A Logical Formulation for Negotiation Among Dishonest Agents

Chiaki Sakama
Computer and Communication Sciences
Wakayama University
Wakayama 640-8510, Japan
sakama@sys.wakayama-u.ac.jp

Tran Cao Son and Enrico Pontelli
Department of Computer Science
New Mexico State University
Las Cruces, NM 88003, USA
tson,epontell@cs.nmsu.edu

ABSTRACT

The paper introduces a logical framework for negotiation among (dis)honest agents. The framework relies on the use of abductive logic programming as a knowledge representation language for agents to deal with incomplete information and preferences. The paper shows how intentionally false or inaccurate information of agents could be encoded in the agents' knowledge bases. Such *disinformation* can be effectively used in the process of negotiation to have desired outcomes by agents. The negotiation processes is formulated under the answer set semantics of abductive logic programming and enables the exploration of various strategies that agents can employ in their negotiation.

Categories and Subject Descriptors

I.2.11 [**Distributed Artificial Intelligence**]: Multi-Agent systems

General Terms

Theory

Keywords

Negotiation, Abductive Logic Programming, Lying, Bullshitting, Preferences

1. INTRODUCTION

Negotiation has been an important research topic in multi-agent research and artificial intelligence. Several formalisms have been developed to model negotiations among agents (e.g., [1, 4, 5, 8, 11, 14, 15, 10]). Some deal with the issue of incomplete information (e.g., [5, 15]) but only few consider both *lying* and *incomplete information* during negotiation [3, 16]. These two issues, together with agents' preferences, are frequently observed during negotiations. The following dialogue between a buyer agent b and a seller agent s illustrates these issues:

b_1 : *"I like a digital camera by the maker C. I want to get one that has good quality at a discount price."*

s_1 : *"The product A is made by C and has good quality. We provide a discounted price to students."*
(In reality, the seller does not know the quality of A.)

b_2 : *"I am not a student."*

s_2 : *"The product B by the maker D is on bargain sale. It has good quality and is provided at a discount price for every customer paying in cash."*
(In reality, the seller knows that B is not of good quality.)

b_3 : *"I do not want products by the maker D at the price."*

s_3 : *"If you join our mailing list, we can provide the product at the lowest price."*

b_4 : *"I'd like to join the list and buy it at the price."*
(In reality, the buyer does not want to join the list and will block every email from the shop with a spam filter.)

In this negotiation, the seller has the goal of selling a product while the buyer wants to buy a product. Although fairly simple, the negotiation highlights all difficulties which need to be addressed for any framework for formalizing negotiation: (*i*) *incomplete information*: the seller does not know whether the buyer is a student or not at the beginning and comes to learn that the buyer is not a student only during the negotiation; (*ii*) *deception*: the seller intentionally lies about the quality of the product B to achieve his goal while the buyer intentionally agrees to join the mailing list only to get the deal which, he thought, is a good one; and (*iii*) *preference and goal change*: the buyer prefers a product made by the maker C but ends up buying a product made by the maker D.

In this paper, we explore the use of *abductive logic programming* in formalizing negotiations with incomplete information, preferences, and disinformation (including lying and bullshitting). Addressing these three issues in a single framework makes our work significantly different from previously developed models of negotiation using logic programming, abduction, and argumentation (e.g., [1, 2, 8, 11, 14, 15, 10]), or using utility-based approaches (e.g., [4, 5]). We represent the knowledge of each agent as a logic program, extended with a set of assumptions, a set of ordered goals, disinformation, and preferences; we employ answer sets as a means to define the basic components of a negotiation process, such as proposals, their acceptability, and responses.

2. ABDUCTIVE LOGIC PROGRAMMING (ALP)

A logic program Π is a set of *rules* of the form

$$c_1; \cdots ; c_k \leftarrow a_1, \ldots, a_m, not\ a_{m+1}, \ldots, not\ a_n \qquad (1)$$

where $0 \leq m \leq n,\ 0 \leq k$, each a_i and c_j is a literal of a propositional language[1] and not represents *negation-as-failure*. Formulae of the form $not\ a$ are called negation as failure literals (naf-literals). For a rule r of the form (1), the left and right hand sides

[1]A rule with variables represents the set of its ground instances.

Cite as: The title of your paper should be written here, Author1, Author2 and Author3, *Proc. of 10th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2011)*, Yolum, Tumer, Stone and Sonenberg (eds.), May, 2–6, 2011, Taipei, Taiwan, pp. XXX–XXX.

of the rule are called the *head* ($H(r)$) and the *body* ($B(r)$) of r, respectively. $head(r)$ denotes the set $\{c_1, \ldots, c_k\}$; and $pos(r)$ and $neg(r)$ denote $\{a_1, \ldots, a_m\}$ and $\{a_{m+1}, \ldots, a_n\}$, respectively. A *non-disjunctive* rule is a rule having a single literal in its head ($k = 1$). A rule with empty head is a *constraint*, while a rule with empty body is a *fact*. A fact $L \leftarrow$ is often written as L.

Let X be a set of ground literals. X is consistent if there is no atom a such that $\{a, \neg a\} \subseteq X$. The body of a rule r of the form (1) is *satisfied* by X if $neg(r) \cap X = \emptyset$ and $pos(r) \subseteq X$. A rule of the form (1) with a non-empty head is satisfied by X if either its body is not satisfied by X or $head(r) \cap X \neq \emptyset$. A constraint is *satisfied* by X if its body is not satisfied by X.

For a consistent set S of ground literals and a program Π, the *reduct* of Π w.r.t. S, denoted by Π^S, is the program obtained from the set of all ground instances of Π by deleting (**i**) each rule that has a naf-literal *not* a in its body with $a \in S$, and (**ii**) all naf-literals in the bodies of the remaining rules.

S is an *answer set (or a stable model)* of Π [7] if it satisfies the following conditions: (**i**) If Π does not contain any naf-literals (i.e., $m = n$ in every rule of Π) then S is a minimal consistent set of literals that satisfies all the rules in Π; and (**ii**) If Π contains some naf-literals (i.e., $m < n$ in some rules of Π), then S is an answer set of Π if S is an answer set of Π^S. Note that Π^S does not contain naf-literals, its answer set is defined in (**i**). A program Π is *consistent* if it has a consistent answer set. Otherwise, it is inconsistent.

An *abductive program* is a pair (P^r, P^a) where P^r and P^a are programs. Every element in P^a is called an *abducible*. An abducible $r \in P^a$ is also called an *abducible rule* (resp. *abducible fact*) if r is a rule (resp. a fact). An abducible containing variables is considered as a shorthand of its ground instantiation. So any instance $\rho(r)$ of an element r from P^a is also an abducible and is written as $\rho \in P^a$. Without loss of generality, we assume that literals in the head of rules of P^a do not occur in the head of rules of P^r. Abducibles are hypothetical rules which are used to account for an observation together with the background knowledge P^r. A set of literals S ($\subset Lit$) is a *belief set* of (P^r, P^a) if S is an answer set of $P^r \cup E$ where $E \subseteq P^a$. An abductive program (P^r, P^a) is *consistent* if it has a belief set; otherwise, it is inconsistent.

