Two Short Notes on Argument: (1) Corrected Specificity (2) Dialectical Refinement

Ronald P. Loui

Springfield, IL and Bay Village, OH
r.p.loui@gmail.com

Abstract. In honor and celebration of Guillermo Simari's milestone two ideas are presented: one old, and one new. The first is the central, revised definition of specificity, from the unpublished *Corrigendum to Poole's Rules and a Lemma of Simari-Loui.* This fulfills a promise made in the footnote of Simari-Loui, in [19]. The second is a new start on formalizing the logic of dialectical dialogue, assuming argument on defeasible reasons as settled work. The phenomena of concern are (1) refinement of predicate sense and (2) refinement of reference, in the face of counter-argument. This is the first attempt to commit these thoughts to paper after the problem was raised during a visit to Universitad Nacional del Sur in Bahía Blanca, Argentina in the early 1990s. [1]

1 Specificity

Poole's rule holds one argument, which he calls a theory, to be more specific than another if there is some way of activating the first without activating the second; and not vice versa.

Simari-Loui [19] made central use of this rule, but as it went to press, an important lemma of the article is wrong. It states that checking specificity is equivalent to checking that the antecedents of the less specific theory can be derived from the antecedents of the more specific theory.

The lemma states that a theory T_1 is more specific than T_2 just in case for every x, an antecedent of a rule used in T_2, x can be defeasibly derived from K_N, the necessary evidence, T_2's rules, and the antecedents of T_1's rules:

[1] The authors on the original report from which the first note is derived were: R. Loui, Supported by NSF R-9008012; J. Norman, Seyfarth, Shaw, Fairweather, and Geraldson, Chicago, IL; also fellow of The Center for Intelligent Computer Systems, Washington University; K. Stiefvater, Supported by NSF CDA-9102090; A. Merrill, Supported by NSF CDA-9123643; A. Costello, Supported by NSF R-9008012 and CDA-9102090; J. Olson, Supported by NSF R-57135A; Department of Computer Science, Washington University, St. Louis; Guillermo Simari had already returned to Argentina when this work began, but was engaged in the discussion

$$(\forall x \in An(T_2))(K_N \cup An(T_1) \cup T_2 \mathrel{|\!\sim} x).$$

In the case of

$$< \{Penguin(O) \ \succ\!\!-\ \neg Flies(O)\}, \neg Flies(O) >$$

versus

$$< \{Bird(O) \ \succ\!\!-\ Flies(O)\}, Flies(O) >,$$

for example, $Bird(O)$ can be defeasibly derived from $Penguin(O)$. And in

$$< \{Cat(G) \ \succ\!\!-\ Aloof(G), \quad Aloof(G) \ \succ\!\!-\ \neg LikesPeople(G)\}, \neg LikesPeople(G) >$$

versus

$$< \{Cat(G) \ \succ\!\!-\ LikesPeople(G)\}, LikesPeople(G) >,$$

$Cat(G)$ defeasibly derives both $Cat(G)$ and $Aloof(G)$.

This last example is a simple counterexample to the lemma. If the lemma were right, then not only is the theory for $LikesPeople(G)$ more specific than the theory for $\neg LikesPeople(G)$ which we have just said is desirable, but also vice versa, which is not desirable. Not only is this intuitively undesirable, but it also violates the antisymmetry of specificity. The required relationship between antecedents is necessary for specificity, but not sufficient.[2]

Correcting this error requires [3] first that the definition of specificity be repaired. As it stands, it is vulnerable to some counter-intuitive behavior. This counter-intuitive behavior plagues Poole's rule generally, not just our use of it. The argument, for example,

$$\begin{array}{rcl} A & -\!\!< & B \wedge C \\ B & -\!\!< & D \\ C & -\!\!< & E \end{array}$$

should be more specific than

$$\begin{array}{rcl} \neg A & -\!\!< & B \\ B & -\!\!< & D \end{array}$$

[2] Where the alleged proof goes wrong is quite easy to see. When there is specificity, every antecedent of the weaker theory can be derived from every antecedent of the stronger theory, but not necessarily from an activator of the stronger theory (an activator is a sentence of the proscribed contingent kinds, which allows a theory's conclusion to be derived, using the necessary evidence, K_N, and the theory's rules). That is, it may be possible to activate the theory without allowing *all* of its antecedents to be defeasibly derived, which is just plain to see.

