The Objective Conception of Context and Its Logic

CHRISTOPHER MENZEL
Department of Philosophy, Texas AhIM University, College Station, Texas 77843-4237 U.S.A.
(email: cmenzel@tamu.edu)

Abstract. In this paper, an “objective” conception of contexts based loosely upon situation theory is developed and formalized. Unlike “subjective” conceptions, which take contexts to be something like sets of beliefs, contexts on the objective conception are taken to be complex, structured pieces of the world that (in general) contain individuals, other contexts, and propositions about them. An extended first-order language for this account is developed. The language contains complex terms for propositions, and the standard predicate ‘*ist*’ that expresses the relation that holds between a context and a proposition just in case the latter is true in the former. The logic for the objective conception features a “global” classical predicate calculus, a “local” logic for reasoning within contexts, and axioms for propositions. The specter of paradox is banished from the logic by allowing ‘*ist*’ to be nonbivalent in problematic cases: it is not in general the case, for any context c and proposition p, that either *ist(c,p)* or *ist*$(c, \neg p)$. An important representational capability of the logic is illustrated by proving an appropriately modified version of an illustrative theorem from McCarthy’s classic Blocks World example.

Key words: context, situation theory, proposition

1. The Objective Conception of Context

Since McCarthy (1987) first emphasized its importance for AI, there has been a rapidly growing body of research in the AI community concerning the notion of context.[1] A review of the literature reveals a wide variety of conceptions of context. However, despite this variety, it seems possible to divide these conceptions into two broad categories: *subjective* and *objective*. On the subjective conception, contexts are something like logical theories, i.e., sets of propositions. For a given proposition to be true in a context on this conception, then, is for it simply to be entailed by it. Thus, that mass is a constant quantity is entailed by, and hence true in, the context of Newtonian mechanics, but false in that of relativistic mechanics. The reason this conception is labeled “subjective” is that its proponents typically tend to think of a context as something like a set of beliefs inside the heads of one or more rational agents. Thus, McCarthy and Buvač note that the context of a conversation consists in the “common assumptions of the participants” (p. 15). Similarly, Giunchiglia (1993) writes that “a context...is inside the reasoning individual” (p. 47). It is “a theory of the world which encodes an individual’s subjective perspective about it” (ibid.).

As on the logical conception, contexts on the objective conception are “first-class citizens.” That is, they are fully-fledged objects that we can refer to and quantify over as adroitly as we do chairs, planets, chromosomes, and numbers.

Minds and Machines **9:** 29–56, 1999.

However, in contrast to the subjective conception, on the objective conception, contexts are *outside* the skull of any reasoning individual. Perhaps the most robust example of this conception is found in situation theory (Barwise and Perry (1986), Devlin (1991)), a powerful, and relatively recent theory of information and information flow (Barwise and Seligman (1997)). On this theory, a situation is a complex but (typically) limited, spatially and temporally extended piece of the real world in which objects – agents, in particular – have properties and stand in relations to other objects. Situations, in this sense, capture one important intuitive conception of context, viz., context as the *setting* in which things like a baseball game, the fabrication of a mechanical part, or a conversation occur. In this sense then, a context provides agents with a shared environment that enables them to make sense of one another's behavior. Notably, the participants in a conversation interpret one another's utterances by connecting them with objects and information available in their shared context. Thus, on the objective conception, the truth of a proposition p in a context c, $ist(c, p)$, is not a *logical* relation between p and other propositions, but a relation of *correspondence* between that proposition and the world: it is a matter of whether the relevant portion of the world – the context c in question – *is* as the proposition *says*.

Not surprisingly, the subjective and objective conceptions lead to rather different logics of context. Insofar as contexts are thought of as sets of beliefs, it is natural to develop a logic that reflects the way in which a limited but (ideally) rational agent in a given context – i.e., a rational agent with a given set of beliefs – reasons in that context. Notable features of such a logic might include the possibility of inconsistent contexts (to allow for the possibility of justified but inconsistent beliefs in "lottery paradox" types of cases) and a nonmonotonic inferential apparatus of some sort. Perhaps the most dramatic difference between the two, however, concerns the overall perspective of the logic. In the influential logic of McCarthy and Buvač (1998), which is strongly subjective in its orientation, it is assumed that every assertion is implicitly contextual; every assertion is embedded in an "outer" context. Thus, for these authors, the basic formulas of the logic are of the form '$c'{:}\psi$' where this includes the notable special case:

(1) $$c' : ist(c, p)$$

indicating "that the proposition p is true in the context c, itself asserted in an outer context c'" (ibid. p. 14). Thought of as sets of beliefs, the outer context framing every assertion reflects two things: first, the situated character of any reasoning agent and, second, the essential capacity of a truly intelligent agent to "transcend" its current context and adopt the broader perspective of an outer context. That context, of course, will be embedded in a new outer context, and so could also be transcended by the agent, as could the new context, and so on. (Context transcendence will be discussed in Section 2.4 below.)

In contrast to the subjective approach, which thinks of assertions in the logic as something like beliefs in the head of a situated agent (or perhaps implications of those beliefs), on the objective conception developed here, assertions in the logic are not *themselves* situated, non-contextual; there is no outer context. However, except for identities, the only noncontextual assertions one is capable of making are assertions about situations and the propositions true in them. More formally, the only atomic formulas in the language are identities and those of the form '$ist(c,p)$'; one cannot, for example, make noncontextual assertions of ordinary atomic propositions such as that Clinton is president; such propositions can only be asserted to be true in a given context, i.e., the expressions denoting such propositions can only occur in the second argument of the '*ist*' predicate. (The language will be defined rigorously in Section 2.2 below.)

The formal difference here is subtle but important. As on the subjective conception, ordinary atomic propositions can only be asserted to be true relative to a context. But unlike the subjective conception, assertions about contexts themselves can be made "absolutely", outside of any context (though they can also be asserted relative to contexts as well). The intuition here is that, on the objective conception, what is true or false in a context is determined by the external world, independent of any agent's beliefs (unless, of course, that information involves such beliefs). Such information is not itself context dependent, and hence the notion of an implicit outer context for every assertion is not appropriate in a logic for the objective conception. At the same time, on this situation theoretically inspired approach, it should be stressed that information about contexts is not, as we might put it, *universal* across contexts. That is, while the information one context carries about another will always be veridical (which is natural on the objective conception where contexts are pieces of the external world), due to the partiality of situations, not all information about a given context is carried by every other context. More formally, while $ist(c,p)$ does follow from $ist(c', ist(c,p))$ on the objective conception, the converse does not hold; from $ist(c,p)$ it does not follow that $ist(c', ist(c,p))$, for any c'.

Now, a logic based on an "absolute" perspective is perhaps of questionable relevance for the classical AI program of building intelligent machines, i.e., at least, machines that can act flexibly and appropriately in changing environments. For this goal, the purpose of a logic is to serve as the basis of a mechanism that, when implemented, enables a machine to reason correctly upon its beliefs and then, on the basis of that reasoning, to choose appropriate actions. The logic is thus itself in the "head" of such a situated agent, and is applied to the agent's salient beliefs, i.e., to its current context in the subjective sense.

However, this is not the vision that motivates a logic based upon the objective conception. Rather, its motivation comes out of the area of knowledge representation (KR), and most specifically the relatively recent branch known as *formal ontology* (see, e.g., Guarino (1995), Guarino (1998), Menzel (1997)). An ontology is a formal representation (typically a set of axioms in a logical language) of the objects and concepts that constitute some (abstract or concrete) real world domain. The purpose of an ontology is to clarify the logic of the domain, to make explicit and clear the relevant concepts one finds there, and thereby to fix the meaning of the vocabulary that is used to describe the domain.

The original motivation for formal ontologies was to fix the meanings of, and logical connections between, the terms in large knowledge bases sufficiently to enable them to be reliably shared and reused across different KR frameworks (Neches et al. (1991), Gruber (1995)), for example, LOOM (MacGregor (1991)) and CLASSIC (Borgida et al. (1989)). For the most part, such knowledge bases originally arose out of large academic Al projects. However, it has become increasingly apparent in recent years that the methods of formal ontology have enormous potential for the area of *enterprise modeling* (cf., e.g., Petrie (1993), Fox and Grüninger (1994)). An enterprise model is a representation of some aspect of a large system, formulated in the language of some enterprise modeling method such as the IDEF methods (Menzel and Mayer (1998)). Enterprise models come in a fairly wide variety of types, corresponding roughly to different types of information: the relatively static information associated with given database schema, the dynamic information of a given manufacturing process, etc.

Models are used for virtually every aspect of enterprise engineering, from the planning and development of the enterprise itself to the definition and maintenance of its processes. It is, therefore, crucial that these models can be shared and reused. However, just as there are many different KR languages in AI, there many different enterprise modeling languages as well. Thus, the same kinds of problems arise for the sharing and reuse of enterprise models as arise for the sharing and reuse of AI knowledge bases. So whatever promise the methods of formal ontology hold for the latter should carry over to the former as well. Notably, the modeling languages themselves need clearly defined logical foundations.

