

Dynamically Ordered Probabilistic Choice Logic Programming

Marina De Vos* and Dirk Vermeir

Dept. of Computer Science
Free University of Brussels, VUB
Pleinlaan 2, Brussels 1050, Belgium
{marinadv,dvermeir}@vub.ac.be
<http://tinf2.vub.ac.be>

Abstract. We present a framework for decision making under uncertainty where the priorities of the alternatives can depend on the situation at hand. We design a logic-programming language, DOP-CLP, that allows the user to specify the static priority of each rule and to declare, dynamically, all the alternatives for the decisions that have to be made. In this paper we focus on a semantics that reflects all possible situations in which the decision maker takes the most rational, possibly probabilistic, decisions given the circumstances. Our model theory, which is a generalization of classical logic-programming model theory, captures uncertainty at the level of total Herbrand interpretations. We also demonstrate that DOP-CLPs can be used to formulate game theoretic concepts.

1 Introduction

Reasoning with priorities and reasoning under uncertainty play an important role in human behavior and knowledge representation. Recent research has been focused on either priorities, [14, 8, 6]¹, or uncertainty, [10, 9, 12, 1] and many others.

We present a framework for decision making under uncertainty where the priorities of the alternatives depend on the different (**probabilistic**) situations. This way we obtain a semantics that reflects all possible situations in which the most rational (probabilistic) decisions are made, given the circumstances. The basic idea for the framework, a logic programming language called “Dynamically Ordered Probabilistic Choice Logic Programming” or DOP-CLP for short, incorporates the intuition behind both ordered logic programs ([8]) and choice logic programs ([4, 5]). The former models the ability of humans to reason with defaults² in a logic programming context, using a static ordering of the rules in the program. This works well, as long as probabilities stay out of the picture, but once they are present something extra is needed to express order. Take the famous “Tweety example” for instance: if you are sure that Tweety is indeed a penguin, you should derive that she cannot fly. But suppose you believe for only 30% that the bird you are holding is indeed a penguin. Is it then sensible to derive that she is

* Wishes to thank the FWO for its support.

¹ [8] uses the word order instead of priority.

² Intuitively, something is true by default unless there is evidence to the contrary.

a non-flying bird?

By also taking into account the probabilities of the antecedents of the rules, in addition to their static order, we can overcome this problem. This leads to a **dynamic ordering** of rules, where the priority of a rule depends on the actual situation.

We aim for a decision making framework that allows decisions to have possibly more than two alternatives, as in the case of ordered logic³. To accomplish this, we turn to a variant of **Choice Logic Programs**[4, 5], in which the possible alternatives for the decisions are described by choice rules. This approach has two nice side effects. First of all, there is not necessarily a partition of the Herbrand base: atoms can belong to more than one decision or to no decision at all. In the former case, there is a probability distribution over the various alternatives. In the latter case, an atom is either true or false, as in classical logic programming. The second advantage of our approach is that we allow a “lazy” evaluation of the alternatives which become active only when they are present in the head of an applicable choice rule.

An interesting application of DOP-CLP is Game Theory. We provide a transformation from strategic games to DOP-CLPs such that a one-to-one mapping is established between the mixed strategy Nash Equilibria of the game and the stable models of its corresponding DOP-CLP.

2 Dynamically Ordered Probabilistic Choice Logic Programs

In this paper, we identify a program with its grounded version, i.e. the set of all ground instances of its clauses. In addition we do not allow function symbols (i.e. we stick to datalog) so the number of literals is finite.

Definition 1. *A Dynamically Ordered Probabilistic Choice Logic Program, or DOP-CLP for short, is a finite set of rules of the form $A \leftarrow^p B$, where A and B are (possibly empty) sets of atoms and $p \in \mathbb{R}^+$. For a rule $r \in P$, the set A is called the **head**, denoted H_r , while the set B is called the **body** of the rule r , denoted B_r . The $p \in \mathbb{R}^+$ denotes the **priority** this rule. A rule without a priority number has an infinite priority. We will denote the priority of rule r as $\rho(r)$. The **Herbrand base** of P , denoted \mathcal{B}_P , is the set of all atoms appearing in P .*

A rule $H \leftarrow^p B$ can be read as:

”The occurrence of the events in B forces a probabilistic decision between the elements $h \in H$ and supports each h with a priority p .”

This means that rules with more than one head atom express no preference among the different alternatives they create.

The priority of a rule r , $\rho(r)$, indicates the maximal impact of the situation, described by the events of the body, on the preference of the head atom over the other alternatives. The dynamic priority of a rule, which we will define later, adjusts this impact according to the probability of the situation described in the body. By combining all the dynamic priorities of rules sharing a common head atom, we obtain an evaluation of the total impact on that atom which can then be used for comparison with other alternatives.

³ In ordered logic the two alternatives are represented using negation.

