Bounded Rational Probabilistic Epistemic Dynamics

(in bilateral bargaining under uncertainty)

Paul Varkey

Multi Agent Systems Group, Department of Computer Science, UIC

4th NASSLLI Jun 20 – 26, 2010
Indiana University Bloomington

Outline

Outline

in the context of NASSLLI'10

in the context of NASSLLI'10

acknowledgements:

- I have had the good fortune of having many great minds (esp. Parikh, van Ditmarsch, Baltag, Smets and Pacuit) expose the key ideas that are foundational to my talk (and this will recognized in the course of my talk)
- In particular, I would like to acknowledge the work that was presented in the Dynamic Epistemic Logic track

in the context of NASSLLI'10

acknowledgements:

- I have had the good fortune of having many great minds (esp. Parikh, van Ditmarsch, Baltag, Smets and Pacuit) expose the key ideas that are foundational to my talk (and this will recognized in the course of my talk)
- In particular, I would like to acknowledge the work that was presented in the Dynamic Epistemic Logic track

in the context of NASSLLI'10

acknowledgements:

- I have had the good fortune of having many great minds (esp. Parikh, van Ditmarsch, Baltag, Smets and Pacuit) expose the key ideas that are foundational to my talk (and this will recognized in the course of my talk)
- In particular, I would like to acknowledge the work that was presented in the Dynamic Epistemic Logic track

in the context of NASSLLI'10

acknowledgements:

- I have had the good fortune of having many great minds (esp. Parikh, van Ditmarsch, Baltag, Smets and Pacuit) expose the key ideas that are foundational to my talk (and this will recognized in the course of my talk)
- In particular, I would like to acknowledge the work that was presented in the Dynamic Epistemic Logic track

terminological caveats:

- As a matter of ensuring terminological consistency, please note that when I say "common knowledge", I mean "common (true) belief" according to Alexandru and Sonja and the wider DEL community – in general, I will follow the wider "probabilistic interactive epistemological" community and properly talk only about beliefs (even when I use the term "knowledge")
- I will use the term "uncertainty" interchangeably with "incomplete information" (for e.g., games under uncertainty OR games under incomplete information)

in the context of NASSLLI'10

acknowledgements:

- I have had the good fortune of having many great minds (esp. Parikh, van Ditmarsch, Baltag, Smets and Pacuit) expose the key ideas that are foundational to my talk (and this will recognized in the course of my talk)
- In particular, I would like to acknowledge the work that was presented in the Dynamic Epistemic Logic track

terminological caveats:

- As a matter of ensuring terminological consistency, please note that when I say "common knowledge", I mean "common (true) belief" according to Alexandru and Sonja and the wider DEL community – in general, I will follow the wider "probabilistic interactive epistemological" community and properly talk only about beliefs (even when I use the term "knowledge")
- I will use the term "uncertainty" interchangeably with "incomplete information" (for e.g., games under uncertainty OR games under incomplete information)

in the context of NASSLLI'10

acknowledgements:

- I have had the good fortune of having many great minds (esp. Parikh, van Ditmarsch, Baltag, Smets and Pacuit) expose the key ideas that are foundational to my talk (and this will recognized in the course of my talk)
- In particular, I would like to acknowledge the work that was presented in the Dynamic Epistemic Logic track

terminological caveats:

- As a matter of ensuring terminological consistency, please note that when I say "common knowledge", I mean "common (true) belief" according to Alexandru and Sonja and the wider DEL community – in general, I will follow the wider "probabilistic interactive epistemological" community and properly talk only about beliefs (even when I use the term "knowledge")
- I will use the term "uncertainty" interchangeably with "incomplete information" (for e.g., games under uncertainty OR games under incomplete information)

in the context of NASSLLI'10

In a conversation that I had with Eric Pacuit, I told him that one of the main reasons I was keen on attending NASSLLI was to learn the **L**ogical approach to multi-agent epistemic modeling & dynamics. I then emphasized that I meant 'Logical' (with a capital 'L'); and not, merely, logical. I then remarked that "hopefully the Logical approach is also the logical approach". He agreed. We both laughed.

keywords & emphases

bilateral-bargaining
bounded-rationality
computational-efficiency
higher-order-uncertainty-and-beliefs
probabilistic-epistemic-dynamics
interactive-epistemology

keywords & emphases

The two main emphases in this talk will be:

- **probabilistic epistemic dynamics**
  - probabilistic representation of beliefs (including higher-order beliefs)
  - probabilistic reasoning (i.e. using Bayesian updating) to realize epistemic dynamics
- **bounded** (or, resource-bounded) **rationality**

keywords & emphases

The two main emphases in this talk will be:

- **probabilistic epistemic dynamics**
  - probabilistic representation of beliefs (including higher-order beliefs)
  - probabilistic reasoning (i.e. using Bayesian updating) to realize epistemic dynamics
- **bounded** (or, resource-bounded) **rationality**

keywords & emphases

The two main emphases in this talk will be:

- **probabilistic epistemic dynamics**
  - probabilistic representation of beliefs (including higher-order beliefs)
  - probabilistic reasoning (i.e. using Bayesian updating) to realize epistemic dynamics
- **bounded** (or, resource-bounded) **rationality**

keywords & emphases

The two main emphases in this talk will be:

- **probabilistic epistemic dynamics**
  - probabilistic representation of beliefs (including higher-order beliefs)
  - probabilistic reasoning (i.e. using Bayesian updating) to realize epistemic dynamics
- **bounded** (or, resource-bounded) **rationality**

keywords & emphases

The two main emphases in this talk will be:

- **probabilistic epistemic dynamics**
  - probabilistic representation of beliefs (including higher-order beliefs)
  - probabilistic reasoning (i.e. using Bayesian updating) to realize epistemic dynamics
- **bounded** (or, resource-bounded) **rationality**

keywords & emphases

The two main emphases in this talk will be:

- **probabilistic epistemic dynamics**
  - probabilistic representation of beliefs (including higher-order beliefs)
  - probabilistic reasoning (i.e. using Bayesian updating) to realize epistemic dynamics
- **bounded** (or, resource-bounded) **rationality**

  The problem that we are considering today (representing and reasoning with higher-order beliefs in the setting of bilateral bargaining) is computationally expensive in nature. In computer science, we are always interested in computational efficiency – an obsession with fast algorithms. Contract algorithms (a close relative of anytime algorithms) are very desirable for particularly hard problems – a contract algorithm for a sophisticated bargaining agent will be presented in today's talk.

the game-theoretic context

Game theory involves the study of strategic interactions among multiple (possibly) rational agents who may (or may not) be completely informed about the strategically relevant features of the interaction, with the ultimate goal of exhibiting (or constructing) equilibria, which are hoped to be indicative of actual rational behavior or, at least, of what it ought to be

- Many real-world scenarios are "informationally rich" and involve significant uncertainty on the part of one or more of the players
- Interactive epistemologies arise naturally in such settings – involving an agent's beliefs, its beliefs about other agents' beliefs, its beliefs about others' beliefs about others' beliefs, etc.
- Uncertainty represented using Aumann's *possible-worlds* knowledge partitions (semantically related to Kripke models) augmented with a (probability) measure
- Interactively nested (i.e. higher order) uncertainty leads to infinite belief hierarchies – first handled by Harsanyi using the clever idea of a recursively defined type space

the game-theoretic context

Game theory involves the study of strategic interactions among multiple (possibly) rational agents who may (or may not) be completely informed about the strategically relevant features of the interaction, with the ultimate goal of exhibiting (or constructing) equilibria, which are hoped to be indicative of actual rational behavior or, at least, of what it ought to be

- Many real-world scenarios are "informationally rich" and involve significant uncertainty on the part of one or more of the players
- Interactive epistemologies arise naturally in such settings – involving an agent's beliefs, its beliefs about other agents' beliefs, its beliefs about others' beliefs about others' beliefs, etc.
- Uncertainty represented using Aumann's *possible-worlds* knowledge partitions (semantically related to Kripke models) augmented with a (probability) measure
- Interactively nested (i.e. higher order) uncertainty leads to infinite belief hierarchies – first handled by Harsanyi using the clever idea of a recursively defined type space

the game-theoretic context

Game theory involves the study of strategic interactions among multiple (possibly) rational agents who may (or may not) be completely informed about the strategically relevant features of the interaction, with the ultimate goal of exhibiting (or constructing) equilibria, which are hoped to be indicative of actual rational behavior or, at least, of what it ought to be

- Many real-world scenarios are "informationally rich" and involve significant uncertainty on the part of one or more of the players
- Interactive epistemologies arise naturally in such settings – involving an agent's beliefs, its beliefs about other agents' beliefs, its beliefs about others' beliefs about others' beliefs, etc.
- Uncertainty represented using Aumann's *possible-worlds* knowledge partitions (semantically related to Kripke models) augmented with a (probability) measure
- Interactively nested (i.e. higher order) uncertainty leads to infinite belief hierarchies – first handled by Harsanyi using the clever idea of a recursively defined type space

the game-theoretic context

Game theory involves the study of strategic interactions among multiple (possibly) rational agents who may (or may not) be completely informed about the strategically relevant features of the interaction, with the ultimate goal of exhibiting (or constructing) equilibria, which are hoped to be indicative of actual rational behavior or, at least, of what it ought to be