For the discussion in the following sections, we associate a name n_r to each rule r and freely use the name to represent the rule.[2]

EXAMPLE 1. *Consider the program* $P = (P^r, P^a)$ *where*

$$\begin{array}{rcl} P^r & = & \{n_1 : \ s \leftarrow, \quad n_2 : \leftarrow \mathit{not}\ p, \mathit{not}\ q\,\}, \\ P^a & = & \{\,n_3 : \ p \leftarrow \mathit{not}\ r, \quad n_4 : \ q \leftarrow \mathit{not}\ r\,\}. \end{array}$$

P^r *is inconsistent.* P *has the belief sets* $\{p, s\}$, $\{q, s\}$, $\{p, q, s\}$, *obtained by adding* $\{n_3\}$, $\{n_4\}$, *and* $\{n_3, n_4\}$ *to* P^r.

3. NEGOTIATION KNOWLEDGE BASES

In this section, we define the notion of a logic programming based negotiation knowledge base (KB). We expect the knowledge base to serve as a means for an agent to create his/her proposals/responses in a negotiation, and to decide whether he/she should accept/reject a proposal. An agent KB must encode his/her beliefs, rules for negotiation with their preferences, possible assumptions about the other agent, and possible information that he/she could lie/bullshit about. We will begin with an extension of abductive logic programming that allows agents to consider preferences.

3.1 ALP with Preferences

For an abductive program (P^r, P^a) and a set of rules X, by $P \cup^r X$ (resp. $P \cup^a X$) we denote the abductive program $(P^r \cup X, P^a)$ (resp. $(P^r, P^a \cup X)$). Given a set of literals S, we denote $S^\neg = \{\,\neg\ell \mid \ell \in S\,\}$.[3]

When multiple sets of abducible rules can be used to generate belief sets of a program, a *preference* relation among abducibles, in the form of $\mathit{prefer}(n_1, n_2)$ can be introduced, allowing us to define a preference relation between belief sets. It is assumed that *prefer* is a transitive and anti-symmetric relation among abducibles of a program. The semantics of abductive programs with such preference relations are defined as follows. First, the relation *prefer* is extended to define a preference relation among sets of abducible rules as follows: for $Q_1, Q_2 \subseteq P^a$, Q_1 is preferred to Q_2 if either (*i*) $Q_1 \subseteq Q_2$ or (*ii*) there exist $n_1 \in Q_1 \setminus Q_2$ and $n_2 \in Q_2 \setminus Q_1$ such that $\mathit{prefer}(n_1, n_2)$ holds. In turn, this provides a mean to compare belief sets of a program $P = (P^r, P^a)$; if S_1 (resp. S_2) is a belief set of P obtained from $P^r \cup Q_1$ (resp. $P^r \cup Q_2$), then S_1 is *preferred* to S_2 if Q_1 is preferred to Q_2. A belief set S of P is the *most preferred* belief set if there is no belief set S' of P that is preferred to S. For example, if we add $\mathit{prefer}(n_3, n_4)$ to P^r in Example 1, then $\{p, s, \mathit{prefer}(n_3, n_4)\}$ is the most preferred belief set of $(P^r \cup \{\mathit{prefer}(n_3, n_4)\}, P^a)$.

3.2 Representing Disinformation

Dishonest agents are those who use intentionally false or inaccurate information. In this paper, we consider the following two cases (p is a propositional sentence):

- Agent a, who believes $\neg p$, utters to another agent b that p is true.
- Agent a, who believes neither p nor $\neg p$, utters to another agent b that p (or $\neg p$) is true.

The first type is called a *lie* [9], while the second type is called a *bullshit* (shortly, *BS*) [6]. In both cases, information p brought to a hearer is false or inaccurate (in contrast to the reality as believed by a speaker). We call p *disinformation*. For instance, consider a salesperson who believes that a product is of poor quality. If he/she utters to a customer that the product has a good quality, it is a lie. If a salesperson ignorant about the quality of a product utters to a customer that the product has a good quality, then this creates BS. In abductive programs, disinformation is defined as follows.

DEFINITION 1 (DISINFORMATION). *Let* $P = (P^r, P^a)$ *be an abductive program. Let L and B be two sets of literals s.t.*

- $\forall l \in L$, $\neg l$ *belongs to every belief set of* P.
- $\forall l \in B$, *neither* l *nor* $\neg l$ *belongs to any belief set of* P.

In this case, (L, B) *is called* disinformation *w.r.t.* P.

Literals in L represent lies, as their opposite facts are included in every belief set of P. Literals in B represent BS, as none of them (or their negations) are present in any answer set. Note that $L \cap B = \emptyset$.

We next consider how disinformation is used by the agents in their reasoning. From the "ethical" viewpoint, we assume that agents try to be honest as much as possible. An agent may use disinformation if he/she cannot achieve a goal without it. The situation is realized using an abductive program as follows.

DEFINITION 2 (ALD-PROGRAM). *Let* $P = (P^r, P^a)$ *be an abductive program and* (L, B) *disinformation w.r.t.* P. *Let*

$$\begin{array}{rcl} I & = & \{\, r \mid r \in P^r \text{ and } head(r) \cap L^\neg \neq \emptyset\}, \\ \Phi & = & \{\, \mathit{prefer}(n_i, n_j) \mid n_i \in I \text{ and } n_j \in P^a\} \ \cup \\ & & \{\, \mathit{prefer}(n_j, n_t) \mid n_j \in P^a \cup I \text{ and } n_t \in (L \cup B)\}. \end{array}$$

[2]We omit the rule names when not needed in the discussion.

[3]We assume $\neg\neg a$ to represent the atom a.

An abductive program with disinformation (ALD-program) (L, B) *w.r.t.* P *is defined as the abductive program:* $\delta(P, L, B) = (P^r \setminus I \cup \Phi, P^a \cup I \cup L \cup B)$.

Because $\delta(P, L, B)$ is also an abductive program, belief sets of $\delta(P, L, B)$ are defined as before. The preference relation between belief sets of $\delta(P, L, B)$ is defined next.

DEFINITION 3 (PREFERRED BELIEF SET). *Given two belief sets* S_1 *and* S_2 *of* $\delta(P, L, B)$, *we say that* S_1 *is preferred to* S_2, *denoted by* $S_2 \ll S_1$, *if*

- $S_1 \cap L \subsetneq S_2 \cap L$; *or*
- $S_1 \cap L = S_2 \cap L$ *and* $S_1 \cap B \subsetneq S_2 \cap B$; *or*
- $S_1 \cap (L \cup B) = S_2 \cap (L \cup B)$ *and the set of abducible rules used in generating* S_1 *is preferred to the set of abducible rules used in generating* S_2.

Both lying and BS are dishonesty, but an agent tries to keep lies as small as possible; lies are considered more sinful than BS [6], that is, lies are wrong beliefs while BS are ungrounded beliefs. For a set of belief sets Σ of $\delta(P, L, B)$, we say that $M \in \Sigma$ is a most preferred belief set of Σ if there exists no $M' \in \Sigma$ such that $M \ll M'$. Observe that an ALD-program $\delta(P, \emptyset, \emptyset)$ reduces to the original abductive program P—in absence of disinformation Def. 3 reduces to the notion of preferred belief set in Sect. 3.1.

PROPOSITION 1. *Let* P *be an abductive program and* (L, B) *be disinformation w.r.t.* P. *If* S *is a belief set of* P, *then* $S \cup \Phi$ *is a belief set of* $\delta(P, L, B)$. *Furthermore, if* S *is a most preferred belief set of* P, *then* $S \cup \Phi$ *is a most preferred belief set of* $\delta(P, L, B)$.