Example 1 (Jakuza). A young member of the Jakuza, the Japanese Mafia, is faced with his first serious job. It will be his duty to control a new victim. Since it is his first job, he goes to his oyabun, head of his clan and mentor, for advice. He tells him that the Jakuza has three methods for controlling its victims: blackmail, intimidation and bribing. The victim can either give in immediately or can put up a stand. In the latter case, she does this by just ignoring the threats of the organization or she can threaten to go to the police. The oyabun is only able to give some information about previous encounters, which still needs to be interpreted in the current situation. So he starts telling about his previous successes. “Every time when I knew that the victim was willing to give in, I resorted to intimidation as this is the easiest technique and each time it worked perfectly. In case we would know that the victim was planning to stand up to us, we looked in to the possibility of bribing. Nine out of ten times, we were successful when we just offered enough money. If you are sure that the victim will run to the police from the moment that you approach him, you have to try to bribe her. Unfortunately this technique worked only 4 times out of 10. When your victim tries to ignore you, you should find something to blackmail her with. However, finding something interesting is not that easy as reflected by a success rate of 3 out of 10 times. So now it is up to you to make a good estimation of the victim’s reaction in order to succeed with your assignment.”

All this information can easily be represented as the next DOP-CLP:

$$\begin{array}{r}
 \text{jakuza} \leftarrow \\
 \text{blackmail} \oplus \text{intimidate} \oplus \text{bribe} \leftarrow^0 \text{jakuza} \\
 \text{stand} - \text{up} \oplus \text{give} - \text{in} \leftarrow^0 \\
 \text{ignore} \oplus \text{police} \leftarrow^0 \text{stand} - \text{up} \\
 \text{intimidate} \leftarrow^{10} \text{give} - \text{in} \\
 \text{bribe} \leftarrow^4 \text{police} \\
 \text{blackmail} \leftarrow^3 \text{ignore} \\
 \text{enough} \oplus \text{more} \leftarrow^0 \text{stand} - \text{up} \\
 \text{bribe} \leftarrow^9 \text{enough}
 \end{array}$$

An interpretation assigns a probability distribution over every state of affairs⁴.

Definition 2. Let P be a DOP-CLP. A (probabilistic) **interpretation** is a probability distribution $\mathbf{I} : 2^{B_P} \rightarrow [0..1]$.

In our examples, we will mention only the probabilities of those states that have a positive probability in the interpretation.

⁴ Each state corresponds to a total interpretation of the choice logic program obtained from P by omitting the priorities. Because we are working with total interpretations we only have to mention the positive part of the interpretation.

Example 2. Recall the Jakuza program of Example 1. The following functions **I**, **J** and **K** are interpretations for this program⁵:

$$\begin{array}{lll}
\mathbf{I}(\{j, br, s, ig, e\}) = \frac{1}{4} & \mathbf{J}(\{j, in, p, e, s\}) = \frac{3}{20} & \mathbf{K}(\{j, bl, s, ig, e\}) = \frac{30}{441} \\
\mathbf{I}(\{j, br, s, ig, m\}) = \frac{5}{12} & \mathbf{J}(\{j, in, p, m, g\}) = \frac{7}{10} & \mathbf{K}(\{j, bl, s, ig, m\}) = \frac{240}{441} \\
\mathbf{I}(\{j, br, s, ig\}) = \frac{1}{6} & \mathbf{J}(\{j, in, p, m, s\}) = \frac{2}{20} & \mathbf{K}(\{j, bl, s, p, e\}) = \frac{12}{441} \\
\mathbf{I}(\{j, br, s, p\}) = \frac{1}{24} & \mathbf{J}(\{j, br, p, m, s\}) = \frac{1}{20} & \mathbf{K}(\{j, bl, s, p, m\}) = \frac{96}{441} \\
\mathbf{I}(\{j, br, s, m\}) = \frac{1}{8} & & \mathbf{K}(\{j, bl, g, ig, e\}) = \frac{5}{441} \\
& & \mathbf{K}(\{j, bl, g, ig, m\}) = \frac{40}{441} \\
& & \mathbf{K}(\{j, bl, g, p, e\}) = \frac{2}{441} \\
& & \mathbf{K}(\{j, bl, g, p, m\}) = \frac{16}{441}
\end{array}$$

Given an interpretation, we can compute the probability of a set of atoms, as the sum of the probabilities assigned to those situations which contain this set of atoms.

Definition 3. Let **I** be an interpretation for a DOP-CLP P . The **probability of set** $A \subseteq \mathcal{B}_P$, denoted $\vartheta_{\mathbf{I}}(A)$ ⁶, is $\vartheta_{\mathbf{I}}(A) = \sum_{A \subseteq Y \subseteq \mathcal{B}_P} \mathbf{I}(Y)$.