- Many real-world scenarios are "informationally rich" and involve significant uncertainty on the part of one or more of the players
- Interactive epistemologies arise naturally in such settings – involving an agent's beliefs, its beliefs about other agents' beliefs, its beliefs about others' beliefs about others' beliefs, etc.
- Uncertainty represented using Aumann's *possible-worlds* knowledge partitions (semantically related to Kripke models) augmented with a (probability) measure
- Interactively nested (i.e. higher order) uncertainty leads to infinite belief hierarchies – first handled by Harsanyi using the clever idea of a recursively defined type space

the game-theoretic context

Game theory involves the study of strategic interactions among multiple (possibly) rational agents who may (or may not) be completely informed about the strategically relevant features of the interaction, with the ultimate goal of exhibiting (or constructing) equilibria, which are hoped to be indicative of actual rational behavior or, at least, of what it ought to be

- Many real-world scenarios are "informationally rich" and involve significant uncertainty on the part of one or more of the players
- Interactive epistemologies arise naturally in such settings – involving an agent's beliefs, its beliefs about other agents' beliefs, its beliefs about others' beliefs about others' beliefs, etc.
- Uncertainty represented using Aumann's *possible-worlds* knowledge partitions (semantically related to Kripke models) augmented with a (probability) measure
- Interactively nested (i.e. higher order) uncertainty leads to infinite belief hierarchies – first handled by Harsanyi using the clever idea of a recursively defined type space

the game-theoretic context

Game theory involves the study of strategic interactions among multiple (possibly) rational agents who may (or may not) be completely informed about the strategically relevant features of the interaction, with the ultimate goal of exhibiting (or constructing) equilibria, which are hoped to be indicative of actual rational behavior or, at least, of what it ought to be

- Many real-world scenarios are "informationally rich" and involve significant uncertainty on the part of one or more of the players
- Interactive epistemologies arise naturally in such settings – involving an agent's beliefs, its beliefs about other agents' beliefs, its beliefs about others' beliefs about others' beliefs, etc.
- Uncertainty represented using Aumann's *possible-worlds* knowledge partitions (semantically related to Kripke models) augmented with a (probability) measure
- Interactively nested (i.e. higher order) uncertainty leads to infinite belief hierarchies – first handled by Harsanyi using the clever idea of a recursively defined type space

"Lots of research is going on in higher-order cognition....there is a limit on how much humans can tolerate in terms of stacks of epistemic operators"
(Hans van Ditmarsch, remark, Jun 23, NASSLLI '10)

the game-theoretic context: bilateral bargaining

A well-studied example is **bilateral bargaining**

the game-theoretic context: bilateral bargaining

Why is the bargaining problem chosen for study?

- It is a very good representative of many real-world strategic interactions and, under the banner of automated negotiation, is being actively studied in the AI community due to the rich array of computational and informational (i.e. interactive epsitemological) problems that it raises as well as the diverse applications that it finds in e-business, electronic security etc.
- Rubinstein, in his seminal 1982 paper, remarked that "Edgeworth presented this problem one hundred years ago, considering it the most fundamental problem in Economics" (referring to Edgeworth's remarkable 1881 book "Mathematical Psychics: An Essay on the Application of Mathematics to the Moral Sciences")
- I have been banging my head against a wall studying this problem for the past $3\frac{1}{2}$ years – I can talk about this problem at length, and I like to talk at length

the game-theoretic context: bilateral bargaining

Why is the bargaining problem chosen for study?

- It is a very good representative of many real-world strategic interactions and, under the banner of automated negotiation, is being actively studied in the AI community due to the rich array of computational and informational (i.e. interactive epsitemological) problems that it raises as well as the diverse applications that it finds in e-business, electronic security etc.
- Rubinstein, in his seminal 1982 paper, remarked that "Edgeworth presented this problem one hundred years ago, considering it the most fundamental problem in Economics" (referring to Edgeworth's remarkable 1881 book "Mathematical Psychics: An Essay on the Application of Mathematics to the Moral Sciences")
- I have been banging my head against a wall studying this problem for the past $3\frac{1}{2}$ years – I can talk about this problem at length, and I like to talk at length

the game-theoretic context: bilateral bargaining

Why is the bargaining problem chosen for study?

- It is a very good representative of many real-world strategic interactions and, under the banner of automated negotiation, is being actively studied in the AI community due to the rich array of computational and informational (i.e. interactive epsitemological) problems that it raises as well as the diverse applications that it finds in e-business, electronic security etc.
- Rubinstein, in his seminal 1982 paper, remarked that "Edgeworth presented this problem one hundred years ago, considering it the most fundamental problem in Economics" (referring to Edgeworth's remarkable 1881 book "Mathematical Psychics: An Essay on the Application of Mathematics to the Moral Sciences")
- I have been banging my head against a wall studying this problem for the past $3\frac{1}{2}$ years – I can talk about this problem at length, and I like to talk at length

the game-theoretic context: bilateral bargaining

Why is the bargaining problem chosen for study?

- It is a very good representative of many real-world strategic interactions and, under the banner of automated negotiation, is being actively studied in the AI community due to the rich array of computational and informational (i.e. interactive epsitemological) problems that it raises as well as the diverse applications that it finds in e-business, electronic security etc.
- Rubinstein, in his seminal 1982 paper, remarked that "Edgeworth presented this problem one hundred years ago, considering it the most fundamental problem in Economics" (referring to Edgeworth's remarkable 1881 book "Mathematical Psychics: An Essay on the Application of Mathematics to the Moral Sciences")
- I have been banging my head against a wall studying this problem for the past $3\frac{1}{2}$ years – I can talk about this problem at length, and I like to talk at length

the game-theoretic context: bilateral bargaining

The Bargaining Problem

A seller and a buyer are negotiating over an item for which they have a valuation of c and v , respectively. How do (should) they split the available profit?

- Offers: seller-offers, alternating offers, etc.
- Delay costs: discounting, fixed costs, etc.
- Horizon: finite, indefinite, etc.
- Information: complete and incomplete (1-sided, 2-sided, higher-order, etc.)

the game-theoretic context: bilateral bargaining

The Bargaining Problem

A seller and a buyer are negotiating over an item for which they have a valuation of c and v , respectively. How do (should) they split the available profit?

- Offers: seller-offers, alternating offers, etc.
- Delay costs: discounting, fixed costs, etc.
- Horizon: finite, indefinite, etc.
- Information: complete and incomplete (1-sided, 2-sided, higher-order, etc.)

the game-theoretic context: bilateral bargaining

The Bargaining Problem

A seller and a buyer are negotiating over an item for which they have a valuation of c and v , respectively. How do (should) they split the available profit?

- Offers: seller-offers, alternating offers, etc.
- Delay costs: discounting, fixed costs, etc.
- Horizon: finite, indefinite, etc.
- Information: complete and incomplete (1-sided, 2-sided, higher-order, etc.)

the game-theoretic context: bilateral bargaining

The Bargaining Problem

A seller and a buyer are negotiating over an item for which they have a valuation of c and v , respectively. How do (should) they split the available profit?

- Offers: seller-offers, alternating offers, etc.
- Delay costs: discounting, fixed costs, etc.
- Horizon: finite, indefinite, etc.
- Information: complete and incomplete (1-sided, 2-sided, higher-order, etc.)

the game-theoretic context: bilateral bargaining

The Bargaining Problem

A seller and a buyer are negotiating over an item for which they have a valuation of c and v , respectively. How do (should) they split the available profit?

- Offers: seller-offers, alternating offers, etc.
- Delay costs: discounting, fixed costs, etc.
- Horizon: finite, indefinite, etc.
- Information: complete and incomplete (1-sided, 2-sided, higher-order, etc.)

the game-theoretic context: bilateral bargaining

The Bargaining Problem

A seller and a buyer are negotiating over an item for which they have a valuation of c and v , respectively. How do (should) they split the available profit?