The transformation of an abductive program to an ALD-program does not affect the constraints of the program:

PROPOSITION 2. *Let* $P = (P^r, P^a)$ *be an abductive program and* (L, B) *disinformation w.r.t.* P. *Let* R *be a set of constraints. If the abductive program* $Q = (P^r \cup R, P^a)$ *is consistent then* (L, B) *is disinformation w.r.t.* Q, *and a set of literals* S *is a belief set of* $\delta(Q, L, B)$ *iff it is a belief set of* $\delta(P, L, B) \cup^r R$.

Observe that an inconsistent abductive program could recover consistency using disinformation.

EXAMPLE 2. *Consider the program* $P = (P^r, P^a)$ *and disinformation* (L, B) *w.r.t.* P *where* $P^a = \emptyset$ *and*

$$\begin{array}{rcl} P^r & = & \{ \neg quality_A \leftarrow product_A, \\ & & \quad product_A \leftarrow, \quad product_B \leftarrow \}, \\ L & = & \{ quality_A \},\ B = \{ quality_B \}. \end{array}$$

In this case, the ALD-program $\delta(P, L, B) = (\delta P^r, \delta P^a)$ *becomes*

$$\begin{array}{rcl} \delta P^r & = & \{ product_A \leftarrow,\ product_B \leftarrow, \\ & & \quad prefer(n_1, n_2) \leftarrow, prefer(n_1, n_3) \leftarrow \}. \\ \delta P^a & = & \{ n_1 : \neg quality_A \leftarrow product_A \\ & & \quad n_2 : quality_A \leftarrow,\ n_3 : quality_B \leftarrow \}. \end{array}$$

Consider $R_1 = \{ \leftarrow not\, quality_A \}$. *The abductive program* $P' = (P^r \cup R_1, P^a)$ *is inconsistent. On the other hand, the ALD-program* $\delta(P, L, B)$ *has the (most preferred) belief set containing* $quality_A$. *Here, a lie has been used in creating the belief set.*

Given $R_2 = \{ \leftarrow not\, quality_B \}$, *we have that* $Q = (P^r \cup R_2, P^a)$ *is also inconsistent, while the program* $\delta(Q, L, B)$ *has the (most preferred) belief set containing* $quality_B$. *Here, information used for constructing the belief set is a BS.*

3.3 Negotiation Knowledge Bases

This section develops the notion of a negotiation knowledge base for negotiation among dishonest agents. For notational convenience, given a set S of literals, let $Goal(S) = \{ \leftarrow not\ \ell \mid \ell \in S \}$. We first define a knowledge base used for negotiation.

DEFINITION 4 (NEGOTIATION KB). *A tuple* $\langle P, L, B, H, N^{\prec} \rangle$ *is a* negotiation knowledge base with disinformation K *(*n-KB*) if:*

- $P = (P^r, P^a)$ *is an abductive program and* (L, B) *is disinformation w.r.t.* P.
- H *is a set of literals (called* assumptions*) such that* $H \cap head(P^r) = \emptyset$ *and* $H \subseteq P^a$.
- $N^{\prec}$ *is a set of literals (called* negotiated conditions*) associated with a strict partial order* $\prec$ *on its elements.*

The n-KB is consistent *if* P *is consistent.*

Intuitively, the abductive program $P = (P^r, P^a)$ is used by an agent during negotiation to achieve his/her own goals, where P^r consists of rules defining the domain-specific knowledge of the agent and P^a consists of abducible rules defining possible negotiation strategies. To add to his/her flexibility during negotiation, the agent could decide to add some disinformation to the program, in this case (L, B). Proposition 1 ensures that every proposal that can be made without the disinformation (w.r.t. P) can be also made w.r.t. $\delta(P, L, B)$, since the notion of a proposal relies on the notion of a belief set (precise definition in the next section).

The set H represents assumptions about features of the other agent that the agent might not have at the beginning of the negotiation; some of this information might become known during the negotiation. $N^{\prec}$ contains literals expressing the desired properties of the outcome for which the n-KB is developed. $\prec$ represents a preference order of the agent with respect to the negotiated conditions; $p \prec q$ means that q is preferred to q. For simplicity, we assume that $\prec$ is extended to an ordering of the subsets of N.

Since the abductive program P serves as a means for the agent to generate arguments or hypotheses in the negotiation, P is assumed to be consistent unless stated otherwise. Prop. 1 shows that if P is consistent then $\delta(P, L, B)$ is also consistent.

We now present two typical n-KBs, one for a seller and one for a buyer, describing the agents discussed in the introduction. Intuitively, the seller agent will have a KB for him/her to negotiate a *sale*, while the buyer agent will have a KB for him/her to negotiate a *purchase*. The buyer wants to get a product for a low price, while the seller wants to sell a product for a high price.

EXAMPLE 3. *Let us consider a seller agent* s, *described by an abductive program* $P_s = (P^r_s, P^a_s)$ *and the n-KB* $K_s = \langle P_s, L_s, B_s, H_s, N^{\prec}_s \rangle$ *to negotiate with customers:*[4]

[4]Arithmetic predicates are written in infix notation.

- P_s^r *consists of the rules:*

$$
\begin{array}{rcl}
senior\_customer & \leftarrow & age \geq 65 \\
student\_customer & \leftarrow & student \\
\neg quality_B & \leftarrow & product_B \\
maker_C & \leftarrow & product_A \\
maker_D & \leftarrow & product_B \\
bargain & \leftarrow & product_B \\
sale & \leftarrow & product_A, \mathbf{price_1} \\
sale & \leftarrow & product_B, \mathbf{price_2} \\
 & \leftarrow & not\ sale \\
product_A & \leftarrow & \\
product_B & \leftarrow & \\
prefer(n_1, n_i) & \leftarrow & (\text{for } i \in \{2,3,4,5\}) \\
prefer(n_i, n_5) & \leftarrow & (\text{for } i \in \{2,3,4\}) \\
 & \leftarrow & high\_pr, low\_pr \\
 & \leftarrow & high\_pr, lowest\_pr \\
 & \leftarrow & low\_pr, lowest\_pr
\end{array}
$$

where $\mathbf{price_1} \in \{\, high\_pr,\ low\_pr \,\}$ *and* $\mathbf{price_2} \in \{\, low\_pr,\ lowest\_pr,\ high\_pr\}$. *Here,* P_s^r *defines various types of customers, and states some features about the available products. It also states the sales conditions and the preferences among the abducible rules of the KB. The last three constraints indicate that the seller sells a product for only one price.*

- H_s *is a set of assumptions that the seller needs to verify about his customers during negotiation. It is given by*

$$H_s = \{\, age \geq 65,\ student,\ cash,\ mail\_list \,\}.$$

- $N_s^{\prec}$ *is the set of possible prices that the seller should be negotiated about and is given by*

$$N_s^{\prec} = \{\, high\_pr,\ low\_pr,\ lowest\_pr \,\}$$

with $\prec = \{\, lowest\_pr \prec low\_pr \prec high\_pr \,\}$. *This indicates that the seller prefers high_pr over low_pr and lowest_pr.*

- $P_s^a = H_s \cup R_s$ *where* R_s *consists of the following rules*

$$
\begin{array}{llcl}
n_1: & high\_pr & \leftarrow & \\
n_2: & low\_pr & \leftarrow & senior\_customer \\
n_3: & low\_pr & \leftarrow & student\_customer \\
n_4: & low\_pr & \leftarrow & bargain,\ cash \\
n_5: & lowest\_pr & \leftarrow & mail\_list,\ cash
\end{array}
$$