A logic of contexts seems particularly well-suited as the logical foundation for enterprise modeling. Because virtually any enterprise is a complex network of connected situations, enterprise models themselves are highly contextual: a database schema typically structures information about one element of a larger enterprise's activities, e.g., vendor and pricing information for raw materials; a process model typically captures one specific enterprise process, e.g., a specific shop floor manufacturing process. A logic of contexts will provide the right sort of framework for keeping the information of many different enterprise models properly contextualized, but in a manner that also allows integration and sharing of that information across different models.

Now, finally, the reason why a logic based upon the objective conception is more appropriate than one based upon the subjective conception is that the information enterprise models typically carry does not usually reflect the perspective of a situated agent. Rather, such information is typically thought of as absolute, a (perhaps temporally extended) snapshot of an enterprise context taken from a "God's eye" point of view. A business systems analyst, for example, strives simply to model how some aspect of a given business works. She is, ideally, a perfect observer, objectively recording data about that aspect of the system in order to generate models that she can subject to analysis. In this “perfect observer” mode, she does not include herself as a part of the modeled system; her standpoint is objective and external to the system and its processes. Again, an industrial engineer developing a manufacturing process plan models the way the intended process is supposed to occur. As its creator, the designer is himself outside of the process being designed, not (typically) in it. As with the systems analyst, his standpoint is objective and external.

In both cases, what is modeled is simply the information carried in each context and the manner in which that information flows between contexts, objectively construed. With its overtly external standpoint and its basis in a modern theory of information, the objective conception of context thus seems to be the right foundation for these modeling activities. (See Menzel et al. (1993) for some initial steps in this direction.)

2. A Language for the Objective Conception

It will be useful at this point to begin formalizing a logic for the objective conception of context. We begin by defining an appropriate language. We will then use the language as a basis for more detailed exposition of the objective conception. That will in turn set the stage for the presentation of the logic.

2.1. Reified propositions and the representation of the *true-in* relation

The fundamental intuitive idea behind any theory of context is that of a proposition being true in a context. The most straightforward accounting of this intuition is that there is a relation *true-in* that propositions bear to contexts. This in turn presupposes that propositions, no less than contexts and ordinary individuals, are first-class citizens: propositions are genuine items in the inventory of the world, and hence can be referred to and quantified over like any other objects. (Propositions so considered are also known as “reified” propositions.) This, however, leaves it open whether we are to use a first-order language or a higher-order language for this purpose. Both options are possible, and both have their advantages and liabilities. However, to deal with some of the semantic subtleties of the *true-in* relation the account here will draw upon techniques that have been developed for first-order languages (see Section 3.4.1), and so we will remain first-order.

To begin with, then, our language for contexts will contain a 2-place predicate ‘*ist*’ indicating the *true-in* relation, and the usual quantificational apparatus of first order logic. Intuitively, quantifiers will range over a single domain containing three kinds of individual objects (our first-class citizens): contexts, propositions, and ordinary objects (concrete and abstract) like people and numbers. We will not introduce any special predicates or sorted variables for any of our three kinds of individuals (though, importantly, it will be possible to define a formula that is true of all and only contexts).

Individual constants will suffice for denoting ordinary objects and contexts, but propositions in general come in a wide variety of logical forms corresponding roughly to the syntactic forms of the sentences in a first-order language. Hence, if we want to capture the richness of these logical forms, we will need something more complex than mere individual constants for denoting propositions. For this reason, some researchers – Buvač, Buvač, and Mason (1995), for instance – simply use first-order sentences to indicate propositions. On this approach, "is true in" is represented not with a predicate but with a modal operator '$\mathtt{ist}$' such that, for any term τ and formula φ, $\neg\mathtt{ist}(\tau, \varphi)$ is a formula as well.

As illustrated in, e.g., McCarthy and Buvač (1998), this approach is adequate for many purposes. Note, however, that, because φ in $\ulcorner\mathtt{ist}(\tau, \varphi)\urcorner$ is not a term, this approach does not treat propositions as first-class citizens. Now, *a priori*, that may in fact turn out to be fine; while we might think of propositions realistically as first-class citizens in the informal situation theoretic picture that motivates our account, it could turn out that we needn't explicitly introduce them into the semantics of the formal theory at all; the definition of the truth of a formula in a context might be all we need.

However, this does not appear to be so. In recent years especially, a wide variety of intuitive, linguistic, and philosophical phenomena have been identified that are, arguably, best explained (or, perhaps, *only* explained) by explicitly introducing reified propositions (Bealer and Mönnich (1989), Zalta (1988)). Best known among these, perhaps, are the logical and semantic phenomena surrounding singular reference and the associated puzzles of *de re* belief (Salmon (1986)). Additionally, a number of linguists have argued that propositions (as well as properties) provide the only satisfactory account of the pervasive phenomenon of nominalization in natural language (see, e.g., Chierchia and Turner (1988), Chierchia, Partee, and Turner (1989)). Furthermore, certain intuitive logical inferences seem to involve reference to, and quantification over, propositions. For example, if Pete and Joyce both believe that the Mariners will win the American League pennant, then there is something that they both believe, viz., *that the Mariners will win the pennant*. Closer to our immediate topic, quantification over propositions also seems to be required if we are to be able say certain things that we might naturally want to say about contexts. For instance, for those of a more philosophical bent wishing to comply with the Quinean dictum "No entity without identity" (Quine (1981), p. 102), a plausible criterion of identity for contexts can be drawn from situation theory: contexts are identical iff the same information holds in both of them, i.e., in the current reconstruction, iff the same propositions are true in them (cf. Barwise (1989), p. 264):

(2) $$c = c' \leftrightarrow (\forall p)(ist(c, p) \leftrightarrow ist(c', p)).$$

Clearly, however, this criterion quantifies over propositions. Therefore, there appear to be good grounds for investigating an alternative to the operator approach that represents "is true in" as a genuine predicate and, consequently, explicitly represents propositions as first-class citizens.

2.2. The Language

Our language $\mathcal{L}$ for a formal logic of contexts, then, will need at least our 2-place predicate '*ist*' and a special class of terms that can be used for referring explicitly to particular propositions like *The Mariners will win the pennant*. As noted above, however, on the objective conception, the only noncontextual propositions we want to be able to express are identities and propositions about contexts. We will do this by restricting the occurrence of predicates expressing ordinary properties and relations so that they only can occur in terms denoting propositions; it will not be possible to use them to construct ordinary atomic formulas.

So, more formally, consider a lexicon consisting of a denumerable set of constants, a denumerable set of variables, countably many n-place predicates, for all n, $0 < n < \omega$, logical operators, $\neg$, $\rightarrow$, $=$, $\forall$, and parentheses (and). We stipulate that the lexicon must include at least the distinguished 2-place predicate '*ist*'. The well-formed expressions of the language $\mathcal{L}$ are given by the following simultaneous recursive definition.

1. Every constant or variable is a *simple term*. Every simple term or propositional term is a *term*.
2. If π is an n-place predicate and $\tau, \tau', \tau_1 \ldots, \tau_n$ are terms, then $\pi(\tau_1, \ldots, \tau_n)$ and $\tau = \tau'$ are (*atomic*) *propositional terms*.
3. If τ and τ' are propositional terms and x any variable, then $\neg\tau$, $\tau \rightarrow \tau'$, and $(\forall x)\tau$ are propositional terms.
4. If τ and τ' are terms and x any variable, then $\neg\tau, \tau \rightarrow \tau'$ and $(\forall x)\tau$ are propositional terms.
5. If τ and τ' are terms, then $ist(\tau, \tau')$ and $\tau = \tau'$ are (*atomic*) *formulas*.
6. If φ and ψ are formulas and x any variable, then $\neg\varphi$, $\varphi \rightarrow \psi$ and $(\forall x)\varphi$ are formulas.
7. Nothing else is a term or formula of $\mathcal{L}$.

Other standard logical operators besides $\forall$ and $\rightarrow$ are defined as usual. Notions of free and bound variables, closure, and of a term being free (or substitutable) for a variable are as usual for formulas and are directly analogous for terms. Square brackets will occasionally be used around propositional terms to aid readability. $(\forall x_1 \ldots x_n)$ will abbreviate $(\forall x_1) \ldots (\forall x_n)$.