- Offers: seller-offers, alternating offers, etc.
- Delay costs: discounting, fixed costs, etc.
- Horizon: finite, indefinite, etc.
- Information: complete and incomplete (1-sided, 2-sided, higher-order, etc.)

the game-theoretic context: bilateral bargaining

The spectrum of informational settings for the bargaining problem

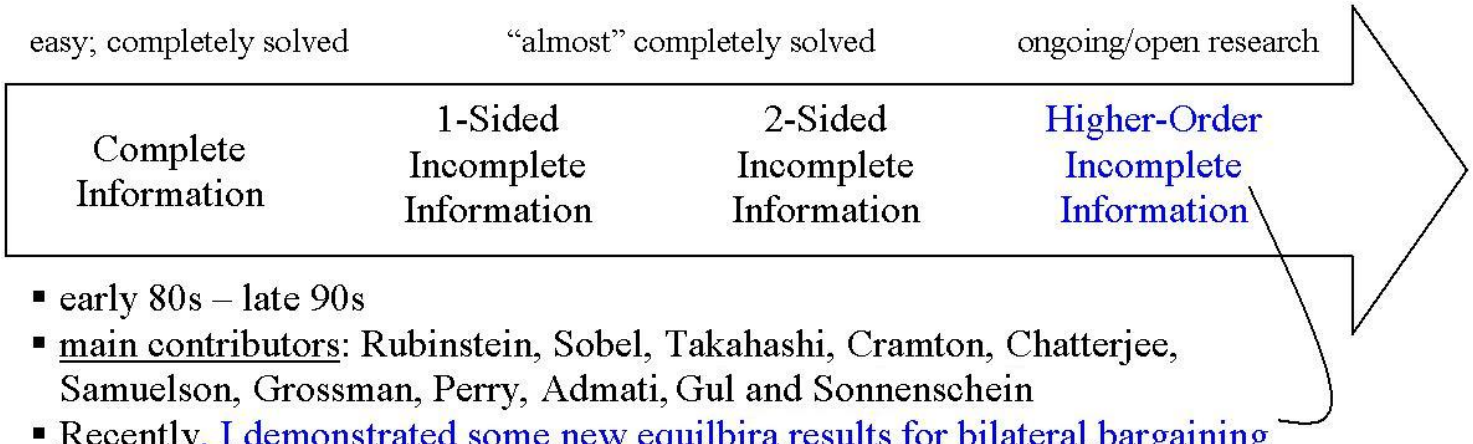

- early 80s – late 90s
- main contributors: Rubinstein, Sobel, Takahashi, Cramton, Chatterjee, Samuelson, Grossman, Perry, Admati, Gul and Sonnenschein
- Recently, I demonstrated some new equilbira results for bilateral bargaining under higher-order uncertainty; presented at MWTD'10; also accepted as short paper at BWGT'10

some limitations of game theory

- multiplicity of equilibria; which one should be chosen? why?
- all notions of equilibria depend strongly on certain common knowledge assumptions – seems to be paradoxical in settings of incomplete information (for e.g. requring the Harsanyi type space to be common knowledge in order to obtain (perfect) Bayesian-Nash equilibria)
- the theory is necessarily silent about a principled approach to belief update in the event of making a completely surprising observation – i.e. when all possibilities are thought to be impossible

some limitations of game theory

- multiplicity of equilibria; which one should be chosen? why?
- all notions of equilibria depend strongly on certain common knowledge assumptions – seems to be paradoxical in settings of incomplete information (for e.g. requring the Harsanyi type space to be common knowledge in order to obtain (perfect) Bayesian-Nash equilibria)
- the theory is necessarily silent about a principled approach to belief update in the event of making a completely surprising observation – i.e. when all possibilities are thought to be impossible

some limitations of game theory

- multiplicity of equilibria; which one should be chosen? why?
- all notions of equilibria depend strongly on certain common knowledge assumptions – seems to be paradoxical in settings of incomplete information (for e.g. requring the Harsanyi type space to be common knowledge in order to obtain (perfect) Bayesian-Nash equilibria)
- the theory is necessarily silent about a principled approach to belief update in the event of making a completely surprising observation – i.e. when all possibilities are thought to be impossible

some limitations of game theory

- multiplicity of equilibria; which one should be chosen? why?
- all notions of equilibria depend strongly on certain common knowledge assumptions – seems to be paradoxical in settings of incomplete information (for e.g. requring the Harsanyi type space to be common knowledge in order to obtain (perfect) Bayesian-Nash equilibria)
- the theory is necessarily silent about a principled approach to belief update in the event of making a completely surprising observation – i.e. when all possibilities are thought to be impossible

Sonja Smets and Alexandru Baltag have communicated the notions of plausibility orderings and soft learning (i.e. belief upgade vs. update) which have been successfully employed in dealing with this issue in the DEL framework and have also pointed out conditions under which one obtains belief stability (and, sometimes, convergence to truth)

decision theory as a paradigm for AI

- Decision theory has been rigorously developed on the foundations of probability theory (Bayesian belief update) and utility theory
  "Ramsey & Savage laid the foundations of personal (subjective) probability elicitation & personal (subjective) preference/utility elicitation"
  (Rohit Parikh, Barwise Invited Lecture, Jun 23, NASSLLI '10)
  - In the same lecture, Parikh also remarked about "the inappropriate precision of the Savage theory", claiming that "humans are not so rational"
  - Admittedly, this is a sharp insight about real human economic and social choice behavior
  - Although, crucially, this is not limiting for the Artifcial Intelligence program, because artificial agents are not limited by human irrationality
- Decision theory can be and has been powerfully operationalized as a control paradigm for AI tasks including reasoning, decision-making, sequential planning, reactive planning and for all of these under settings of uncertainty
  - Bilateral bargaining is an example of a sequential decision-making and planning problem under uncertainty

decision theory as a paradigm for AI

- Decision theory has been rigorously developed on the foundations of probability theory (Bayesian belief update) and utility theory
  "Ramsey & Savage laid the foundations of personal (subjective) probability elicitation & personal (subjective) preference/utility elicitation"
  (Rohit Parikh, Barwise Invited Lecture, Jun 23, NASSLLI '10)
  - In the same lecture, Parikh also remarked about "the inappropriate precision of the Savage theory", claiming that "humans are not so rational"
  - Admittedly, this is a sharp insight about real human economic and social choice behavior
  - Although, crucially, this is not limiting for the Artifcial Intelligence program, because artificial agents are not limited by human irrationality
- Decision theory can be and has been powerfully operationalized as a control paradigm for AI tasks including reasoning, decision-making, sequential planning, reactive planning and for all of these under settings of uncertainty
  - Bilateral bargaining is an example of a sequential decision-making and planning problem under uncertainty

decision theory as a paradigm for AI

- Decision theory has been rigorously developed on the foundations of probability theory (Bayesian belief update) and utility theory
  "Ramsey & Savage laid the foundations of personal (subjective) probability elicitation & personal (subjective) preference/utility elicitation"
  (Rohit Parikh, Barwise Invited Lecture, Jun 23, NASSLLI '10)
  - In the same lecture, Parikh also remarked about "the inappropriate precision of the Savage theory", claiming that "humans are not so rational"
  - Admittedly, this is a sharp insight about real human economic and social choice behavior
  - Although, crucially, this is not limiting for the Artifcial Intelligence program, because artificial agents are not limited by human irrationality
- Decision theory can be and has been powerfully operationalized as a control paradigm for AI tasks including reasoning, decision-making, sequential planning, reactive planning and for all of these under settings of uncertainty
  - Bilateral bargaining is an example of a sequential decision-making and planning problem under uncertainty

decision theory as a paradigm for AI

- Decision theory has been rigorously developed on the foundations of probability theory (Bayesian belief update) and utility theory
  "Ramsey & Savage laid the foundations of personal (subjective) probability elicitation & personal (subjective) preference/utility elicitation"
  (Rohit Parikh, Barwise Invited Lecture, Jun 23, NASSLLI '10)
  - In the same lecture, Parikh also remarked about "the inappropriate precision of the Savage theory", claiming that "humans are not so rational"
  - Admittedly, this is a sharp insight about real human economic and social choice behavior
  - Although, crucially, this is not limiting for the Artifcial Intelligence program, because artificial agents are not limited by human irrationality
- Decision theory can be and has been powerfully operationalized as a control paradigm for AI tasks including reasoning, decision-making, sequential planning, reactive planning and for all of these under settings of uncertainty
  - Bilateral bargaining is an example of a sequential decision-making and planning problem under uncertainty

decision theory as a paradigm for AI

- Decision theory has been rigorously developed on the foundations of probability theory (Bayesian belief update) and utility theory
  "Ramsey & Savage laid the foundations of personal (subjective) probability elicitation & personal (subjective) preference/utility elicitation"
  (Rohit Parikh, Barwise Invited Lecture, Jun 23, NASSLLI '10)
  - In the same lecture, Parikh also remarked about "the inappropriate precision of the Savage theory", claiming that "humans are not so rational"
  - Admittedly, this is a sharp insight about real human economic and social choice behavior
  - Although, crucially, this is not limiting for the Artifcial Intelligence program, because artificial agents are not limited by human irrationality
- Decision theory can be and has been powerfully operationalized as a control paradigm for AI tasks including reasoning, decision-making, sequential planning, reactive planning and for all of these under settings of uncertainty
  - Bilateral bargaining is an example of a sequential decision-making and planning problem under uncertainty

decision theory as a paradigm for AI

- Decision theory has been rigorously developed on the foundations of probability theory (Bayesian belief update) and utility theory
  "Ramsey & Savage laid the foundations of personal (subjective) probability elicitation & personal (subjective) preference/utility elicitation"
  (Rohit Parikh, Barwise Invited Lecture, Jun 23, NASSLLI '10)
  - In the same lecture, Parikh also remarked about "the inappropriate precision of the Savage theory", claiming that "humans are not so rational"
  - Admittedly, this is a sharp insight about real human economic and social choice behavior
  - Although, crucially, this is not limiting for the Artifcial Intelligence program, because artificial agents are not limited by human irrationality
- Decision theory can be and has been powerfully operationalized as a control paradigm for AI tasks including reasoning, decision-making, sequential planning, reactive planning and for all of these under settings of uncertainty
  - Bilateral bargaining is an example of a sequential decision-making and planning problem under uncertainty

