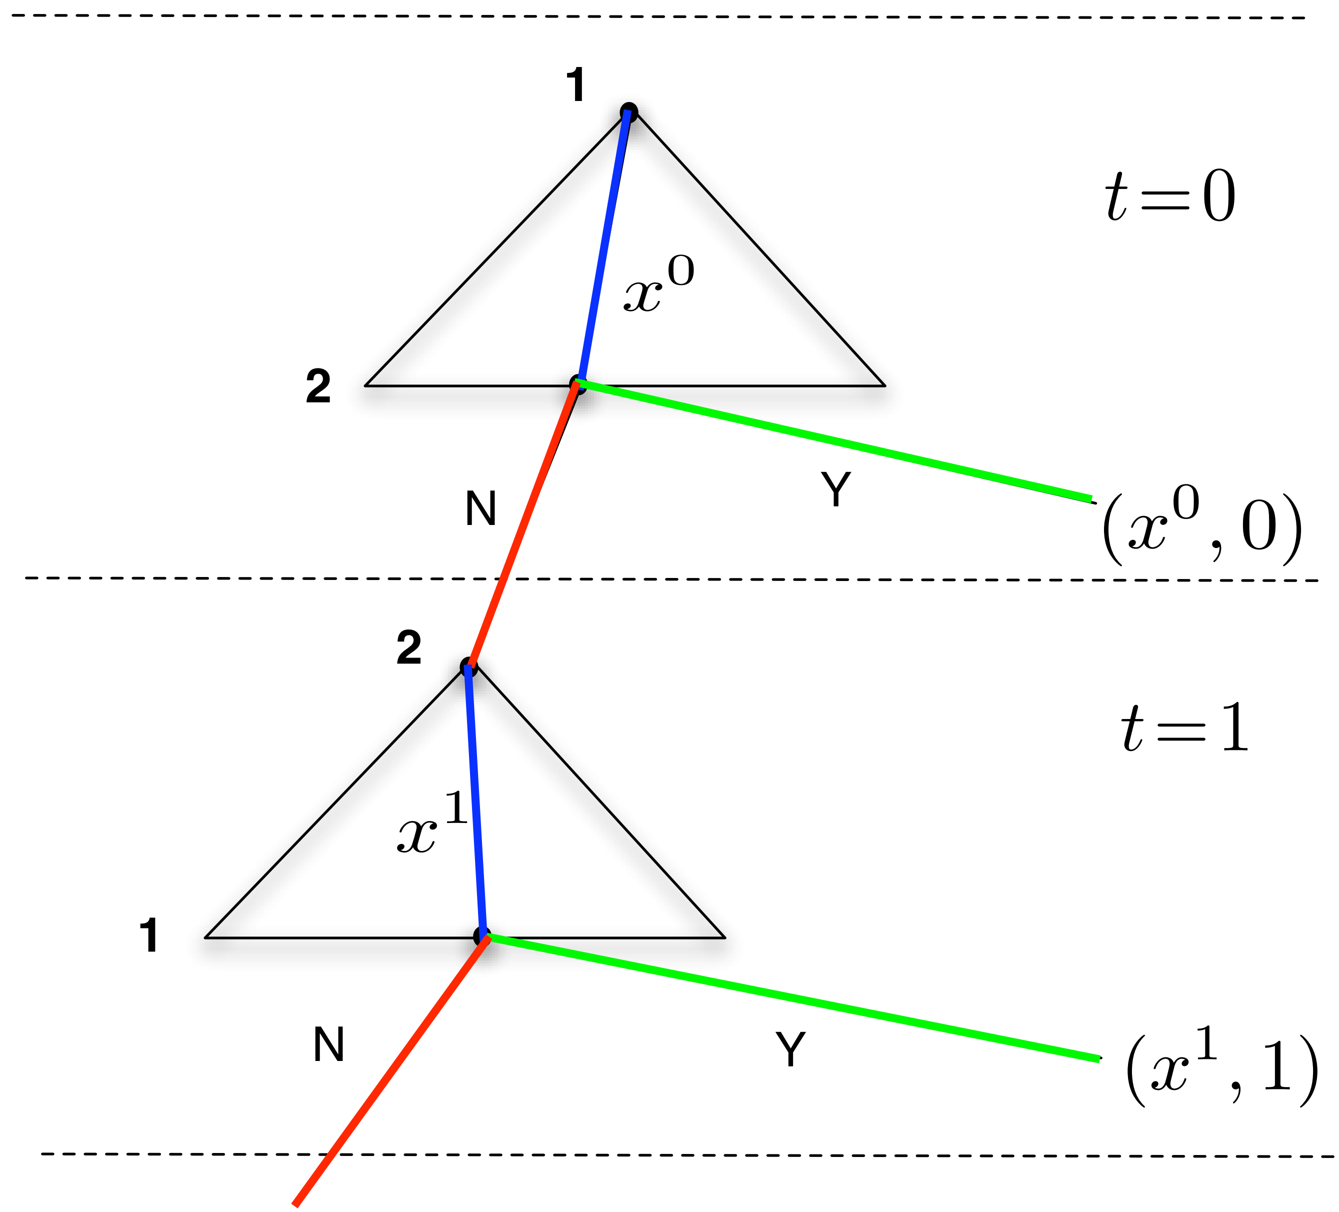

1
$t=0$
x^0
2
N
Y
$(x^0, 0)$
2
$t=1$
x^1
1
N
Y
$(x^1, 1)$

Alternating offers game

(Rubinstein)