$\mathcal{L}$ is an instance of a broad class of extended first-order intensional languages that take propositions (and, more generally, properties and relations as well) to be first-order objects (see esp. Bealer (1982) and Turner (1987)).[2] But because these languages are not widely used in AI, it is worth commenting on a few of $\mathcal{L}$'s features. First, assuming $\mathcal{L}$ is endowed with a rich stock of constants and predicates beyond '*ist*', the language will contain a correspondingly rich variety of propositional terms. For example, suppose we let 'm' stand for the Seattle Mariners and 'W' for the property of winning the American League pennant. Then there is the propositional term '$W(m)$' denoting (intuitively) the proposition that the Mariners win the pennant. Notice, however, that this expression is *not* a formula of $\mathcal{L}$: since the proposition it expresses is essentially contextual, the term cannot be asserted by itself, outside of any context. Rather, it can only be asserted as part of an atomic formula of the form '*ist(c,W(m))*' to the effect that it is true in some context c (or in identities like '$p = W(m)$'). This, of course, reflects the objective conception, which adopts a standpoint outside of all contexts. Hence, being a "God's eye" view of the space of contexts and not itself another context, the only information that is available from that standpoint – the only noncontextual information that can be asserted – consists of identities and statements to the effect that a given proposition is true in a given context (and boolean and quantified elaborations thereof).[3]

In contrast to expressions like '$W(m)$', '*ist*$(c, W(m))$' is *both* a term *and* a formula. As such, in addition to being assertible as is (qua formula), it can also be asserted (qua term) as an argument to a predicate, the 'ist' predicate in particular. This permits us to express the nesting of contextual information that is so essential to any theory of contexts (see especially Section 2.4 below). Note that this is completely kosher. All that is required of terms in a language is that they can be assigned a unique denotation. All that is required of formulas is that they can be assigned a unique truth value. That some expressions are assigned both a denotation (qua term) and a truth value (qua formula) is perfectly consistent with these requirements.[4]

Second, although propositions are often called "higher-order" objects, the fact that there are terms denoting propositions does not of itself make $\mathcal{L}$ a higher-order language. Being higher-order is a *semantic*, or *model theoretic*, property of a language; as Shapiro (1991) notes in his excellent study (p. 13), "languages ..., by themselves, are neither first- nor higher-order." (See also Enderton (1972), ch. 4). A language + model theory is truly higher-order, roughly, if and only if there are quantifiable variables in the language that range over the elements of some sort of power-set construction over the domain of individuals (or some other primitive semantic domain such as worlds or times). An object in the model theory is said to be higher-order if it is an element of such a construction. Properties and propositions in standard possible world semantics, for example, are higher-order in this sense: the class of propositions is defined to be the set of all functions from possible worlds to truth values (or, equivalently, the power set of the set of possible worlds), and properties the set of all functions from possible worlds into the the power set of the set of all possible individuals (see e.g., Dowty, Wall, and Peters (1981)). In the model theory for $\mathcal{L}$, however (which is only presented informally here), there is just a single sort of variable ranging over a single domain. The propositions constitute a subset of that domain. There are no higher-order objects in the technical sense of "higher-order". In particular, although the class of propositions is stipulated to be closed under certain logical operations (see Section 3.3 below), it is not a power set construction of any sort. Semantically, propositions are just first-order individuals, and $\mathcal{L}$, with its model theory, is a first-order language in the truest sense.

2.3. Localism and the Evaluative Standpoint

On the objective conception, a context provides a sort of "model" that determines the truth values of propositions. However, because of the limited nature of contexts, propositions that are true in one context may fail to be true in another. Call this phenomenon *localism*. The ground of localism is the simple fact that which propositions hold in a context is a function of how things stand with regard to the objects that are present in the context. Since contexts can differ both with regard to the objects present in them and with regard to how things stand with those objects, localism follows immediately. There are two notable manifestations of localism that we will want to capture in our logic. First, basic atomic propositions concerning the properties of, and relations among, objects that are not present in a context are simply not true there, regardless of whether or not they are true in the contexts in which they *are* present. The tallest building in Melbourne, Australia, for example, is not present in the context of a city league softball game in College Station, Texas; hence, the proposition that it is, say, over 300 meters high does not hold in that context; for nothing in the context makes it true. Thus, a context carries "positive" information (i.e., information in the form of an atomic proposition) only about those objects that it "knows about," i.e., which are present in it; there is just no positive information available within a context regarding things outside the context, i.e., no information about what properties they have or what relations they bear to other things. Call this *atomic localism*.[5]

Second, quantified propositions (see Section 3.3) are evaluated in a context relative to the set of objects present in the context. Thus, an existentially quantified proposition that holds in one context might well fail in a broader context. Call this *quantificational localism*. Lewis (1986) illustrates quantificational localism with the proposition *there is no beer*, which turns out – to the horror of the neighborhood frat house – to be true in the context determined by a salient refrigerator. However, relative to the broader context of the neighborhood in which the refrigerator is situated, and in particular relative to the context circumscribed by the corner mini-mart, the proposition is – to the frat brothers' great relief – false.

2.3.1. Localism and Partiality

Let b be the tallest building in Melbourne, $300m$ the property of being 300 meters tall, and c^* the context of a city league softball game in College Station. By atomic localism, the proposition $300m(b)$ is not true in the context c^*, that is $\neg ist(c^*, 300m(b))$. However, atomic localism does not determine whether or not we should count the *negated* atomic proposition $\neg 300m(b)$ that it is not the case that the tallest building b in Melbourne is over 300 meters high as true in c^*. The option we will adopt represents a fairly strong departure from standard situation theory.

Say that a context c is *weakly partial* with regard to information if, for some object e, there is no positive information about e in c.[6] By atomic localism, then (assuming that every individual has some property or stands in some relation to other individuals), every context in which fewer than all individuals are present is weakly partial. Say that c is *strongly partial* with regard to information if there is some atomic proposition p such that neither p nor its complement $\neg p$ is true in c. Strong partiality, then, is essentially the failure of bivalence in general with regard to the *true-in* relation. In standard standard situation theory, every situation in which fewer than all individuals are present (i.e., every, or nearly every, situation) is strongly partial with regard to information. A situation can carry no negative information – indeed, no information at all about objects that are not present in it; the situation is simply silent with regard to such information. In particular, in situation theory, neither $300m(b)$ nor $\neg 300m(b)$ is true in c^*.

As we will see in Section 3.4.1, it will be possible in the logic here for a context to be strongly partial with regard to information, but only under rare conditions in which logical paradox threatens. Such pathological cases aside, on the conception developed here, a context *does* carry negative information about objects that are not present in it. In particular, but for the problematic cases just noted, a negated atomic proposition $\neg p$ is deemed true in a context just in case p simply is not true there. Thus, the proposition $\neg 300m(b)$ will be taken to be true in the context of a softball game in College Station, despite the distinct absence of the building in question.

It seems to me that both options regarding the evaluation of negated propositions are compatible with the objective conception. The difference reflects a choice regarding what we might call the *evaluative standpoint*. Say that a piece of information – a proposition – p is *generated* in a context c just in case all of the objects that p is about are present in c. In standard situation theory, the evaluative standpoint is *internal* – one, in a sense, puts oneself "inside" the situation and evaluates the propositions that are generated there, as those are the only ones that one can "see" from within the context. Those propositions that are not generated in the context, from the internal standpoint, are simply not available for evaluation. In contrast to standard situation theory – and perhaps somewhat more in the spirit of the objective conception – the evaluative standpoint adopted here is *external*, outside all particular contexts. From this perspective, any proposition generated in any context can be evaluated with respect to any other context, even contexts in which the proposition is not generated. In particular, from this standpoint, a typical negated atomic proposition $\neg P(a_1, \ldots, a_n)$ will be evaluated as true in a context c just in case $P(a_1, \ldots, a_n)$ is not true in c, i.e., just in case the objects $a_1, \ldots, a_n$ fail to stand in the relation P in c, regardless of whether or not they are present in c. Most contexts, that is to say, while weakly partial with respect to information, are not strongly partial.

It is worth emphasizing that, from the external evaluative standpoint adopted here, there are two ways in which a negated atomic proposition $\neg Foo(a)$ can be true in a context. It can be true because a is present in c and fails to Foo there, or it can be true because a – which may well be Foo in other contexts is simply absent from c; there is just no (positive) information about a and its properties available in the context to validate the proposition $Foo(a)$. We must therefore distinguish between having the property *non-Foo* in a context – which, by localism, entails presence in c – from not being Foo at c. Properly defined (in light of localism), the property of being *non-Foo* in a context entails presence in the context:

(3) $$\textit{ist}(c, \textit{non-Foo}(x)) \equiv_{df} \textit{ist}(c, [(\exists y) y = x \wedge \neg Foo(x)]).$$

It then follows that, while, being *non-Foo* at c — $\textit{ist}(c, \textit{non-Foo}(x))$ – entails not being Foo there – $\textit{ist}(c, \neg Foo(x))$ – the converse is not true.[7]

2.4. Subsumption and context transcendence

As McCarthy has repeatedly emphasized in his work on context, a significant capacity underlying intelligent problem solving is the ability to *transcend* one's current context and reason about broader contexts in which the original context is embedded. Indeed, to represent this phenomenon robustly might be thought of as a strong condition of adequacy on any logic of context. However, subjective and objective conceptions differ in regard to *what* is being represented. An adequate logic for the subjective conception will represent context transcendence from within the head of an intelligent agent. An adequate logic for the objective conception, by contrast, will represent the logical properties of contexts that make transcendence possible. Notably, it will be capable of representing the *subsumption* of one context by another in which information about what holds in one context c', $\textit{ist}(c', p)$, is itself information that holds in a broader context c, $\textit{ist}(c, \textit{ist}(c', p))$. (Subsumption is defined and axiomatized in Section 3.4.) An agent in c' that apprehends this relation can transcend that context and situate itself in the broader context c, from which it can formulate strategies for action and problem solving that were not possible relative to the original context c'. So, for instance, by transcending the context of the house refrigerator to the broader neighborhood context, the frat brother in charge of refreshments is able to form a feasible plan for solving the pressing problem of beer depletion.