decision theory as a paradigm for AI

- Decision theory has been rigorously developed on the foundations of probability theory (Bayesian belief update) and utility theory
  "Ramsey & Savage laid the foundations of personal (subjective) probability elicitation & personal (subjective) preference/utility elicitation"
  (Rohit Parikh, Barwise Invited Lecture, Jun 23, NASSLLI '10)
  - In the same lecture, Parikh also remarked about "the inappropriate precision of the Savage theory", claiming that "humans are not so rational"
  - Admittedly, this is a sharp insight about real human economic and social choice behavior
  - Although, crucially, this is not limiting for the Artifcial Intelligence program, because artificial agents are not limited by human irrationality
- Decision theory can be and has been powerfully operationalized as a control paradigm for AI tasks including reasoning, decision-making, sequential planning, reactive planning and for all of these under settings of uncertainty
  - Bilateral bargaining is an example of a sequential decision-making and planning problem under uncertainty

a review of decision theory

- (subjective) uncertainty about the world represented using a probability distribution
- (subjective) preferences over outcomes (technically, require preferences over lotterries over outcomes; along with satisfaction of "rational preference" axioms – transitivity & completeness)
- actions/decisions/choices change the state of the world in a known (possibly stochastic) manner
- decisions made using the **M**aximum **E**xpected **U**tility principle

a review of decision theory

- (subjective) uncertainty about the world represented using a probability distribution
- (subjective) preferences over outcomes (technically, require preferences over lotterries over outcomes; along with satisfaction of "rational preference" axioms – transitivity & completeness)
- actions/decisions/choices change the state of the world in a known (possibly stochastic) manner
- decisions made using the **M**aximum **E**xpected **U**tility principle

a review of decision theory

- (subjective) uncertainty about the world represented using a probability distribution
- (subjective) preferences over outcomes (technically, require preferences over lotterries over outcomes; along with satisfaction of "rational preference" axioms – transitivity & completeness)
- actions/decisions/choices change the state of the world in a known (possibly stochastic) manner
- decisions made using the **M**aximum **E**xpected **U**tility principle

a review of decision theory

- (subjective) uncertainty about the world represented using a probability distribution
- (subjective) preferences over outcomes (technically, require preferences over lotterries over outcomes; along with satisfaction of "rational preference" axioms – transitivity & completeness)
- actions/decisions/choices change the state of the world in a known (possibly stochastic) manner
- decisions made using the **M**aximum **E**xpected **U**tility principle

a review of decision theory

- (subjective) uncertainty about the world represented using a probability distribution
- (subjective) preferences over outcomes (technically, require preferences over lotterries over outcomes; along with satisfaction of "rational preference" axioms – transitivity & completeness)
- actions/decisions/choices change the state of the world in a known (possibly stochastic) manner
- decisions made using the **M**aximum **E**xpected **U**tility principle

Example 1. Should I eat a peanut butter sandwich? p(enjoy it) $= 0.99999$, utility is very high; p(suddenly develop peanut allergy) $= 0.00001$; outcome – at best (extreme discomfort; hefty emergency room bill), at worst (??), utility is very low.
The decision is USUALLY to have the sandwich

a review of decision theory

- (subjective) uncertainty about the world represented using a probability distribution
- (subjective) preferences over outcomes (technically, require preferences over lotterries over outcomes; along with satisfaction of "rational preference" axioms – transitivity & completeness)
- actions/decisions/choices change the state of the world in a known (possibly stochastic) manner
- decisions made using the **M**aximum **E**xpected **U**tility principle

Example 2. Should I cross the street today?

- approaching driver is law-abiding and careful
  probability: 0.9999; outcome: safely cross street, attend NASSLLI, try ethnic food on Indiana Ave., share a few drinks with new friends at The Upstairs Pub, enlarge academic network; utility: very, very high
- drunk student driving home after excessive partying all night
  probability: 0.0001; outcome: death; utility: very, very low

a review of decision theory

- (subjective) uncertainty about the world represented using a probability distribution
- (subjective) preferences over outcomes (technically, require preferences over lotterries over outcomes; along with satisfaction of "rational preference" axioms – transitivity & completeness)
- actions/decisions/choices change the state of the world in a known (possibly stochastic) manner
- decisions made using the **M**aximum **E**xpected **U**tility principle

Example 2. Should I cross the street today?

- approaching driver is law-abiding and careful
  probability: 0.9999; outcome: safely cross street, attend NASSLLI, try ethnic food on Indiana Ave., share a few drinks with new friends at The Upstairs Pub, enlarge academic network; utility: very, very high
- drunk student driving home after excessive partying all night
  probability: 0.0001; outcome: death; utility: very, very low

what about higher-order uncertainty?

How do we represent higher-order uncertainty – i.e. beliefs about beliefs, beliefs about beliefs, etc?

what about higher-order uncertainty?

How do we represent higher-order uncertainty – i.e. beliefs about beliefs, beliefs about beliefs, etc?

using probability distributions over probability distributions!

Outline

bargaining in the decision-theoretic framework: some intuitive epistemological insights

Analyze the game's epistemology at various levels:

- Both agents (seller & buyer) have first-order uncertainty about the condition of the car: if the seller believes the car is in poor condition (similarly, if the buyer believes the car is in good condition, or, desires this particular car), it finds itself in a weak bargaining position
- Say the seller starts with a low offer (or, reduces prices drastically during negotiation) – this is an indication to the buyer that the seller believes there is something wrong with the car; which is then (or, should be) reflected in the buyer's second-order beliefs
- Also consider the incentive to decieve/mis-lead/bluff/posture – even if car is bad, seller starts with high price to make it seem like car is good
- Say the seller's first-order belief about the car is that it is in poor condition, but that its second-order belief is that the buyer believes that the car is in good condition – in this case, the seller believes that it is in a better bargaining position than it really is and, therefore, can start with a higher offer than it would have started with had it only considered its first order belief

bargaining in the decision-theoretic framework: some intuitive epistemological insights

Analyze the game's epistemology at various levels:

- Both agents (seller & buyer) have first-order uncertainty about the condition of the car: if the seller believes the car is in poor condition (similarly, if the buyer believes the car is in good condition, or, desires this particular car), it finds itself in a weak bargaining position
- Say the seller starts with a low offer (or, reduces prices drastically during negotiation) – this is an indication to the buyer that the seller believes there is something wrong with the car; which is then (or, should be) reflected in the buyer's second-order beliefs
- Also consider the incentive to decieve/mis-lead/bluff/posture – even if car is bad, seller starts with high price to make it seem like car is good
- Say the seller's first-order belief about the car is that it is in poor condition, but that its second-order belief is that the buyer believes that the car is in good condition – in this case, the seller believes that it is in a better bargaining position than it really is and, therefore, can start with a higher offer than it would have started with had it only considered its first order belief

bargaining in the decision-theoretic framework: some intuitive epistemological insights

Analyze the game's epistemology at various levels:

- Both agents (seller & buyer) have first-order uncertainty about the condition of the car: if the seller believes the car is in poor condition (similarly, if the buyer believes the car is in good condition, or, desires this particular car), it finds itself in a weak bargaining position
- Say the seller starts with a low offer (or, reduces prices drastically during negotiation) – this is an indication to the buyer that the seller believes there is something wrong with the car; which is then (or, should be) reflected in the buyer's second-order beliefs
- Also consider the incentive to decieve/mis-lead/bluff/posture – even if car is bad, seller starts with high price to make it seem like car is good
- Say the seller's first-order belief about the car is that it is in poor condition, but that its second-order belief is that the buyer believes that the car is in good condition – in this case, the seller believes that it is in a better bargaining position than it really is and, therefore, can start with a higher offer than it would have started with had it only considered its first order belief

bargaining in the decision-theoretic framework: some intuitive epistemological insights

Analyze the game's epistemology at various levels:

- Both agents (seller & buyer) have first-order uncertainty about the condition of the car: if the seller believes the car is in poor condition (similarly, if the buyer believes the car is in good condition, or, desires this particular car), it finds itself in a weak bargaining position
- Say the seller starts with a low offer (or, reduces prices drastically during negotiation) – this is an indication to the buyer that the seller believes there is something wrong with the car; which is then (or, should be) reflected in the buyer's second-order beliefs
- Also consider the incentive to decieve/mis-lead/bluff/posture – even if car is bad, seller starts with high price to make it seem like car is good
- Say the seller's first-order belief about the car is that it is in poor condition, but that its second-order belief is that the buyer believes that the car is in good condition – in this case, the seller believes that it is in a better bargaining position than it really is and, therefore, can start with a higher offer than it would have started with had it only considered its first order belief

bargaining in the decision-theoretic framework: some intuitive epistemological insights

Analyze the game's epistemology at various levels:

- Both agents (seller & buyer) have first-order uncertainty about the condition of the car: if the seller believes the car is in poor condition (similarly, if the buyer believes the car is in good condition, or, desires this particular car), it finds itself in a weak bargaining position
- Say the seller starts with a low offer (or, reduces prices drastically during negotiation) – this is an indication to the buyer that the seller believes there is something wrong with the car; which is then (or, should be) reflected in the buyer's second-order beliefs
- Also consider the incentive to decieve/mis-lead/bluff/posture – even if car is bad, seller starts with high price to make it seem like car is good
- Say the seller's first-order belief about the car is that it is in poor condition, but that its second-order belief is that the buyer believes that the car is in good condition – in this case, the seller believes that it is in a better bargaining position than it really is and, therefore, can start with a higher offer than it would have started with had it only considered its first order belief

general insights on interactive epistemological modeling in a decision-theoretic framework

An agent's beliefs ought to contain not just first-order beliefs but second-order beliefs also

general insights on interactive epistemological modeling in a decision-theoretic framework

An agent's beliefs ought to contain not just first-order beliefs but second-order beliefs also

But, what about third-order beliefs? and fourth-order beliefs?

general insights on interactive epistemological modeling in a decision-theoretic framework

An agent's beliefs ought to contain not just first-order beliefs but second-order beliefs also

But, what about third-order beliefs? and fourth-order beliefs?