formally an Alt. Offers Game is given by:

$$G = \langle\!\langle \{1,2\}, X \cup \{D\}, (\succeq_i) \rangle\!\rangle$$

where preferences are time-dependent

$\succeq_i$ is defined over $(X \times T) \cup \{D\}$

histories are of type

$(x^0, N, x^1, N \ldots, X^t)$ non-terminal

$(x^0, N, x^1, N \ldots, X^t, Y)$ terminal

Alternating offers preferences

$\succeq_i$ must fulfils some "basic" constraints

i- disagreement is the worst possible outcome

$$(x \times t) \succeq_i D$$

ii- pie is desirable

$$(x \times t) \succ_i (y \times t) \iff x_i > y_i$$

iii- time is valuable

$$(x \times t) \succ_i (x \times s) \text{ if } t < s$$

Alternating offers: equilibria

Given an Alter. Offers game

$$G = \langle \{1,2\}, X \cup \{D\}, (\succeq_i) \rangle$$

PROPERTY: there are infinite Nash Equilibria

Equilibria example

strategy: players keep asking the whole pie until time t=n then they ask x^* and each player will accept only x^*

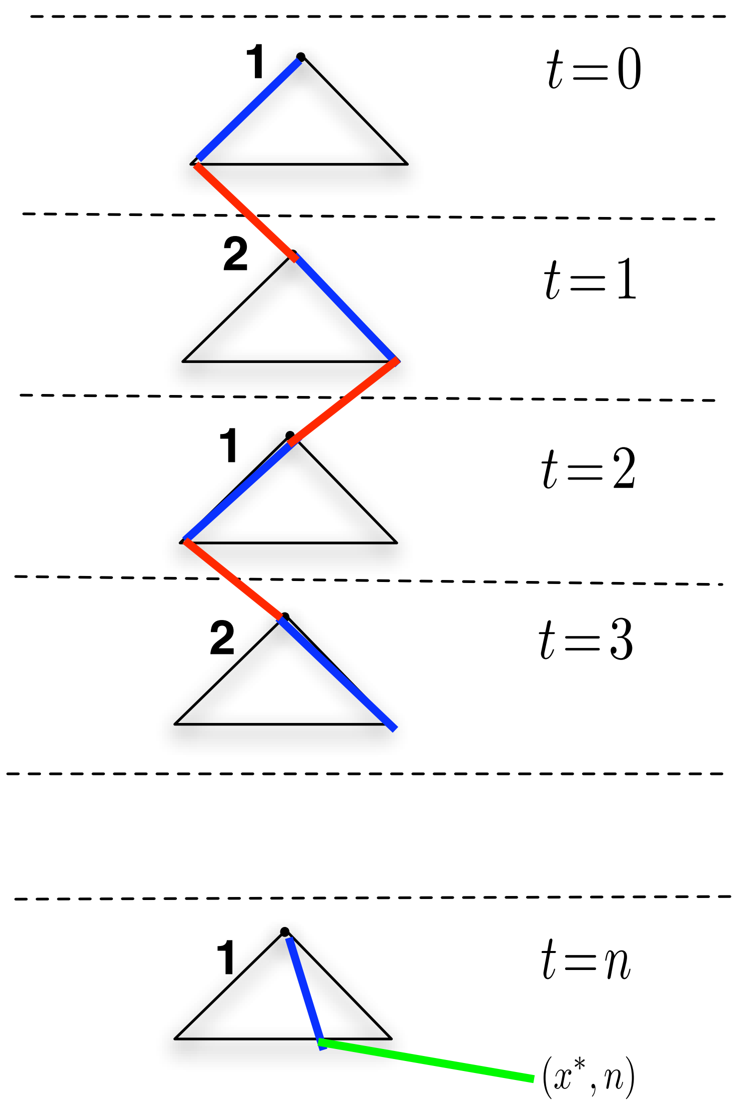

Preferences: more constraints

iv- stationarity

$$(x \times t) \succ_i (y \times t+1) \text{ iff } (x \times 0) \succ_i (y \times 1)$$

v- increasing loss to delay

$$x_i - v_i(x^i, 1) \text{ increasing function of } x_i$$

Alternating offers: equilibria

THEOREM: if $\succeq_i$ fulfils all constraints i-v then there exists a unique strategy profile (σ^*, δ^*) which is a Nash Equilibria

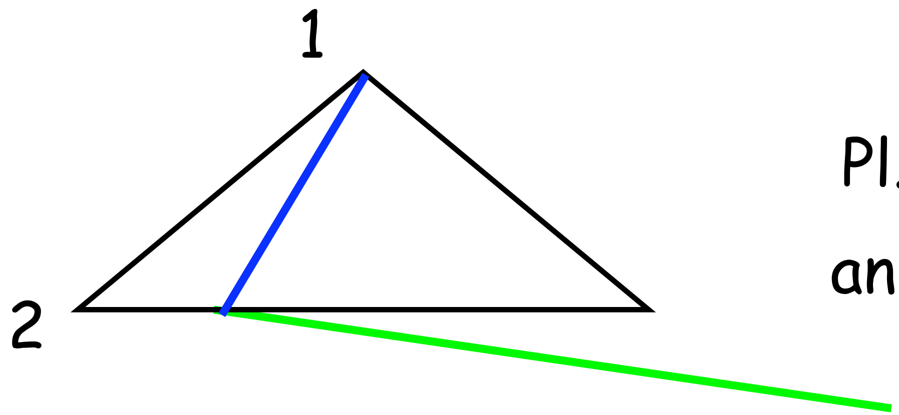

Equilibria

Pl. 1 proposes (x_1^*, x_2^*)

and Pl. 2 accepts straight away

$((x_1^*, x_2^*), 1)$

(x_1^*, x_2^*) is depends on both $\succeq_1$ and $\succeq_2$

Imperfect information: strategies

Imperf.Inf. assumption: players may have only partial info on past actions.