It is instructive to see how the language $\mathcal{L}$ represents the facts about subsumption on the basis of which such an agent might transcend from one context into another. Roughly speaking, one transcends a context c when, after reasoning from within c, one begins to reason from within a broader context c', a context within which, in particular, one can reason about the original context c. Thus, to return to the example above, upon discovering the fraternity's current plight, the brother in charge of refreshments hastens to the local mini-mart, purchases an ample supply of Sam Adams, and returns triumphantly home. The reasoning underlying his actions can be explicated (in part, at least) in terms of his transcending the context of his immediate beerless environment,

(4) $$ist(f, \neg(\exists x)Beer(x)),$$

or, using the often more perspicuous notation of McCarthy (1993) noted above,[8]

(5) $$f : \neg(\exists x)Beer(x)$$

to the broader context of his local neighborhood, which both subsumes the original beerless context and contains the source of a remedy,

(6) $$ist(n, ist(f, \neg(\exists x)Beer(x)) \wedge ist(m, (\exists x)Beer(x))),$$

i.e.,

(7) $$n : ist(f, \neg(\exists x)Beer(x)) \wedge ist(m, ((\exists x)Beer(x))).$$

That is, in the context of the broader neighborhood n in which our hero is situated, while there is no beer in the fridge, there is at the local mini-mart. Thus, by transcending the context of his refrigerator to the broader neighborhood context he is able to form a feasible plan for replenishing the frat house's depleted supply.

The key here, again, is the ability to express information of the form $ist(c, ist(c', p))$ to the effect that information about what holds in one context – $ist(c', p)$ – can itself be information that holds in a broader context. This is possible in $\mathcal{L}$ by virtue of the fact that every formula φ of $\mathcal{L}$ is also a term (though not vice versa), and hence capable of appearing in the second argument place of '*ist*'.

3. A Logic for the Objective Conception

With the language defined and some intuitive underpinnings laid, we can now begin constructing the logic – CL – for the objective conception of context. Let φ, ψ, and θ be any formulas, τ, τ', and τ'' any propositional terms, x, y, k, k', and p any distinct variables, and let ϵ^x_τ be the result of replacing every free occurrence of a variable x in the expression ϵ with τ. (Any of these may occur with subscripts.) Then the universal closures of any instance of the following schemas L1 – L29 is an axiom.[9]

L1 $\varphi \rightarrow (\psi \rightarrow \varphi)$

L2 $(\varphi \rightarrow (\psi \rightarrow \theta)) \rightarrow ((\varphi \rightarrow \psi) \rightarrow (\varphi \rightarrow \theta))$.

L3 $(\neg\varphi \rightarrow \psi) \rightarrow ((\neg\varphi \rightarrow \neg\psi) \rightarrow \varphi)$.

L4 $(\forall x)\varphi \rightarrow \varphi^x_\sigma$ for any term σ free for x in φ.

L5 $(\forall x)(\varphi \rightarrow \psi) \rightarrow (\varphi \rightarrow (\forall x)\psi)$, where x does not occur free in φ.

L6 $x = x$.

L7 $x = y \rightarrow (\varphi \rightarrow \varphi')$, where y is substitutable for x in φ, and φ' is the result of replacing one or more occurrences of x in φ with y.

Rule of Inference

MP ψ follows from φ and $\varphi \rightarrow \psi$.

As usual, a *derivation* in CL is a sequence $S = \langle\varphi, \ldots, \varphi_n\rangle$ such that each φ_i is either an axiom of CL or follows from earlier elements of the sequence by **MP**. S is a *derivation of* φ in CL if $\varphi = \varphi_n$. We write $\vdash \varphi$ if there is a derivation of φ in CL. φ is derivable from Γ in CL – $\Gamma \vdash \varphi$ – if, for some finite subset Γ' of Γ, $\bigwedge \Gamma' \rightarrow \varphi$, where $\bigwedge \Gamma'$ is any conjunction of the members of Γ'.

A useful derived rule of inference that will be used below is Generalization on Constants:[10]

GC $(\forall x)\varphi$ follows from φ^x_a, for any constant a.

3.2 "LOCAL" PREDICATE LOGIC

Now we need to identify the logical principles that govern matters within contexts. Note first it is possible to define the notion of context in terms of the *true-in* relation: an entity is a context just in case something is true in it: $\mathit{Context}(x) \equiv_{df} (\exists y)\mathit{ist}(x, y)$. This is important, because most of the remaining axioms concern contexts, not all things generally, and hence are typically of the form $\ulcorner \mathit{Context}(x) \rightarrow \varphi \urcorner$. To facilitate readability, however, we will assume the existence of a distinguished set of sorted variables that range over contexts, and we let c and c' be any two such variables.[11] To begin, then, we note that classical propositional logic carries over into contexts directly:

L8 $\mathit{ist}(c, [\tau \rightarrow (\tau' \rightarrow \tau)])$.

L9 $\mathit{ist}(c, [(\tau \rightarrow (\tau' \rightarrow \tau'')) \rightarrow ((\tau \rightarrow \tau') \rightarrow (\tau \rightarrow \tau'')])$.

L10 $\mathit{ist}(c, [(\neg\tau \rightarrow \tau') \rightarrow ((\neg\tau \rightarrow \neg\tau') \rightarrow \tau)])$.

Although the propositional logic of contexts is classical, the logic of quantification within contexts is basically a free logic. The reason for this, of course, is quantificational localism: the things present in a typical context constitute only a proper subset of all the things there are. Thus, what may be true of everything in the context, may not be true of everything in general. Hence, since we want to interpret quantifiers locally, we cannot preserve truth by allowing arbitrary universal instantiations. More explicitly, suppose everything in context c has property Foo, that is, interpreting quantifiers locally,

(8) $$ist(c, (\forall x)Foo(x)),$$

and suppose that b is an object that is not in c, and which is also not Foo in any context. Then, in particular, it is not the case that b is Foo at c, i.e.,

(9) $$\neg ist(c, Foo(b)).$$

Hence, on pain of contradiction, we do not want to be able to infer

(10) $$ist(c, Foo(b))$$

from (8), and so we do not want local universal instantiation, viz.,

(11) $$ist(c, (\forall x)\tau \rightarrow \tau^x_\sigma), \sigma \text{ free for } x,$$

as a general logical principle. Rather, we want to restrict the terms we can instantiate for the bound variable to those denoting objects present in the context c. We can accomplish this by introducing a notational definition:

DEFINITION 1. $E!(x) \equiv_{df} (\exists y) y = x$.

Note that since $\ulcorner(\exists y)y = x\urcorner$ is also a term, it (hence its abbreviation $\ulcorner E!(x)\urcorner$) can appear as an argument to 'ist'. For readability, we define $\ulcorner E!(\tau_1, \ldots, \tau_n)\urcorner$ to be the conjunction $\ulcorner E!(\tau_1) \wedge \ldots \wedge E!(\tau_n)\urcorner$. Given this, we introduce the axioms for local quantification:

L11 $ist(c, (\forall x)\tau \rightarrow (E!(\sigma) \rightarrow \tau^x_\sigma))$, for any term σ that is free for x in τ.

L12 $ist(c, (\forall x)(\tau \rightarrow \tau') \rightarrow (\tau \rightarrow (\forall x \tau')))$, where x does not occur free in τ.

Local identity axioms, as well as local versions of all the axioms below, follow from a "reflection" principle to the effect that propositions provable in the global logic are true in every context:

L13 If $\vdash \varphi$, then $\vdash ist(c, \varphi)$.

Complementing the local schemas above we have schemas that, in essence, give us MP and Generalization for local inferences:

L14 $ist(c, (\tau \rightarrow \tau')) \rightarrow (ist(c, \tau) \rightarrow ist(c, \tau'))$.

L15 $(\forall x)ist(c, \tau) \rightarrow ist(c, (\forall x)\tau)$.

That is, if τ is true in c for everything whatever, then $(\forall x)\tau$ is true in c as well. Making matters more explicit, L14 and LiS yield the following metatheorems:

M1 If $\vdash ist(c, \sigma)$, and $\Gamma \vdash ist(c, \sigma \rightarrow \tau)$, then $\Gamma \vdash ist(c, \tau)$.
M2 If $\Gamma \vdash ist(c, \tau^x_k)$, and k does not occur in any member of Γ, then $\Gamma \vdash ist(c, (\forall x)\tau)$.

For reasons having to do with negation and the truth-at relation discussed below, we cannot prove what is *almost* the contrapositive of L15, and hence we assume it as a further principle:

L16 $ist(c, (\exists x)\tau) \rightarrow (\exists x)ist(c, \tau)$.

Atomic localism – that only atomic propositions about objects present in a context can be true in it – is expressed by the following schema:

L17 $ist(c, [\pi(x_1, \ldots, x_n) \rightarrow E!(x_1, \ldots, x_n)])$.

Note that, because '=' is categorized as an operator, not a predicate, L17 does not apply to identities, hence we get no contradiction from the fact that $ist(c, b = b)$ even if b is not present in c, i.e., even if $ist(c, \neg E!(b))$. The justification for this is that identities are true objectively, outside of all contexts, as well as across all contexts, regardless of how things stand among the objects present in those contexts. We express this noncontextual character of identity propositions in the following principle:

L18 $ist(c, x = y) \rightarrow x = y$.

That is, any identity that is true in a context is true simpliciter. The converse is a consequence of L6 and L13.