We realize that all higher orders of beliefs are relevant (possibly to varying degrees) to the decision-maker

general insights on interactive epistemological modeling in a decision-theoretic framework

An agent's beliefs ought to contain not just first-order beliefs but second-order beliefs also

But, what about third-order beliefs? and fourth-order beliefs?

We realize that all higher orders of beliefs are relevant (possibly to varying degrees) to the decision-maker

Problem: in general, these infinite hierarchies of beliefs cannot be represented on a finite computer

general insights on interactive epistemological modeling in a decision-theoretic framework

An agent's beliefs ought to contain not just first-order beliefs but second-order beliefs also

But, what about third-order beliefs? and fourth-order beliefs?

We realize that all higher orders of beliefs are relevant (possibly to varying degrees) to the decision-maker

Problem: in general, these infinite hierarchies of beliefs cannot be represented on a finite computer

WHAT TO DO?

decision theory with finite interactive epistemologies

decision theory with finite interactive epistemologies

Agents **represent higher-order uncertainty only up to a finite level** of the infinite interactive epistemology

decision theory with finite interactive epistemologies

Agents **represent higher-order uncertainty only up to a finite level** of the infinite interactive epistemology

Agents that can represent (and, therefore, reason with) more (i.e. deeper) levels of the interactive belief hierarchy are said to be more 'sophisticated' (or, said to have a higher 'strategy level')

decision theory with finite interactive epistemologies

Agents **represent higher-order uncertainty only up to a finite level** of the infinite interactive epistemology

Agents that can represent (and, therefore, reason with) more (i.e. deeper) levels of the interactive belief hierarchy are said to be more 'sophisticated' (or, said to have a higher 'strategy level')

These modeling assumptions can be philosophically defended on the basis of bounded (or, limited) rationality (for humans as well as machines)

decision theory with finite interactive epistemologies

Agents **represent higher-order uncertainty only up to a finite level** of the infinite interactive epistemology

Agents that can represent (and, therefore, reason with) more (i.e. deeper) levels of the interactive belief hierarchy are said to be more 'sophisticated' (or, said to have a higher 'strategy level')

These modeling assumptions can be philosophically defended on the basis of bounded (or, limited) rationality (for humans as well as machines)

These insights and the development of the formal model to represent and reason in this manner (known as the I-POMDP) are due to my advisor Piotr Gmytrasiewicz and his former student, Prashant Doshi

bilateral bargaining using decision theory with finite epistemologies

The elegance of this model inspired my own studies/research about how decision theoretic agents can participate in bilateral bargaining using such finite representations of interactive beliefs

bilateral bargaining using decision theory with finite epistemologies

The elegance of this model inspired my own studies/research about how decision theoretic agents can participate in bilateral bargaining using such finite representations of interactive beliefs

Why did I chose the bargaining problem?

- Interested in the applicability of the finitely-nested decision-theoretic model to non-trivial problems
- Interested in the usefulness of the finitely-nested decision-theoretic model from a computational complexity/efficiency perspective

bilateral bargaining using decision theory with finite epistemologies

The elegance of this model inspired my own studies/research about how decision theoretic agents can participate in bilateral bargaining using such finite representations of interactive beliefs

Why did I chose the bargaining problem?

- Interested in the applicability of the finitely-nested decision-theoretic model to non-trivial problems
- Interested in the usefulness of the finitely-nested decision-theoretic model from a computational complexity/efficiency perspective

the **strategy level** of the agent models

Level-0 agents (denoted as L-0Seller and L-0Buyer) only model the world (i.e. they maintain beliefs about the world, for e.g. the condition of the car); they do not model other agents' beliefs (i.e. they do not maintain higher-order beliefs about other agents' beliefs)

the **strategy level** of the agent models

Level-0 agents (denoted as L-0Seller and L-0Buyer) only model the world (i.e. they maintain beliefs about the world, for e.g. the condition of the car); they do not model other agents' beliefs (i.e. they do not maintain higher-order beliefs about other agents' beliefs)

Level-1 agents model the world as well as the opponents as being some Level-0 type agents; for e.g. an L-1Buyer can model the seller as being one of many possible L-0Seller type agents (each ascribed some probability) – therfore, we say that Level-1 agents are capable of maintaining second-order beliefs

the **strategy level** of the agent models

Level-0 agents (denoted as L-0Seller and L-0Buyer) only model the world (i.e. they maintain beliefs about the world, for e.g. the condition of the car); they do not model other agents' beliefs (i.e. they do not maintain higher-order beliefs about other agents' beliefs)

Level-1 agents model the world as well as the opponents as being some Level-0 type agents; for e.g. an L-1Buyer can model the seller as being one of many possible L-0Seller type agents (each ascribed some probability) – therfore, we say that Level-1 agents are capable of maintaining second-order beliefs

Level-2 agents model the world as well as the opponents as being some Level-0 **OR** Level-1 type agents – therfore, we say that Level-2 agents are capable of maintaining third-order beliefs

the **strategy level** of the agent models

Level-0 agents (denoted as L-0Seller and L-0Buyer) only model the world (i.e. they maintain beliefs about the world, for e.g. the condition of the car); they do not model other agents' beliefs (i.e. they do not maintain higher-order beliefs about other agents' beliefs)

Level-1 agents model the world as well as the opponents as being some Level-0 type agents; for e.g. an L-1Buyer can model the seller as being one of many possible L-0Seller type agents (each ascribed some probability) – therfore, we say that Level-1 agents are capable of maintaining second-order beliefs

Level-2 agents model the world as well as the opponents as being some Level-0 **OR** Level-1 type agents – therfore, we say that Level-2 agents are capable of maintaining third-order beliefs

and so on.....

Level-0 agents

L-0Buyer accepts an offer x if it gains some minimum profit (mpd) from doing so (i.e. accept outstanding offer x if $v - x \geq mpd$) – an L-0Buyer is completely characterized by its valuation v and its mpd

Level-0 agents

L-0Buyer accepts an offer x if it gains some minimum profit (mpd) from doing so (i.e. accept outstanding offer x if $v - x \geq mpd$) – an L-0Buyer is completely characterized by its valuation v and its mpd

L-0Seller starts with any arbitrary offer; then, if it is rejected, decrements it by some fixed amount or some random amount – an L-0Seller is, therefore, characterized by its "decrement policy"

L-1Buyer

L-1Buyer is characterized by both its valuation as well as its beliefs about the seller – which it considers as being one of fintely many possible L-0Seller types (each assigned some initial (i.e. prior) non-zero probability)

L-1Buyer

L-1Buyer is characterized by both its valuation as well as its beliefs about the seller – which it considers as being one of fintely many possible L-0Seller types (each assigned some initial (i.e. prior) non-zero probability)

An example of an L-1Buyer's beliefs, belief updates and best-response computation

Example 2. Suppose that the buyer, assumed to have a *valuation* of 0.7, believes, with probabilities 0.5, 0.4 and 0.1, respectively, that the seller is one of three possible types – a *subintentional* L0-Seller with a fixed *schedule* of offers, say $\{1.0, 0.9, 0.7, 0.4, 0.3, 0.2, 0.1\}$, or, an L0-Seller that decrements the most recently rejected offer by a fixed value, say 0.1, or, an L0-Seller that decrements the most recently rejected offer by a random value (this is the *random* model). Let these models be called *Model 1*, *Model 2* and *Model 3* respectively. Now, suppose that the opposing seller agent makes offers according to the following schedule $\{1.0, 0.9, 0.8, 0.7, 0.5, 0.3, 0.1\}$. Assume that the discount factor at every step is 0.7.

Table 2 shows the buyer's belief updates and decision-making at every stage.

