A Simple Computational Model for Nonmonotonic and Adversarial Legal Reasoning

Giovanni Sartor
CIRFID, University of Bologna, Via Galliera 3, 40121, Bologna
IDG-CNR, Via Panciatichi 56/16, 50127 Firenze
e-mail SARTOR@CIRFID.UNIBO.IT

Keywords: Knowledge representation, nonmonotonic reasoning, arguments, rules and exceptions, legal system

In many commonsense contexts only incoherent and conflicting information is available. In such contexts reasonable conclusions must be derived from inconsistent sets of premises. This is especially the case in legal reasoning: legal norms can be issued by different authorities, in different times, to reach incompatible socio-political objectives, and the meaning of those norms can be semantically indeterminate.

Logic deduction alone is insufficient to derive justified conclusions out of inconsistent legal premises, since in the most popular logical systems (such as classical or intuitionistic logic) everything can be deduced from any contradiction. Nevertheless, much research now underway shows that formal methods can be developed for reasoning with conflicting information. The possibility of obtaining justified conclusions from an inconsistent set of premises increases when an ordering is defined over that set, since the ordering of the premises can be translated into an ordering of the competing arguments. This fact is particularly relevant for legal reasoning, since lawyers effectively solve normative conflicts by using ordering relations.

In the following pages, a model for reasoning with ordered defaults, interpreted as unidirectional inference rules, is proposed: a language for representing (possibly) contradictory rules is introduced, a notion of argument is defined, and types of arguments are distinguished. A simple interpreter in Prolog able to develop those arguments is also illustrated. Finally, the significance of the proposed model (and, more generally, of the acceptance of inconsistency) for the formal analysis of legal systems is discussed.

1. INTRODUCTION

In legal contexts it is frequently necessary to reason from *prima facie* incompatible premises. This is due to the following main reasons:

a. The defeasibility of legal norms. As other commonsense rules, legal norms can be applied only by default, *i.e.*, unless there are prevailing reasons to the contrary. Defeasibility in law and morality has a specific motivation: each concrete situation may express conflicting interests and values, and this fact is reflected in the possibility of ascribing to a single concrete act a plurality of possibly incompatible qualifications. For example, a work contract establishing the accessory obligation to accomplish certain political, religious, or sexual activities represents the exercise of contractual autonomy, but also an attempt on personal liberty; the killing of the assaulter is a homicide, but also the defence of the life of the assaulted person; the act of a soldier disobeying an order violates military discipline, but may protect fundamental constitutional values. In such cases we must distinguish the different, and possibly incompatible, *prima facie* evaluations of the concerned situation — each one separately considering certain aspects of that situation — and the overall judgement resulting from the comparison and balancing of all those evaluations[1]. The co-ordination of conflicting profiles of legal relevance is often accomplished by establishing preference relations, asserting that one norm prevails over the others, generally or under specific circumstances. Those relations qualify the prevailing norm as an exception, capable of blocking the application of the weaker norms (the rules).

Techniques have been developed to consistently represent sets of rules and exceptions (for example, by means of abnormality predicates), but exceptions can be more naturally expressed as statements contradicting the corresponding rules or denying their applicability. This representation is adopted, for example, in default logic (Reiter and Criscuolo [1981] 1987, 144f), in argumentation models based on specificity (on preferring the most specific argument, cf. Poole 1985), in logic programs of rules and exceptions (Kowalski and Sadri 1991), in abduction (Poole 1988), and in some experiences in computer science and law, such as, in particular, the OBLOG project (Gordon 1987).

b. The dynamics of normative systems. Law provides institutionalized procedures for the production of new legal norms. Those norms can (and normally do, since law has to adapt itself to changes in social and political situations) contradict norms already in force. The relation between more recent and older regulation is codified in the so called chronological principle, establishing the predominance of more recent regulations.

In dealing with dynamic normative systems a drastic solution is possible: old norms are to be deleted (abrogated) from the legal systems when conflicting with the new ones. Nevertheless, a softer strategy is also available: inconsistent norms issued through time can be preserved in the legal system, under the

[1] This aspect is well conceptualised by W.D. Ross (1930; 1939), who introduces the notion of *prima facie duty*: "I suggest '*prima facie* duty' or 'conditional duty' as a brief way of referring to the characteristic (quite distinct from that of being a proper duty) which an act has, in virtue of being of a certain kind (e.g. the keeping of a promise), of being an act that would be a duty proper if it were not at the same time of another kind which is morally significant" (W.D. Ross 1930, 19). According to this author "Moral intuitions are not principles by the immediate application of which our duty in particular circumstances can be deduced. They state [...] *prima facie* obligations. [...] We are *not* obliged to do that which is only *prima facie* obligatory. We are only bound to do that act whose *prima facie* obligatoriness in those respects in which it is prima facie obligatory most outweighs its *prima facie* disobligatoriness in those aspects in which it is *prima facie* disobligatory" (W.D. Ross 1939, 84-85). In recent legal theory the *prima facie* character of legal qualifications has been wrongly presented as deriving from the application of specific categories of norms (as the "principles" of Dworkin 1977). It concerns, instead, every legal and moral prescription.

condition that precedence is given to the most recent ones. This last solution is normally to be preferred when the conditions for the application of the old rules are not fully subsumable into the conditions for the application of the new ones: these last norms can then be considered as exceptions.

c. The concurrence of multiple legal sources. In modern legal systems, democracy and pluralism on the one hand, and the need for administrative regulation on the other, require that normative powers are distributed over a number of bodies and authorities. The authorities holding these powers may issue contradictory prescriptions. Those conflicts are normally solved according to the hierarchical criterion, which affirms the predominance of the prescriptions issued by higher authorities (*rectius*, the prescriptions issued in the exercise of higher normative competences).

Also in source conflicts a drastic solution is available: inhibiting the entrance into the normative systems to norms coming from inferior authorities or eliminating those norms, when they are incompatible with higher sources. Again, in most cases a different strategy seems preferable: preserving all conflicting prescriptions and giving the precedence to those stemming from higher sources.

d. Semantic indeterminacy. Legal language often leaves a space of semantic indeterminacy[2]. Some legal theorists postulate that in those contexts there is always just one right interpretation and only this interpretation is to be included in the legal system. This notion of legal system (the legal system as the set of the "right" interpretations) is of very little use for modelling legal reasoning and representing legal knowledge, since it simply assumes the result of legal reasoning, without considering the interpretative choices in the normative contexts and reasoning patterns in which they take place.

Also in dealing with alternative interpretations another solution is possible: preserving all possible interpretations in the legal system and giving the precedence the most reliable ones, according to the preferences (grades of plausibility) attributed by the interpreter.

To deal with all factors of inconsistency just considered — *i.e.*, to derive reasonable consequences out of a legal system including rules and exception, conflicting norms issued in subsequent times, prescriptions of different authorities, and alternative interpretations — we need inference procedures taking into account an ordering relation. This ordering can be built by assuming that higher legal sources are preferred to lower ones (hierarchical criterion), that subsequent norms are preferred to preceding ones (chronological criterion), that exceptions are preferred to rules (speciality criterion), that more plausible interpretations are preferred to less plausible ones (hermeneutic criterion)[3].