- *Let us assume that the seller intends to claim that both products* A *and* B *are of good quality, if needed. In other words, he/she would lie about* $quality_B$ *and BS about* $quality_A$*; this is described by the disinformation:*

$$(L_s, B_s) = (\{quality_B\}, \{quality_A\}).$$

Rules n_1*-*n_5 *specify different pricing scenarios. Intuitively,* n_1 *says that any customer who agrees to buy the product with high_pr is unconditionally accepted. Senior citizens (*n_2*) and students (*n_3*) are entitled to low_pr. low_pr is also applied to bargain products and for every customer paying in cash (*n_4*). A special discount lowest_pr is applied to a customer who subscribes to the shop's mailing list and purchases the product in cash (*n_5*).*

Observe that, with the introduction of the disinformation (L_s, B_s)*, the abductive program* (P_s^r, P_s^a) *in the n-KB constructs an ALD-program* $\delta(P_s, L_s, B_s) = (\delta P_s^r, \delta P_s^a)$*, with* δP_s^a *equal to* $P_s^a \cup L_s \cup B_s$ *plus the rule:*

$$n_6: \quad \neg quality_B \quad \leftarrow \quad product_B$$

and $\delta P_s^r = P_s^r \setminus \{\neg quality_B \leftarrow product_B\}$ *plus the preferences* $prefer(n_6, n_i)$ *for* $i = 1, \ldots, 5$ *and* $prefer(n_j, n_t)$ *for* $j = 1, \ldots, 6$ *and* n_t *are the names of the facts in* (L_s, B_s)*. It is easy to see that* $\delta(P_s, L_s, B_s)$ *is consistent and thus* K_s *is consistent. Furthermore, none of its belief set contains two distinct prices, i.e., guaranteeing that the choice of a price is unique for the seller.*

EXAMPLE 4. *An n-KB* $K_b = \langle P_b, L_b, B_b, H_b, N_b^{\prec} \rangle$ *with* $P_b = (P_b^r, P_b^a)$ *for the buyer b is given by*

- P_b^r *consists of*

$$
\begin{array}{rcl}
purchase & \leftarrow & product_X,\ quality_X,\ \mathbf{price_3} \\
 & \leftarrow & not\ purchase \\
\neg student & \leftarrow & \\
cash & \leftarrow & \\
prefer(n_1, n_2) & \leftarrow & \\
prefer(n_2, n_3) & \leftarrow & \\
prefer(n_1, n_3) & \leftarrow & \\
 & \leftarrow & low\_pr, lowest\_pr
\end{array}
$$

where $X \in \{A, B\}$ *and* $\mathbf{price_3} \in \{\, low\_pr,\ lowest\_pr \,\}$.

- $H_b = \{\, quality_A,\ quality_B,\ maker_C,\ maker_D,\ product_A,\ product_B \}$.
- $N_b^{\prec} = \{\, low\_pr,\ lowest\_pr \,\}$ *with* $\prec = \{\, low\_pr \prec lowest\_pr \,\}$.
- $P_b^a = H_b \cup R_b$ *where* R_b *consists of*

$$
\begin{array}{llcl}
n_1: & lowest\_pr & \leftarrow & maker_C \\
n_2: & lowest\_pr & \leftarrow & maker_D \\
n_3: & low\_pr & \leftarrow & maker_C
\end{array}
$$

The set H_b *represents properties of products that the buyer needs to check. The set* $N_b^{\prec}$ *specifies negotiated conditions and the buyer prefers to pay the lowest price.*

Suppose that the buyer does not care about the mailing list of the seller but could pretend to join it if it works to his/her advantage. He/she decides to use the disinformation $(L_b, B_b) = (\emptyset, \{mail\_list\})$ *w.r.t.* P_b*. In this case, he/she will use the ALD-program* $\delta(P_b, L_b, B_b)$ *in his/her negotiation. It is easy to check that* K_b *is also consistent.*

4. PROPOSALS AND ACCEPTABILITY

In this section we explore the notion of proposal. In generating a proposal, an agent has a goal and can make assumptions about the receiver of the proposal. The agent may also decide to reveal some information about his/her state-of-belief, rendering some conditions on the feasibility of the proposal.

DEFINITION 5 (PROPOSAL). *Let us consider an n-KB* $K = \langle P, L, B, H, N^{\prec} \rangle$ *of an agent a, and a set of literals* $G \subseteq N^{\prec}$*. A tuple* $\gamma = \langle G, S, R \rangle$ *is a* proposal *for* G *w.r.t.* K *if* $\delta(P, L, B) \cup^r Goal(G)$ *has a belief set* M *such that*

- $S = M \cap H$*, and*
- $R \subseteq M \setminus H$*.*

We refer to G*,* S*,* R*, and* M *as the* goal, assumptions, conditions, *and* support *of* $\langle G, S, R \rangle$*, respectively.*

The proposal is honest *if* $M \cap (L \cup B) = \emptyset$*; it is* deceptive *if* $M \cap L \neq \emptyset$*; and it is* unreliable *if* $M \cap B \neq \emptyset$*.*

$\alpha(K, G)$ *denotes the set of all possible proposals for* G *w.r.t.* K*.*

Intuitively, a proposal $\langle G, S, R \rangle$ states that the goal of a is to negotiate for the objective G. The reason that a puts forward the proposal is that he/she has a support for it (the belief set M). By making the proposal, a indicates assumptions that he/she has made about the receiver of the proposal (the set S). In addition, a also reveals additional information supporting the goal G (the set R), which informs the receiver of the proposal that the information in R should not be violated.

EXAMPLE 5. *In Example 3,* $\langle\{high\_pr\}, \emptyset, \{product_A\}\rangle$ *and* $\langle\{low\_pr\}, \{student\}, \{product_B\}\rangle$ *are two possible proposals by the seller w.r.t.* K_s*. The former indicates that he/she can sell* $product_A$ *for the high price. The latter states that he/she can sell* $product_B$ *for the low price but the customer must be a student. Both of these proposals are honest. A deceptive proposal by the seller is* $\langle\{low\_pr\}, \{student\}, \{product_B, quality_B\}\rangle$*, indicating that he/she can sell the product B, which has good quality, for low price but he/she assumes that the buyer is a student. The proposal is deceptive because the seller has the fact* $\neg quality_B$ *in* K_s*, which shows that* $quality_B$ *is a lie. Similarly, we can see that* $\langle\{high\_pr\}, \emptyset, \{product_A, quality_A\}\rangle$ *is an unreliable proposal w.r.t.* K_s*.*

In Example 4, $\langle\{low\_pr\}, \{product_A, maker_C, quality_A\}, \emptyset\rangle$ *and* $\langle\{lowest\_pr\}, \{product_B, maker_D, quality_B\}, \{\neg student\}\rangle$ *are two possible proposals by the buyer. The first states that the buyer is willing to buy* $product_A$ *with the low price, assuming that it is made by maker C and it is of good quality; the second is similar, for buying* $product_B$ *with the lowest price, assuming that it is made by the maker D and it is of good quality. In addition, the buyer indicates that he/she is not a student in the second proposal. These proposals by the buyer are honest.*

Given an agent a and a proposal $\gamma = \langle G, S, R\rangle$ from another agent b, we can see one of the following cases:

- a accepts γ: This means that the goal of γ matches the goal of a, or at least as good as what a wants. The assumptions stated in S are acceptable to a and are consistent with his/her knowledge. Furthermore, a must see whether the conditions stated in R could be accepted.
- a rejects γ: This means that there is no possible way that a can view γ as his/her proposal.
- a sees some alternative proposals for the goal of γ, yet γ is not suitable for a, i.e., a considers γ a negotiable proposal.