Offer	Belief Update	Expected Next Offer	Expected Profit On Reject	Accept/Reject
1.0	All models survive	$0.5 \times 0.9 + 0.4 \times 0.9 + 0.1 \times 0.5 = 0.86$	< 0	Reject
0.9	All models survive	$0.5 \times 0.7 + 0.4 \times 0.8 + 0.1 \times 0.45 = 0.715$	< 0	Reject
0.8	Model 2 and Model 3 survive	$0.8 \times 0.7 + 0.2 \times 0.4 = 0.64$	$0.7 \times (0.7 - 0.64) = 0.042$	Reject
0.7	Model 2 and Model 3 survive	$0.8 \times 0.6 + 0.2 \times 0.35 = 0.55$	$0.7 \times (0.7 - 0.55) = 0.105$	Reject
0.5	Only Model 3 survive	$1 \times 0.25 = 0.25$	$0.7 \times (0.7 - 0.25) = 0.315$	Reject
0.3	Model 3 survives	$1 \times 0.15 = 0.15$	$0.7 \times (0.7 - 0.15) = 0.385$	Accept

L-2Seller

L-2Seller is characterized by its beliefs about the buyer – which it considers as being one of fintely many possible L-0Buyer OR L-1Buyer types (each assigned some initial (i.e. prior) non-zero probability)

L-2Seller

L-2Seller is characterized by its beliefs about the buyer – which it considers as being one of fintely many possible L-0Buyer OR L-1Buyer types (each assigned some initial (i.e. prior) non-zero probability)

All the key insights regarding higher-order probabilistic belief dynamics & bounded rationality will be presented in the context of the L-2Seller

L-2Seller: decision-theoretic optimal sequential planning with finite interactive epistemology

L-2Seller: decision-theoretic optimal sequential planning with finite interactive epistemology

The seller evaluates "all" alternative conditional plans

- Based on its prior beliefs, it can evaluate the expected profit it obtains when a particular offer, say x, is either accepted or rejected – if it is accepted, the immediate income is just $(c-x)$; if the offer is rejected, the agent updates its beliefs (and, higher-order beliefs) and then computes the discounted expected profit from the next stage (i.e. offer)
- In this manner, it evaluates all possible conditonal plans of the following form: "If I offer $x1$ and if it is rejected, then I revise my beliefs (using Bayes' rule) and then (if I offer $x1-2$...), ...),....), which gives me an expected profit of $Pr(x1)$; if instead, I offer $x2$,....." and, thereby, computes an optimal conditional plan

L-2Seller: decision-theoretic optimal sequential planning with finite interactive epistemology

The seller evaluates "all" alternative conditional plans

- Based on its prior beliefs, it can evaluate the expected profit it obtains when a particular offer, say x, is either accepted or rejected – if it is accepted, the immediate income is just $(c-x)$; if the offer is rejected, the agent updates its beliefs (and, higher-order beliefs) and then computes the discounted expected profit from the next stage (i.e. offer)
- In this manner, it evaluates all possible conditonal plans of the following form: "If I offer $x1$ and if it is rejected, then I revise my beliefs (using Bayes' rule) and then (if I offer $x1-2$...), ...),....), which gives me an expected profit of $Pr(x1)$; if instead, I offer $x2$,....." and, thereby, computes an optimal conditional plan

L-2Seller: decision-theoretic optimal sequential planning with finite interactive epistemology

The seller evaluates "all" alternative conditional plans

- Based on its prior beliefs, it can evaluate the expected profit it obtains when a particular offer, say x, is either accepted or rejected – if it is accepted, the immediate income is just $(c-x)$; if the offer is rejected, the agent updates its beliefs (and, higher-order beliefs) and then computes the discounted expected profit from the next stage (i.e. offer)
- In this manner, it evaluates all possible conditonal plans of the following form: "If I offer $x1$ and if it is rejected, then I revise my beliefs (using Bayes' rule) and then (if I offer $x1-2$...), ...),....), which gives me an expected profit of $Pr(x1)$; if instead, I offer $x2$,....." and, thereby, computes an optimal conditional plan

L-2Seller: decision-theoretic optimal sequential planning with finite interactive epistemology (illustrated)

L-2Seller: decision-theoretic optimal sequential planning with finite interactive epistemology (illustrated)

The seller's prior beliefs:

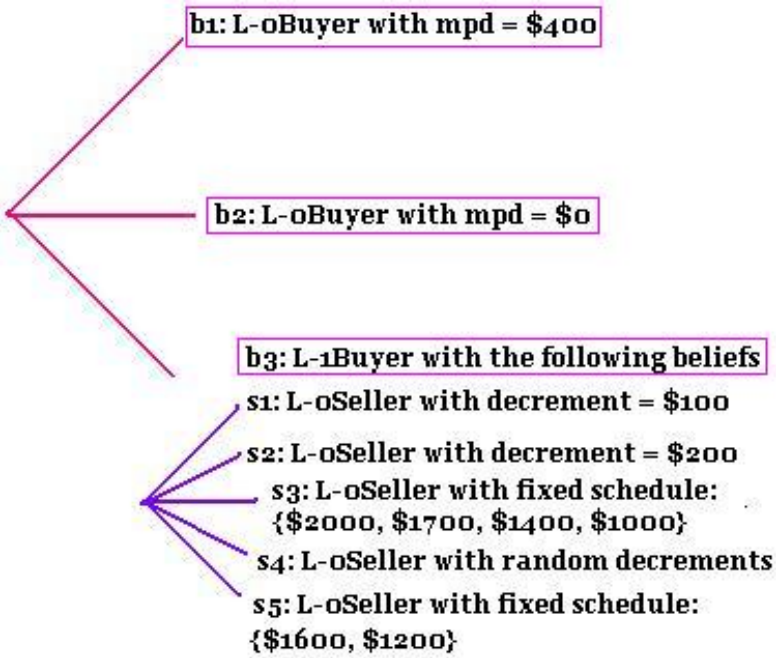

L-2Seller: decision-theoretic optimal sequential planning with finite interactive epistemology (illustrated)

The seller's prior beliefs:

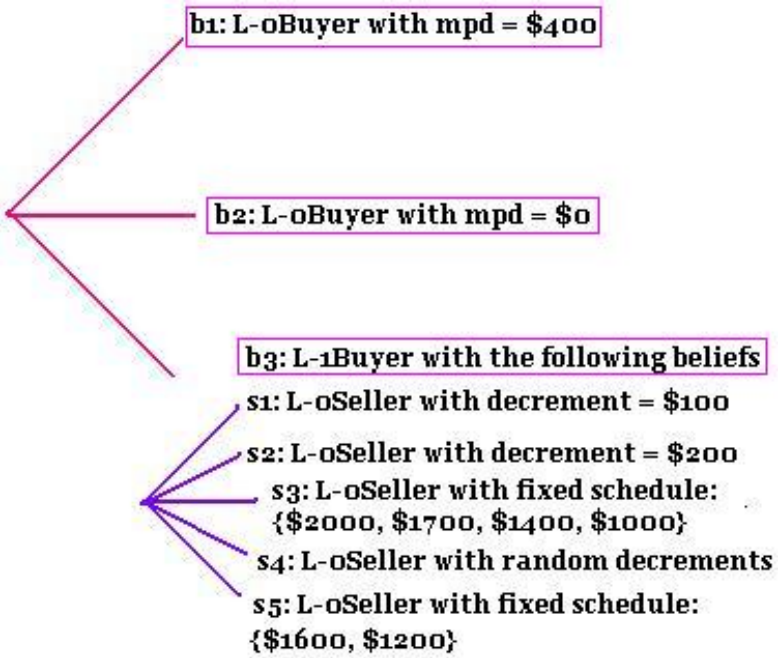

The seller's prior beliefs (abbreviated):

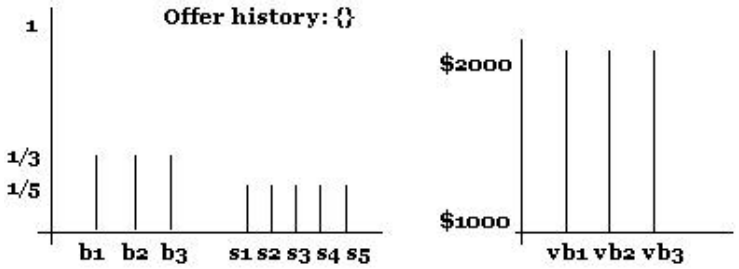

L-2Seller: decision-theoretic optimal sequential planning with finite interactive epistemology (illustrated)

The seller's deliberation tree:

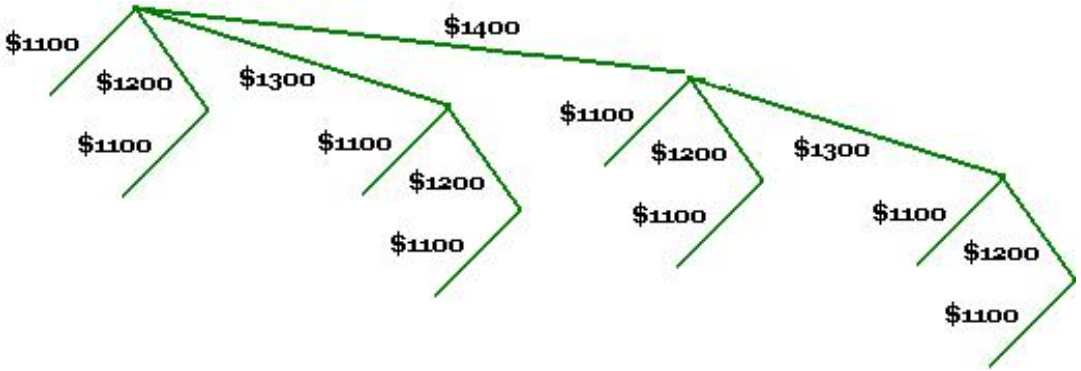

$\{\}$

The seller's beliefs (abbreviated):