as a result some actions are determined by chance

$$G = \langle N, H, P, f_c, (\mathcal{I}_i)(\succeq_i) \rangle$$

$P(h) = c$ the next action for history h is determined by the lottery $f_c(h)$

a strategy in an Ext. Game of imperfect information is a function that assign to each non-terminal history a lottery over possible actions

preferences are over (induced) lotteries on the set of terminal histories

Markovian model of Negotiation

Markov processes are suitable for modelling past-independent behaviours

(hence imperfect-information games)

we consider the imperfect-information variant of the alternating offers game

which is: we assume players actions being state-dependent, rather than path-dependent

Markovian model of Negotiation

the imperfect-info alternating offer game can be naturally encoded as a DTMC

(players decision is a lottery over the possible actions)

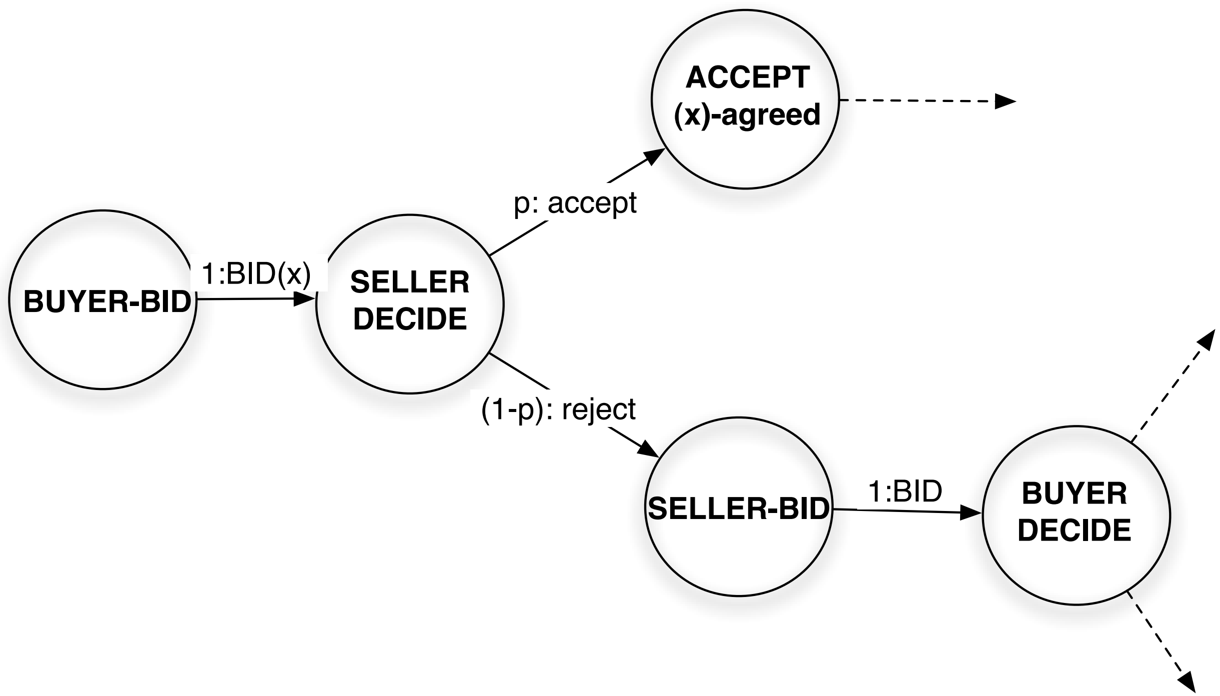