Two main approaches to deal with inconsistent ordered sets of premises can be distinguished:

i. Assumption-based approach. This approach tries to identify best or preferred subsets of the set of premises under consideration, *i.e.*, to identify one (or some) consistent subset(s) of assumptions allowing all justified consequences being derived. Here the accent is on the whole of the knowledge base, and a "definitive" maximal selection is pursued, from which every justified conclusion can be logically deduced (although a recomputation of that selection may be necessary after changes in the knowledge base).

ii. Argument-based approach. This approach, instead, looks for preferred arguments, *i.e.*, arguments leading to justified consequences. Here the accent goes on single inferences, to wit on minimal sets of premises implying the desired conclusion. The argument construction is to be performed each time a consequence is derived.

In the first approach we can mention the proposals of Alchourrón and Makinson (1981), Alchourrón (1986), and Brewka (1991a, 1991b), whose application to legal contexts is considered in Sartor (1992a).

In the argument based approach, we can recall Simari and Loui (1992), who use specificity as the ordering criterion (according to Poole 1985), and Prakken (1992), who also considers explicit orderings. These authors, like default logic (cf. Reiter [1980] 1987), refuse contrapositive inferences: defeasible statements are considered unidirectional inference rules. Prakken (1992) translates the rejection of contrapositive inferences into a technique for comparing arguments of explicitly ordered premises: when confronting arguments leading to contradictory conclusions p and $\neg p$, only the rules at the top of the inference chains for p and $\neg p$ must be considered.

In fact, the reasons that in some commonsense contexts block contrapositive inferences also prevent us from questioning a rule just because it causes an inconsistency indirectly (by realising a condition for other rules being applied). For example, let us consider the arguments $A1 = \{w \leftarrow q;\ q \leftarrow p;\ p\}$ for w and $A2 = \{\neg w \leftarrow s;\ s \leftarrow u,\ u\}$ for $\neg w$. Only the rules $w \leftarrow q$ and $\neg w \leftarrow s$ should be compared since intermediate rules are not responsible for the conflict: their consequents are consistent with the competing argument and can be contradicted only by reasoning contrapositively (for deriving $\neg q$, and putting into question $q \leftarrow p$ we need to reason contrapositively from $w \leftarrow q$ and $\neg w$).

In the following pages we will develop a model for reasoning with inconsistent ordered premises in a logic programming framework, model that is both argument-based and inspired by the rejection of contrapositive inferences: we will consider how to build and compare arguments, to derive plausible and justified conclusions. This model is especially inspired to Prakken (1991a, 1991b, 1992), but is characterised by the following features:

— Thanks to a limitation of the language (corresponding to logic programming approaches) a simple and effective notion of argument is obtained and translated into inferential procedures (par. 2 and 3) implemented into an interpreter in Prolog (par. 5).

— The language can express exceptions concerning specific rules (par. 4).

— Arguments concerning preference relations between norms can be imbedded into the reasoning in which the legal qualification established by those norms are derived (par. 7.2).

2. THE LANGUAGE

We use simple rules of the form

$$\boldsymbol{n}: p0 \leftarrow p1 \wedge \ldots \wedge pn$$

where $\boldsymbol{n}$ is the rule name, and each p_i is a literal. A literal is a formula q or $\neg q$, where q is an atom, and $\neg$ is interpreted as classical negation. Note that negative literals can occur in both the consequent (head) and the antecedent (body) of those rules.

[2] Cf., among the others, Alexy (1981a, 190 f), who distinguishes three types of semantic indeterminacy, ambiguity, vagueness, and evaluative openness.

[3] On those criteria and their harmonisation, cf. Sartor (1992a).

We also admit degenerate rules with an empty body, that we shall call *facts*. A rule name ***n*** is a label of the form

$r(X_1, \ldots, X_n)$,

where r is a new function symbol, and $X_1, \ldots, X_n$ are the free variables in the rule. By substituting $X_1, \ldots, X_n$ with appropriate terms $t_1, \ldots, t_n$ we obtain names $r(t_1, \ldots, t_n)$ for all instances of the rules. We write $\mathbf{n2} \succ \mathbf{n1}$ to mean that the rule $\mathbf{n2}$ is preferred to the rule $\mathbf{n1}$. The $\succ$ relation is transitive and antisymmetric (a strict partial order).

Let us consider a simple example concerning the protection of privacy:

r1(X): permitted(publishing(X)) ←
of_public_interest(X);
r2(X): ¬permitted(publishing(X)) ←
private_information(X);
f1: of_public_interest(mary's_photo);
f2: private_information(mary's_photo).

Let us first assume that the ordering $\succ$ is empty. In such a situation we have two competing arguments:

A_1={**r1**(mary's_photo):
permitted(publishing(mary's_photo)) ←
of_public_interest(mary's_photo);
f1: of_public_interest(mary's_photo)}

for permitted(publishing(mary's_photo)), and

A_2={**r2**(mary's_photo):
¬permitted(publishing(mary's_photo)) ←
private_information(mary's_photo);
f2: private_information(mary's_photo)}

for ¬permitted(publishing(mary's_photo)).

Using rule names we can represent those arguments simply as A_1 = {**r1**(mary's_photo); **f1**}, and A_2 = {**r2**(mary's_photo); **f2**}.

Neither argument is binding (justifying) since there is no preference relation between the adversary rules **r1** and **r2**. In such a situation we can draw only plausible, or uncertain, conclusions: both the permission and the prohibition of publishing the photo are merely plausible consequences of the premises set.

Let us extend the order, making privacy more important than liberty of communication, *i.e.*, let us add the relation **r2**(X) $\succ$ **r1**(X). Then, according to intuition, ¬permitted(publishing(mary's_photo)) becomes the justified consequence. In fact:

— **r2** prevails over the adversary rule **r1**, and therefore the argument A_2 for ¬permitted(publishing(mary's_photo), consequent of **r2**(mary's_photo) defeats the argument A_1 for permitted(publishing(mary's_photo)), consequent of **r1**(mary's_photo);

— no direct counterargument can be raised against the remaining part of A_2, *i.e.*, against {**f2**}.

3. NOTIONS OF INFERENCE

Let us try to build a formal framework to model the type of defeasible reasoning exemplified above (cf. Sartor 1992b). We distinguish different notions of consequence of an inconsistent set Π of ground rules (we consider each open rule as the set of its ground instances, *i.e.*, its instances not containing variables):

— *Logical consequence.* Since Π is inconsistent, every statement is a logical consequence of Π.

— *Grounded consequence.* We say that p a grounded consequence of Π iff p is derivable from a consistent subset of Π, using the rules in Π as unidirectional inference rules[4]. Not every logical consequence is also a grounded consequence (consequences derivable only by the *falsum* inference rule are surely not).