In choice logic programs, the basis of DOP-CLP, a rule is applicable when the body is true, and is applied when both the body and a single head atom are true. This situation becomes more tricky when probabilities come into play. Applicability is achieved when the body has a non-zero probability. In order for a rule to be applied it must be applicable. In addition, we demand that at least one head element has a chance of happening and that no two of them can happen simultaneously.

Definition 4. Let **I** be an interpretation for a DOP-CLP P .

1. A rule $r \in P$ is called **applicable** iff $\vartheta_{\mathbf{I}}(B_r) > 0$.
2. An applicable rule $r \in P$ is **applied** iff $\exists a \in H_r \cdot \vartheta_{\mathbf{I}}(a) > 0$ and $\forall S \in 2^{H_r}$ with $|S| > 1 : \vartheta_{\mathbf{I}}(S) = 0$.

We have been referring to alternatives of decisions without actually defining them. Two atoms are alternatives if they appear together in the head of an applicable choice rule. Alternatives are thus dynamic, since the applicability of the rules depends on the interpretation.

Definition 5. Let **I** be an interpretation for a DOP-CLP P .

- Two atoms $a, b \in \mathcal{B}_P$ are **alternatives** wrt **I** iff \exists applicable $r \in P \cdot \{a, b\} \subseteq H_r$.
- The set of all alternatives of an atom $a \in \mathcal{B}_P$ wrt **I** is denoted $\Omega_{\mathbf{I}}(a)$ ⁷.
- A set $D \subseteq \mathcal{B}_P$ is a **maximal alternative set** wrt **I** iff $\forall a, b \in D \cdot a$ and b are alternatives and $\forall c \notin D \cdot \exists a \in D \cdot a$ and c are no alternatives.
- $\Delta_{\mathbf{I}}$ is the set of all maximal alternative sets wrt **I** .
- An atom $a \in \mathcal{B}_P$ is called **single** iff $\Omega_{\mathbf{I}}(a) = a$.

⁵ For brevity, the names of the atoms are abbreviated.

⁶ When the set A contains just one element a we omit the brackets and write $\vartheta_{\mathbf{I}}(a)$.

⁷ Notice that $a \in \Omega_{\mathbf{I}}(a)$. The set $\Omega_{\mathbf{I}}(a) \setminus \{a\}$, denoted $\Omega_{\mathbf{I}}^-(a)$, is the set of all alternatives of a excluding itself.

A naive approach to defining a probability distribution is to insist that the sum of probabilities of the multiple elements in the head of a choice rule must be one. This approach fails in situations of the following kind:

$$\begin{array}{l} a \oplus b \oplus c \leftarrow^0 \dots \\ a \oplus b \leftarrow^4 \dots \\ \dots \end{array}$$

In this situation, the atom c would not stand a chance of obtaining a positive probability, although this might be the most favorable alternative.

To overcome this problem, we introduced maximum alternative sets. They group all the atoms that have an alternative relation with each other. It is those sets that will be used for the probability distribution. In the next definition we call an interpretation total if it defines a probability distribution in which the probabilities of the elements of any maximal alternative set add up to one. Furthermore, for all decisions that need to be made, an alternative is selected for every possible outcome.

Definition 6. A interpretation \mathbf{I} for a DOP-CLP P is called **total** iff $\forall D \in \Delta_{\mathbf{I}}$,

- $\sum_{a \in D} \vartheta_{\mathbf{I}}(a) = 1$, and
- $\forall A \subseteq \mathcal{B}_P$ such that $\vartheta_{\mathbf{I}}(A) > 0 \cdot |A \cap D| = 1$.

Example 3. Reconsider the Jakuza program of Example 1 and the interpretations of Example 2. The interpretation \mathbf{I} is not total. Indeed, consider the maximal alternative set $\{ig, p\}$. We have $\{ig, p\} \cap \{br, j, s, m\} = \emptyset$, while $\mathbf{I}(\{br, j, s, m\}) > 0$, and $\vartheta_{\mathbf{I}}(ig) + \vartheta_{\mathbf{I}}(p) = 5/6 + 1/12 \neq 1$. The interpretations \mathbf{J} and \mathbf{K} are total.

As we mentioned earlier, the dynamic priority of a rule adjusts the (static) preference of the rule to the probability that this situation might actually occur. It does this by giving the maximal contribution of the body atoms to the general preference of the head atoms. The dynamic priority of an atom is obtained by taking into account every real contribution of any situation that provides a choice for this atom.

Definition 7. Let \mathbf{I} be an interpretation for a DOP-CLP P . The **dynamic priority** of a rule $r \in P$, denoted $\varrho_{\mathbf{I}}(r)$, equals $\varrho_{\mathbf{I}}(r) = \rho(r) * \vartheta_{\mathbf{I}}(B_r)$.