Markovian model of Negotiation

players' strategies depend on 2 parameters

i)- the Offer proposal function $p_{a \to \hat{a}}(t)$

IP^b initial price proposed by player b

RP^b reserved price of player b

T^b time-deadline of player b

ii)- the Acceptance Probability function

$S\_AP(x)$ for the Seller

$B\_AP(x)$ for the Buyer

Offer Function families

Conceder: player concedes a lot in early stage of negotiation

Boulware: player concedes a lot only close to deadline

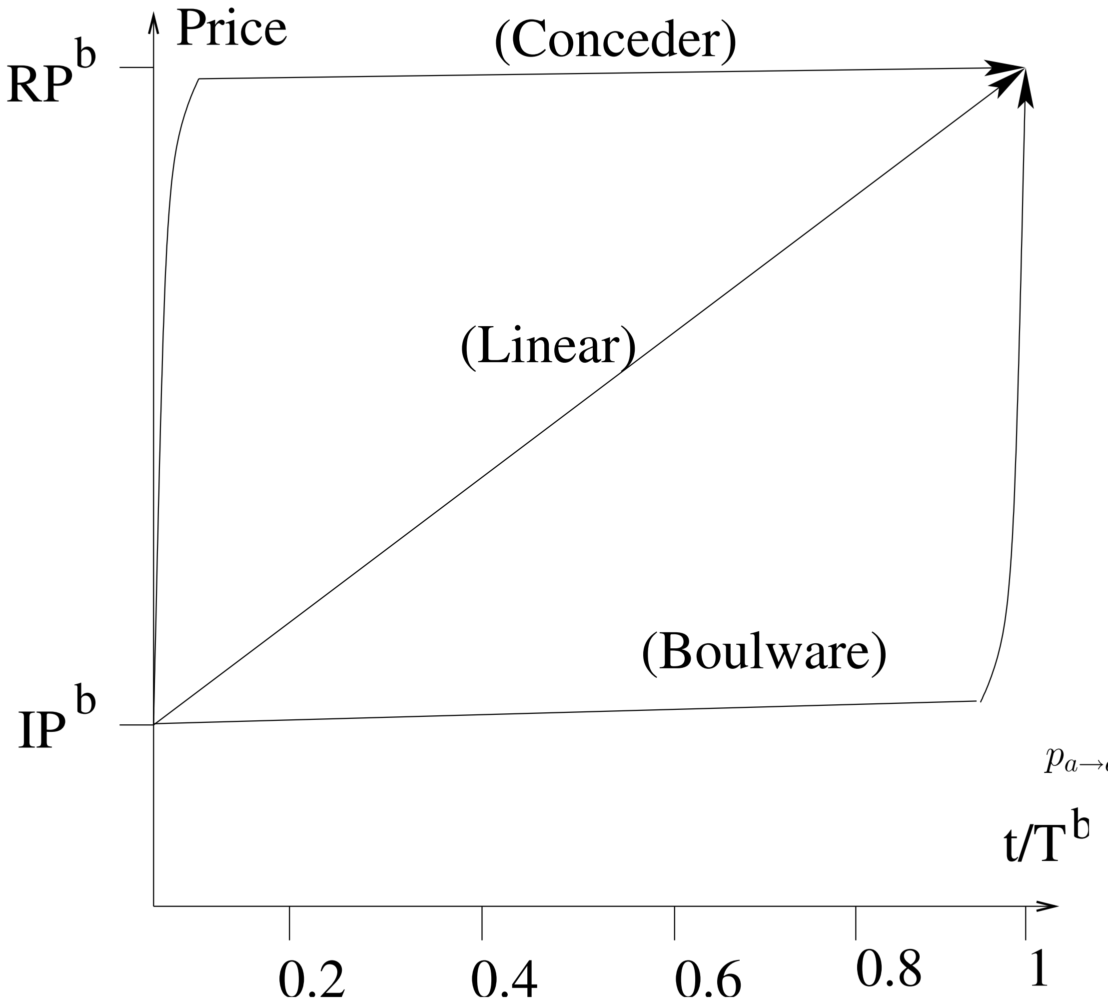

Buyer

$$p_{a\rightarrow \hat{a}}(t) = \begin{cases} IP^a + \phi^a(t)(RP^a - IP^a) & \textbf{for } a=b \text{ buyer,} \\ RP^a + (1-\phi^a(t))(IP^a - RP^a) & \textbf{for } a=s, \text{ seller} \end{cases}$$

$$\phi^a(t) = k^a + (1-k^a)\left(\frac{t}{T^a}\right)^{\frac{1}{\psi}}$$

Offer Function approximation

with the PRISM model-checker we are forced to use two-segments linear approximation of non-linear Offer Functions

