CLASSIFICATION WITHOUT ASSERTION*

ERIC MCCREADY
Aoyama Gakuin University

1 Introduction

Numeral classifiers are found in many languages with bare nominals. In such languages, nominals without (overt) determiners are grammatical in general. Bare nominals can be interpreted as definite or indefinite, and as singular or plural, as in (1). However, in cases where something is to be indicated about a particular number of objects, it is impossible (again, in general) to use a numeral quantifier alone, as shown by (2). Instead, it is necessary to use a phrase of the form *Num+Classifier(+Genitive)*, where the *Classifier* is a term that limits the counting to a specific kind of object. The genitive is required in certain contexts in some classifier languages, such as Japanese, which will be the focus of this paper. Sentences of this form are interpreted as indefinite, or at minimum, as partitive:

(1) zasshi-o katta
magazine-ACC bought

'(I) bought a magazine/the magazine/some magazines/the magazines.'

(2) * san zasshi-o katta
three magazine-ACC bought

(intended) 'I bought three magazines.'

(3) san-satsu-no zasshi-o katta
three-CL.printed.matter-GEN magazine-ACC bought

'I bought three magazines.'

In this paper I consider only non-measure classifiers. This domain is already broad: It contains classifiers that pick out various idiosyncratic domains. Some 'vanilla' examples are *hon*, which is used with nominals that denote long, thin objects; *satsu*, used with nominals denoting various

*Thanks to Pranav Anand, David Beaver, Chris Davis, Bart Geurts, Christopher Potts, and audiences at Kyoto University, FAJL 5, ICCS 2008, and the Tokyo Dynamic Semantics Meeting.

kinds of printed and bound materials; *hiki*, for small animals; *too*, for large animals, and so on. More idiosyncratic and perhaps peculiar examples include *hashira*, for gods; *shoo*, for wins (in e.g. sports), and the interesting *kaku*, for cup-saucer pairs, and *zen*, for pairs of chopsticks. These last two will play a significant role in the present paper, for the reason that they count pluralities rather than singularities.

The main question addressed by the present paper is: What is the meaning of these classifiers? The meaning of the combined numeral quantifier and classifier is intuitively obvious: To give a count of the number of objects of which the sentence is to hold, i.e. to serve as a quantifier. In the following, for example, *wu ben* indicates that the otherwise unspecified number of books bought by the speaker is (at least) five.

(4) wo mai le wu ben shu
I buy ASP five CL book

'I bought five books.' (Chinese)

(5) $\exists X[|X| = 5 \wedge book^*(X) \wedge bought^*(s,X)]$

But this content seems to come solely from the numeral *wu*. What is the content of *ben*? Does it even have content? Could we think of it merely as a piece of grammatically induced morphology, i.e. as something like agreement? Perhaps this is indeed the source of the morphology, at least in some sense. However, this does not mean that the classifier carries no content, as we will see. In particular, I will argue that classifiers bear a kind of pragmatic content: conventional implicature.

Conventional implicatures are a kind of pragmatic content that is in some respects very close to semantic content. The difference between the two can be clarified by considering examples with nominal appositives (Potts, 2003, 2005) such as (6). Here, it is asserted by the sentence that John enjoys the outdoors; it is also indicated, in some way distinct from standard assertions, that John is a cyclist. This latter type of content is conventionally implicated. That this indication is indeed distinct from assertion can be seen by considering the effect of denying (6) by e.g. using (7): By doing so, one denies that John is a nature lover, but not that he likes bicycles.

(6) John, a cyclist, enjoys the outdoors.

(7) That's not true.

Conventional implicatures have received substantial recent attention in semantics, pragmatics and philosophy of language. There is still controversy about the correct way to characterize their properties, or even if conventional implicatures as a category are truly independent of other kinds of pragmatic meaning such as presupposition (c.f. Schlenker (2007), as well as other papers in that volume). But it is common ground that conventional implicatures (at least of the type exemplified by nominal appositives) are independent of so-called 'at-issue' content. This idea has been given a formal implementation embodied in the logic $\mathcal{L}_{CI}$ (Potts, 2005): According to this logic, conventionally implicated content and at-issue content are completely disjoint in type and do not interact at all.

However, not all such content is fully separate, e.g. pejoratives (Williamson, 2009). In (8), *Limey* plays a dual role:[1] It conveys a negative attitude toward a group of people and it contributes a predication to compositional semantics.

[1]I hope no one will be offended by the use of this relatively mild pejorative. I myself do not subscribe to the relevant attitudes. In any case, the pejorative here is mentioned, not used.

(8) Jimmy is a Limey.