3.3. Axioms for propositions

It is widely agreed that the phenomena of intentionality and singular reference noted above in Section 2 are best explained in terms of propositions that are "fine-grained" in the sense that their identity conditions are determined by their internal structure rather than by their extensions, i.e., their mere truth or falsity in and across contexts (see, e.g., Bealer (1982), Chierchia and Turner (1988), Soames (1987), Zalta (1988)). The intuitive idea here is that complex propositions are "built up" from their simpler components by means of a collection of logical functions. Initially, we assume the existence of all the atomic propositions that can be built up from a given n-place relation and n arguments. We then postulate that atomic propositions are identical if and only if they contain the same relation and that relation has the same arguments in the propositions. The closest we can get to capturing this without quantifying over properties and relations is captured in the following principle (for brevity, let σ be $\pi(\tau_1, \ldots, \tau_n)$, σ' be $\rho(\tau'_1 \ldots, \tau'_n)$, σ^* be $\pi(x_1 \ldots, x_n)$, and $\sigma' *$ be $\rho(x_1, \ldots, x_n)$):

L19 $\sigma = \sigma' \rightarrow ((\tau_1 = \tau'_1 \wedge \ldots \wedge \tau_n = \tau'_n) \wedge (\forall c x_1 \ldots x_n)(ist(c, \sigma^*) \leftrightarrow ist(c, \sigma'^*)))$.

It is also assumed that atomic propositions built up from relations of different arity are distinct:

L20 $\pi(\tau_1, \ldots, \tau_n) \neq \rho(\tau'_1, \ldots, \tau'_m)$, if $n \neq m$.

Negated propositions are constructed from given propositions by means of a logical negation function *neg*, conditional propositions by means of a function *cond*, and universally quantified propositions by means of logical functions *u-quant*.[12] Fine-grainedness is ensured by requiring that the logical functions have disjoint ranges and are all one-to-one – or at least nearly so; *u-quant* will be assumed to be one-to-one *up to similarity*, where, roughly, p and q are similar if their internal structure differs only with regard to the individuals they contain. (The model theory for $\mathcal{L}$, of course, contains a more precise definition.) The one-to-one-ness of *neg* and *cond* is captured in schemas L21 and L22, and that of *u-quant* by L23:

L21 $\neg\tau = \neg\tau' \rightarrow \tau = \tau'$

L22 $[\tau \rightarrow \tau'] = [\sigma \rightarrow \sigma'] \rightarrow (\tau = \sigma \wedge \tau' = \sigma')$.

L23 $[(\forall x)\tau] = [(\forall x)\tau'] \rightarrow \tau = \tau'$.

That the ranges of the logical functions are pairwise disjoint is expressed in the following schema:

L24 $\tau = \tau'$, where τ and τ' are propositional terms of different syntactic categories.

Finally, note that these requirements on the logical functions ensure that each proposition has a unique *decomposition* – analogous to the decomposition of a formula into its subformulas and, ultimately, its primitive terms and predicates – that reveals how it was built up from simpler entities. Given this, we will assume in addition that every proposition is *structurally well-founded* in the sense that no proposition can contain itself in its own decomposition. Structural well-foundedness is captured in the following schema (cf. Bealer (1982), p. 65):

L25 $k = \tau - k' = \tau'$, where τ is a propositional term containing k', and τ' is a propositional term containing k.

3.4. THE LOGIC OF SUBSUMPTION

As noted, it is an essential feature of $\mathcal{L}$ that it can represent the properties of contexts that lie at the heart of context transcendence. Most notably, as discussed in Section 2.4, $\mathcal{L}$ is able to represent the subsumption of one context by another. The final plank in our logic of context concerns the logic of subsumption. Formally, we define subsumption as follows:

DEFINITION 2. $c \preccurlyeq c \equiv_{df} \forall p(ist(c', p) \rightarrow ist(c, ist(c', p)),$

that is, loosely, c subsumes c' if, whenever p is true in the context c', the context c carries the information that p is true in c'.

Note that, by atomic localism L17, it follows the fact that c subsumes c' that the narrower context c' is itself present in the broader context c, i.e.,

(12) $$c' \preccurlyeq c \rightarrow ist(c, E!(c'))$$

This makes good sense on the objective conception: for, by localism, the information that holds in a given context is a function of the objects that are present in the context. Since $ist(c', p)$ is itself information in c, and since c' is a fully-fledged object, it follows that c' is present in c as well.

What about the converse of (12)? Though not a theorem of our principles to this point, it does seem intuitively valid. For suppose c' is present in c, and suppose that two objects a and b stand in some relation R in c' (and hence, by atomic localism, are present in c'). Intuitively, if contexts on the objective conception are simply pieces of the real world, it seems that the only way to make sense of the idea of one context c' being present in another c is to think of c' as a *subcontext* of the c, i.e., as a smaller context embedded in a larger one – as the seventh inning is embedded in a given baseball game, for example. If so, however, it seems reasonable that every object present in the subcontext c' must be present in the larger context c, the objects a and b in particular. Now, since a, b, and c' are all present in c, and since a bears R to b in c', intuitively once again, the information *that* a bears R to b in c' is carried by c. More generally, the larger context c "*sees*" everything that is true in the subcontext c', i.e., the converse of (12) is intuitively valid. We thus add it as an axiom:

L26 $ist(c, E!(c')) \rightarrow c' \preccurlyeq c,$

By the definition of subsumption, if c subsumes c', then if p is true in c', c carries the information that p is true in c'. What about the converse of this latter proposition? That is, if c carries the information that p is true in (a subcontext) c', is that information veridical? It is an important – and intuitive – part of the objective conception that it is:

L27 $ist(c, ist(c', p)) \rightarrow ist(c', p).$

This is perhaps the point where the objective conception parts ways most dramatically from the subjective conception. For if contexts are thought of as constituted by an agent's beliefs, then there is no *a priori* reason to suppose that what is true in one context according to another – i.e., what one agent believes according to the beliefs of another agent – is going to be veridical. The objective conception, however, pretty much demands veridicality: the information a context carries is determined by how things stand in the context. But what is true in its subcontexts is part of how things stand in a context. Hence, such information will always be veridical.[13]

Similar considerations yield a concomitant principle, viz., that no contradiction is true in any context, a reasonable principle on the objective conception insofar as we assume that the world itself (hence any piece thereof) is consistent. This is expressed in the following axiom:

L28 $\neg ist(c, \tau \wedge \neg\tau)$.

3.4.1. Self-aware contexts, self-referential propositions, and the specter of paradox

The converse of L27, of course, is invalid, since not every context carries information about what is true in every other context; that is, not every context subsumes every other context. However, a qualified form of the converse does follow directly from the definition of subsumption, viz., that if p is true in a given context c', then the information that p is true in c' is carried by every context that subsumes c':

(13) $$ist(c', p) \rightarrow (c' \preccurlyeq c \rightarrow ist(c, ist(c', p))).$$

But what about information about what *isn't* true in a given context? For example, does it follow from $\neg ist(c', p)$ and $c' \preccurlyeq c$ that $ist(c', q)$, where q is the negation of p? In fact, this will follow from (13) if the truth-in relation is bivalent, i.e., if

$$(*) \quad ist(c, \tau) \vee ist(c, \neg\tau).$$

Given that both CL's global and local propositional logics are classical, (∗) is not an unreasonable assumption. Nevertheless, (∗) is highly problematic. As noted, the language of CL contains terms for propositions that are first-class citizens capable of a certain sort of self-reference. Moreover, the language contains a (relativized) truth predicate. Under such circumstances, the specter of paradox looms large. The most immediate threat, viz., that of an overtly self-referential liar proposition $p = \neg ist(c, p)$ that says of itself that it is not true in some context c is averted in virtue of the structural well-foundedness of propositions. However, this by no means frees the theory from the threat of paradox. Rather, given (∗) and a couple of the intuitive properties of contexts and propositions made explicit in the above principles, it is possible to construct "empirical" paradoxes of the sort brought to prominence in Kripke (1975) that do not involve overtly self-referential propositions.

Consider, in particular, the proposition $[(\forall x)(Rx \rightarrow \neg ist(c, x))]$, i.e., the proposition that everything that is R is false at c. (Let 'λ' serve as an abbreviation for this proposition.) Suppose that λ itself is the only thing present in c that is R, i.e., that $ist(c, (\forall x)(R(x) \leftrightarrow x = \lambda))$, and also that c subsumes itself, $c \preccurlyeq c$.[14] Now, by our global propositional logic, either $ist(c, \lambda)$ or $\neg ist(c, \lambda)$. So suppose $ist(c, \lambda)$. Then $ist(c, R(\lambda) \rightarrow \neg ist(c, \lambda))$ and, hence, since $ist(c, R(\lambda))$ (which follows from $ist(c, (\forall x)(R(x) \leftrightarrow x = \lambda)))$, we have $ist(c, \neg ist(c, \lambda))$. However, since $c \preccurlyeq c$, then by our assumption that $ist(c, \lambda)$ and (13) we have $ist(c, ist(c, \lambda))$. It follows from our local propositional logic that $ist(c, [ist(c, \lambda) \wedge \neg ist(c, \lambda)])$, contradicting L28. Hence, $\neg ist(c, \lambda)$. So suppose instead that $\neg ist(c, \lambda)$. By $(*)$ it follows that $ist(c, \neg\lambda)$, i.e., $ist(c, [\neg(\forall x)(Rx \rightarrow ist(c, x))])$. By definition of $\exists$ and $\wedge$ we have $ist(c, [\exists x(R(x) \wedge ist(c, x))])$. But since $ist(c, (\forall x)(R(x) \leftrightarrow x = \lambda))$, it follows that $ist(c, [R(\lambda) \wedge ist(c, \lambda)])$ and hence $ist(c, ist(c, \lambda))$. So by L27, $ist(c, \lambda)$. Contradiction!