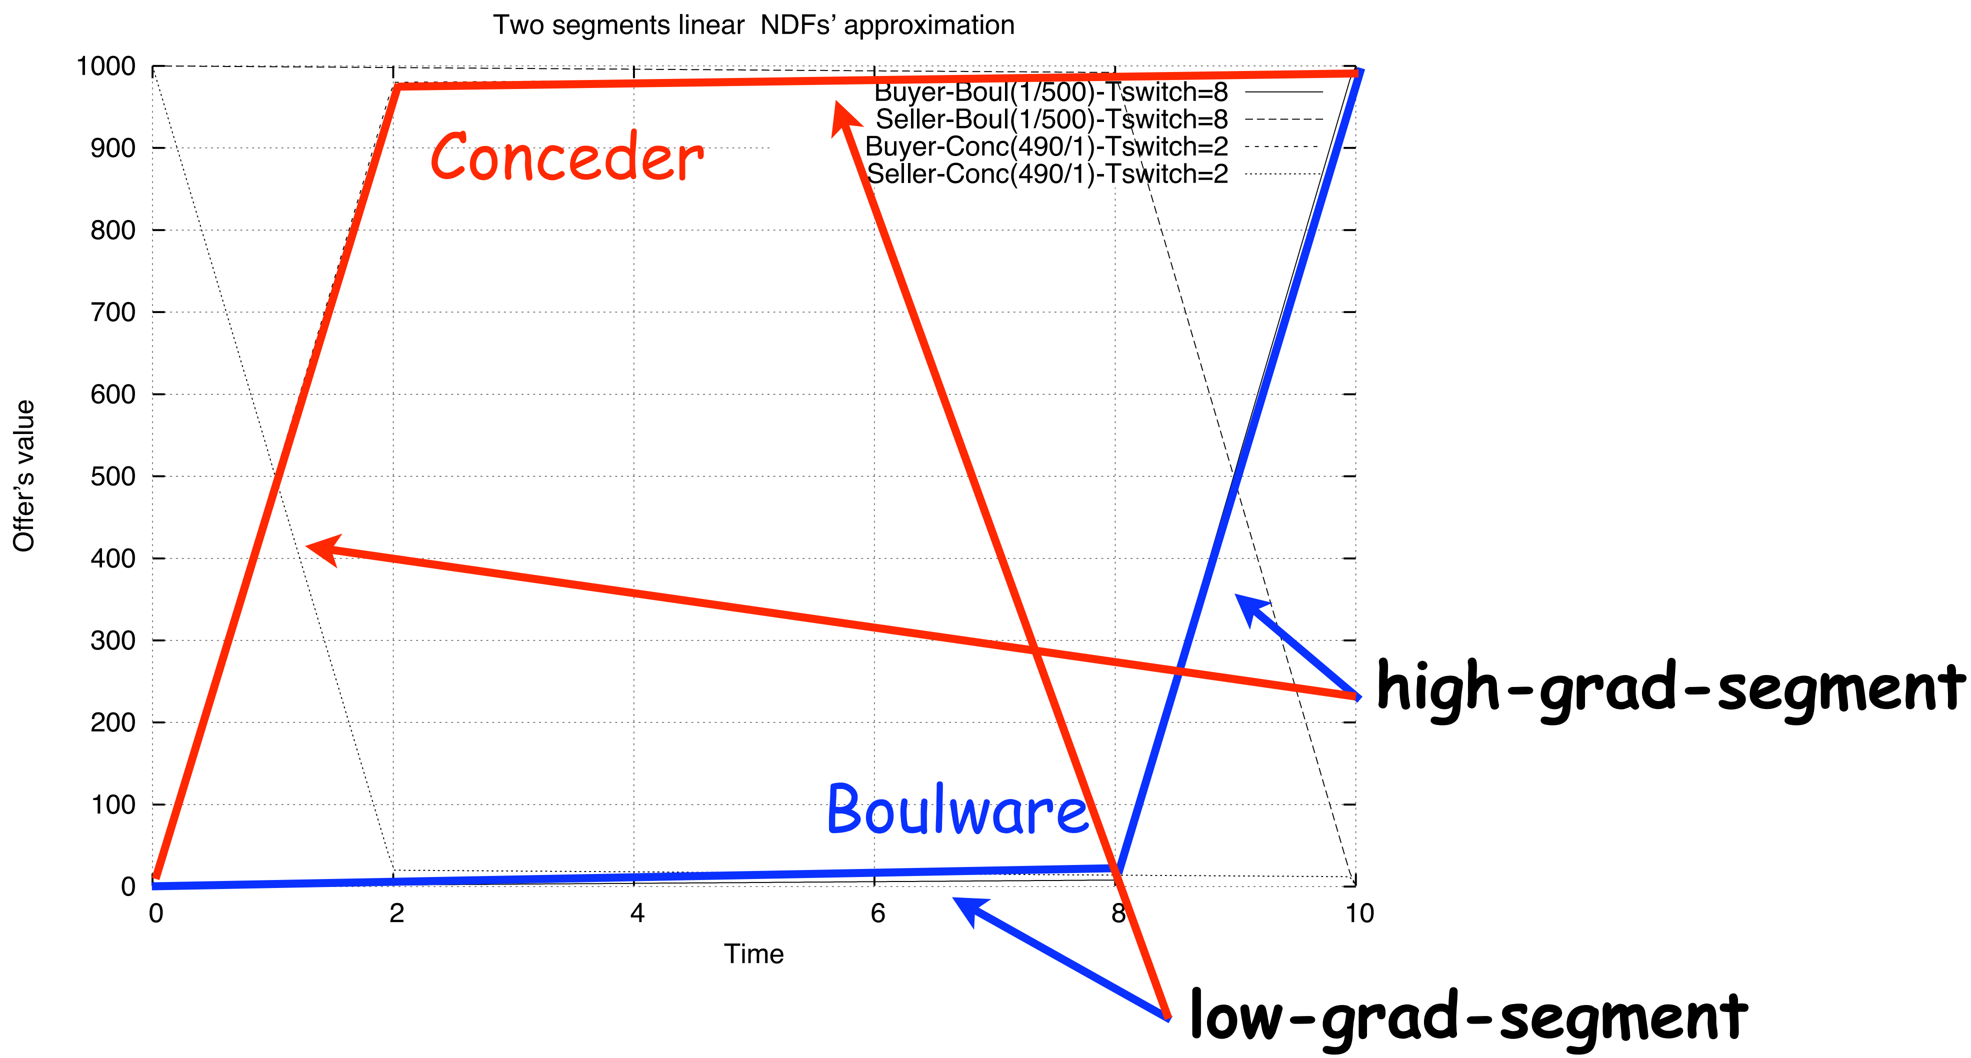

Acceptance Probability functions

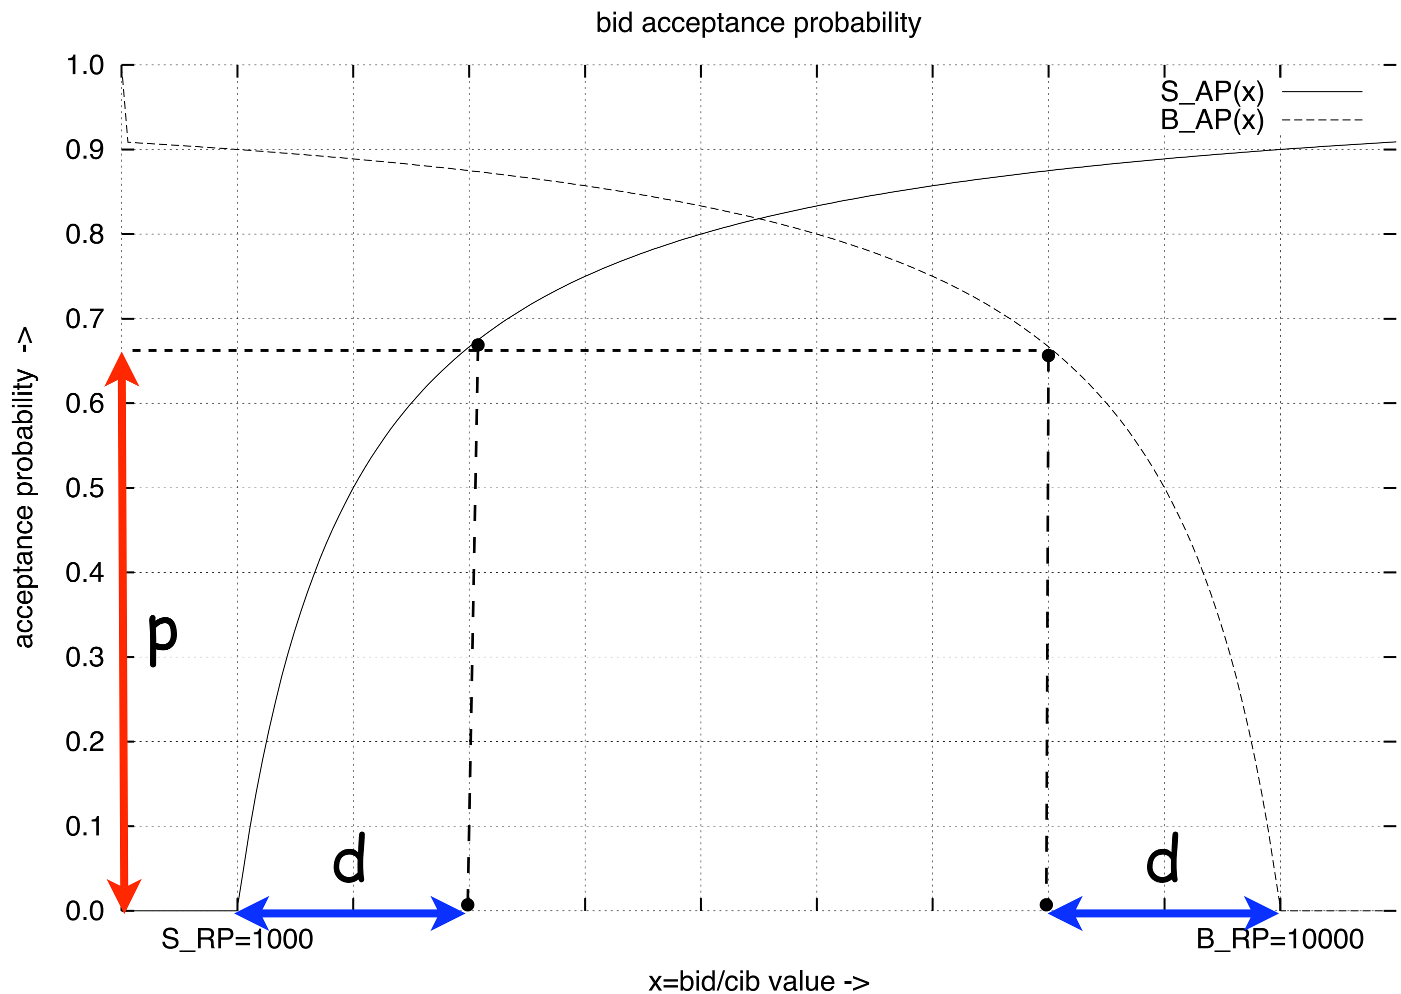

$$S\_AP(x,t) = \begin{cases} 0 & \textbf{if } (x \leq S\_RP) \wedge (t < T^s) \\ 1 - \frac{S\_RP}{x} & \textbf{if } (x > S\_RP) \wedge (t < T^s) \\ 1 & \textbf{if } (t \geq T^s) \end{cases}$$

$$B\_AP(x,t) = \begin{cases} 1 & \textbf{if } (x <= 0) \vee (t \geq T^b) \\ 1 + \frac{S\_RP}{x - (B\_RP + S\_RP)} & \textbf{if } (S\_RP < x < B\_RP) \wedge (t < T^b) \\ 0 & \textbf{if } (x > B\_RP) \wedge (t < T^b) \end{cases}$$

PCTL Model-Checking

probabilistic extension of CTL for referring to Discrete Time Markov Chains

PCTL syntax

$$\phi ::= \; tt \mid a \mid \phi \wedge \phi \mid \neg\phi \mid \mathcal{P}_{\trianglelefteq p}(\varphi)$$
$$\varphi ::= \phi \, U^{I} \, \phi$$