— *Plausible consequence.* We say that p is a plausible (or undefeated) consequence of Π if an argument in favour of p is not worse than any argument to the contrary.

— *Justified consequence.* We say that p is a justified consequence of Π, if an argument in favour of p is better than every argument to the contrary.

To formalise the notions just introduced we need to specify the concepts of argument and counterargument, and distinguish types of arguments.

a. Argument and counterargument. As we have said above our rules are unidirectional inference rules, which do not admit contrapositive inferences. The set of the statements derivable from a rule set Σ — *i.e.*, the extension of Σ, denoted as $E(\Sigma)$ — can therefore be obtained by applying repeatedly the rules in Σ to the result of the previous application of those rules (starting from the empty set), until no new consequence can be obtained in this way (a fix point has been reached). So we can also say that $E(\Sigma)$ is the minimum set H such that for every (ground) rule ***n***: $p_0 \leftarrow p_1, \ldots, p_n$ in Σ, if $p_1, \ldots, p_n$ are in H, then also p_0 is in H (cf. Gelfond and Lifschitz 1990).

We say that A is an *argument* for p in Π iff A is minimal among the consistent subsets Γ of Π such that $p \in E(\Gamma)$. Therefore, an argument A for p is a consistent minimal set of (ground) rules allowing p to be derived. For any $q \in E(A)$, A includes a subargument for q, *i.e.*, an $A_1 \subseteq A$ such that A_1 is an argument for q. The consistency requirement for arguments can be translated into the demand that no argument contains (ground) rules with complementary heads. In fact, if an argument contained two rules with contradictory heads, it would be either inconsistent (if the antecedents of both rules were satisfied) or non minimal (if one rule were not satisfied). An argument B for q is a *counterargument* to an argument A iff A includes a subargument A_1 for $\overline{q}$ (where $\overline{q}$ denotes the complement of q). A counterargument puts into question the corresponding arguments, since it culminates in the negation of a statement whose derivation is necessary to reach the conclusion of the argument.

b. Defeated argument. We say that an argument A *is defeated* iff a subargument $A_i \subseteq A$ is *directly defeated*. An argument A_i for q is *directly defeated* iff there exists an argument B for $\overline{q}$ such that:

i. B includes a rule **r2**: $\overline{q} \leftarrow w_1 \wedge \ldots \wedge w_n$ such that **r2** $\succ$ **r1**, where **r1**: $q \leftarrow s$ is the rule for q in A_i (being an

[4] Our rules behave in this regard as to the already mentioned inference rules of default logic. They can be also understood as a formalisation of the warrants of Toulmin (1969, 97ff).

argument, A_i must be minimal, and therefore can contain only one of such rules).

ii. All strict subarguments of B (all arguments B_i such that $B_i \subset B$) are justifying, *i.e.*, B includes justifying arguments for $w_1, \ldots, w_n$.

When any subargument A_i of an argument A is directly defeated, A is invalidated: it can be discarded since it cannot lead to new plausible or justified conclusions. In such a case we also say that B is a *defeating counterargument* to A.

c. Plausible argument. We say that an argument A for p is plausible (or undefeated) iff A has no defeating counterargument B. We can equivalently say that an argument A is plausible iff:

i. A for p is not directly defeated;

ii. A includes plausible arguments for $s_1, \ldots, s_n$, where $\mathbf{r1}$: $p \leftarrow s_1 \wedge \ldots \wedge s_n$ is the rule for p in A.

A plausible argument is therefore an argument that is not defeated: no inference step necessary to reach the conclusion of the argument clashes against a prevailing counterargument. This does not mean that the argument is able to justify its conclusion: it may have a counterargument not inferior to it. In this last case the argument would be *merely plausible.*

d. Justifying argument. We say that an argument A for p is justifying iff A has no plausible counterargument B. We can equivalently say that an argument A is plausible iff

i. There is no plausible argument for $\overline{p}$;

ii. A includes justifying arguments for $s_1, \ldots, s_n$, where r: $p \leftarrow s_1 \wedge \ldots \wedge s_n$ is the rule for p in A.

A justifying argument is therefore an argument that is not questioned by any valid (undefeated) objection.

The notions of argument just introduced allow us to specify different types of consequence:

— *Grounded consequence:* p is a grounded consequence of Π iff there is an argument $A \subseteq \Pi$ for p.

— *Plausible consequence:* p is a plausible consequence of Π iff there is a plausible argument $A \subseteq \Pi$ for p.

— *Justified consequence:* p is a justified consequence of Π iff there is a justifying argument $A \subseteq \Pi$ for p.

We also say that p is a merely grounded consequence if p is grounded but not plausible, and that p is a merely plausible consequence, iff p is plausible and not justified.

Let us now apply the reasoning patterns just sketched to our privacy example. Let us first consider the premises set

Π_1={**r1**(X): permitted(publishing(X)) ← of_public_interest(X);
r2(X): ¬permitted(publishing(X)) ←
private_information(X);
f1: of_public_interest(mary's_photo)}

with an empty $\succ$ relation. Π_1 includes a justifying argument for permitted(publishing(X)) to wit A_1 = {**r1**(mary's_photo); **f1**}, and no argument for ¬permitted(publishing(mary's_photo)). Let us add the fact **f2**: private_information(mary's_photo) to Π_1, so to obtain the rule set Π_2 = {**r1**(mary's_photo); **r2**(mary's_photo); **f1**; **f2**}.

In Π_2, A_1 is no longer justifying, since it has the plausible counterargument A_2 = {**r2**(mary's_photo); **f2**} for ¬permitted(publishing(mary's_photo))}. Both A_1 and A_2 are plausible non justifying arguments, so that both permitted(publishing(mary's_photo)) and ¬permitted(publishing(mary's_photo)) are merely plausible consequences.

Let us now extend the $\succ$ relation with **r2**(X) $\succ$ **r1**(X). Not surprisingly, we obtain the result that A_2 becomes justifying — and its conclusion ¬permitted(publishing(mary's_photo)), justified — while A_1 becomes merely grounded (grounded and not plausible). In fact, A_1 (the only counterargument to A_2) is directly defeated by A_2 and therefore loses its plausibility, so that A_2 becomes able to justify its conclusion.

4. APPLICABILITY PREDICATES AND CROSS REFERENCES

In many common sense contexts it is useful to express rules asserting that other rules are — or are not — applicable given certain conditions. This is especially relevant when we want to express "selective exceptions", *i.e.*, exceptions not intended to exclude the derivation of a certain predicate, but only meant to prevent the use of certain rules. Selective exceptions are especially useful for formulating second level exceptions, stating that determinate exceptions are not applicable under certain conditions. For example, in Italian law a contract is invalid if stipulated by a person under age (first level exception), but this provision is inapplicable if the under age person lied about his age (second level exceptions). The second level exception cannot be interpreted as the assertion that every contract stipulated by somebody falsely affirming to be in full age is *ipso facto* valid: there are many causes of contract invalidity — *e.g.*, having been forced to stipulate the contract, the absence of the prescribed form, etc. — for which a lie on age is totally irrelevant.