Henry Prakken [15] has noticed that Theorem 4.16 of the paper is also in error. An example can be given in which two arguments defeat each other. Specificity is antisymmetric, but not defeat. The proof goes sour at "the same is true for defeat."

[3] Prakken (in [15]) corrects the error in an equally intuitive way. He suggests restricting the search for activators to those that activate theories in exactly the same way that the evidence does (this must be done carefully if there are multiple derivations). This fix for Poole also appears to allow a pruning lemma; however, it

But it is not more specific according to the unadulterated Poole definition, because of two separate flaws in the definition.

The first reason is that $E \wedge (B \vee \neg C)$ allows the former theory to be activated without activating the latter. This prevents the desired conclusion that the former theory is more specific. What is basically wrong here is that the weaker theory should be allowed to use the defeasible rule $E \ \succ\!\!-\ C$, in order to derive (defeasibly) B.

To see the second flaw, consider $E \wedge (A \vee \neg C)$, which is also disjunctive, but this time uses the disjunction to derive the theory's ultimate conclusion. This is essentially a side-stepping of the non-triviality condition for activators. The non-triviality condition should be strengthened.

Fixing the flaws in Poole's rule is important because this kind of comparison is the comparison used in the Yale Shooting Problem arguments (Hanks-McDermott [6]), as exhibited among the examples in the paper by Simari-Loui [19].

The rule also suffers in an example reminiscent of Royal Elephants (Simonet [20]), consider

$$\begin{array}{rcl} D & -\!\!< & B \wedge C \\ B \wedge C & -\!\!< & A \end{array}$$

compared with what should be an inferior theory:

$$\begin{array}{rcl} \neg D & -\!\!< & B \\ B & -\!\!< & E \end{array}$$

Neither is more specific by Poole's rule. The example appears to require right-weakening of rules: allowing rules to be derived from rules by weakening the consequent.[4]

But right-weakening is notoriously problematic. [5]

[4] For example,

$$\text{If } A \ \succ\!\!-\ B \wedge C, \text{ then } A \ \succ\!\!-\ C;$$

then the theory

$$\begin{array}{rcl} D & -\!\!< & B \wedge C \\ B & -\!\!< & E \\ C & -\!\!< & A \end{array}$$

would be the defeater of the weaker theory. The first theory does not defeat the third, but with right-weakening, it allows the construction of a third theory which directly accounts for the considerations used in the weaker theory, and which ought to be more specific.

[5] When there are competing arguments, such as

$$B \ -\!\!<\ A$$

versus

$$\neg B \ -\!\!<\ A,$$

a rule can be right-weakened with an arbitrary dilution, such as

$$A \ \succ\!\!-\ B, \text{ therefore, } A \ \succ\!\!-\ B \vee C,$$

allowing an argument for C. This cannot be done without right-weakening, since

The proposed development for specificity, which fixes both kinds of counterintuitive behavior due to disjunction, which allows proper treatment of the last example, and which allows a proper pruning lemma, is as follows.

Let $\Delta^{\downarrow} \subseteq L^2$, $K_N \subseteq L$, and $K_C \subseteq L$ be, respectively, instantiated defeasible rules, necessary (background) knowledge, and contingent knowledge (evidence), referring to a first-order language L.

An argument, $< T, h >$ has $T \subseteq \Delta^{\downarrow}$ and h derivable from $T \cup K_N \cup K_C$, where derivation may use rules of FOL, and a modus ponens for defeasible rules, i.e.

$$\frac{\vdash p}{\mid\!\sim p}$$

$$\frac{\mid\!\sim p, \quad p \succ q}{\mid\!\sim q}$$

Also, T is minimal; no proper subset of T allows derivation of h.

A rule, R, is a *top rule* of argument $< T, h >$ just in case its consequent, $Con(R)$, is not needed for the derivation of anything but the argument's conclusion. Because the rules used in arguments are a minimal set, that is equivalent to saying that the antecedent of any rule can be derived (from evidence) using rules other than this top rule.[6]

Definition 1. **Top**$(R, < T, h >)$ *iff for every r in T, $An(r)$ can be defeasibly derived from $K_N \cup K_C$ using $T - \{R\}$.*

Example 1. $B \succ C$ is a top rule in the argument from A to C, using: $A \succ B$, and $B \succ C$.