This leads to the following definition.

DEFINITION 6 (PROPOSAL CLASSIFICATION). *Consider an n-KB (of an agent a)* $K = \langle P, L, B, H, N^{\prec}\rangle$ *and a proposal (from another agent b)* $\gamma = \langle G, S, R\rangle$*. Let* $\delta Q = \delta(P, L, B) \cup^r Goal(G)$*. Then,*

- γ *is* acceptable *w.r.t.* K *if* δQ *has a belief set* M *such that* $S \subseteq M$ *and* $M \cap H \subseteq R \cap H$*. We say that* γ *is acceptable* without disinformation *if* $M \cap (L \cup B) = \emptyset$*,* with disinformation*, otherwise.*
- γ *is* rejectable *if* δQ *is inconsistent.*
- γ *is* negotiable*, otherwise.*

Intuitively, δQ encodes the set of possible proposals for G by the agent with the n-KB K. If δQ is inconsistent then the proposal is rejectable. γ is acceptable if Q has a belief set M satisfying the following conditions:

- it is compatible with the assumptions made about him/herself, i.e., it must include S; and
- if there are assumptions made by the receiving agent (a) about the proposer (b), then these must be compatible with the information revealed by the proposer, i.e., $M \cap H \subseteq R \cap H$.

The first condition is needed, since a negotiated goal is acceptable to both parties only if their supports agree. The second condition implies that a proposal is based on the same set of shared assumptions. A proposal is negotiable if it is neither acceptable nor rejectable. Note that when an agent considers a proposal acceptable or negotiable, he/she may use disinformation included in his/her knowledge base.

EXAMPLE 6. *For* K_s *and* K_b *from Examples 3 and 4,*

- $\langle\{high\_pr\}, \{product_A\}, \emptyset\rangle$ *is acceptable without disinformation w.r.t.* K_s *as* $\delta(P_s \cup^r Goal(\{high\_pr\}), L_s, B_s)$ *has a belief set* M *containing* $high\_pr$ *and* $product_A$ *and no disinformation.*
- $\langle\{high\_pr\}, \{product_A, quality_A\}, \emptyset\rangle$ *is acceptable with disinformation w.r.t.* K_s *as* $\delta(P_s \cup^r Goal(\{high\_pr\}), L_s, B_s)$ *has a belief set* M *containing* $high\_pr$*,* $product_A$*, and* $quality_A$*. This is because any belief set that allows the seller to accept this proposal needs to contain* $quality_A$*, which is disinformation.*
- $\langle\{low\_pr\}, \{product_B, maker_D, quality_B\}, \emptyset\rangle$ *is a negotiable proposal w.r.t.* K_s *since* $\delta(P_s \cup^r Goal(\{low\_pr\}), L_s, B_s)$ *has a belief set containing its assumptions but requires at least one of the sets* $\{student\}$*,* $\{age \geq 65\}$*, or* $\{cash\}$*.*
- $\langle\{high\_pr\}, \emptyset, \{product_A, maker_C, quality_A\}\rangle$ *is a rejectable proposal w.r.t.* K_b *because* $\delta(P_b \cup^r Goal(\{high\_pr\}), L_b, B_b)$ *has no belief set containing* $high\_pr$*.*

An agent can employ a stronger condition for accepting a proposal in Definition 6, e.g., by restricting the belief set M to be $M \cap (L \cup B) = \emptyset$. This may lead an agent to consider γ acceptable only if it is acceptable without disinformation. With this stronger condition, the proposal $\langle\{high\_pr\}, \{product_A, maker_C, quality_A\}, \emptyset\rangle$ is rejectable for the seller if he/she wants to be honest. On the other hand, it is negotiable if the seller uses disinformation. We will return to this issue in Section 6.

Let K be an n-KB and $\Gamma_a(K)$, $\Gamma_n(K)$, and $\Gamma_r(K)$ be the set of proposals that are acceptable, negotiable, and rejectable w.r.t. K.

PROPOSITION 3. *Let* K *be an arbitrary n-KB. Then,* $\Gamma_a(K)$*,* $\Gamma_u(K)$*, and* $\Gamma_r(K)$ *are pairwise disjoint. Furthermore, for every proposal* γ *we have that* $\gamma \in \Gamma_a(K) \cup \Gamma_u(K) \cup \Gamma_r(K)$*.* □

5. NEGOTIATION USING N-KBS

We will now present a model of negotiation between two agents a and b who use n-KBs K_a and K_b respectively in their negotiations. We assume that K_a and K_b share the same language. Furthermore, we will assume that the set of assumptions in K_a is disjoint from the set of assumptions in K_b. We envision a negotiation will contain several rounds, each represented in Figure 1. In each round, an agent, called A, puts forward a proposal that includes the goal, the assumptions that A made about his opponent B ($Assumptions_{A>B}$), and the information about his/herself ($Information_A$). The second agent, B, will respond with a proposal with the same structure.

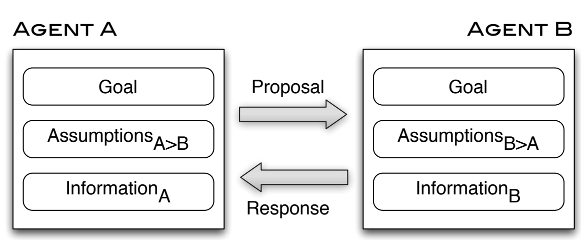

Figure 1: Proposal and Response

We will begin with the definition of a response to a proposal. In our framework, a response could be an arbitrary proposal, an acceptance or a rejection of the current proposal. The only restriction is that an acceptance can only be made if the current proposal is acceptable. This leads to the following definition.

Definition 7 (Response). *Let $K_a = \langle P, L, B, H, N^{\prec} \rangle$ be an n-KB of an agent a and $\gamma_b = \langle G, S, R \rangle$ be a proposal by b w.r.t. its n-KB K_b. A* response *to γ_b by a is*

(i) *A proposal $\gamma_a = \langle G', S', R' \rangle$;* **or**

(ii) *$\langle \top, \emptyset, \emptyset \rangle$, denoting* acceptance of the proposal *if γ_b is acceptable w.r.t. K_a.* **or**

(iii) *$\langle \bot, \emptyset, \emptyset \rangle$, denoting* rejection of the proposal.

Let us now define the notion of a negotiation. Intuitively, a negotiation is a series of responses between two agents, who, in alternation, take into consideration the other agent's response and put forward a new response; this can be either accept, reject, or a new proposal that may involve explanations of why the latest proposal (of the other agent) was not acceptable.