Selective exceptions do not necessary need a metalanguage: they can be formulated in an object level language including rule names (cf. Brewka 1986; Poole 1987). It is sufficient to translate every rule of the form:

$r(X_1, \ldots, X_n)$: $p \leftarrow q$

into the pair:

$r(X_1, \ldots, X_n)$: $p \leftarrow q \wedge$ applicable($r(X_1, \ldots, X_n)$);
$r(X_1, \ldots, X_n)$: applicable($r(X_1, \ldots, X_n)$)

where the fact applicable($r(X_1, \ldots, X_n)$) asserts that all instances of the rule $r(X_1, \ldots, X_n)$ are applicable, and the derivation of the consequent of the original rule is conditioned to the applicability of that rule. From now on, we will assume that each rule has been translated into this form, and interpret original rules (and their names) as abbreviations for their translation.

Applicability predicates constitute a powerful and flexible linguistic feature. For example, we can have the following representation of the legal context just considered:

r1(X): valid(X) ← contract(X);
r2(Y, X): ¬ valid(X) ← under_age_for(Y, X);
r3(Y, X): ¬ applicable(r2(Y, X)) ← declared_full_age(Y);
f1: contract(sale1);
f2: under_age_for(mark, sale1);
f3: declared_full_age(mark).

Let us assume the priority relations **r3** $\succ$ **r2** $\succ$ **r1**. The set of rules above is to be translated into the following:

r1(X): valid(X) ← contract(X) ∧ applicable(r1(X));

r1(X): applicable(r1(X));
r2(Y, X): ¬ valid(X) ← under_age_for(Y, X) ∧
applicable(r2(Y, X));
r2(Y, X): applicable(r2(Y, X));
r3(Y, X): ¬applicable(r2(Y, X)) ← declared_full_age(Y) ∧
applicable(r3(Y, X));
r3(Y, X): applicable(r3(Y, X));
f1: contract(sale1) ← applicable(f1);
f1: applicable(f1);
f2: under_age_for(mark, sale1) ← applicable(f2);
f2: applicable(f2);
f3: declared_full_age(mark) ← applicable(f3);
f3: applicable(f3).

In this extended set we have a justifying argument for valid(sale1):

A_1={**r1**(sale1): valid(sale1) ← contract(sale1) ∧
applicable(r1(sale1));
r1(sale1): applicable (r1(sale1));
f1: contract(sale1) ← applicable(f1);
f1: applicable (f1)}.

There is a counterargument to A_1:

A_2={**r2**(mark, sale1): ¬ valid(sale1) ←
under_age_for(mark, sale1) ∧
applicable(r2(mark, sale1));
r2(mark, sale1): applicable(r2(mark, sale1));
f2: under_age_for(mark, sale1) ← applicable(f2);
f2: applicable(f2)}.

A_2 could block A_1 since **r2** ≻ **r1**. Nevertheless this is not the case, since A_2 is not plausible, being defeated by the argument

A_3={**r3**(mark, sale1): ¬applicable(r2(mark, sale1)) ←
declared_full_age(mark) ∧
applicable(r3(mark, sale1));
r3(mark, sale1): applicable(r3(mark, sale1));
f3: declared_full_age(mark) ← applicable(f3);
f3: applicable(f3)}.

A_3 directly defeats the subargument {**r2**(mark, sale1): applicable(r2(mark, sale1))} of A_2.

5. A SIMPLE INTERPRETER IN PROLOG

Let us now develop an interpreter for our language. The rules are represented by unit clauses of the form:

n: $p_0 \leftarrow p_1 \wedge \ldots \wedge p_m$

where ":", ← and ∧ are infix operators, the first having the lowest priority and being therefore the predicate symbol. Each rule is translated into the corresponding pair

n: $p_0 \leftarrow p_1 \wedge \ldots \wedge p_m \wedge$ applicable(*n*);
n: applicable(*n*)

by the predicate rule(*Name, TranslatedRule*), defined by two clauses, the first extending each rule with its applicability condition, and the latter establishing that the rule is applicable.

```
rule(Name: Head ← Body ∧ applicable(Name)):-
    Name: Head ← Body.
rule(Name: applicable(Name) ← true) :- Name: _.
```

The ≻ relation is represented by the predicate $Name_1$ preferred_to $Name_2$, defined in terms of the predicate $Name_1$ stronger_than $Name_2$:

```
X preferred_to Y :- X stronger_than Y.
X preferred_to Y :- X stronger_than Z, Z preferred_to Y.
```

Let us now define the notion of grounded consequence. The predicate grounded_consequence(*Consequence*, A_1, A_2) states that *Consequence* is a grounded consequence of the argument A_2, built by consistently extending the partial argument A_1.

```
grounded_consequence(L, A1, A2) :-
    rule(Name: L ← Body),
    consistent(L, A1),
    grounded_consequence(Body,
        [Name: L ← Body | A1], A2).
grounded_consequence(L1 ∧ L2, A1, A3) :-
    not var(L1),
    grounded_consequence(L1, A1, A2),
    grounded_consequence(L2, A2, A3).
grounded_consequence(true, A, A).
```

Let us go to the notion of plausible consequence, *i.e.*, of undefeated grounded consequence. The predicate plausible_consequence(*L*, A_1, A_2) states that the literal *L* is a plausible consequence of the argument A_2, built by consistently extending the partial argument A_1.

```
plausible_consequence(L, A1, A2) :-
    rule(Name: L ← Body),
    consistent(L, A1),
    not directly_defeated(Name, L),
    plausible_consequence(Body,
        [Name: L ← Body | A1], A2).
plausible_consequence(L1 ∧ L2, A1, A3):-
    not var(L1),
    plausible_consequence(L1, A1, A2),
    plausible_consequence(L2, A2, A3).
plausible_consequence(true, A, A).
```

The predicate directly_defeated($Name_2$, *L*) states that any argument deriving the consequence *L* with the rule $Name_1$ is directly defeated. This happens when the complement of *L* can be derived by using a rule $Name_2$ that is preferred to $Name_1$, and the body of $Name_2$ is a justified consequence:

```
directly_defeated(Name1, L) :-
    complement(L, ComplL),
    rule(Name2: ComplL ← Body),
    Name2 preferred_to Name1,
    justified_consequence(Body,
        [Name: L ← Body], A2).
```

Finally, let us define the notion of justified consequence. The predicate justified_consequence(*L*, A_1, A_2) states that the literal *L* is a justified consequence of the argument A_2, built by consistently extending the partial argument A_1.

```
justified_consequence(L, A1, A2):-
    rule(Name: L ← Body),
    consistent(L, A1),
```

```
    complement(L, ComplL),
    not plausible_consequence(ComplL, [], _),
    justified_consequence(Body,
        [Name: L ← Body | A1], A2).
justified_consequence(L1 ∧ L2, A1, A3):-
    not(var(L1)),
    justified_consequence(L1, A1, A2),
    justified_consequence(L2, A2, A3).
justified_consequence(true, A, A).
```

This simple interpreter is integrated with definitions for auxiliary predicates, such as complement(*L*, *ComplementOfL*), which relates each literal L to its complement $\overline{L}$, and consistent(*L*, *PartialArgument*), which checks if a rule with consequent L can be consistently added to *PartialArgument*. The consistency check is performed by examining if *PartialArgument* contains a rule with consequent $\overline{L}$ (if a rule for $\overline{L}$ is in *PartialArgument*, the addition of the rule for L would violate either consistency or minimality of the argument being constructed). This consistency check is sufficient for ground rules, but more general solutions are available according to established techniques. The interpreter is capable of determining the grounded, plausible, and justified consequences of any set of rules.