The **dynamic priority** of an atom $a \in \mathcal{B}_P$, denoted $\varrho_{\mathbf{I}}(a)$, is $\varrho_{\mathbf{I}}(a) = \sum_{r \in P: a \in H_r} \varrho_{\mathbf{I}}(r)$.

The dynamic priority will be used to determine which alternatives of a decision are eligible candidates and which ones are not. An atom is said to be blocked if there exists an alternative which has higher dynamic priority. Preferred atoms are those that block every other alternative. The competitors of an atom are those alternatives which are not blocked by this atom. Their dynamic priority is thus at least as high as that of the atom.

Definition 8. Let \mathbf{I} be an interpretation for a DOP-CLP P . An atom $a \in \mathcal{B}_P$ is **blocks** by $b \in \Omega_{\mathbf{I}}^{-}(a)$ w.r.t. \mathbf{I} iff $\varrho_{\mathbf{I}}(b) > \varrho_{\mathbf{I}}(a)$.

An atom $a \in \mathcal{B}_P$ is called **preferred** in \mathbf{I} iff $\forall b \in \Omega_{\mathbf{I}}^{-}(a) \cdot \varrho_{\mathbf{I}}(a) > \varrho_{\mathbf{I}}(b)$.

The atom a is a **competitor** of the atom $b \in \Omega_{\mathbf{I}}^{-}(a)$ w.r.t. \mathbf{I} if b does not block a w.r.t. \mathbf{I} .

In standard logic programming an interpretation is a model if every rule is either not applicable or applied. When priorities are involved, in order for an interpretation to become a model, it must be possible to assign a zero-probability to atoms which have a more favorable alternative with non-zero probability.

Definition 9. Let \mathbf{I} be an interpretation for the DOP-CLP P . \mathbf{I} is a **model** for P iff $\forall r \in P$:

- $\vartheta_{\mathbf{I}}(B_r) = 0$, i.e. r is not applicable, or
- r is applied, or
- $\forall a \in H_r \cdot \exists b$ competitor of a w.r.t. $\mathbf{I} \cdot \vartheta_{\mathbf{I}}(b) > 0$.

Example 4. Consider again the Jakuza program of Example 1 and its interpretations of Example 2. The interpretation \mathbf{I} is not a model, since the rule $blackmail \leftarrow^{10} ignore$ does not satisfy any of the above conditions. This rule is applicable, since $\vartheta_{\mathbf{I}}(ignore) = 5/6$; it is not applied, since $\vartheta_{\mathbf{I}}(blackmail) = 0$; it does not have any competitors since $\varrho_{\mathbf{I}}(blackmail) = 5/2$ while $\varrho_{\mathbf{I}}(intimidate) = 0$ and $\varrho_{\mathbf{I}}(bribe) = 29/12$. The interpretations \mathbf{J} and \mathbf{K} are both models.

Proposition 1. Let P be a DOP-CLP and let \mathbf{I} be a model for it. If $a \in \mathcal{B}_P$ is a preferred atom then $\exists r \in P : a \in H_r \cdot r$ is applied .

In some cases, atoms receive a probability which they actually do not deserve. This happens when there is some better qualified alternative (i.e., an alternative that has a higher dynamic priority) that should obtain this probability. Such atoms are called assumptions, since they were just "assumed" to have a chance of happening.

Definition 10. Let \mathbf{I} be an interpretation for a DOP-CLP P . An atom $a \in \mathcal{B}_P$ is called an **assumption** w.r.t. \mathbf{I} iff $\vartheta_{\mathbf{I}}(a) > 0$ when either a is blocked or a is single and $\varrho_{\mathbf{I}}(a) = 0$. \mathbf{I} is **assumption-free** iff it contains no assumptions.

Example 5. Consider once more the Jakuza program of Example 1 and its interpretations in Example 2. The interpretation \mathbf{J} is not assumption-free, as $\vartheta_{\mathbf{J}}(intimidate) > 0$ and the alternative $bribe$ blocks $intimidate$, since $\varrho_{\mathbf{J}}(bribe) = 4 + 27/20 > 3.5 = \varrho_{\mathbf{J}}(intimidate)$. Intuitively, because bribing is more successful than intimidation, intimidation should not be considered at all. The interpretation \mathbf{K} is assumption-free.

Proposition 2. Let P be a DOP-CLP and let \mathbf{I} be a total assumption-free interpretation for P . If $a \in \mathcal{B}_P$ is a preferred atom then $\exists r \in P : a \in H_r \cdot r$ is applied and $\vartheta_{\mathbf{I}}(a) = 1$.