$(*)$, then, is unsound. The *true-in* relation is "gappy"; certain problematic propositions are such that neither they nor their negations are true in some contexts. To cash the notion of "problematic" here semantically, the model theory for CL builds on the "rule of revision" semantics developed independently by Gupta (1982) and Herzberger (1982). Very roughly, the idea is that the understanding of most concepts involves knowing some kind of rule or procedure for *applying* the concept. By contrast, understanding the concept of truth involves knowing a *revision* rule, a procedure for improving any given candidate for the extension of the truth predicate. The procedure works like this. Suppose the extensions of all names and nonlogical predicates in your language are fixed, and that you are given an arbitrary initial candidate extension for the truth predicate. Now evaluate all the sentences in your language so interpreted by the ordinary Tarskian semantic rules. All the sentences that come out true constitute a new candidate extension for the truth predicate. By continuing in this fashion one gradually eliminates the arbitrariness and error present in one's initial candidate and approaches a model that exemplifies the Tarskian ideal wherein any given sentence is in the extension of the truth predicate if and only if it is true. At certain "good" limit stages in the procedure, things stabilize as much as they are going to, and the only sentences that foil the Tarskian ideal – the only problematic sentences – are paradoxical sentences like the liar, which forever oscillate in truth value through the stages of the revision procedure, and hence are appropriately designated unstable.

In our adaptation, an analogous revision procedure is used to build up the extension of the true-in i-elation (at each context). At "good" stages – which we use as models of $\mathcal{L}$ – only pairs $\langle c, p\rangle$ where, roughly, p is both true and stable in c are included in the extension of the true-in relation. Pairs in which p is a paradoxical proposition like 1 that never stabilize at c are left out. But if λ is unstable at c, so is its negation $\neg\lambda$. Hence, given that only stable propositions can be true in a context, we have neither $ist(c, \lambda)$ nor $ist(c, \neg\lambda)$. This fact makes for a useful definition:

DEFINITION 3. $Stable(c, \tau) \equiv_{df} ist(c, \tau) \vee ist(c, \neg\tau)$,

that is, a proposition is stable at c iff either it or its negation is true (i.e., stably true) in c.

Now, to return to the question that prompted this discussion, in models of $\mathcal{L}$, the sentence '$\neg ist(c, \lambda)$' turns out true (assuming the truth of '$ist(c, \forall x(R(x) \leftrightarrow x = \lambda))'$) because λ is unstable in c. However, its reified counterpart – the proposition $\neg ist(c, \lambda)$ – is also unstable in c, and so we also have $\neg ist(c, \neg ist(c, \lambda))$. Hence, recalling our example above, we cannot in general infer $ist(c, \neg ist(c', \tau))$ from $\neg ist(c', t)$ and $c' \preccurlyeq c$. Rather, we need also to know that τ is stable at c. That is, we have as a theorem of CL:

(14) $$(\neg ist(c', \tau) \wedge Stable(c', \tau) \wedge c' \preccurlyeq c) \rightarrow ist(c, \neg ist(c', \tau)).$$

Conversely, the only way that $\neg ist(c', \tau)$ can be (stably) true in a broader context c is if τ is stable at c'. But this in turn implies that $\neg\tau$ is (stably) *true* in c'. This, in fact, is the last principle of CL:[15]

L29 $(ist(c, \neg ist(c', \tau)) \wedge c' \preccurlyeq c) \rightarrow ist(c', \neg\tau)$.

4. Lifting one Type of Context into Another: The Blocks World Example

Many research areas in philosophy, cognitive science, and AI are defined by a key set of intuitive data: well-known problems, puzzles, examples, phenomena, etc. some subset of which, at least, any new approach to the area must address. By addressing one or more of them successfully, the approach meets a sort of minimal criterion of adequacy that, at the least, demonstrates that the approach has a certain measure of promise. In the area of context, McCarthy's Blocks World example functions in this way. The example consists chiefly of a theorem that shows how what is true in a given context can be "lifted" into a subsuming context in a slightly different form, and then used to derive new information about that context. It provides a good test of the basic logic in an account of context, notably its account of subsumption. Consequently, in this penultimate section I want to use CL to set up and derive $\mathcal{L}$'s version of the well known theorem found in McCarthy's example (cp. Akman and Surav (1996b)). The proof will, very roughly, follow that of McCarthy (1993) (see also McCarthy and Buvač (1998)).

4.1. THE RELEVANCE OF STABILITY

The proof of the Blocks World theorem will be instructive for several reasons. Notable among these is the fact that McCarthy's proof, transcribed into CL, assumes a principle that is in fact invalid in CL, namely:[16]

(15) $$(ist(c, \tau) \rightarrow ist(c, \tau')) \rightarrow ist(c, \tau \rightarrow \tau').$$

For (15) to follow, one must assume $Stable(c, \tau \rightarrow \tau')$.[17] There is an important point here. Those interested only in practical applications of logic in AI might want to think of the previous section as an irrelevant, unilluminating technical exercise that gratuitously opens the problematic Pandora's box of logical paradox. To the contrary, however, as this example shows, the notion of stability (or something of the sort) is needed in any account that wishes consistently to attribute to contexts features assumed in many applications of the notion both inside and outside AI. Specifically, these features are: (i) that contexts and propositions (or sentences) are first-class citizens; (ii) that one context can subsume another; (iii) that propositions have a "first-order" form (which, given (i), enables the possibility of self-reference); and (iv) that there is a true-in relation between contexts and propositions.

4.2. The Blocks World Example in CL

The differences between the subjective and the objective conceptions force one to recast the Blocks World example somewhat. Notably, since contexts on the objective conception are (typically fleeting) individual things rather than (eternal) logical theories, for the example to have any interest we must talk not about particular contexts, but about context *types*. More exactly, we must talk generally about contexts in which the "axioms" of the original example – axioms embodying the logical conception of context – are true.

That noted, then, in the Blocks World example, one starts with a context type *above-theory* that expresses several simple principles governing the relations *on* and *above*, viz., that *on(x, y)* implies *above(x, y)*, and that the above relation is transitive. (As in the original example, we allow for a finite number of additional principles as well.) We express the *above-theory* type in the language of CL as follows:

$$
\begin{aligned}
& AT(c) \equiv_{df} \\
& ist(c, [(\forall xy)(on(x, y) \rightarrow above(x, y))]) \wedge \\
(16) \quad & ist(c, [(\forall xyz)((above(x, y) \wedge above(y, z)) \rightarrow above(x, z))]) \wedge \\
& [\dots \text{ additional axioms } \dots].
\end{aligned}
$$

A second context type *blocks* is then introduced that includes the propositions of (some version of) the situation calculus (where contexts are taken to be situations) as well as axioms (unstated here) about the Blocks World. Accordingly, *blocks* contexts are not interested in simple propositions about what is on or above what *simpliciter*, but rather in propositions about what is on or above what in what contexts. Thus, unlike *above-theory* contexts, *blocks* contexts also involve 3-place *on* and *above* relations.

Therefore, it is essential that *blocks* contexts be able to relate their 3-place *on* and *above* relations to their 2-place counterparts in *above-theory* contexts. Consequently, it entails two propositions to this effect. The *blocks* type can then be represented in the language of CL as follows:[18]

$$\begin{aligned}(17)\quad & Bl(c) \equiv_{df} \\ & ist(c, [(\forall xyz)(on(x, y, z) \leftrightarrow ist(z, on(x, y)))]) \wedge \\ & ist(c, [(\forall xyz)(above(x, y, z) \leftrightarrow ist(z, above(x, y)))]) \wedge \\ & [\text{... situation calculus axioms ...}] \wedge \\ & [\text{... Blocks World axioms ...}].\end{aligned}$$

The *above-theory* context is now "lifted" into the blocks context by asserting that, in any *blocks* context, every context is an *above-theory* context, i.e., a context in which the facts coded into the *AT* predicate hold. This is expressed by the following sentence:

(18) $$(\forall x)(Bl(x) \rightarrow ist(x, (\forall y)(Context(y) \rightarrow AT(y)))).$$

As noted, the proof requires the following CL lemma:

(19) $$(\forall x)(Stable(x, \tau \rightarrow \tau') \rightarrow ((ist(x, \tau) \rightarrow ist(x, \tau')) \rightarrow ist(x, \tau \rightarrow \tau'))).$$

All we need to assume to apply this proposition is the reasonable thesis that atomic propositions involving the 3-place *on* and *above* relations are stable at all contexts:

(20) $$(\forall wxyz)(Stable(w, on(x, y, z)) \wedge Stable(w, above(x, y, z)))$$

which yields

(21) $$(\forall wxyz)(Stable(w, on(x, y, z) \rightarrow above(x, y, z))).$$

We now turn to the theorem (McCarthy and Buvač (1998), p. 21).