PCTL Model-Checking

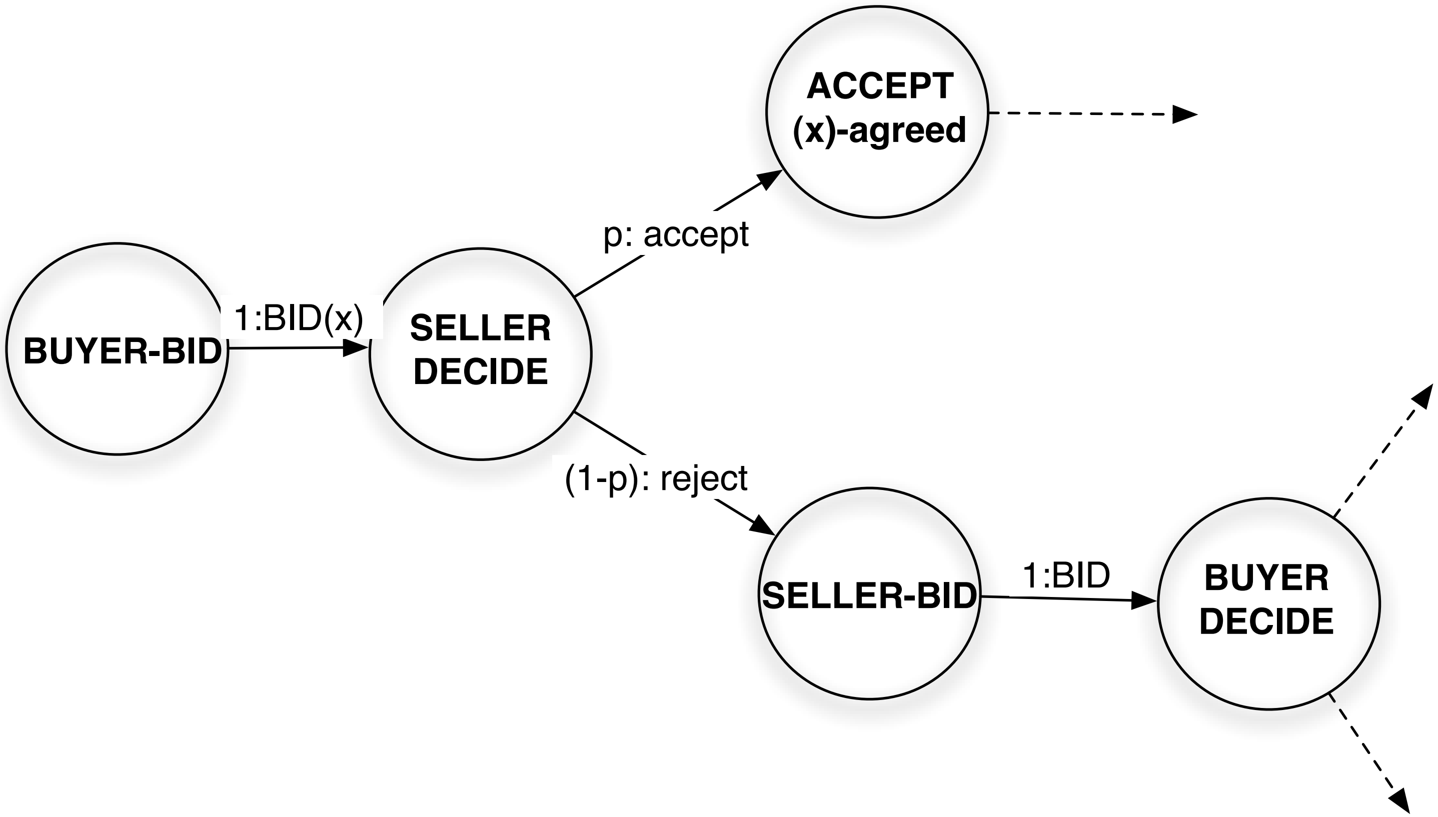

$$\phi_1 \equiv P_{\geq 0.8}[\diamond(agreed = 100)]$$

$$\phi_x \equiv P_{?}[\diamond(agreed = x)]$$

Model Verification

by verifying $\phi_x \equiv P_{?}[\diamond(agreed = x)]$

we devise the distribution of probability over the set of possible agreements, hence the expected utility

by comparing a number of strategy profiles we devise how strategy parameters affect the expected outcome of negotiation