Interpretations evaluate the likelihood of every possible outcome, by assigning a probability distribution to every situation. These probabilities are influenced by the atoms which are present in each situation. In order to quantify this influence one must know whether the events which occur in any such interpretation are independent of each other. An interpretation which assumes that there is no inter-dependence between atoms, is said to be "independent", as follows:

Definition 11. Let \mathbf{I} be an interpretation for a DOP-CLP P . We say that \mathbf{I} is **independent** iff $\forall A \subseteq \mathcal{B}_P$:

- $\vartheta_{\mathbf{I}}(A) = 0$, or
- $\forall D \in \Delta_{\mathbf{I}} \text{ s.t. } |D \cap A| \leq 1 \cdot \vartheta_{\mathbf{I}}(A) = \prod_{a \in A} \vartheta_{\mathbf{I}}(a)$.

Example 6. Consider the interpretations \mathbf{J} and \mathbf{K} of Example 2. The interpretation \mathbf{J} is not independent as $\vartheta_{\mathbf{J}}(\{y, in, p, e, s\}) = 3/20 \neq \vartheta_{\mathbf{J}}(y) * \vartheta_{\mathbf{J}}(in) * \vartheta_{\mathbf{J}}(p) * \vartheta_{\mathbf{J}}(e) * \vartheta_{\mathbf{J}}(s) = 1 * 19/20 * 1 * 3/20 * 8/20$. The interpretation \mathbf{K} is independent.

Definition 12. Let P be a DOP-CLP. A total independent assumption-free model is said to be **stable**. A stable model is **crisp** if it assigns probability one to a single subset of the Herbrand base.

Example 7. For the last time we return to the Jakuza example and its three interpretations \mathbf{I}, \mathbf{J} and \mathbf{K} from Example 2. Combining the results from Examples 3, 4, 5 and 6, we can conclude that \mathbf{K} is the only stable model of the three.

A stable model for the Jakuza example represents a rational choice where the probability of the action is consistent with the estimates on the victim's reactions. In general, stable models reveal all possible situations in which the decisions are made rationally, considering the likelihood of the events that would force such decisions.

3 An Application of DOP-CLPs: Equilibria of Strategic Games

3.1 Strategic Games

A strategic game models a situation where several agents (called players) independently choose which action they should take, out of a limited set of possibilities. The result of the actions is determined by the combined effect of the choices made by each of the players. Players have a preference for certain outcomes over others. Often, preferences are modeled indirectly using the concept of *payoff* where players are assumed to prefer outcomes where they receive a higher payoff.

Example 8 (Bach or Stravinsky). Two people wish to go out together to a music concert. They have a choice between a Bach or Stravinsky concert. Their main concern is to be together, but one person prefers Bach and the other prefers Stravinsky. If they both choose Bach then the person who preferred Bach gets a payoff of 2 and the other a payoff of 1. If both go for Stravinsky, it is the other way around. If they pick different concerts, they both get a payoff of zero.

The game is represented in Fig. 1. One player's actions are identified with the rows and the other player's with the columns. The two numbers in the box formed by row r and column c are the players' payoffs when the row player chooses r and the column player chooses c . The first of the two numbers is the payoff of the row player.

Definition 13 ([11]). A **strategic game** is a tuple $\langle N, (A_i), (u_i) \rangle$ where

- N is a finite set of **players**;
 - for each player $i \in N$, A_i is a nonempty set of **actions** that are available to her⁸
- and,

⁸ We assume that $A_i \cap A_j = \emptyset$ whenever $i \neq j$.

	<i>Bach</i>	<i>Stravinsky</i>
<i>Bach</i>	2, 1	0, 0
<i>Stravinsky</i>	0, 0	1, 2

Fig. 1. Bach or Stravinsky (BoS)

	<i>Head</i>	<i>Tail</i>
<i>Head</i>	1, 0	0, 1
<i>Tail</i>	0, 1	1, 0

Fig. 2. Matching Pennies (Example 9).

– for each player $i \in N$, $u_i : A = \times_{j \in N} A_j \rightarrow \mathbb{R}$ is a **utility function** which describes the players’ preferences.

An element $\mathbf{a} \in A$ is called a **profile**. For a profile \mathbf{a} we use \mathbf{a}_i to denote the component of \mathbf{a} in A_i . For any player $i \in N$, we define $A_{-i} = \times_{j \in N \setminus \{i\}} A_j$. Similarly, an element of A_{-i} will often be denoted as \mathbf{a}_{-i} . For any $\mathbf{a}_{-i} \in A_{-i}$ and $a_i \in A_i$, (\mathbf{a}_{-i}, a_i) is the profile $\mathbf{a}' \in A$ in which $\mathbf{a}'_i = a_i$ and $\mathbf{a}'_j = \mathbf{a}_j$ for all $j \neq i$.

A game $\langle N, (A_i), (u_i) \rangle$ is played when each player $i \in N$ selects a single action from the set A_i of actions available to her. Since players are thought to be rational, it is assumed that a player will select an action that, to the best of her knowledge, leads to a “preferred” profile. Of course, this is limited by the fact that she must decide without knowing what the other players will choose.

The notion of Nash equilibrium shows that, in many cases, it is possible to limit the possible outcomes (profiles) of the game.