Let Δ be a set of defeasible rules, let h be a sentence in the language, L, and let A be a set of sentences in L.

A finite sequence of sentences, $< B_1, \ldots, B_n >$ is a *consistent defeasible derivation (CD-derivation)* of h from A using rules Δ just in case h is derived by A's activating ground instances of rules, and the set of all intermediate sentences is consistent in L.

---

bad, since the argument

$$\begin{array}{rcl} C & \prec & (B \vee C) \wedge \neg B \\ B \vee C & \prec & A \\ \neg B & \prec & A \end{array}$$

has counterargument

$$B \prec A,$$

but if $A \wedge D \succ \neg B$, the arbitrary dilution can actually be supported. Instead, fix the rule for specificity so that it treats the problem properly without requiring right-weakening.

[6] It is not sufficient to say that a rule is top just in case it participates in eliminating a literal from the goal clause: consider $< \{Q \succ R, S \succ T\}, R \wedge T >$ which is an argument for $R \wedge T$; only the latter is a top rule; otherwise, the argument

Definition 2. *$< B_1, \ldots, B_n >$ is a* **CD-derivation of** *h* **from** *A iff*

1. *$B_n = h$;*
2. *For each B_i, either*
   a. *$\{B_j \mid j < i\} \cup A \vdash B_i$; or*
   b. *for some ground instance of a rule R in Δ, $An(R) = B_j$ for some $j < i$ and $B_i = Con(R)$;*
3. *$\{B_i \mid i \leq n\} \nvdash \bot$.*

Example 2 (continued). $< A, B, C >$ is a CD-derivation of C from A using: $A \succ B$ and $B \succ C$.

Definition 3. *There is a* **CD-derivation** *of h from A* **using Δ with all top rules** *of $< T, h >$ just in case*

1. *there is a CD-derivation of h from A using Δ;*

2. *for each x such that $Top(x, < T, h >)$, there is no CD-derivation of h from A using $\Delta - \{x\}$.*

$< T_1, h_1 >$ is more specific than $< T_2, h_2 >$ just in case some legitimate sentence activates T_2 for h_2 without activating T_1 for h_1, using CD-derivations from *the two theories' combined set of rules*, using *every top rule* of T_2; and there is no such *asymmetric activator* of T_1 for h_1 that does not also activate T_2 for h_2.

The requirement to use top rules is just a strengthening of Poole's nontriviality condition that asymmetric activators do not activate the theory simply by FOL rules, side-stepping the defeasible rules. The combination of theories is more profound. It signals the importance of pairwise comparison as opposed to an n-wise, holistic evaluation of merit (which is what Geffner-Pearl tends toward [5]) on one extreme, or a conception of specificity as intrinsic, perhaps even measurable (which is what the algebra of Simari-Loui suggests), on the other extreme. Transitivity no longer holds of specificity.

That is,

Definition 4. *e is an* **asymmetric activator** *of $< T_1, h_1 >$* **but not** *$< T_2, h_2 >$ just in case*

1. *there is some CD-derivation of h_1 from $K_N \cup \{e\}$ using $T_1 \cup T_2$ with all top rules of $< T_1, h_1 >$;*

2. *there is no CD-derivation of h_2 from $K_N \cup \{e\}$ using $T_1 \cup T_2$ with all top rules of $< T_2, h_2 >$.*

Let eAA_{inotj} symbolize that e is an asymmetric activator of $< T_i, h_i >$ but not $< T_j, h_j >$

Definition 5. $< T_1, h_1 >$ *is* **more specific** *than* $< T_2, h_2 >$ *just in case*

1. there is some asymmetric activator in S_C *of* $< T_2, h_2 >$ *but not of* $< T_1, h_1 >$*;*
i.e., there is some $e \in S_C$ *s.t.* eAA_{2not1}*; and*

2. there is no asymmetric activator in S_C *of* $< T_1, h_1 >$ *but not* $< T_2, h_2 >$*;*
i.e., there is no $e \in S_C$ *s.t.* eAA_{1not2}.