6. NONMONOTONIC AND ADVERSARIAL REASONING

Nonmonotonic and adversarial reasoning are strictly related in our model: by adding new premises to the available knowledge base, not only new arguments become possible, but old arguments are defeated or reinstated, so that the sets of justified and plausible consequences do not grow monotonically. Let us consider some rules and exceptions from Italian tort law:

fault_liability(*X*, *F*):
X is liable for the harmful fact *F* ←
  X accomplished culpably the fact *F*;
incapacity(*X*, *F*):
¬(*X* is liable for the harmful fact *F*) ←
  X was incapable during the fact *F*;
incapacity_exc1(*X*, *F*):
¬applicable(incapacity(*X*, *F*)) ←
  X's incapacity during the fact *F* was due to his fault;

over which the following ordering is established: **incapacity_exc1**(*X*, *F*) ≻ **incapacity**(*X*, *F*) ≻ **fault_liability**(*X*, *F*).

Let us assume that Mary wants compensation from John, saying that he crashed her fence with his car. She has been able to prove the following facts:

f1: *john* accomplished culpably the harmful fact *fence_crash*.

In the premises set Π_1 = {**fault_liability**; **incapacity**; **incapacity_exc1**; **f1**}, the argument A_1 = {**fault_liability**(*john*, *fence_crash*); **f1**} justifies the conclusion that John is liable for the destruction of the fence.

John can free himself by satisfying the condition of an exception, hence causing a contradiction with the rules establishing liability. So, for example, if he proves that he was incapable at the time of the accident, *i.e.*, that

f2: *john* was incapable during the fact *fence_crash*,

we obtain the premises set Π_2 = {**fault_liability**; **incapacity**; **incapacity_exc1**; **f1**; **f2**}. In Π_2, John's liability is excluded by the argument A_2 = {**incapacity**(*john*, *fence_crash*); **f2**} for ¬(*john* is liable for the harmful fact *fence_crash*). A_1 is directly defeated by A_2, since **incapacity**(*X*, *F*) ≻ **fault_liability**(*X*, *F*), and the only subargument of A_2 ({**f2**}) is justifying.

Mary can prove that John's incapability derived from his fault (for example, being caused by his drunkenness):

f3: *john's* incapacity during the fact *fence_crash* was due to his fault.

In the premises set Π_3 = {**fault_liability**; **incapacity**; **incapacity_exc1**; **f1**; **f2**; **f3**}, the argument A_3 = {**incapacity_exc1**(*john*, *fence_crash*); **f3**} for ¬applicable(incapacity(*john*, *fence_crash*)) invalidates A_2, by directly defeating the subargument {**incapacity**(*John*, *fence_crash*): applicable(incapacity(*john*, *fence_crash*)} of A_2. Therefore, A_1 becomes reinstated, and John appears liable again (by a justifying argument).

Let us assume that the incapacity of John was not due to drunkenness, but to alcoholism, so being a permanent state. In such a case, it would be dubious whether he could be said to be faulty for his incapacity, since although he caused his alcoholism, once he became an alcoholic, he could not stop being an alcoholist by an act of will. So we would have two competing rules

alcoholism1(*X*, *F*):
X's incapacity during the fact *F* was due to his fault ←
  X's incapacity during the fact *F* was due to alcoholism;
alcoholism2(*X*, *F*):
¬(*X*'s incapacity during the fact *F* was due to his fault) ←
  X's incapacity during the fact *F* was due to alcoholism;

and fact **f3** would be substituted by the following:

f3': *john*'s incapacity during the fact *fence_crash* was due to alcoholism.

From the premises set Π_4 = {**fault_liability**; **incapacity**; **incapacity_exc1**: **alcoholism1**; **alcoholism2**; **f1**; **f2**; **f3'**} we have no justified conclusion about John's responsibility. Clearly, if we choose **alcoholism1** (if the ≻ relation is extended with **alcoholism1** ≻ **alcoholism2**) we can derive that John is liable; if we choose **alcoholism2** (if the ≻ relation is extended with **alcoholism2** ≻ **alcoholism1**) we can derive that he is not liable.

In fact, Π_4 includes the merely plausible argument A_1 = {**fault_liability**(*John*, *fence_crash*); **f1**} for John's liability. A_1 is undefeated, since it has the counterargument A_2 = {**incapacity**(*john*, *fence_crash*); **f2**}, but A_2 is not defeating, being itself merely plausible: it has the counterargument A_4 = {**incapacity_exc1**(*john*, *fence_crash*); **alcoholism1**(*john*, *fence_crash*); **f3'**} for ¬applicable(incapacity(*john*, *fence_crash*)). A_4 also is merely plausible: it has the merely plausible counterargument A_5 = {**alcoholism2**(*john*, *fence_crash*), **f3'**}, which is contradicted by the merely plausible argument A_6 = {**alcoholism1**(*john*, *fence_crash*); **f3'**}.

The possibility of distinguishing plausible and justified arguments is very important in law: frequently legal sources are susceptible to alternative interpretations and the choice among

those interpretations may involve an evaluative (and subjective) element. Therefore, every representation of a legal context by a consistent set of premises represents a drastic simplification of a wider decisional environment. Our model preserves the original complexity of that environment, by distinguishing the controversial (and therefore adversarial) aspects of the legal problem domain, expressed by conflicting clauses, and the evaluative choices, formulated by extending the $\succ$ relation.

7. TWO EXTENSIONS OF THE PROPOSED MODEL

The model here proposed can be extended to meet some significant demands.

7.1. The formalisation of meaning postulates

Our approach is based on the rejection of contrapositive inferences, a very intuitive choice for defeasible rules. Nevertheless this choice is inadequate for meaning postulates. To clarify this point, let us reconsider the following example.

r1(X): X shall be put in prison $\leftarrow$ X killed somebody;
r2(X): $\neg$ (X shall punished) $\leftarrow$ X acted for self defence;
f1: *peter* killed somebody;
f2: *peter* acted for self defence.

Let us also assume the preference relation **r2** $\succ$ **r1** and add an additional rule **r3**, more reliable than **r2** (**r3** $\succ$ **r2**), and stating that being put in prison is a punishment:

r3(X): X shall be punished $\leftarrow$ X shall be put in prison.