Definition 8 (Negotiation). *Let a and b be two agents and K_a and K_b be their n-KBs respectively. A* negotiation *between a and b, starting with a, is a possible infinite sequence of proposals $\omega_1, \ldots, \omega_i, \ldots$ where $\omega_i = \langle G_i, S_i, F_i \rangle$ and*

- *ω_{2k+1} is a proposal w.r.t. K_a ($k \geq 0$)*
- *ω_{2k} is a proposal w.r.t. K_b ($k \geq 1$)*
- *ω_{i+1} is a response to ω_i for every $i \geq 1$.*

A negotiation ends *at i if $\omega_i = \langle \top, \emptyset, \emptyset \rangle$ or $\omega_i = \langle \bot, \emptyset, \emptyset \rangle$.*

When $G_i \neq G_{i+2}$, we say that a goal change *has occurred for the agent who proposes ω_i.*

We classify negotiations as follows.

Definition 9 (Un/Successful Negotiation). *A negotiation is* successful *(resp.* unsuccessful*) if it is finite and ends with $\omega_i = \langle \top, \emptyset, \emptyset \rangle$ (resp. $\omega_i = \langle \bot, \emptyset, \emptyset \rangle$). We call ω_{i-1} the* accepted *(resp.* rejected*) proposal of the negotiation.*

Example 7. *Consider the seller s and the buyer b agents (Examples 3 and 4).*

$b_1 : \langle \{low\_pr\}, \{product_A, quality_A, maker_C\}, \emptyset \rangle$
$s_1 : \langle \{low\_pr\}, \{student\}, \{product_A, quality_A, maker_C\} \rangle$
$b_2 : \langle \{low\_pr\}, \{product_A, quality_A, maker_C\}, \{\neg student\} \rangle$
$s_2 : \langle \{low\_pr\}, \{cash\}, \{product_B, maker_D, quality_B\} \rangle$
$b_3 : \langle \{lowest\_pr\}, \{product_B, maker_D, quality_B\}, \{cash\} \rangle$
$s_3 : \langle \{lowest\_pr\}, \{cash, mail\_list\}, \{product_B, maker_D, quality_B\} \rangle$
$b_4 : \langle \top, \emptyset, \emptyset \rangle$.

The seller bullshits in s_1 and lies in s_2. The buyer lies in b_4. A goal change has occurred at b_3 (for the buyer) and s_3 (for the seller).

We next define the notion of a constructive negotiation.

Definition 10. *A negotiation $\langle G_i, S_i, R_i \rangle_{i=1}^{\infty}$ is* constructive *if there are no two indices $1 \leq i < j$ such that $\langle G_i, S_i, R_i \rangle = \langle G_j, S_j, R_j \rangle$ and $\langle G_{i+1}, S_{i+1}, R_{i+1} \rangle = \langle G_{j+1}, S_{j+1}, R_{j+1} \rangle$.*

Intuitively, in a constructive negotiation, agents do not repeat their answers. Thus, constructive negotiations lead to finiteness.

Theorem 1. *Every constructive negotiation is finite.*

Proof. *(Sketch) Each agent has only finitely many proposals and each agent can, at each point, terminate the negotiation either by accepting or rejecting the proposal.* □

A negotiation represents one possible way for two agents to reach an agreement (or disagreement). In the course of reaching an agreement, two agents might have different alternatives. The *negotiation tree* accounts for all possible negotiations between two agents. By the *level* of a node in a tree we mean the number of links lying on the path connecting the root to such node. Also, for a proposal $\omega = \langle G, S, F \rangle$ and an n-KB K, by $\beta(K, \omega)$ we denote the set of all responses to ω w.r.t. K.

Definition 11 (Negotiation Tree). *Let a and b be two agents with the n-KBs K_a and K_b. A* negotiation tree *between a and b, starting with a, is a labeled tree $T_{a,b}$ where*

- *The root of $T_{a,b}$ is □;*
- *Each child of □ has the label of the form $(K_a, \langle G, S, R \rangle, K_b)$ where $\langle G, S, R \rangle$ is a proposal w.r.t. K_a;*
- *If $\eta = (K, \omega, K')$ is a node at level $i \geq 1$, then every child of η has the label of the form (K', ω', K) where $\omega' \in \beta(K', \omega)$; and*
- *Nodes labeled by $(K, \langle \top, \emptyset, \emptyset \rangle, K')$ or $(K, \langle \bot, \emptyset, \emptyset \rangle, K')$ have no children.*

Intuitively, a negotiation tree represents all possible negotiations between a and b, starting with a; the tree allows us to predict the possible results of a negotiation between the two agents. It is easy to see that each path from a node of level 1 of $T_{a,b}$ to a leaf is a negotiation between a and b. Observe also that a negotiation tree is analogous to the notion of a protocol as used in [8].

Definition 12 (Classification of Negotiation Trees). *A negotiation tree is* finite *if it has a finite number of nodes. A finite tree is* successful *if it has a leaf whose label is of the form $(K, \langle \top, \emptyset, \emptyset \rangle, K')$; it is* unsuccessful *if all of its leaves have a label of the form $(K, \langle \bot, \emptyset, \emptyset \rangle, K')$.*

The idea of constructive negotiation is extended to negotiation trees.

Definition 13 (Constructive Negotiation Tree). *A negotiation tree is* constructive *if every branch of the tree encodes a constructive negotiation.*

The next theorem follows immediately from Theorem 1 and the definition of a negotiation tree.

Theorem 2. *Every constructive negotiation tree is finite.* □

6. NEGOTIATION STRATEGIES

The previous sections provide the basic definitions for modeling negotiation. Given two agents a and b, the negotiation tree can be used to predict whether a negotiation between them will succeed or fail. In practice, agents engaging in a negotiation commonly employ their own strategies. We formalize this notion as follows.

Definition 14 (Strategy). *Given an agent a with the n-KB K, a* negotiation strategy *for a is a function F that maps each proposal ω and negotiation $\vec{h}$ to a proposal that satisfies the following properties:*

- *$F(\omega, \vec{h}) \in \beta(K, \omega)$, and*
- *$\vec{h}, F(\omega, \vec{h})$ is a negotiation.*

Given two agents a and b with the strategies F_a and F_b, respectively, the outcome of a negotiation between them now depends on F_a and F_b. We characterize this in the next definition.

Definition 15 (Negotiation with Strategy). *Given the agents a and b with the strategies F_a and F_b, respectively, a negotiation $\omega_1, \omega_2, \ldots$ between a and b, started by a, is said to be (F_a, F_b)-induced if for every $i > 0$, $\omega_{2i+1} = F_a(\omega_{2i}, \sigma_{2i})$ and $\omega_{2i} = F_b(\omega_{2i-1}, \sigma_{2i-1})$, where $\sigma_j = \omega_1, \ldots, \omega_j$ for $j > 0$.*

Agents are interested in different types of strategies. For example, strategies that guarantee the termination of a negotiation, strategies that do not use disinformation, strategies that guarantee the success of a negotiation, etc. We will discuss some of these next. For two negotiations $\vec{h}$ and $\vec{t}$, we write $\vec{h} \sqsubset \vec{t}$ if $\vec{h}$ is a proper prefix of $\vec{t}$.

DEFINITION 16 (OBSERVANT STRATEGY). *A strategy F of an agent a is* observant *if $F(\omega, \vec{h}) \neq F(\omega, \vec{t})$ for every pair of negotiations $\vec{h}$ and $\vec{t}$ such that $\vec{h} \sqsubset \vec{t}$.*

The observant strategy says an agent does not repeat the same response to the same proposal ω in a negotiation. If at least one agent's strategy is observant then their negotiation will terminate.