Given $< T_1, h_1 >$ and $< T_2, h_2 >$, use the symbolization $e \mapsto_{T_i} f$ to assert the existence of a CD-derivation of f from $K_N \cup e$ using $T_1 \cup T_2$ with all top rules of $< T_i, h_i >$. $e \mapsto f$ if there is a CD-derivation at all from $K_N \cup e$ using $T_1 \cup T_2$, not requiring use of top rules. Note that $\mapsto$ and $\mapsto_{T_i}$ are defined only for a pair of theories being compared.

Note also that

eAA_{1not2} just in case

1. $e \mapsto_{T_1} h_1$;
2. $e \not\mapsto_{T_2} h_2$;

The new pruning lemma makes use of both the top rule restriction and the union of theories when checking activation.

To find whether there is an asymmetric activator of $< T_1, h_1 >$ but not $< T_2, h_2 >$ it is usually sufficient to check whether the conjoined antecedents of the top rules of $< T_1, h_1 >$ is an asymmetric activator of $< T_1, h_1 >$ but not $< T_2, h_2 >$. For simplicity, first assume that h_1 can be derived from the conjoined consequents of top rules in $< T_1, h_1 >$, the last step in deriving the theory's conclusion uses just K_N and the consequents of top rules; that is, intermediate conclusions from consequents of non-top rules are used only to derive antecedents of later rules.

Lemma 1 (restricted pruning). *For any arguments* $< T_1, h_1 >$ *and* $< T_2, h_2 >$*, there exists an asymmetric activator of* $< T_1, h_1 >$ *but not* $< T_2, h_2 >$ *just in case the following has the property* AA_{1not2}*:*
$Conjoin_\Gamma(An(R_i))$,
where $\Gamma = \{R_i : Top(R_i, < T_1, h_1 >)\}$*;*
under the assumption that h_1 *can be derived from the conjoined consequents of top rules in* $< T_1, h_1 >$*, i.e.,*
$Conjoin_\Gamma(Con(R_i)) \mapsto h_1$.

Proof. 1. First consider the case where $< T_1, h_1 >$ has a single top rule. Suppose there is an e such that eAA_{1not2}. That is, $e \mapsto_{T_1} h_1$ and $e \not\mapsto_{T_2} h_2$. Let $e' = An(R)$, where by assumption, R is the only rule in T_1 such that $Top(R, < T_1, h_1 >)$.
Clearly, $e' \mapsto_{T_1} h_1$, since we assume $Conjoin_\Gamma(Con(R_i)) \mapsto h_1$. So

Assume to the contrary that $e' \mapsto_{T_2} h_2$. Recall that $e \mapsto_{T_1} h_1$, so $e \mapsto e'$, since e' is just the antecedent of the top rule which must be used. Chain this with the assumption that $e' \mapsto_{T_2} h_2$, and get $e \mapsto_{T_2} h_2$. But e is supposed to be an asymmetric activator of $< T_1, h_1 >$ but not $< T_2, h_2 >$. This is a contradiction. So e' must be an asymmetric activator.
2. Next, consider the case where $< T_1, h_1 >$has multiple top rules. Let e' be $Conjoin_\Gamma(An(R_i))$. The same argument applies, but it is no longer obvious that $e \mapsto e'$.
Assume $e \mapsto_{T_1} h_1$. Consider any rule, $R \in \Gamma$, i.e., any R such that $Top(R, < T_1, h_1 >)$. Show that $e \mapsto An(R)$. This suffices to show that $e \mapsto e'$.
Assume to the contrary that $e \not\mapsto An(R)$. This is a contradiction, because $e \mapsto_{T_1} h_1$ requires that any CD-derivation of h_1 from $K_N \cup e$ using $T_1 \cup T_2$ use all of the top rules of $< T_1, h_1 >$, including R, and it is impossible to use R without deriving $An(R)$.

Now relax the assumption regarding the derivability of h_1 from consequents of top rules. The full lemma is that whenever there is an asymmetric activator of $< T_1, h_1 >$ but not $< T_2, h_2 >$, an asymmetric activator is formed by conjoining (1) the antecedents of top rules with (2) the conditional whose antecedent (2a) conjoins the consequents of non-top rules, and whose consequent (2b) is the claim supported by the argument.