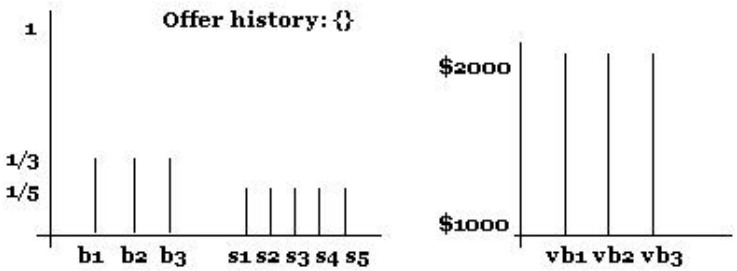

L-2Seller: decision-theoretic optimal sequential planning with finite interactive epistemology (illustrated)

The seller's deliberation tree:

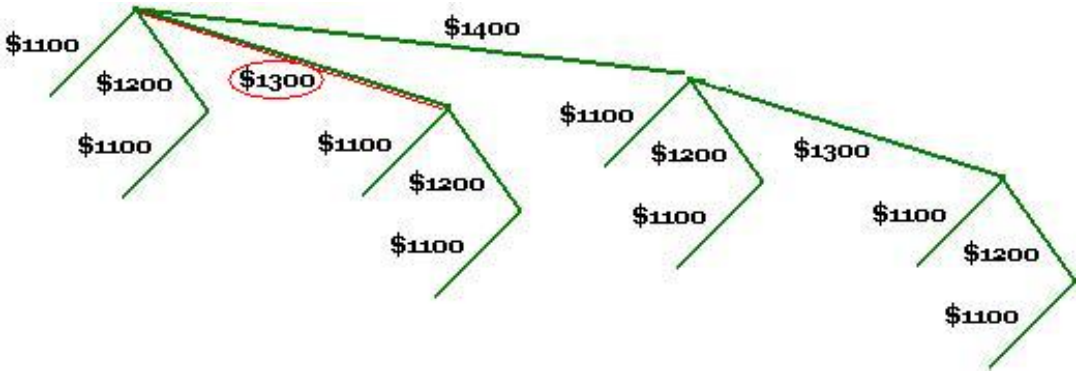

$\{\} \rightarrow \{\$1300\}$

The seller's beliefs (abbreviated):

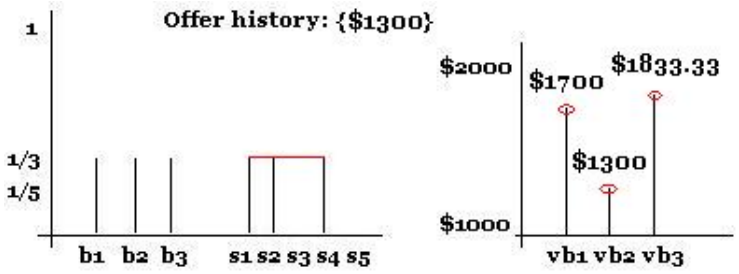

L-2Seller: decision-theoretic optimal sequential planning with finite interactive epistemology (illustrated)

The seller's deliberation tree:

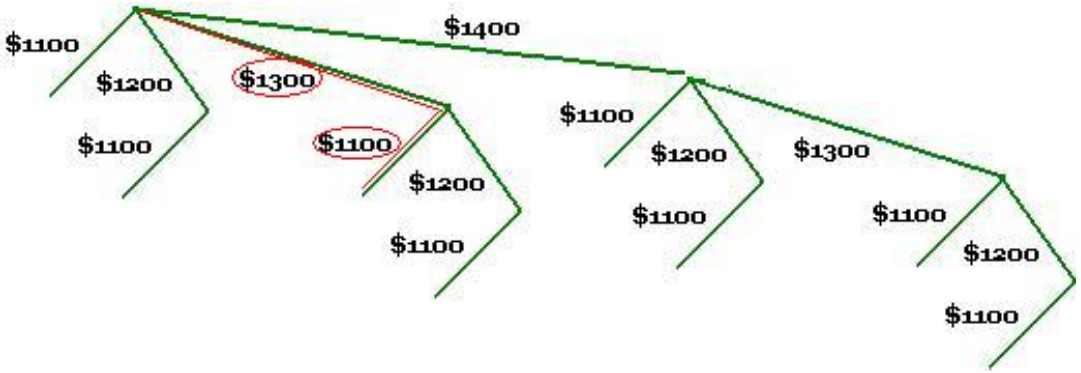

$$\{\} \rightarrow \{\$1300\} \rightarrow \{\$1300, \$1100\}$$

The seller's beliefs (abbreviated):

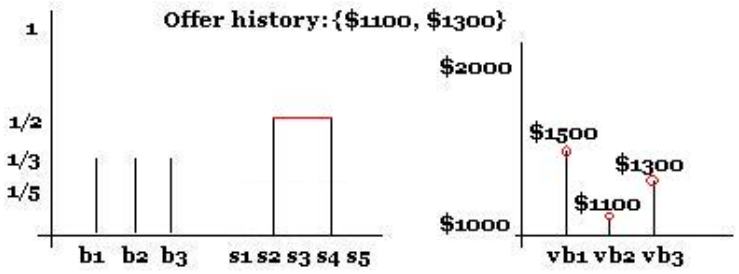

L2Seller: complexity analysis & constant factor improvements through disk-based memoization

The seller's deliberation tree:

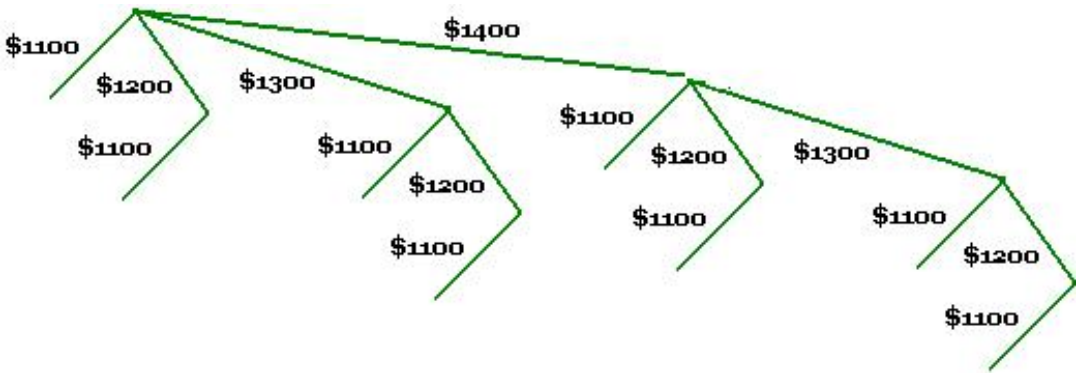

L2Seller: complexity analysis & constant factor improvements through disk-based memoization

The seller's deliberation tree:

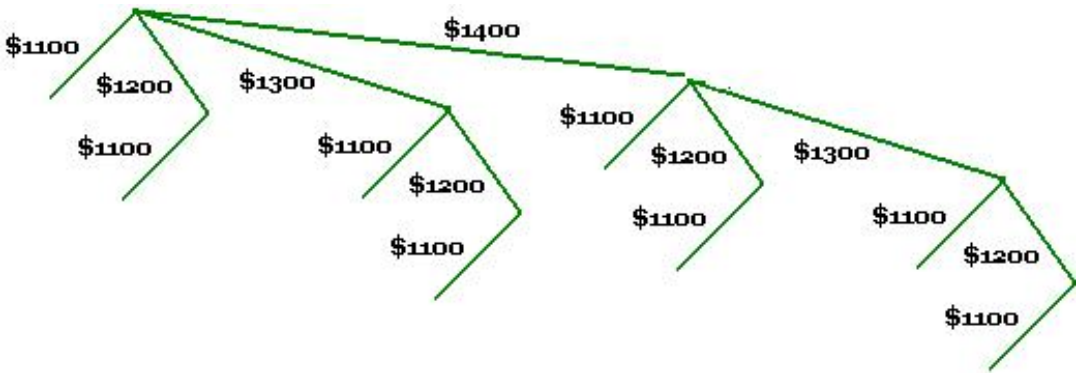

The seller's action space (set of possible offers) is uncountable – discretization of which is then necessary to make it amenable to a discrete algorithm

L2Seller: complexity analysis & constant factor improvements through disk-based memoization

The seller's deliberation tree:

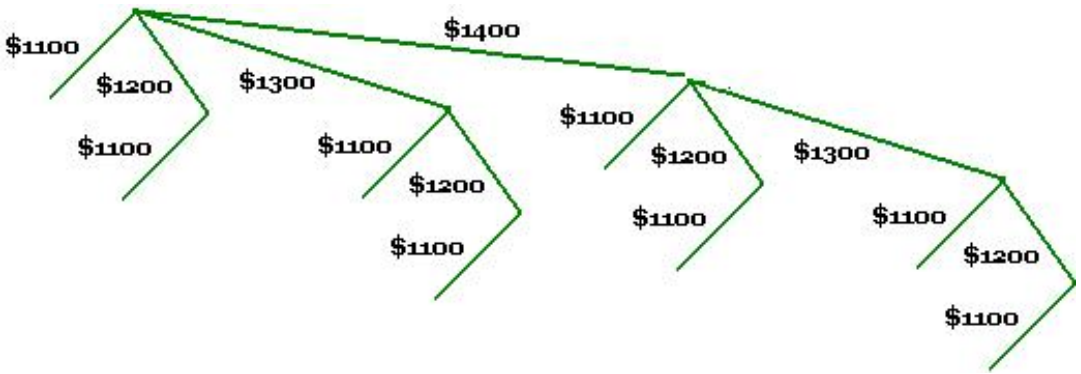

The seller's action space (set of possible offers) is uncountable – discretization of which is then necessary to make it amenable to a discrete algorithm