Definition 14 ([11]). A **Nash equilibrium** of a strategic game $\langle N, (A_i), (u_i) \rangle$ is a profile \mathbf{a}^* satisfying $\forall a_i \in A_i \cdot (\mathbf{a}_{-i}^*, \mathbf{a}_i^*) \geq_i (\mathbf{a}_{-i}^*, a_i)$.

Intuitively, a profile \mathbf{a}^* is a Nash equilibrium if no player can unilaterally improve upon his choice. This means that, given the other players’ actions \mathbf{a}_{-i}^* , \mathbf{a}_i^* is the best player i can do⁹.

Although the notion of Nash equilibrium is intuitive, it does not provide a solution to every game. Take for example the Matching Pennies game.

Example 9 (Matching Pennies). Two people toss a coin. Each of them has to choose head or tail. If the choices differ, person 1 pays person 2 a Euro; if they are the same, person 2 pays person 1 a Euro. Each person cares only about the amount of money that she receives. The game modeling this situation is depicted in Fig. 2. This game does not have a Nash equilibrium.

The intuitive strategy to choose head or tail with 50% frequency each (yielding a profit in 25% of the cases) corresponds with a mixed strategy Nash equilibrium where agents assign a probability distribution over their actions.

⁹ Note that the actions of the other players are not known to i .

Definition 15 ([11]). The *mixed extension* of the strategic game $\langle N, (A_i), (u_i) \rangle$ is the strategic game $\langle N, (\Delta(A_i)), (U_i) \rangle$ in which $\Delta(A_i)$ is the set of probability distributions over A_i , and $U_i : \times_{j \in N} \Delta(A_j) \rightarrow \mathbb{R}$ assigns to each $\alpha \in \times_{j \in N} \Delta(A_j)$ the expected value under u_i of the lottery over A that is induced by α (so that $U_i(\alpha) = \sum_{\mathbf{a} \in A} (\prod_{j \in N} \alpha_j(\mathbf{a})) u_i(\mathbf{a})$ if A is finite).

Note that $U_i(\alpha) = \sum_{a_i \in A_i} \alpha_i(a_i) U_i(\alpha_{-i}, e(a_i))$, for any mixed strategy profile α , where $e(a_i)$ is the degenerate mixed strategy of player i that attaches probability one to $a_i \in A_i$. This because we are working with finite sets of actions (e.g. A_i).

Definition 16 (Mixed Strategy Nash Equilibrium). A *mixed strategy Nash equilibrium* of a strategic game is a Nash equilibrium of its mixed extension.

Example 10. Although the matching pennies game (Example 9) does not have a Nash equilibrium, it has the single mixed strategy Nash equilibrium $\{\{Head : 1/2, Tail : 1/2\}, \{Head : 1/2, Tail : 1/2\}\}$, which corresponds to how humans would reason. Apart from its two Nash equilibria, the Bach and Stravinsky game (Example 8) also has the extra mixed strategy Nash equilibrium $\{\{Bach : 2/3, Stravinsky : 1/3\}, \{Bach : 1/3, Stravinsky : 2/3\}\}$.

Each strategic game has at least one mixed strategy Nash equilibrium. Furthermore, each Nash equilibrium is also a mixed strategy Nash equilibrium and every crisp mixed strategy Nash equilibrium (where all the probabilities are either 0 or 1) responds to a Nash equilibrium.

3.2 Transforming Strategic Games to DOP-CLPs

In this subsection, we combine propose an intuitive transformation from strategic games to DOP-CLPs such that the stable models of the former correspond with the mixed strategy Nash equilibria of the latter.

Definition 17. Let $\langle N, (A_i), (u_i) \rangle$ be a strategic game. The corresponding DOP-CLP P equals $P = \{A_i \leftarrow^0 \mid \forall i \in N\} \cup \{\mathbf{a}_i \leftarrow^{u_i(\mathbf{a})} \mathbf{a}_{-i} \mid \mathbf{a} \in A, \forall i \in N\}$.

The corresponding DOP-CLP contains two types of rules. First, there are the real choice rules which represent, for each player, the actions she can choose from. The zero priority assures that the choice itself does not contribute to the decision making process. Rules of the second type represent all the decisions a player can make (the heads) according to the situations that the other players can create (the bodies). A rule's priority corresponds with the payoff that the deciding player would receive for the pure strategy profile corresponding to the head and body of the rule.