Lemma 2. . *For any arguments* $< T_1, h_1 >$ *and* $< T_2, h_2 >$*, there exists an asymmetric activator of* $< T_1, h_1 >$ *but not* $< T_2, h_2 >$ *just in case the following has the property* AA_{1not2}*:*

$$Conjoin_\Gamma(An(R_i)) \wedge (Conjoin_\Gamma(Con(R_i)) \supset h_1)$$

where Γ *is as before.*

Proof. The proof again begins: suppose some eAA_{1not2}; i.e., $e \mapsto_{T_1} h_1$, $e \not\mapsto_{T_2} h_2$. e' is the conjunction of top rule antecedents, and $e''$is e' conjoined with the material conditional as above. We want $e''AA_{1not2}$. Clearly $e'' \mapsto_{T_1} h_1$. Show $e'' \not\mapsto_{T_2} h_2$. Suppose it did, i.e., $e'' \mapsto_{T_2} h_2$: $e \mapsto_{T_1} h_1$, so $e \mapsto e''$ since all top rules must be used and e'' is the weakest sentence that activates all top rules and also derives h_1. This property of being weakest is the key observation. Chaining, $e \mapsto_{T_2} h_2$, but this is a contradiction. So $e''AA_{1not2}$ if for any e, eAA_{1not2}.

The revised rules have been implemented twice. A version with an underlying first-order logic, on which this section focuses, was implemented in C and is quick; it is primarily limited by the underlying resolution theorem-prover.[7]

[7] The main programmers were Adam Costello, Andrew Merrill, and Ronald Loui; the 92k, 3300 lines of source code contain a dedication to the late computer scientist, Eugene Nathan Johnson, c. 1944 – 1984; the program is called "nathan".

The second version is in LISP with a restricted propositional language (just negation of atomic formulae and conjunction), with provision for analogical (case-based) reasoning, and with additional features peculiar to certain forms of legal reasoning.[8]

2 Refinement

One of the clear purposes of dialogical dialectic is the clarification of reference and restriction of predication. A mathematical model of such dialogue is sketched here.

Twenty-five years after Simari-Loui and related works, the AI logic community has capably formalized the dynamics of arguments built upon a given set of claims and rules. This was the inherited framework, mainly from Doyle [4] and Reiter [18]. Specificity, priority, undercutting and reinstatement occupied much of the attention of research for decades. First-order predicate quantification was also a hindrance. It took some time to gain widespread acceptance of the dialectical pro-con process, its ampliativity and potential non-determinism, its dependence on search, and the non-monotonicity of process that was distinct from non-monotonicity of syntax. ([11], [9], see e.g., Baroni, Cerutti, Giacomin, and Simari, [2])

The phenomena that are well researched concern the rules 'if p then q', 'if p and r then defeasibly not q', 'if s then not(if p then defeasibly q)', represented in various ways.

Even as these investigations were gaining speed, this author was puzzling over what seemed to be a more common phenomenon in dialogical argument. That phenomenon was the clarification of claims, the refinement of parts of claims, through argument.

There are two distinct kinds of clarification considered next. A third kind is notationally a bit farther afield.

Of the two we can easily address, first is the clarification of reference.

If a person makes a claim '$P(a)$', or 'All $x \in A$ $P(x)$', or 'Asians do well at Harvard', another person may respond in dialogue that 'A' is not referentially specific (we can quibble over whether this is reference or antecedent predication, but anyone reading this note should have already read Quine on this subject, [17]). Or a person might respond that 'P' is not sufficiently qualified or restricted in extent. The idea is that if 'Asians' were more carefully tailored, or 'does well at Harvard' were more carefully tailored, then the claim could be provisionally accepted and dialogue could continue. If not refined, however, the dialogue would enter an argument subdialogue.

The purpose of adversarial engagement is to improve the claim. It is a language game on the semantics, essentially the translation of the shorthand

[8] The main programmer was Jon Olson; its 52k, 2000 lines of source code are called "lmnop" after the initials of the last names of its designers. See Loui-Norman et al.[12]

natural language into formal logical symbols. David Lewis is associated with this kind of dialogical referential refinement, e.g., in [8], and it is related to the non-dialectical desire in natural language dialogue to agree on anaphora (consider Webber [21] and the literature that followed on the subject).