THEOREM: $(\forall w)(Bl(w) \rightarrow ist(w, (\forall xyz)(on(x, y, z) \rightarrow above(x, y, x))))$.

Proof:[19] Suppose

(22) $$Bl(c)$$

and that

(23) $$ist(c, on(a, b, c')).$$

From (17) and our local propositional logic, we have

(24) $$ist(c, [(\forall xyz)(on(x, y, z) \rightarrow ist(c', on(x, y)))]).$$

From (23), by L17 (and our local MP L14), it follows that

(25) $$ist(c, E!(a, b, c')),$$

so by L11 we have

(26) $$ist(c, [on(a, b, c') \to ist(c', on(a, b))]).$$

Thus, by (23) and (26) we have

(27) $$ist(c, ist(c', on(a, b)))$$

and, by L27, that

(28) $$ist(c', on(a, b)).$$

By (18), (25), and L11,

(29) $$ist(c, AT(c')).$$

So, by the definition (16) of AT,

(30) $$ist(c, ist(c', [(\forall xy)(on(x, y) \to above(x, y))])).$$

Thus, by L27 again,

(31) $$ist(c', [(\forall xy)(on(x, y) \to above(x, y))]).$$

By localism L17 and (28) it follows that

(32) $$ist(c', E!(a, b))$$

and so by a couple applications of L11 and L14, (31) and (32) yield

(33) $$ist(c', above(a, b)).$$

From (17) we have

(34) $$ist(c, [(\forall xyc')(ist(c', above(x, y)) \to above(x, y, c'))]),$$

and so from (25), (34), L11 and L14,

(35) $$ist(c, [ist(c', above(a, b)) \to above(a, b, c')]).$$

By (25), we have $ist(c, E!(c'))$, i.e.,

(36) $$c' \preccurlyeq c,$$

and so that fact, together with (33) and (13), entail

(37) $$ist(c, ist(c', above(a, b))).$$

Hence, by (35), (37), and L14), we have

(38) $$ist(c', above(a, b, c')),$$

so by the deduction theorem, it follows from $Bl(c)$ alone that

(39) $$ist(c, on(a, b, c')) \rightarrow ist(c, above(a, b, c'))$$

Thus, given (19) and (21), (39) yields

(40) $$ist(c, [on(a, b, c') \rightarrow above(a, b, c')]),$$

and so by M2

(41) $$ist(c, (\forall xyz)(on(x, y, z) \rightarrow above(x, y, z))),$$

and so, by the deduction theorem again, we have

(42) $$Bl(c) \rightarrow ist(c, (\forall xyz)(on(x, y, z) \rightarrow above(x, y, z))),$$

and so by GC we have our theorem:

(43) $$(\forall w)(Bl(w) \rightarrow ist(w, (\forall xyz)(on(x, y, z) \rightarrow above(x, y, z)))).$$

5. Conclusion

In this paper I've attempted to clarify the objective conception of context and develop a logic for it. On this conception, contexts are taken to be (in general) limited pieces of the real world, whose denizens include not only ordinary individuals, but contexts and propositions as well. The objective conception stands in contrast to the subjective conception, which understands contexts to be (something like) sets of propositions believed by rational agents. The objective conception is, of course, not a philosophical *competitor* to the subjective conception. Both conceptions are valid, stemming from different intuitions, and both, hopefully, will prove useful in appropriate domains. In particular, a logic for the objective conception appears to be the appropriate formal foundation for enterprise modeling and, more generally, for knowledge representation applications in which the standpoint of the representing agent is considered to be outside the space of contexts being represented.

Acknowledgements

I am indebted to one of the anonymous reviewers of this paper for forcing me to be clearer about the motivations for the objective conception, and to another for indicating the need for further exposition regarding the language $\mathcal{L}$. The basics of the logic CL as well as its semantics were developed when I was a visitor in the Division of Manufacturing Technology at the Commonwealth Scientific and Industrial Research Organization (CSIRO), Melbourne Laboratory. My thanks to Dr. Laszlo Nemes for his great hospitality and for making my stay possible. Thanks also to Saša Buvač, Varol Akman, and my colleague Colin Allen for their comments and encouragement. An earlier version of this paper was presented at the AAAI Fall 1997 Symposium on Context in Knowledge Representation and Natural Language.

Notes

[1]A particularly good, up-to-date list of references is found in McCarthy and Buvač (1998). See also the useful overview by Akman and Surav (1996a).

[2]$\mathcal{L}$ is counted as an intensional logic because it contains terms that denote (in $\mathcal{L}$'s intended model theory) entities – viz., propositions – for which an standard extensional substitutivity principle fails, specifically, the principle

$$ist(c, p) \leftrightarrow ist(c', p') \rightarrow ist(c, p) = ist(c', p')$$

that is, the principle that equivalent propositions are identical. Because it fails, it follows that propositions are not truth values, and hence are not themselves extensional.

[3]I am assuming a "rigid" account of names, so that identities are assumed to be absolute across contexts. (The *locus classicus* of this view, of course, is Kripke (1972).) This seems to be the most appropriate account of names for the objective conception.

[4]Strictly speaking, neither the language of Bealer (1982) nor that of Turner (1987) has expressions that are both terms and formulas. But the reason for this is simply that these two theories are theories not only of propositions but of properties and relations as well, and the terms for the latter entities in general require a variable binding apparatus (to capture their adicity) that causes them to differ structurally from formulas. However, were these two theories to restrict their focus to propositions, the variable binding apparatus would be unnecessary and any syntactic differences between terms and formulas that remained (eg., square brackets or Turner's bare λ operator) could be eliminated.

[5]Atomic localism is the situation theoretic analog of the thesis of *serious actualism* in the philosophy of modality, according to which objects have properties only at those possible worlds in which they exist (cf., eg, Adams (1981), Menzel (1990), Menzel (1991), and Plantinga (1983)).

[6]More formally, c is weakly partial if there is some object e such that for no $n + 1$-place relation R ($n \geq 0$) and individuals $e_1, \ldots, e_n$, is it the case that the proposition $R(e_1, \ldots e, \ldots e_n)$ is true in c.

[7]The distinction between internal and external perspectives has an exact parallel in the semantics of quantified modal logic in the distinction between "internalism" and "perspectivalism". This is hardly surprising; formally, very little separates a context from a possible world. Prior (1957) is still the most cogent defense of internalism. Perspectivalism is defended in Menzel (1993); see also Adams (1981) for a sort of intermediate position.

[8]I emphasize that this notation is being used only as an abbreviation, unlike McCarthy (1993), where reference to an "outer context" in each formula is essential. Note also that I am not exactly sure of the nature of the ':' predicate in McCarthy's (and Buvač's) work. It usually appears to be an object language predicate, but, for example, it is not part of the lexicon from which Buvač (1994) builds his language, and hence appears there to be a metalinguistic predicate.

[9]A universal closure of a formula φ is the result of affixing universal quantifiers $\forall v$ to φ, for at least every variable v occurring free in φ. Note that the stipulation that axioms be universal closures of instances of the schemas means that all theorems are all sentences, i.e., they contain no free variables. The reason for this is the special semantics for variables that is used in the model theory for $\mathcal{L}$ to make the theory of propositions work smoothly.

[10]Note that the definition of $\Gamma \vdash \varphi$ prevents one from using GC to derive, e.g., $(\forall x)P(x)$ from $P(a)$.

[11]To make this an official extension of the language, we would add the axiom $\ulcorner(\forall x)\varphi^x_c \leftrightarrow (\forall x)(Context(x) \rightarrow \varphi)\urcorner$ for any such variable c.

[12]Philosophically it is natural to postulate logical functions for all the boolean operators and the existential quantifier, but they are formally unnecessary, and so are omitted for purposes here.

[13]Interestingly, L27 is a valid principle of Buvač (1994), which is presented as an attempt to formalize the ideas behind McCarthy and Buvač (1998). In fairness, there are sections of McCarthy and Buvač (1998) (eg, Section 6) where the authors seem clearly to have more of an objective conception in mind for the given application.

[14]There is nothing in the logic that prevents such "self-aware" contexts, nor should there be; see esp. Barwise (1989), ch. 8 for numerous examples.

[15]It should be noted that recent work by Antonelli (1994) and Kremer (1993) suggest that CL will be essentially incomplete. Though undesirable, perhaps, since the purpose of the logic at this point is for representation and conceptual clarification rather than reasoning, I do not consider this a serious liability, but rather simply a reflection of the theory's expressive power.

[16]For a counterexample, let $\mathcal{M}$ be a model and suppose τ and τ' are both unstable relative to a context c in $\mathcal{M}$ and that they are out of phase with one another relative to c, i.e., at any level in the construction that generates $\mathcal{M}$, τ is true at c iff τ' is false at c. Then (I abuse notation slightly here) the conditional $ist(c, \tau) \rightarrow ist(c, \tau')$ is true in $\mathcal{M}$ because the antecedent is false (since τ is unstable at c) but $ist(c, \tau \rightarrow \tau')$ is false in $\mathcal{M}$. For when τ and τ' are both unstable at c and out of phase, the conditional proposition $\tau \rightarrow \tau'$ is unstable at c as well (it will be true at c in stages where τ is false, and false at c in stages where τ is true).