We would expect that Peter is not to be put in prison, for the following reasons: he acted for self defence, rule **r2** is preferred to rule **r1**, and the two rules are incompatible in the framework of the available knowledge, where prison turns out to be a punishment. Instead, we derive the surprising consequence that Peter shall be put in prison, but this does not constitute a punishment. This conclusion derives from the fact that **r3** has been used as a default, to wit as a unidirectional inference rule. Therefore rule **r3** determines the rejection of rule **r2** contradicting its consequent, rather than rule **r1** satisfying its antecedent.

The meaning postulates should be used, instead, as additional axioms for checking the compatibility between defaults. In other words, when establishing the compatibility of a default rule **r1**: $p \leftarrow s$ with a partial argument A, the inconsistency check should be applied to the set $\Sigma = \{p\} \cup A \cup M$ where M contains the relevant meaning postulates. In a logic programming model, the meaning postulates could be represented as constraints.

7.2. Modelling comparative evaluations (reasoning about preference relations)

The model sketched above takes into account preference relations, but is not able to reason about them. Preferences are represented as categorical assertions: universal facts including free variables — those facts establishing that each instance of a rule **r1**(X_1, ..., X_n) is preferred to every instance of a competing rule **r2**(Y_1, ..., Y_m) —, or concrete facts, obtained by substituting constants to the variables.

In many legal contexts, instead, the precedence accorded to certain rules is justified only in specific circumstances. This is especially the case when there is a dynamic tension between competing values or interests. Such contexts require that the patterns of reasoning introduced above for deriving legal consequences are also used for establishing preferences.

This aspect of legal reasoning can be represented in our model by modifying the notion of defeat: a justifying argument must establish that the top rule of the defeating counterargument is stronger than the top rule of the defeated argument. This solution can be implemented by a slight modification of our interpreter: we give up transitivity of the preference relation (by deleting the definition of the predicate preferred_to) and modify as follows the definition of the predicate directly_defeated (where Name_1 is the top rule of the defeated argument and Name_2 is the top rule of the defeating counterargument):

directly_defeated(Name_1, L) :-
  complement(L, ComplL),
  rule(Name_2: ComplL $\leftarrow$ Body),
  justified_consequence(Name_2 stronger_than Name_1,
    [Name_2: L $\leftarrow$ Body], A1),
  justified_consequence(Body, A1, A_2).

To illustrate this extension of our model[5], let us now formalise a judgement of the German constitutional Court, the "Lebach Urteil", illustrated in Alexy (1980). This decision concerned a television documentary about a grave crime, in which the names of the participants to the crime were mentioned and their photos shown. One of those participants (who had a minor role in the fact) affirmed that this programme violated his privacy and compromised the chances of his resocialisation, and therefore violated his personality right (right to the free development of the personality) protected by the German Constitution.

The German constitutional Court, as Alexy observes, develops his decision in three steps:

a. The Court admits that in the present case a comparative evaluation is necessary since two conflicting constitutional provisions are applicable: the right to privacy, excluding the publication of private information, and the right to communication, granting the liberty of propagating information. Moreover, it observes that none of these rights is to be unconditionally preferred to the other: only the particular circumstances of the single case allow a choice to be made.

b. It affirms that usually the interest to be informed by television about crimes overcomes the grave aggression to privacy regularly determined by transmissions concerning those crimes. This predominance, nevertheless, is not to be recognised for transmissions taking place when the fact is no more actual.

c. It concludes that, under this last circumstance (as in the Lebach case), the representation of a documentary concerning a criminal fact is not admissible if it can determine a new or additional prejudice to the author of the fact.

Our approach allows a straightforward representation of the argument of the Court:

---

[5] In addition, we need a constraint stating that the stronger_than relation is antisymmetric: it cannot be true that two rules R_1 and R_2 are such that both R_1 stronger_than R_2 and R_2 stronger_than R_1. If we do not have a mechanism for reasoning with constraints, we can establish that in general any rule stating that R_1 is stronger that R_2 has a corresponding rule stating that R_2 is not stronger than R_1. This can be simply obtained by adding a new clause to the definition of the predicate "rule":
rule(Name: $\neg R_2$ stronger_than R_1 $\leftarrow$ Body $\wedge$ applicable(Name)) :-
  Name: R_1 stronger_than R_2 $\leftarrow$ Body.

a. Incompatible prima facie qualifications. In the first step the Court observes that *prima facie*, two alternative conclusions are derivable in the framework of the constitutional rules. We can represent this normative context by the following rules:

r1(*X*): ¬(*X* is permitted) ← *X* violates privacy

r2(*X*): *X* is permitted ← *X* is a form of communication.

r3(*X*): *X* violates privacy ←
X is the television transmission of a documentary about a criminal fact indicating the authors of that fact

r4(*X*) *X* is a form of communication ←
X is the television transmission of a documentary about a criminal fact indicating the authors of that fact.

Rules **r1** and **r2** express the constitutional evaluation of privacy and communication; rule **r3** and **r4** specify (concretise) the constitutional rules in relation to the facts of the case: the transmission of a television documentary about a criminal fact, which indicates the authors of this fact, constitutes both a communication and a violation of privacy.

Let us add also the basic fact of the case:

f1: *lebach_transmission* is the television transmission of a documentary about a criminal fact indicating the authors of that fact

Given this fact, and no preference relation among the rules, we are able to derive, as the Court did, two merely plausible conclusions: the transmission is permitted, by the argument $A1$ = {**r1**(*lebach_transmission*); **r3**(*lebach_transmission*); **f1**}, while it is not permitted by the argument $A2$ = {**r2**(*lebach_transmission*); **r4**(*lebach_transmission*); **f1**}.

b. First comparative evaluation. The instances of rule **r2** (protecting the liberty of communication) that concern a television transmission of the indicated type are usually preferred to the corresponding instances of rule **r1** (protecting privacy)

p1(*X*): r2(*X*) stronger_than r1(*X*) ←
X is the television transmission of a documentary about a criminal fact mentioning the authors of that fact

The preference rule **p1** strengthens argument $A2$. More exactly **p1** permits the construction of a preference argument $P1$ for r2(*Lebach*) stronger_than r1(*Lebach*). This argument, together with $A2$, defeats $A1$, and therefore the conclusion of $A2$ — the permission of the transmission — becomes justified. Nevertheless, the judges also affirm that this preference only exists for the transmissions of actual facts, *i.e.*, that the rule **p1** is not applicable to subsequent transmissions:

$\mathbf{e_{p1}}$(*X*): ¬applicable(p1(*X*)) ← X is subsequent to the fact.

Obviously, the exception $\mathbf{e_{p1}}$ prevails over **p1**.

p2: e_{p1}(*X*) stronger_than p1(*X*).

Since in the present case it holds that

f2: *lebach_transmission* is subsequent to the fact

the exception $\mathbf{e_{p1}}$ blocks the preference relation **p2** from being applied to the present case, and therefore the conflict between **r1** and **r2** remains undecided: both contradictory legal qualifications of the Lebach transmission remain merely plausible.

c. Second comparative evaluation. To reach a justified conclusion, a second comparative evaluation is needed, establishing that privacy prevails under the conditions of the case:

p3(X): r1(X) stronger_than r2(X) ←
X is the television transmission of a documentary about a criminal fact mentioning the authors of that fact ∧
X is subsequent to the fact ∧
X causes a new violation of privacy.