Note that $A(x) \succ\!\!- P(x)$ can be refined simply by qualification in existing rule-based frameworks, $A(x) \wedge Q(x) \succ\!\!- P(x)$, with "Asians" now eliding the qualified "Asians who are American-born" though $P(x)$ is not as easily refined with an antecedent qualifier, e.g. "Asians do well at Harvard academically". Apparently, one simply has to rewrite the predicate, $A(x) \succ\!\!- Q(x)$, and $P(x) \succ\!\!- Q(x)$, perhaps even $Q(x) \succ\!\!- P(x)$.

Although this appears at first to be a problem of language rather than logic, it may provide a model of one of the more elusive aspects of logic: namely, its interface with language. First, the initial expression of claims may be limited by the finiteness of locution or representation. HP Grice comes to mind here, as there may be a limit to what one can say during one's turn. But it may simply be an idealization that all formal expression of claims be maximally precise at initial claiming. Second, there may be a legitimate logical dynamics in the interpersonal agreement over semantics when two persons, or more, enter into the meeting of minds. Predicate refinement, or referential refinement, may be an important part of that meeting, whether as revision of representation or revision of the content of claims. Third, there is the concept of open-texture, what AI might call underspecification and ex-post learning, where the assumption is that revision of initial representation will take place, even if the initial representation leaves no gap between natural language and formal symbolization. This may be because of finite expression, finite envisionment, or the limit of multi-party agreement. For the latter, it may be that underspecification is necessary in order to reach initial agreement, i.e., the basis of the claim from which it derives its truth-like authority or assertability.

In tandem with the development of formal defeasible dialectical argument in AI, the AI and Law community has investigated the nature of open texture of predicates. There the semantics of a linguistic element ("vehicle in the park") undergoes revision as cases are decided in time that affect what is considered a 'vehicle' and what is considered to be 'in the park'. (Hart [7]) A good example of this is the original desire to prohibit vehicles from the park, with the explicit exception of parade vehicles, hence, 'no vehicles in the park' with some exceptions. But over time, vehicles may include drones, and being in the park may implicate the airspace over the park at some lower altitudes. Further specification is part of the semantics of the edict, and the semantics forsees a process of revision over time as the world evolves, not just revision through argument, as hard cases are considered with the originally existing concepts.

Regardless of how one wants to divide the responsibility between logic and linguistics, there is clearly a challenge to find a framework for thinking about dialectical dialogue that results in reference and predication refinement.

The basic addition is a sequential index to a predicate and a term, so P becomes P_1, and A becomes A_1, a becomes a_1. The initial claim $P_1(A_1)$. So P', P'', P''' can simply be P_1, P_2, P_3.

But refinement is forced by the adversarial dialogue. In the case of "Asians do well at Harvard",[9] one counter move would be "Foreign-born Asians do not always do well at Harvard." $not(A(x) \wedge FB(x) \succ P(x))$. This could be an observation that may not be disputed, or can be subject to further dispute. If it is sufficiently justified or jointly presumed, it forces revision of the original claim. Note that a single counterexample might not force refinement, since the claim is presumably an expression of a defeasible rule.

In any case, $A_2 = A_1 - FB$. We can express A_2 alternately as $A_{(-FB)}$.

A slightly different refinement would be restriction to a subclass, i.e., "US-born Asians do well at Harvard" in which case we have $A(x) \wedge US(x) \succ P(x)$, where we could write the refinement as $A_{(\wedge US)}$.

If the counter move is that "Asians do well at Harvard only in terms of academics", the refinement is on the predicate in the consequent. The move from $A(x) \succ P(x)$ to $A(x) \succ Q(x)$ is not interesting except insofar as $P(x)$ and $Q(x)$ were so closely related that they could have been mistaken in the initial claim.

A predicate P is refined to a **polysemic correlate** Q, when $P(x) \succ Q(x)$ and $Q(x) \succ P(x)$.