The algorithm is exponential in the action space

L2Seller: complexity analysis & constant factor improvements through disk-based memoization

The seller's deliberation tree:

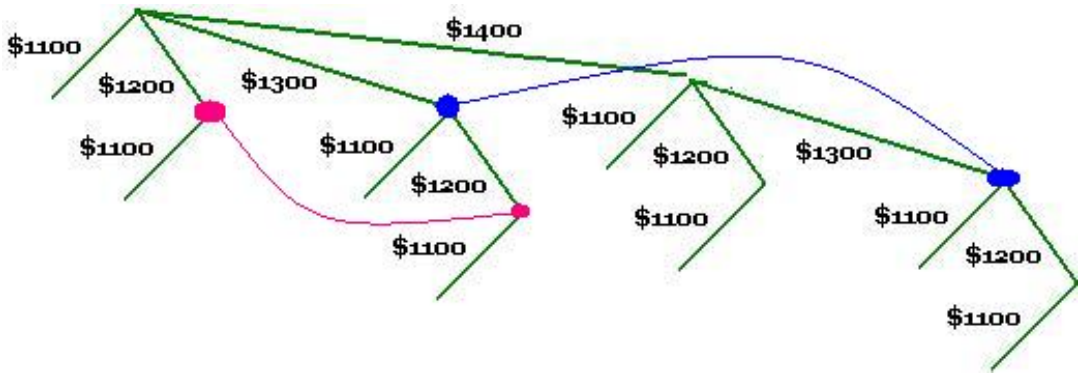

The seller's action space (set of possible offers) is uncountable – discretization of which is then necessary to make it amenable to a discrete algorithm

The algorithm is exponential in the action space

A simple improvement: The complexity can be slighlty mitigated through the use of a disk-based caching mechanism that stores belief states that have been solved to avoid its recomputation – this is essentially memoized dynamic programming

L2Seller: complexity analysis & constant factor improvements through disk-based memoization

The seller's deliberation tree:

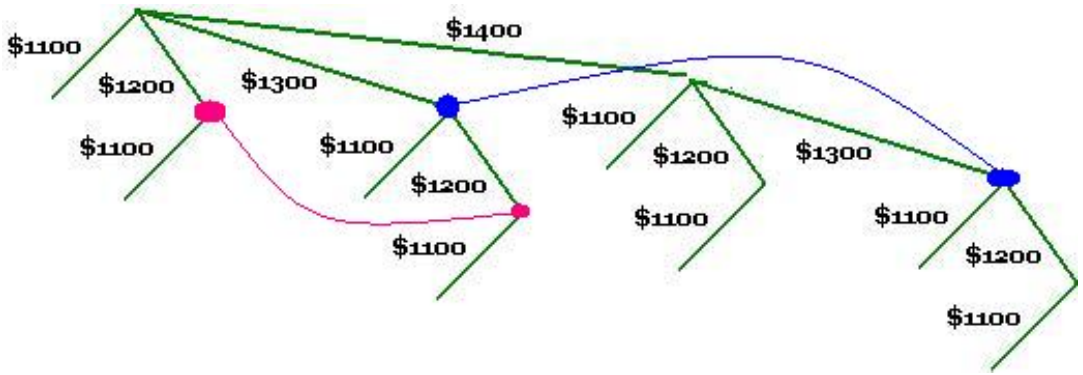

The seller's action space (set of possible offers) is uncountable – discretization of which is then necessary to make it amenable to a discrete algorithm

The algorithm is exponential in the action space

A simple improvement: The complexity can be slighlty mitigated through the use of a disk-based caching mechanism that stores belief states that have been solved to avoid its recomputation – this is essentially memoized dynamic programming

The order of the complexity is still exponential – we only achieve a constant factor improvement (at best!)

bounded rationality realized as **trade-off** between **optimality** and **efficiency!**

bounded rationality realized as **trade-off** between **optimality** and **efficiency!**

Recall again that the complexity is exponential in the dimension of the action space

bounded rationality realized as **trade-off** between **optimality** and **efficiency!**

Recall again that the complexity is exponential in the dimension of the action space

The dimension of the action space is just the discretization resolution

bounded rationality realized as **trade-off** between **optimality** and **efficiency!**

Recall again that the complexity is exponential in the dimension of the action space

The dimension of the action space is just the discretization resolution

Higher (i.e. finer) resolution corresponds to a more fine-grained search through the space of conditional plans and, therefore, corresponds to a solution that is closer to the optimal

bounded rationality realized as **trade-off** between **optimality** and **efficiency!**

Recall again that the complexity is exponential in the dimension of the action space

The dimension of the action space is just the discretization resolution

Higher (i.e. finer) resolution corresponds to a more fine-grained search through the space of conditional plans and, therefore, corresponds to a solution that is closer to the optimal

We immediately see a tradeoff that can be obtained between optimality and efficiency (speed) based on the setting of the discretization resolution – since we have an upper-bound of the running-time, we can use this to implement this optimal planning algorithm as a contract algorithm

empirical evidence

This table provides empirical evidence for both the memoization enhancement as well as the achievement of bounded rationality!

d	τ	optimal schedule	θ	$\left(c, 2^{\lfloor \frac{0.9}{d}+1 \rfloor}\right)$	$\hat{\mu}$
0.02	inf	0.48→0.28→0.2→0.14→0.12→0.1	0.216036	174, 2^{46}	17.861, 17.826
0.02	7	0.48→0.28→0.2→0.14→0.12→0.1	0.216036	610, 2^{46}	64.328, 62.477
0.02	6	0.48→0.28→0.2→0.14→0.12→0.1	0.216036	513, 2^{46}	52.626, 52.501
0.02	5	0.48→0.28→0.2→0.14→0.1	0.215916	408, 2^{46}	41.153, 41.418
0.02	4	0.48→0.28→0.18→0.12	0.214824	295, 2^{46}	29.369, 29.228
0.02	3	0.46→0.26→0.18	0.212658	174, 2^{46}	16.94, 16.426
0.02	2	0.48→0.28	0.21224	45, 2^{46}	4.512, 4.768
0.05	inf	0.55→0.3→0.2→0.15→0.1	0.205083	66, 2^{19}	6.536, 6.252
0.05	6	0.55→0.3→0.2→0.15→0.1	0.205083	162, 2^{19}	14.958, 14.774
0.05	5	0.55→0.3→0.2→0.15→0.1	0.205083	138, 2^{19}	12.759, 13.378
0.05	4	0.55→0.3→0.2→0.15	0.203733	106, 2^{19}	9.575, 9.763
0.05	3	0.55→0.3→0.2	0.201405	66, 2^{19}	6.111, 5.988
0.05	2	0.6→0.35	0.202639	18, 2^{19}	2.029, 1.774
0.1	inf	0.5→0.3→0.2→0.1	0.201527	30, 2^{10}	3.09, 2.916
0.1	5	0.5→0.3→0.2→0.1	0.201527	49, 2^{10}	4.729, 4.5
0.1	4	0.5→0.3→0.2→0.1	0.201527	43, 2^{10}	3.916, 3.891
0.1	3	0.5→0.3→0.2	0.199453	30, 2^{10}	2.9, 2.8
0.1	2	0.6→0.4	0.2	9, 2^{10}	1.264, 1.192
0.2	inf	0.5→0.3→0.1	0.18025	10, 2^{5}	1.118, 1.242
0.2	4	0.5→0.3→0.1	0.18025	11, 2^{5}	1.41, 1.277
0.2	3	0.7→0.5→0.3	0.190684	10, 2^{5}	1.287, 1.125
0.2	2	0.5→0.3	0.197568	4, 2^{5}	0.658, 0.741

Outline

conclusions & future work

conclusions & future work

Finite epistemological decision theoretic modeling using probabilistic beliefs and higher-order beliefs are useful for sequential planning problems in multi-agent settings under uncertainty

conclusions & future work

Finite epistemological decision theoretic modeling using probabilistic beliefs and higher-order beliefs are useful for sequential planning problems in multi-agent settings under uncertainty

Action space discretization provides a natural realization of bounded rationality as a tradeoff between optimality and efficiency

conclusions & future work

Finite epistemological decision theoretic modeling using probabilistic beliefs and higher-order beliefs are useful for sequential planning problems in multi-agent settings under uncertainty

Action space discretization provides a natural realization of bounded rationality as a tradeoff between optimality and efficiency

Memoization (i.e. caching) is helpful in practice to mitigate the exponentiality of the space of conditional plans

Thank You