Let us add the last fact:

f3: *lebach_transmission* causes a new violation of privacy.

The preference relation **p3**, allows to extend $A1$ into a defeating counterargument against $A2$, and therefore makes $A1$ justifying and its conclusion — the non permissibility of the transmission of documentary on Lebach — justified. Note that our formalisation immediately reflects the structure of the argument of the German Court, in a precise formal framework.

8. CONCLUSION: DEFEASIBLE ARGUMENTS AND THE LEGAL SYSTEM

The possibility to develop reasonable arguments from inconsistent premises (as in the model here illustrated) has considerable implications for the formal analysis of legal systems.

For our purposes, the legal system — and its representation in computable structures, *i.e.*, a legal knowledge base — can be considered as a set of normative premises used to justify (to derive) legal decisions. The conceptual framework here proposed allows two qualifications of normative systems being distinguished: consistency and determinacy (on determinacy, cf. also Gordon 1991):

— A legal system Σ can be said to be (*unconditionally*) *consistent* iff, in every possible case C[6], the set $\Sigma \cup C$ is logically consistent. We can also say that a rule set Σ is (unconditionally) consistent iff in no possible case C, $\Sigma \cup C$ has merely grounded consequences. Consistency is not a reasonable objective in the representation of legal knowledge: given the causes of inconsistency seen in par. 1 above, a consistent formalisation of legal contents presupposes a drastic manipulation of the original sources and is unable to model essential aspects of legal knowledge and legal reasoning, such as dealing with rules and exceptions[7].

— A legal system Σ instead, can be said to be (*unconditionally*) *determinate* iff, in no possible case C the set $\Sigma \cup C$ has merely plausible consequences. Every indeterminate legal system is also inconsistent, but the converse is not always true: preference relations, as we have seen in some examples above, may establish determinacy in inconsistent sets. Determinacy is a very important quality for legal systems, or

[6] A case can be considered, in our framework, as a consistent set of factual literals, *i.e.*, of literals that do not contain predicates appearing in the consequents of legal norms (cf. Sartor 1992a).

[7] The treatment of inconsistencies is therefore also a contribution to the development of isomorphic models of law. On isomorphism in the representation of legal knowledge cf., among the others, Gordon (1988), Karpf (1991), Bench-Capon and Coenen (1992).

better is a fundamental ideal for both legislation and legal argumentation. It would be wrong (mystifying) to interpret determinacy as an objective quality of every legal system, so hiding the creative nature of the activities intended to rationalise the legal system in relation to the ideal of determinacy. To understand the nature of those activities, two attitudes towards the legal system must be distinguished:

a. An external (or realistic) point of view, that looks at law as the set of criteria *de facto* used by legal decision makers. The external view is adopted by the "observer" who intends the describe the normative contexts in which legal decisions take place, or to anticipate those decisions.

b. An internal (prescriptivistic) point of view that considers law as the set of criteria that should be used (according to some normative model of legal reasoning) in legal decision-making. This view is adopted by the "participant" who intends to take, justify, or suggest a legal decision.

As far as the (*a*) aspect is concerned, an indeterminacy arises when alternative legal premises (alternative legal ideologies, cf. A. Ross 1958) are effective in society, or may anyway seem susceptible of being accepted by legal decision-makers. In such a situation, certainty of law is compromised, since the citizen is not able to anticipate the legal decisions concerning his behaviour. Legislators, judges, and legal scientists should try to correct this situation (issuing new norms, stating new decisions, proposing new interpretative arguments), but certainly indeterminacy cannot be eliminated by simply (and falsely) "postulating" determinacy.

As far as the (*b*) aspect is concerned, instead, an indeterminacy arises when a person involved in a legal evaluation (typically a judge, but also a legal scientist) is perplexed about the premises he should use to justify its decision or about their ordering. Also this situation can be represented by a knowledge base including incompatible premises and letting the derivation of alternative possible solutions. Nevertheless, if the indeterminacy is relevant for the case at hand the perplexed legal decision-maker cannot suspend his judgement nor adopt arbitrarily one of the possible solutions. This last solution would be against the universalisability principle (cf. Hare 1962; Alexy 1978, 250ff), requiring that equal cases are solved in the same way: different cases, equal in all relevant aspects, could be treated differently by choosing alternative solutions.

Therefore, determinacy of the internal point of view must be pursued, by adding new premises or by establishing new preference relations, until no merely plausible consequences can be derived, as far as the case at hand is concerned (this aspect is well represented in the Lebach example above). This does not mean that there is always a unique right legal solution, but rather that the legal decision-maker should always try to reach determinacy by making the best possible argument — in the framework of the socio-legal-political ideology he endorses, this ideology being under critical examination.

The external and internal views of the legal system can be related with two types of knowledge based systems (cf. Sartor 1993, 57ff.):

a. Systems intended for predicting use, *i.e.*, intended to anticipate future legal decision, or to suggest argument susceptible of being accepted by the legal decision makers.

b. Systems intended for a decisional use, *i.e.*, intended to adopt, or suggest the right legal decision accordingly to the normative ideology of their responsible (this is the case for systems used in administrative decisions).

The reasoning model above described can be used to develop both types of systems:

a. In a predicting system indeterminate knowledge bases model disagreement. Justified consequences derived from such a knowledge base indicate conclusions that will probably be adopted (or accepted) by all the decision-makers whose opinions are faithfully represented in the knowledge base, while merely plausible consequences will be adopted (or accepted) only by some of those decision-makers (they represent alternative possible outcomes and strategies).

b. In a decisional system indeterminate knowledge bases model perplexity. Justified consequences derived from such a knowledge base indicate conclusions upon which the systems responsible should have no doubt (according to his own legal ideology), while merely plausible consequences point to open problems, in which a choice is required between equally justified alternatives. In these last cases the system responsible should extend the knowledge base or decide according to non formalised criteria.

To reach more sophisticate performances, models such as the one here described could be supplemented with tools for suggesting new premises, modifying the old ones, helping in the assessment of preference relations, etc. Nevertheless all those aspects (whose computational handling raises many difficult or unsolved problems) can be treated quite independently from the subject here considered, only concerning the derivation of the consequences of the available premises.

Let us conclude with a short remark concerning the understanding of permissive norms that becomes possible once that inconsistency is accepted in legal systems. In legal theory there is still a lively discussion concerning the role of permissive norms: (*a*) are permissions simply the negation of prohibitions (as in standard deontic logic, where permitted means not forbidden $Pp = \neg O \neg p$), and in this case (*b*) what new content can a permissive norm bring into a consistent legal system? In fact, if p is not forbidden in the system (it holds already that $\neg O \neg p$), then the permissive norm Pp does not say anything new; while is p is forbidden (it holds already that $O \neg p$), than Pp creates an inconsistency, so violating the postulate of consistency. Some authors have also affirmed that permissions simply abrogate (eliminate from the system) pre-existing forbidding norms, so avoiding inconsistency. Nevertheless, according to this last opinion Pp would become redundant immediately after eliminating $O \neg p$, and would have no further effect on the dynamics of the legal system.