This could be statistical, but as we know in defeasible reasoning, not all defeasible rule connections are adopted on statistical grounds, nor are all statistically strong correlations accepted as defeasible rules in an argument game. We can express P_2 in this case as $P_{(>Q)}$.

A related counter move is negative rather than positive, e.g., "Asians do not do well at Harvard socially". Here, the relation between P and Q contains a negative assertion, $P(x) \succ Q(x)$where $Q(x) = R(x) \wedge S(x)$, but $A(x) \succ R(x)$, $not(A(x) \succ S(x))$. We could express this as $P_{(-S)}$, but there is a more general phenomenon here.

An **implicature of the consequent** is a property (predicate) immediately and solely (defeasibly) derivable from the consequent of a defeasible rule.

There are senses of "doing well at Harvard" that are normally implied, and are not correlates, but are essential parts of each sense of the concept or phrase that may be in play. One can do well musically. One can do well relative to predicted performance. One can do well in the sense of doing good, i.e., doing charitable works, which may be the way that one does well in an ethics-based assessment of performance. The logician of course simply requires that the correct sense be attached to a predicate Q^*, where Q^* is ideally the result of the dialectic that forces refinements $Q_1, Q_2, Q_3 \ldots Q^*$. But such a requirement does not permit discussion of how one gets from Q_1 to Q^*.

For our $Q(x)$, there may be several implicatures of the consequent:

$$Q(x) \succ T(x),$$

[9] This example is chosen with careful consideration of the author's own experience.

$$Q(x) \succ U(x),$$
$$Q(x) \succ V(x),$$
etc.

The adversarial argument may sever each individually: "Asians at Harvard do not generally ski well", "Asians at Harvard do not normally write senior theses well", etc. The response really does depend on dialectic, or at least a process of statistical argument, because it may in fact be defensible that Asians at Harvard write senior theses well (numerically, stereotypically, or even axiomatically[10]).

In any case, the refinement from these kinds of move is $Q_{(-T)(-U)(-V)}$, where we choose not to distinguish between a refinement that loses a logical entailment and a refinement that loses implicature. The refinement of the consequent in this way also could have an effect on the set of rules upon which arguments are constructed.

One way to do a revision would be to retain the symbol Q, but remove the severed implicatures from the rule-base. But a more literal and manipulable notation would add undercutters to the refined predicate:

$$\begin{aligned} P(x) &\succ Q_{(-T)(-U)(-V)}(x) \text{ is accepted, so long as} \\ & not(Q_{(-T)(-U)(-V)}(x) \succ T(x)) \\ & not(Q_{(-T)(-U)(-V)}(x) \succ U(x)) \\ & not(Q_{(-T)(-U)(-V)}(x) \succ V(x)). \end{aligned}$$

are added to the rule base.

A third kind of refinement is refinement of a referential term, rather than predicate. So a particular Asian at Harvard, e.g. *JeremyLin*, may be asserted to have done well at Harvard. If the objection is that there are a few people named 'Jeremy Lin', that may not be an interesting argument. If the objection is that we know Mr. Lin excelled on the basketball court at Harvard, but we have no idea what his academics were like, then we can write $P(x) \succ Q_{(-T)}(x)$, with $not(Q_{(-T)}(x) \succ T(x))$, a removal of the sense of doing well normally associated. But if the objection is that *JeremyLin*, the freshman, was different from *JeremyLin*, the senior, we have a different sort of problem. What is needed here is the indexicals that are further associated with context refinement. To provide a notation that would support such indexicals, we might need to represent individuals as fluents (e.g., McCarthy [13], McCarthy and Hayes [14], Baker [1]), then represent the restriction to slices of fluents.

For decades, it has bothered this author that most of what dialectic seems to focus on, in real life conversation, bears no relation to the logic of argument celebrated in AI knowledge representation and reasoning, and in AI and Law research.

[10] Axiomatically, for example, if one wanted to take the definition of doing well at Harvard to be the specific identifiable way that Asians do: a jarring but not uncommon maneuver in the web of temporarily constructed dialogical meaning.