PROPOSITION 4. *Given two agents a and b with strategies F_a and F_b, respectively. If either F_a or F_b is observant then every negotiation between a and b is constructive.*

So far, our discussion focused on the development of a general framework for negotiation and does not distinguish the type of information that an agent uses in achieving his/her goals. In practice, an agent might prefer to be honest before he/she starts using disinformation. We will now discuss some strategies that differentiate between honesty and dishonesty. In particular, strategies that guarantee that an agent lies or BS only if s/he has no alternative can be built using the preference relation between belief sets (Def. 3), since such relation favors belief sets without disinformation. For a proposal ω, let $\Sigma(\omega) = \{M \mid M \text{ supports some } \gamma \in \beta(K, \omega)\}$.

DEFINITION 17 (DELIBERATE STRATEGY). *A strategy F of an agent a with the n-KB K is* deliberate *if*

- *$F(\omega, \vec{h})$ is supported by a most preferred belief set of $\Sigma(\omega)$ whenever $\Sigma(\omega) \neq \emptyset$;*
- *$F(\omega, \vec{h}) = \langle \bot, \emptyset, \emptyset \rangle$, otherwise.*

A deliberate strategy, which is also observant, is called a best-practice *strategy.*

A deliberate strategy does not guarantee termination of a negotiation. However, a best-practice strategy does. Furthermore, we can observe that an agent with a best-practice strategy may accept a less preferred outcome of a negotiation even though he/she might obtain a more preferred outcome had he/she used disinformation. Similarly, he/she may sometimes reject a proposal even though this might be negotiable and further negotiation might yield a preferred outcome, had he/she lied or BSed. This can be seen in the negotiation in Example 7: a deliberate seller will respond to b_1 with $s'_1 = \langle \{low\_pr\}, \{student\}, \{product_A, maker_C\} \rangle$ rather than s_1 since he/she has a belief set without disinformation ($quality_A$) that supports s'_1 and belief set without disinformation is preferred to belief set with BS (Def. 3); similarly, a deliberate buyer, on the other hand, will not accept the proposal s_3 since there is no belief set without disinformation supporting its acceptance; the buyer could have respond with the proposal $b'_4 = \langle \{lowest\_pr\}, \{cash\}, \{product_B, quality_B, maker_D\} \rangle$ which is supported by a belief set without disinformation.

Other strategies can be developed to further characterize agents, e.g., an extremely-selfish agent will likely want to put forward one of his/her proposal for his/her best goal regardless of the response from the other agent; a somewhat self-centered agent will likely want to put forward a proposal for the best goal that he/she could be achieved given the information obtained from the other agent; a greedy agent would attempt to satisfy a proposal at the cost of being dishonest; etc. As we have seen, however, such strategies might not guarantee termination of the negotiation if they are not observant. One disadvantage of observant strategies lies in that they require the agent to memorize the full history of the negotiation. We will next discuss a possible way to address this issue.

We start with a more refined definition of responses. Intuitively, a response to a proposal $\langle G, S, R \rangle$, by an agent who has an acceptable behavior and the willingness to complete the negotiation, needs to take into consideration:

- *The assumptions that have been made by the proposer:* the response should identify and make explicit those assumptions that are wrong about him/herself, as far as he/she will not lie/BS on those assumptions—e.g., if the seller assumes that the buyer is a student but the buyer is not, then the buyer should identify this and inform the seller, as far as if he/she does not disguise him/herself as a student;
- *The information that the proposer reveals about him/herself:* the response needs to conform to this information—e.g., if the seller says that s/he does not have the product A, then the the buyer should not assume that $product_A$ is available, even though $product_A$ is a viable assumption in his/her KB.

DEFINITION 18 (CONSCIOUS RESPONSE). *Let a be an agent with n-KB $K_a = \langle P, L, B, H, N^{\prec} \rangle$ and let $\gamma_b = \langle G, S, R \rangle$ be a proposal by b w.r.t. its n-KB K_b. A* conscious response *to γ_b by a is*

(i) *A proposal $\gamma_a = \langle G', S', R' \rangle$ w.r.t. K_a with a support M such that $G \preceq G'$, $R \cap H \subseteq S'$, and $S^{\neg} \cap M \subseteq R'$ where M is the support of γ_a, if γ_b is not rejectable w.r.t. K_a;* **or**

(ii) *A proposal $\gamma_a = \langle G', S', R' \rangle$ w.r.t. K_a with a support M such that $G \not\preceq G'$ and $S^{\neg} \cap M \subseteq R'$ where M is a support of γ_a if γ_b is rejectable w.r.t. K_a;* **or**

(iii) *$\langle \top, \emptyset, \emptyset \rangle$, denoting* acceptance of the proposal*, if γ_b is acceptable w.r.t. K_a;* **or**

(iv) *$\langle \bot, \emptyset, \emptyset \rangle$, denoting* rejection of the proposal.

Intuitively, given that the proposal $\langle G, S, R \rangle$ is acceptable to an agent, he/she could accept it (case (*iii*)) or attempt to negotiate for some better options (case (*i*)). If the proposal is negotiable, he/she could continue and attempt to get a better option (case (*i*)). If the proposal is rejectable, he/she could try to negotiate for something that is not as good as the current goal (case (*ii*)). In any cases, the agent can stop with a rejection (case (*iv*)). An agent should generate a new proposal whose goal depends on the goal of the given proposal, whose assumptions cover the conditions stated in the original proposal ($R \cap H \subseteq S'$), and whose conditions identify all incorrect assumptions made in the original proposal ($S^{\neg} \cap M \subseteq R'$). An agent, who decides to consider preferable proposals, would require that the support for the new proposal must be preferred to any support for accepting γ_b.

EXAMPLE 8. *Let us consider the proposal $\gamma_1 = \langle \{low\_pr\}, \{cash\}, \{product_A, quality_A, maker_C\} \rangle$ (Read as "[s]: I can sell you the $product_A$, made by $maker_C$, and has good quality for low_pr if you pay in cash"). We can check that γ_1 is acceptable w.r.t. K_b. As such, the buyer could accept this proposal. However, the buyer could respond with the proposal $\gamma'_1 = \langle \{lowest\_pr\}, \{product_A, quality_A, maker_C\}, \{cash\} \rangle$. (Read as "[b]: Can I get the $lowest\_pr$?")*

Let $\gamma_2 = \langle \{low\_pr\}, \{product_A, maker_C\}, \emptyset \rangle$ (Read as "[b]: Can I have $product_A$ from $maker_C$ for low_pr?"). It is easy to see that γ_2 is not acceptable but negotiable w.r.t. K_s since any rule for concluding low_pr requires additional assumption (e.g., student or age ≥ 65). The seller can respond with the proposal $\gamma'_2 = \langle \{low\_pr\}, \{student\}, \{product_A, maker_C\} \rangle$. (Read as "[s]: Yes, if you are a student.").

Let $\gamma_3 = \langle \{high\_pr\}, \emptyset, \{product_A, quality_A, maker_C\} \rangle$ (Read as "[s]: I can sell you the $product_A$, made by $maker_C$, and has good quality for $high\_pr$"). γ_3 is rejectable w.r.t. K_b because no rule in K_b can be used to derive $high\_pr$. As such, the buyer could reject this proposal. He could also weaken the goal by responding with $\gamma'_3 = \langle \{low\_pr\}, \{product_A, quality_A, maker_C\}, \emptyset \rangle$. (Read as "[b]: Can I get it for low_pr?")