Example 11. The Bach and Stravinsky game (Example 8) can be mapped to the DOP-CLP P :

$$\begin{array}{ll} b_1 \oplus s_1 \leftarrow^0 & b_2 \oplus s_2 \leftarrow^0 \\ b_1 \leftarrow^2 b_2 & b_2 \leftarrow^1 b_1 \\ b_1 \leftarrow^0 s_2 & b_2 \leftarrow^0 s_1 \\ s_1 \leftarrow^0 b_2 & s_2 \leftarrow^0 b_1 \\ s_1 \leftarrow^1 s_2 & s_2 \leftarrow^2 s_1 \end{array}$$

This program has three stable models:

$$\begin{array}{lll}
\mathbf{I}_1(b_1, b_2) = 2/9 & \mathbf{I}_2(b_1, b_2) = 1 & \mathbf{I}_3(b_1, b_2) = 0 \\
\mathbf{I}_1(b_1, s_2) = 4/9 & \mathbf{I}_2(b_1, s_2) = 0 & \mathbf{I}_3(b_1, s_2) = 0 \\
\mathbf{I}_1(s_1, b_2) = 1/9 & \mathbf{I}_2(s_1, b_2) = 0 & \mathbf{I}_3(s_1, b_2) = 0 \\
\mathbf{I}_1(s_1, s_2) = 2/9 & \mathbf{I}_2(s_1, s_2) = 0 & \mathbf{I}_3(s_1, s_2) = 1
\end{array}$$

In this example, the probabilities of the actions correspond with the one given for mixed strategy Nash equilibria. The following theorem demonstrates that this is generally true.

Theorem 1. *Let $\langle N, (A_i), (u_i) \rangle$ be a strategic game and let P be its corresponding DOP-CLP and let \mathbf{I} and α^* be respectively an interpretation for P and a mixed strategy profile for $\langle N, (A_i), (u_i) \rangle$ such that $\forall \mathbf{a} \in A, \forall i \in N, \alpha_i(\mathbf{a}) = \vartheta_{\mathbf{I}}(\mathbf{a}_i)$. Then, \mathbf{I} is a stable model iff α^* is a mixed strategy Nash equilibrium.*

4 Relationships to Other Approaches

4.1 Logic Programming

It is easy to see that positive logic programs are a subclass of the dynamically ordered choice logic programs, and that the stable models for both systems coincide. All necessary properties follow immediately from the way we handle single atoms. With the current semantics it is impossible to have a mapping between the stable models of a choice logic program ([4]) and the crisp stable models of the corresponding DOP-CLP. Indeed, our system is more credulous, since it allows a pure choice (probability 1) when two alternatives are equally preferred. However, we have that every stable model of a CLP is also a crisp stable model of the corresponding DOP-CLP.

4.2 Priorities

The logic programming language using priorities that corresponds best to our approach is dynamically ordered choice logic programming (OCLP) introduced in [6]. Although OCLP does not work with probabilities, these two systems have a common approach to and a similar notion of alternatives, in the sense that alternatives appear in the head of an applicable choice rule. OCLP also requires that this choice rule has a higher priority than the rule for which one computes the head atoms' alternatives. So, the main difference with our approach is the way that OCLP uses priority to create alternatives. Ordered logic programs ([8]) can easily be transformed to DOP-CLPs in such a way that the credulous stable models of the former correspond with the crisp stable models of the latter. For the same reason that we mentioned for CLPs, it is not yet possible to represent the skeptical stable model semantics for ordered logic programs. In [3], preference in extensive disjunctive logic programming is considered. As far as overriding is concerned the technique corresponds rather well with skeptical defeating of [6], but alternatives are fixed as an atom and its (classical) negation. Dynamic preference in extended logic programs is introduced in [2] in order to obtain a better suited well-founded semantics. Preferences/priorities are incorporated here as rules in the program.

While alternatives make our system dynamic, [2] introduces the dynamics via a stability criterion that overrules preference information but the alternatives remain static. A totally different approach is proposed in [14]. Here the preferences are defined between atoms without references to the program. After defining models in the usual way, one then uses preferences to filter out the less preferred models.

4.3 Uncertainty

A lot of researchers [1, 10, 9, 12, 13] have tackled the problem of bringing probabilities into logic programming. The probabilities used can be divided into two categories depending on the type of knowledge symbolized: statistical or belief. [1] concentrates on the first type while [10, 12, 13] are more interested in the latter. [9] is one of the few that is able to handle both types. Our formalism focuses mainly on knowledge of belief although it is possible to use statistical knowledge for defining the static priorities. An other difference between the various systems is the way they introduce probabilities and handle conjunctions. For example, [9] works with probability intervals and then uses the rules of probability to compute the probability of formulae. In this respect, we adopt the possible world/model theory of [10, 12]. However, we introduce probabilities at the level of interpretations, while they hard-code the alternatives by means of disjoint declarations together with probabilities, and the other atoms are computed by means of the minimal models of the logic program.