One (fairly trivial) indication

The less a player concedes the higher his expected utility is going to be

Conceder-Conceder

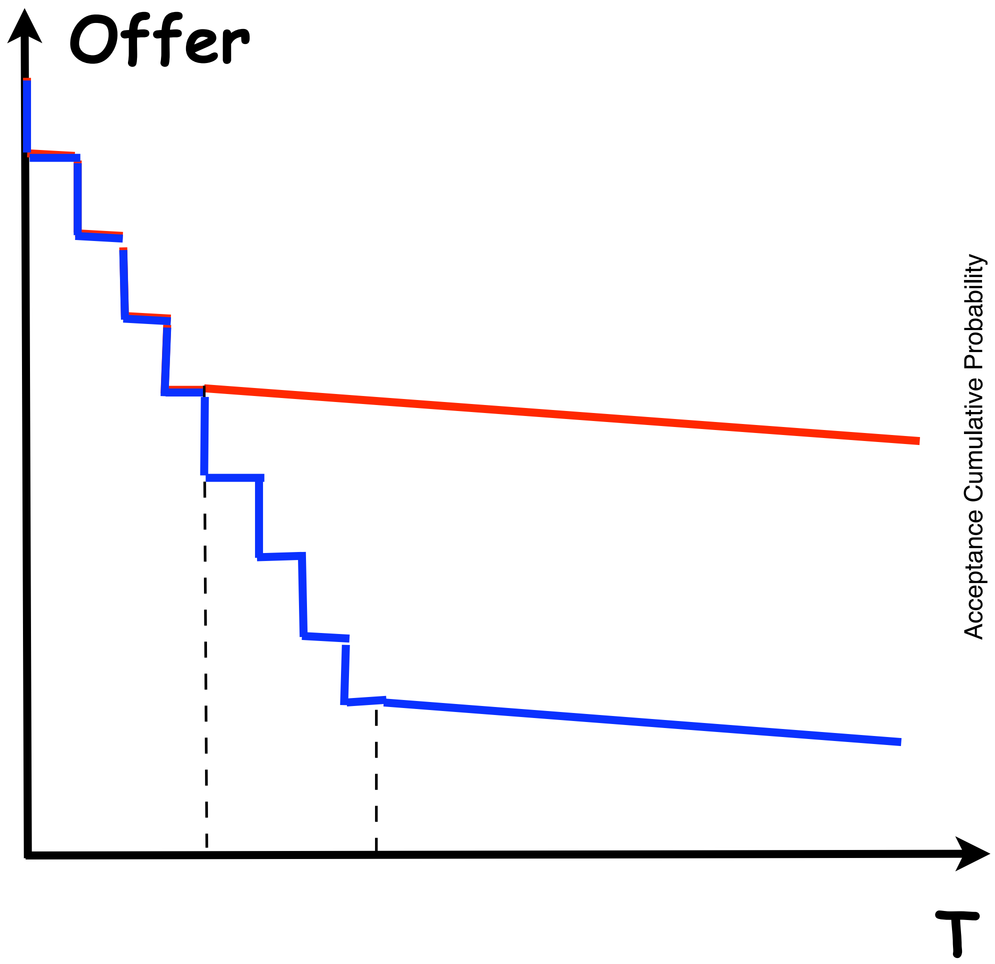

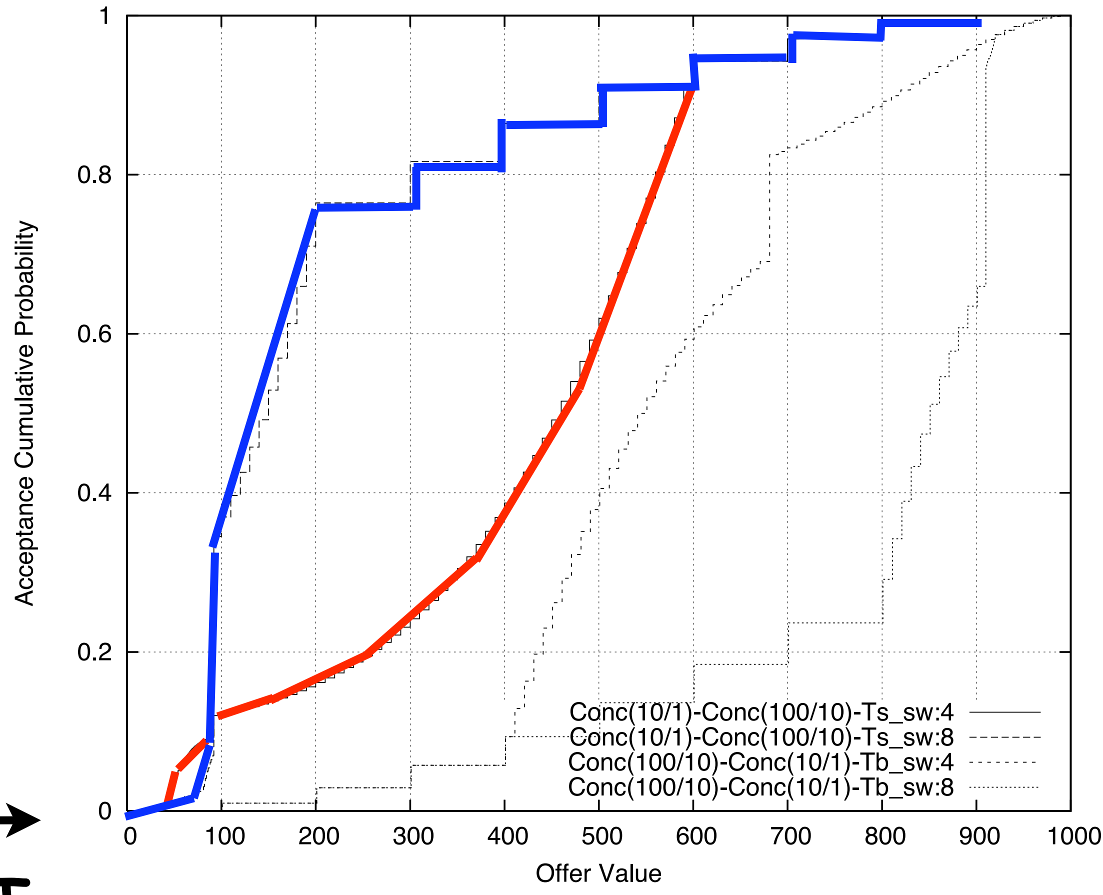

<u>**Seller Offer Function**</u>

<u>**Cumulative Acceptance Prob**</u>

Conclusion

we have shown that:

under certain assumption a game of imperfect information can be encoded into a discrete-time Markovian model

PCTL model-checking can be used to verify such a model

model-checking allows for comparing of strategy profiles

such an approach differ from both classical game-theory analysis and from simulative analysis

issue:

can we perform a deeper analysis through model-checking?
how about Nash-Equilibria analysis through model-checking ?