Our approach allows us to give a positive answer to both questions above: permissive norms are simply the negation of obligations, and they contribute positively to the content of the legal system. The positive contribution of permissive norms consists in the fact that they contradict forbidding norms, and so prevent the application of the latter (if those norms are inferior to the permissions). For example, constitutional permissions (the so called liberty rights) block subsequent (but inferior) forbidding norms established by legislation. The contradicting role of permissions also explains the relevance of the distinction between weak and strong permission (Alchourrón 1969): a behaviour p is weakly permitted, in a certain normative system Π, if $O \neg p$ is not deducible from Π; p is strongly permitted, in a normative system Π, if Pp is deducible from Π. Strong permission has a different relevance from weak permission since it presupposes a permissive norm, that may block the application of existing, or future, forbidding norms.

9. REFERENCES

Alchourrón, C.E. 1969. Logic of Norms and Logic of Normative Propositions. *Logique et Analyse:* 242-268.

Alchourrón, C.E. 1986. Conditionality and the Representation of Legal Norms. In *Automated Analysis of Legal Texts*, edited by A.A. Martino and F. Socci, 175-186. Amsterdam: North Holland.

Alchourrón, C.E., and D. Makinson. 1981. Hierarchies of Regulations and Their Logic. 1971. In *New Studies on Deontic Logic*, edited by R. Hilpinen, 123-148. Dordrecht: Reidel.

Alexy, R. 1978. *Theorie der juristischen Argumentation.* Frankfurt: Suhrkamp.

Alexy, R., 1980. Die logische Analyse juristischer Entscheidungen. In *Argumentation und Recht*, edited by W. Hassemer, A. Kaufmann, and U. Neumann, 181-212, ARSP, Beiheft 14. Wiesbaden: Steiner.

Bench-Capon, T.J.M., and F.P. Coenen. 1992. Isomorphism and Legal Knowledge Based Systems. *Artificial Intelligence and Law* 1: 65-86.

Brewka, G. 1986 Tweety - Still Flying: Some Remarks on Abnormal Birds, Applicable Rules and Default Prover. In *Proceedings AAAI-86*, Philadelphia (PA), 8-12.

Brewka, G. 1991a. Belief Revision in a Framework for Default Reasoning. In *The Logic of Theory Change. Workshop, Konstanz, Germany, October 13-15, 1989. Proceedings*, edited by A. Fuhrmann, and M. Morreau, 206-222. Berlin: Springer.

Brewka, G. 1991b. *Nonmonotonic Reasoning. Logical Foundations of Commonsense.* Cambridge: Cambridge. University Press.

Dworkin, R.M. 1977. *Taking Rights Seriously.* London: Duckworth.

Gelfond, M., and V. Lifschitz. 1990. Logical Programs with Classical Negation. In *Proceedings of the Seventh International Logic Programming Conference,* 579-597. New York: MIT Press.

Gordon, T.F. 1987. OBLOG-2: A Hybrid Knowledge Representation System for Defeasible Reasoning. In *The First International Conference on Artificial Intelligence and Law. Proceedings of the Conference*, 231-239. New York: ACM Press.

Gordon, T.F. 1988. The Importance of Nonmontonicity for Legal Reasoning. In *Expert Systems in Law: Impacts on Legal Theory and Computer Law*, edited by H. Fiedler, F. Haft, and R. Traunmüller, 111-126. Tübingen: Attempto.

Gordon, T.F. 1989. Issue Spotting in a System for Searching Interpretation Spaces. In *The Second International Conference on Artificial Intelligence and Law: Proceedings of the Conference,* 157-164. New York: ACM Press.

Gordon, T.F. 1991. An Abductive Theory of Legal Issues. *International Journal of Man-Machine Studies* 35: 95-118.

Hare R.M. 1962. *Freedom and Reason.* Oxford: Claredon Press.

Karpf, J. 1991. Isomorphic Modelling of Statutory Law. *Informatica e diritto* 12: 89-111.

Kowalski, R. A., and F. Sadri. 1990. Logic Programming with Exceptions. In *Proceedings of the Seventh International Logic Programming Conference,* 598-613. New York: MIT Press.

Poole, D.L. 1985. On the Comparison of Theories: Preferring the Most Specific Explanation. In *Proceedings IJCAI 1985,* 144-147.

Poole, D.L. 1987. Variables in Hypotheses. In *Proceedings IJCAI 1987.*

Poole, D.L. 1988. A Logical Framework for Default Reasoning. *Artificial Intelligence* 36: 27-47.

Prakken, H. 1991a. A Tool to Model Disagreement in Law: Preferring the Most Specific Argument. In *The Third International Conference on Artificial Intelligence and Law. Proceedings of the Conference*, 165-174. New York: ACM Press.

Prakken, H. 1991b. *A Formal Theory about Preferring the Most Specific Argument.* Rapport IR-259. Amsterdam: Vrije Universiteit, Faculteit der Wiskunde en Informatica.

Prakken, H. 1992. An Argumentation Framework in Default Logic. *Annals of Mathematics and Artificial Intelligence.* To appear.

Reiter, R. [1980] 1987. A Logic for Default Reasoning. In Read*ings in Nonmonotonic Reasoning*, edited by M.L. Ginsberg, 68-93. Los Altos (California): Morgan Kaufmann. First published in *Artificial Intelligence* 13: 81-132.

Reiter, R., and G. Criscuolo. [1981] 1987. On interacting defaults. In *Readings in Nonmonotonic Reasoning,* edited by M.L. Ginsberg, 94-100. Los Altos (California): Morgan Kaufmann. First published in *Proceedings of the Seventh International Joint Conference on Artificial Intelligence - IJCAI 1981,* Vancouver, 270-276.

Ross, A. 1958. *On Law and Justice.* London: Steven and Sons.

Ross, W.D. 1930. *The Right and the Good.* Oxford: Claredon Press.

Ross, W.D. 1939. *Foundations of Ethics.* Oxford: Claredon Press.

Sartor, G. 1992a. Normative Conflicts in Legal Reasoning. *Artificial Intelligence and Law.* To appear.

Sartor, G. 1992b. *Defeasible Arguments from Inconsistent Knowledge bases: A Logic Programming Approach.* Technical Report CIOC - 90. Bologna: CIOC.

Sartor, G. 1993. *Artificial Intelligence in Law.* Oslo: Tano.

Simari, G.R., and R.P. Loui. 1992. A Mathematical Treatment of Defeasible Reasoning and its Implementation. *Artificial Intelligence* 53: 125-157.

Toulmin, S. 1958. *The Uses of Argument.* Cambridge (U.K.): Cambridge University Press.