I should note that McCarthy and Buvač (1998) (p. 22) actually use a lifting axiom in their proof (based on their principle **discharge**): $ist(assuming(c, p), q) \rightarrow ist(c, p \rightarrow q)$, of which (15) seems an accurate reconstruction in CL.

[17]Or at least, that appears to be the only way to prove it. Suppose that $Stable(c, \tau \rightarrow \tau')$ and that $ist(c, \tau) \rightarrow ist(c, \tau')$, but that $\neg ist(c, \tau \rightarrow \tau')$. Then by stability, we have $ist(c, \neg(\tau \rightarrow \tau'))$, i.e. $ist(c, \tau \wedge \neg\tau')$. It follows from our local propositional logic that $ist(c, \tau)$ and $ist(c, \neg\tau')$, but we assumed $ist(c, \tau) \rightarrow ist(c, \tau')$, hence we also have $ist(c, \tau')$, and hence, by local propositional logic again, $ist(c, \tau' \wedge \neg\tau')$, contradicting L28. So $ist(c, \tau \rightarrow \tau')$.

[18]That we are limited to finitely axiomatizable context types on this approach is probably no great source of concern for AI researchers, but in fact it is no intrinsic limitation, as the language and semantics of CL can be extended to allow for special context predicates that can be used to build atomic formulas (not just terms), and which could then, if necessary, be axiomatized with countably many axioms. (Indeed, such predicates might be required in any case if we wish to develop a more robust theory of contexts based on, say situation theory or the situation calculus.) For example, if there were denumerably many Blocks World axioms, then if 'Bl' were a special context predicate we could have an axiom of the form $Bl(c) \rightarrow \varphi$ for each such axiom φ.

[19]Given the way the logic is set up, it is easiest here to suppose now that c and c' are individual constants.

References

Adams, R. (1981), 'Actualism and Thisness', *Synthese* 49, 3–41.

Akman, V. and Surav, M. (1996), 'Steps toward Formalizing Context', *AI Magazine* 17(3), 55–72.

Akman, V. and Surav, M. (1996), 'The Use of Situation Theory in Context Modeling', *Computational Intelligence* 12(4), 427–438.

Antonelli, A. (1994), 'The Complexity of Revision', *Notre Dame Journal of Formal Logic* 35, 67–72.

Barwise, J. (1989), *The Situation in Logic*, CSLI Lecture Notes, Number 17, Stanford, California: CSLI Publications.

Barwise, J. and Perry, J. (1986), *Situations and Attitudes*, Cambridge, Massachusetts: Bradford Books/MIT Press.

Barwise, J. and Seligman, J. (1997). '*Information Flow: The Logic of Distributed Systems*', Cambridge: Cambridge University Press.

Bealer, G. (1982), *Quality and Concept*, Oxford: Clarendon Press.

Bealer, C., and U. Mönnich (1989), 'Property Theories', in D. Gabbay and F. Guenthner (eds.), *Handbook of Philosophical Logic*, vol. 4, Dordrecht: D. Reidel.

Borgida, A., Brachman, R., McGuinness, D, and Resnick, L. (1989), 'CLASSIC: A Structural Data Model for Objects', in *Proceedings of the 1989 ACM SIGMOD International Conference on Management of Data*, 59–67.

Buvač, S. (1994), 'Quantificational Logic of Context', in *Proceedings of the 13th National Conference on Artificial Intelligence.*

Buvač, S., V. Buvač, and I. Mason (1995), 'Metamathematics of Context', *Fundamenta Informaticae* 23(3), 263–301.

Chierchia, G., Partee, B. and Turner, R. (eds.) (1989), *Properties, Types, and Meaning, Volume I, Foundational Issues*, Dordrecht: Kluwer Academic Publishers.

Chierchia G. and Turner, R. (1988), 'Semantics and Property Theory', *Linguistics and Philosophy* 11, 261–302.

Devlin, K. (1991), *Logic and Information*, Cambridge, UK: Cambridge University Press.

Dowty, D., Wall, R., and Peters, S.. (1981), *Introduction to Montague Semantics*, Dordrecht: D. Reidel Publishing Company.

Enderton, H. (1972), *A Mathematical Introduction to Logic*, New York: Academic Press.

Fox, M. S., and Grüninger, M. (1994), 'Ontologies for Enterprise Integration', in *Proceedings of the 2nd Conference on Cooperative Information Systems*

Giunchiglia, F. (1993), 'Contextual Reasoning', *Epistemologia* 16, 345–364.

Gruber, T.R. (1995), 'Toward Principles for the Design of Ontologies Used for Knowledge Sharing', International Journal of Human and Computer Studies 43(5/6), 907–928.

Guarino, N. (1995), 'Formal Ontology, Conceptual Analysis and Knowledge Representation', *International Journal of Human and Computer Studies* 43(5/6), 625–640.

Guarino, N. (ed.) (1998), *Formal Ontology in Information Systems: Proceedings of the 1st International Conference*, Amsterdam: IOS Press.

Gupta, A. (1982), 'Truth and Paradox', *Journal of Philosophical Logic* 11, 1–60.

Herzberger, H. (1982), 'Notes on Naive Semantics', *Journal of Philosophical Logic* 11, 61–102.

Kremer, P. (1993), 'The Gupta-Belnap systems S# and S* are not Axiomatisable', *Notre Dame Journal of Formal Logic* 34, 583–596.

Kripke, 5. (1972), *Naming and Necessity*, Cambridge, Massachusetts: Harvard University Press.

Kripke, S. (1975), 'Outline of a Theory of Truth', *The Journal of Philosophy* 72, 690–716.

Lewis, D. (1986), *On the Plurality of Worlds*, Oxford: Basil Blackwell.

McCarthy, J. (1979), 'First-order Theories of Individual Concepts and Propositions', in Michie, D. (ed.) Machine Intelligence, Edinburgh: Edinburgh University Press.

McCarthy, J. (1987), 'Generality in Artificial Intelligence', Communications of the ACM 30(12), 1030–1035.

McCarthy, J. (1993), 'Notes on Formalizing Context', in *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*.

McCarthy, J., and Buvač, S. (1998), 'Formalizing Context (Expanded Notes)', in Aliseda A., R. van Glabeek, and D. Westersthhl (eds), *Computing Natural Language*, CSLI Lecture Notes, vol. 81, Stanford University: CSLI Publications.

MacGregor, M. (1991), 'Using a Description Classifier to Enhance Deductive Inference', *Proceedings of the Seventh IEEE Conference on AI Applications*, 141–147.

Menzel, C. (1990), 'Actualism, Ontological Commitment, and Possible World Semantics', *Synthese* 85, 355–389.

Menzel, C. (1991), 'The True Modal Logic', *Journal of Philosophical Logic* 20, 331–374.

Menzel, C. (1993), 'Singular Propositions and Modal Logic', *Philosophical Topics* 21, 113–148.

Menzel, C. (1997), 'Modeling Method Ontologies: A Foundation for Enterprise Model Integration', in Grüninger, M. and Farquhar, A. (eds.), *Ontological Engineering: Papers from the 1997 AAAI Spring Symposium*, Menlo Park: AAAI Press, Technical Report SS-97-06, 73–83.

Menzel, C., and Mayer, R. (1998), 'The IDEF Family of Languages', in Bernus, P. Mertins, K., and Schmidt, G. (eds.), *Handbook on Architectures for Information Systems*, New York: Springer-Verlag.

Menzel, C. Mayer, R., and Sanders, L. (1993), 'Representation, Information Flow, and Model Integration', in Petrie (1993), 131–141.

Neches, R., Fikes, R., Finin, T., Gruber, T., Patil, R., Senator, T., and Swartout, W.R. (1991), 'Enabling Technology for Knowledge Sharing', *AI Magazine* 12(3), 36–56.

Petrie, A. C. (ed.) (1993), *Enterprise Model Integration: Proceedings of the First International Conference*, Cambridge, Massachusetts: MIT Press.

Plantinga, A. (1983), 'On Existentialism', *Philosophical Studies* 44, 1–20.

Prior, A.N. (1957), *Time and Modality*, Oxford: Oxford University Press.

Quine, W.V.O. (1981), 'On the Individuation of Attributes', reprinted in *Theories and Things*, Cambridge, Massachusetts.: Harvard University Press, 100–112.

Salmon, N. (1986), *Frege's Puzzle*, Cambridge, Massachusetts: MIT Press/Bradford Books.

Shapiro, S. (1991), *Foundations without Foundationalism: A Case for Second-order Logic*, Oxford: Clarendon Press.

Soames, S. (1987), 'Direct Reference, Propositional Attitudes, and Semantic Content', *Philosophical Topics* 15, 47–87.

Turner, R. (1987), 'A Theory of Properties,' *Journal of Symbolic Logic* 52, 455–472.

Zalta, E. (1988), *Intensional Logic and the Metaphysics of Intentionality*, Cambridge, Massachusetts: MIT Press/Bradford Books.