4.4 Games and Logic Programming

The logical foundations of game theory have been studied for a long time in epistemic logic. Only recently, researchers have become interested in the relationships between game theory and logic programming. The first to do so was [7]. It was shown that n -person games or coalition games can be transformed into an argumentation framework such that the NM-solutions of the game correspond with the stable extensions of the corresponding argumentation framework. [7] illustrated also that every argumentation framework can be transformed into a logic program such that the stable extensions of the former coincide with the stable models of the latter. In [4] it was demonstrated that each strategic game could be transformed into a CLP such that the Nash equilibria of the former correspond with the stable models of the latter. [6] shows that OCLPs can be used for an elegant representation of extensive games with perfect information such that, depending on the transformation, either the Nash or the subgame perfect equilibria of the game correspond with the stable models of the program. Concerning mixed strategy Nash equilibria of strategic games, the approach which is the most related to ours is the Independent Choice Logic of [13]. [13] uses (acyclic) logic programs to deterministically model the *consequences* of choices made by agents. Since choices are external to the logic program, [13] restricts the programs further, not only to be deterministic (i.e. each choice leads to a unique stable model) but also to be independent, in the sense that literals representing alternatives may not influence each other, e.g. they may not appear in the head of rules. ICL is further extended to reconstruct much of classical game theory and other related fields. The main difference with our approach is that we do not go outside of the realm of logic programming to recover the notion of

equilibrium. The basis of his formalism does not contain probabilities but works with selector functions over the hypotheses and then works with the (unique) stable model that comes from the program itself. This way one creates a possible world semantics. Our transformation makes sure that every atom is an alternative of a choice/decision for which a probability can be computed.

References

1. Fahiem Bacchus. **LP**, a Logic for Representing and Reasoning with Statistical Knowledge. *Computational Intelligence*, 6:209–231, 1990.
2. Gerhard Brewka. Well-Founded Semantics for Extended Logic Programs with Dynamic Preferences. *Journal of Artificial Intelligence Research*, 4:19–36, 1996.
3. Francesco Buccafurri, Wolfgang Faber, and Nicola Leone. Disjunctive Logic Programs with Inheritance. In Danny De Schreye, editor, *International Conference on Logic Programming (ICLP)*, pages 79–93, Las Cruces, New Mexico, USA, 1999. The MIT Press.
4. Marina De Vos and Dirk Vermeir. Choice Logic Programs and Nash Equilibria in Strategic Games. In Jörg Flum and Mario Rodríguez-Artalejo, editors, *Computer Science Logic (CSL'99)*, volume 1683 of *Lecture Notes in Computer Science*, pages 266–276, Madrid, Spain, 1999. Springer Verslag.
5. Marina De Vos and Dirk Vermeir. On the Role of Negation in Choice Logic Programs. In Michael Gelfond, Nicola Leone, and Gerald Pfeifer, editors, *Logic Programming and Non-Monotonic Reasoning Conference (LPNMR'99)*, volume 1730 of *Lecture Notes in Artificial Intelligence*, pages 236–246, El Paso, Texas, USA, 1999. Springer Verslag.
6. Marina De Vos and Dirk Vermeir. A Logic for Modelling Decision Making with Dynamic Preferences. Accepted at Jelita 2000. *Lecture Notes in Artificial Intelligence*. Springer Verslag.
7. Phan Minh Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77(2):321–358, 1995.
8. D. Gabbay, E. Laenens, and D. Vermeir. Credulous vs. Sceptical Semantics for Ordered Logic Programs. In J. Allen, R. Fikes, and E. Sandewall, editors, *Proceedings of the 2nd International Conference on Principles of Knowledge Representation and Reasoning*, pages 208–217, Cambridge, Mass, 1991. Morgan Kaufmann.
9. Raymond Ng and V.S. Subrahmanian. A semantical framework for supporting subjective and conditional probabilities in deductive databases. In Koichi Furukawa, editor, *Proceedings of the 8th International Conference on Logic Programming*, pages 565–580. MIT, June 1991.
10. Liem Ngo. Probabilistic Disjunctive Logic Programming. In Eric Horvitz and Finn Jensen, editors, *Proceedings of the 12th Conference on Uncertainty in Artificial Intelligence (AI-96)*, pages 387–404, San Francisco, aug 1996. Morgan Kaufmann Publishers.
11. Martin J. Osborne and Ariel Rubinstein. *A Course in Game Theory*. The MIT Press, Cambridge, Massachusetts, London, England, third edition, 1996.
12. David Poole. Logic programming, abduction and probability. In Institute for New Generation Computer Technology (ICOT), editor, *Proceedings for the International Conference on Fifth Generation Computer Systems*, pages 530–538. IOS Press, 1992.
13. David Poole. The independent choice logic for modeling multiple agents under uncertainty. *Artificial Intelligence*, 94(1–2):7–56, 1997.
14. Chiaki Sakama and Katsumi Inoue. Representing Priorities in Logic Programs. In Michael Maher, editor, *Proceedings of the 1996 Joint International Conference and Symposium on Logic Programming*, pages 82–96, Cambridge, September 2–6 1996. MIT Press.