There has always been work on argument and dialectic in AI and natural language, as well as in the wider world of informal logic. The question of semantics is ably approached through the competing analogies and disanalogies that formalize reasoning from precedent on open-textured predicates (see Loui-Norman et al.[12], Loui-Norman [10], Praken-Sartor [16], Loui [3]). But there, the dominant dynamics is from $P(x) \succ\!\!- Q(x)$ to $not(P(x) \succ\!\!- Q(x))$ with no representation, as far as this author has seen, of the refinement that results from dialectic. The impact of adversarial dialogue on claims may be constructive: it often makes the claims sharper, the concepts finer, and the references more specific. It is not just rule qualification, though some of it could be represented that way. It is not just unrelated symbol substitution. Because the symbol substituted bears, at the very least, correlate and implicature relations to the original symbol. One can look at the failure to sustain $P(x) \succ\!\!- Q(x)$ and try a new argument based on $P(x) \succ\!\!- Q'(x)$, i.e., $Q_{(-T)}(x)$, but understanding Q's relation to T may be helpful in seeing how the whole dialogue fits together.[11]

References

1. Andrew B Baker. Nonmonotonic reasoning in the framework of situation calculus. *Artificial Intelligence*, 49(1-3):5–23, 1991.
2. Pietro Baroni, Federico Cerutti, Massimiliano Giacomin, and Guillermo Ricardo Simari. *Computational Models of Argument: Proceedings of COMMA 2010*, volume 216. Ios Press, 2010.
3. Peter Boltuc, Ronald P. Loui, Felmon Davis, and D. E. Wittkower. Scientific and legal theory formation in an era of machine learning: remembering background rules, coherence, and cogency in induction. 2015.
4. Jon Doyle. A truth maintenance system. *Artificial intelligence*, 12(3):231–272, 1979.
5. Hector Geffner and Judea Pearl. Conditional entailment: Bridging two approaches to default reasoning. *Artificial Intelligence*, 53(2-3):209–244, 1992.
6. Steve Hanks and Drew McDermott. Nonmonotonic logic and temporal projection. *Artificial intelligence*, 33(3):379–412, 1987.
7. Herbert Lionel Adolphus Hart, Herbert Lionel Adolphus Hart, and Leslie Green. *The concept of law.* Oxford University Press, 2012.
8. David Lewis. Convention cambridge. *Mass.: Harvard UP*, 1969.
9. Ronald P. Loui. Process and policy: Resource-bounded nondemonstrative reasoning. *Computational Intelligence*, 14(1):1–38.
10. Ronald P Loui and Jeff Norman. Eliding the arguments of cases. 1997.
11. Ronald Prescott Loui. Defeat among arguments: a system of defeasible inference. *Computational Intelligence*, 3:100–106, 1987.
12. Ronald Prescott Loui, Jeff Norman, Jon Olson, and Andrew Merrill. A design for reasoning with policies, precedents, and rationales. In *ICAIL*, 1993.
13. J Mac Carthy. Situations and actions and causal laws. standford artificial intelligence project memo 2, 1963.

[11] As David Makinson once told me in a cafe in Paris named after St. Louis's own Josephine Baker, the important step is to invent the notation.

14. John McCarthy and Patrick J Hayes. Some philosophical problems from the standpoint of artificial intelligence. In *Readings in nonmonotonic reasoning*, pages 26–45. Morgan Kaufmann Publishers Inc., 1987.
15. Henry Prakken. An argumentation framework in default logic. *Annals of Mathematics and Artificial Intelligence*, 9:93–132, 1993.
16. Henry Prakken and Giovanni Sartor. Modelling reasoning with precedents in a formal dialogue game. In *Judicial applications of artificial intelligence*, pages 127–183. Springer, 1998.
17. Willard Van Orman Quine et al. Logic and the reification of universals. *From a logical point of view*, 6, 1953.
18. Raymond Reiter. A logic for default reasoning. *Artif. Intell.*, 13:81–132, 1980.
19. Guillermo Ricardo Simari and Ronald Prescott Loui. A mathematical treatment of defeasible reasoning and its implementation. *Artif. Intell.*, 53:125–157, 1992.
20. Geneviève Simonet. Nonmonotonic inference rules for multiple inheritance with exceptions. *Proceedings of the IEEE*, 74:1345–1353, 1986.
21. Bonnie Lynn Webber. *A formal approach to discourse anaphora.* Routledge, 2016.