DEFINITION 19 (ADAPTIVE AGENT). *An agent a with the n-KB $K = \langle P, L, B, H, N^{\prec} \rangle$ and a strategy F is said to be* adaptive *if for every proposal $\langle G, S, R \rangle$*

- $F(\langle G, S, R \rangle, \vec{h}) = \langle G', S', R' \rangle$ *such that*
  - $\langle G', S', R' \rangle$ *is a conscious response to* $\langle G, S, R \rangle$*;*
  - *if* $\langle G, S, R \rangle$ *is acceptable w.r.t.* K *then* G' *is preferred to* G *or* $\langle G', S', R' \rangle = \langle \top, \emptyset, \emptyset \rangle$.
- a *changes his/her n-KB to* $K' = \langle P \cup^r (R \cap H), L, B, H, N^{\prec} \rangle$ *after his/her response to a proposal.*

Intuitively, an agent is adaptive if it imports the information received during the negotiation into his/her n-KB and keeps this information for the next round of negotiation. Furthermore, an adaptive agent prefers to accept a proposal if a better outcome cannot be achieved. We can show that if both agents are adaptive then their negotiations will terminate.

PROPOSITION 5. *Every negotiation between adaptive agents will terminate.*

7. RELATED WORK

Logic programming and abductive logic programming have been used to formulate negotiation by many researchers, e.g., [2, 11, 14, 12, 15]. Our work differs from these proposals in that we introduce disinformation, while none of those proposals do.

Our work has similarities to [2], in that the notion of belief set is used as a means for exchanges between agents. However, the main goal in [2] is to coordinate belief sets of two agents. The framework does not have a mechanism for generating new proposals. We use abductive programs and specify a way for computing new proposals as part of the response to a given one.

Our work is also similar to [11, 14], in that an abductive logic programming based framework is used to model negotiation. In our framework, the assumptions can be used in conjunction with predefined strategies, represented in the abductive rules, and disinformation to generate proposals, whereas [11, 14] use only assumptions. The system in [12] uses induction to construct proposals but does not consider preferences, while our approach does not use induction and considers preferences.

The proposal in this paper can be seen as an extension of [15], with the introduction of disinformation. Also, we use abductive logic programming instead of logic programming with consistency restoring rules (CR-Prolog). Unlike CR-Prolog, abductive programs do not require the minimality of assumptions used in the generation of belief sets. Furthermore, our formalization does not distinguish between a proposal or an extended proposal, making it simpler than what presented in [15].

Our work is in the same spirit as the approaches to argumentation-based negotiation (ABN) [1, 8, 10], in that it considers explanations as a part of a proposal/response. The main difference between our work and ABN lies in our use of abductive logic programs, a non-monotonic logic, while ABN's logic is monotonic. Our framework does not compute explanations for accepting/rejecting a proposal in advance as in [1], and it allows negotiators to non-monotonically modify their beliefs using incoming information. [8] introduces priorities over arguments and uses abduction to seek conditions to support arguments. But they do not integrate abduction and preferences as done in this paper.

There have been a number of proposals to formalize negotiation with multiple issues or incomplete information (e.g., [4, 5]). The key difference between these approaches and ours is that they rely on the use of utility functions and deadlines in the construction a counter-offer (or a response). Our approach does not consider deadlines. It provides agents with a way to construct their responses which, together with their strategies, can take into consideration the agents' preferences and disinformation. Neither [4] nor [5] consider disinformation.

Note also that negotiation with disinformation (deception) has been considered in [16], although this work focuses on multi-agents in a dynamic environment, where agents act and interact to achieve their individual or cooperative goals. The issue has also been discussed in the context of game theory (e.g., [3]). These approaches differ from most of the works we have discussed so far, including our own.

Finally, let us remark that [13] introduced preference over abducibles to specify preferred explanations. They realize it in the context of prioritized logic programming, where priorities are specified outside of a program. We use abducible rules and preferences for specifying negotiation strategies and reasoning with disinformation, which is completely new in this paper.

8. CONCLUSION

In this paper, we extend abductive logic programming with preferences and disinformation (ALD-programs) and use ALD-programs to formalize negotiation by defining the basic concepts of negotiation using the belief sets of ALD-programs. The main features of our formalism are that it (*a*) includes support (as explanation) in a proposal/response; (*b*) can deal with incomplete information, preference, and changes in goals; (*c*) can take into consideration the disinformation that an agent is willing to use; (*d*) allows an agent to compute proposals/responses (and their support) on a case-by-case basis; and (*e*) allows an agent to develop his/her own strategy and employ it in the negotiation. In this work, we focus on the development of the basic notions for modeling negotiation and the termination of negotiations. The proposed framework is general and it allows us to address several important issues, such as how to achieve an optimal result, how to obtain a shortest (in terms of rounds) negotiation, etc. The investigation of these issues and the development of a system for automated negotiation are our most immediate future goals.

9. REFERENCES

[1] L. Amgoud et al. A unified and general framework for argumentation-based negotiation. In *Proc. 6th AAMAS*, pages 1018–1025. ACM Press, 2006.

[2] W. Chen, M. Zhang, and N. Foo. Repeated negotiation of logic programs. In *Proc. 7th NRAC.* , 2006.

[3] D. Ettinger, P. Jehiel. A Theory of Deception. *American Econ. Journal: Macroeconomics*, 2(1):1–20, 2010.

[4] S. S. Fatima et al. An agenda-based framework for multi-issue negotiation. *Artif. Intell.*, 152(1):1–45, 2004.

[5] S. S. Fatima, M. Wooldridge, and N. R. Jennings. Bargaining with incomplete information. *Ann. Math. Artif. Intell.*, 44(3):207–232, 2005.

[6] H. G. Frankfurt. *On Bullshit*. Princeton Univ. Press, 2005.

[7] M. Gelfond and V. Lifschitz. Logic programs with classical negation. *ICLP*, pages 579–597, 1990.

[8] A. Kakas and P. Moraitis. Adaptive agent negotiation via argumentation. In *AAMAS*, pages 384–391, 2006. ACM.

[9] J. E. Mahon. Two definitions of lying. *J. Applied Philosophy*, 22(2):211–230, 2008.

[10] I. Rahwan, S. D. Ramchurn, N. R. Jennings, P. Mcburney, S. Parsons, and L. Sonenberg. Argumentation-based negotiation. *The Knowledge Engineering Review*, 18:343–375, 2003.

[11] F. Sadri, F. Toni, and P. Torroni. An abductive logic programming architecture for negotiating agents. In *JELIA*, pages 419–431. Springer Verlag, 2002.

[12] C. Sakama. Inductive Negotiation in Answer Set Programming. In *DALT'08*, pages 143–160. Springer Verlag, 2008.

[13] C. Sakama and K. Inoue. Prioritized Logic Programming and its Application to Commonsense Reasoning. *Artificial Intelligence*, 123(1-2):185–222, 2000.

[14] C. Sakama and K. Inoue. Negotiation by abduction and relaxation. In *Proc. 6th AAMAS*, pages 1018–1025. ACM Press, 2007.

[15] T. C. Son and C. Sakama. Negotiation using logic programming with consistency restoring rules. In *IJCAI'09*, pages 930–935, 2009.

[16] G. Zlotkin and J. S. Rosenschein. Incomplete information and deception in multi-agent negotiation. In *IJCAI'91*, pages 225–231, San Francisco, CA, USA, 1991. Morgan Kaufmann Publishers Inc.