I will argue that classifiers also introduce *mixed content* of this type. Specifically, classifiers conventionally implicate that the quantification carried out by the numeral they associate with takes place over a particular domain; this amounts to a temporary restriction of the domain of discourse (Gauker, 1998). At the same time, they put the content denoted by the nominal into a form appropriate for quantifying over, roughly in the way indicated by Chierchia (1998). Since $\mathcal{L}_{\mathrm{CI}}$ cannot handle this kind of mixed content, formalizing this picture requires a modification of the logic. I will make use of the extension of $\mathcal{L}_{\mathrm{CI}}$ to $\mathcal{L}_{\mathrm{CI}}{}^{+}$ proposed by McCready (2010).

The structure of the paper is as follows. The next section, section 2, will make the case that classifiers do indeed bear mixed content. Section 3 will give a formal implementation of the idea within $\mathcal{L}_{\mathrm{CI}}{}^{+}$, which will also be outlined there. Section 4 sums up and indicates some possible extensions of the proposed analysis. It also considers possible alternative formulations of the general view expressed in the paper.

2 Classifier Meaning as Conventionally Implicated

The aim of this section of the paper is to determine just what the content of numeral-classifier combinations amounts to. The basic picture seems clear enough: The combination functions as a quantifier. In a sentence like (4), five books are said to have been bought. But this content obviously comes from the numeral. The question therefore is: What is the contribution of the classifier? The claim that I will make is that its contribution is dual. It functions to restrict the domain of quantification, as stated above: This part of its meaning is, I will claim, conventionally implicated. However, classifiers also work to individuate the domain of objects for quantification, a claim made by many authors. The usual evidence for this is that numeral quantification without classifiers is not grammatical. Here, I will show that choice of classifier can also impact truth conditions in some cases. If it is indeed correct that classifiers bear content in both conventionally implicated and truth-conditional domains, it means that they are mixed content bearers; this observation motivates the treatment of them using $\mathcal{L}_{\mathrm{CI}}{}^{+}$ provided in section 3.

The first question to be addressed is the semantic status of the classifier meaning: Is it truth-conditional, or meaning of some other kind? It seems clear that the numeral introduces truth-conditional content, for if the number of individuals does not match the numeral, the sentence is just false. If one uses (4) in a situation where only four books have been bought, the sentence makes a false claim. For quantification to take place, however, it is necessary to individuate the denotation of the nominal; I will turn to this issue in a moment. Before that, I want to consider the meaning of the classifier itself. I claimed in the previous paragraph that the classifier itself carries mixed content, one part of which is related to individuation, and one which has to do with domain restriction. As I will show when I discuss the process of individuation, this content must necessarily be viewed as truth-conditional.

The domain-restricting aspect of the classifier content, however, is not so easily classifiable. To see this, consider what happens if the classifier type does not match the object quantified over. For example, suppose we use the classifier *ken* 'buildings-CL' (Japanese) to count dogs in some sentence. If the 'matching' part of classifier content is part of the truth-conditions, that sentence should be false. The reason is that if *ken* functions to restrict the domain to those objects that satisfy the predicate *building*, then the domain will contain only objects which are buildings; since no building is a dog, there will be no witness for any existential statement made, and falsehood will result. This is the case even if the quantification is universal, given that universal quantifiers in natural language carry existential presuppositions (cf. von Fintel 1994).

Is this prediction correct? It seems that it is not. In such instances the sentence is *not* false; instead it is inappropriate. It is wrong to use classifiers to quantify over objects belonging to inappropriate domains. But this wrongness does not induce real falsity. (9) shows that this is the case.

(9) # otoko-ga 2-satsu haittekita
man-NOM 2-CL(books) entered

Quantificational content: '2 men entered.'
Classifier-content: the objects quantified over have the form of printed, bound matter

This looks highly incoherent, but it is not false—just wrong. But, if we assume that the domain restriction induced by classifiers takes place in the truth-conditional domain, the sentence should be false: If the objects quantified over are required to satisfy *printed.bound.matter*, there will be no object in the domain of which the sentence is true, since $man \cap printed.bound.matter = \varnothing$; therefore there is no available witness for the truth of (9), and the sentence is false, indeed necessarily false on the assumption that no possible object is both a man and a book. But this is not the intuition. This means that it cannot be the case that the domain restriction of the classifier is truth-conditional.

It is therefore necessary to explore the possibility of the classifier bearing other types of meaning, in particular the various kinds of pragmatic meaning. There seem to be three main options: Presupposition, conversational implicature, and conventional implicature. Could the content of the classifier be presuppositional? Assuming that it is would be in line with common assumptions about 'grammatical' meanings such as the content of number and gender features, which are often claimed to be presuppositional in nature (cf. Sauerland 2008). Since it is possible to view the problem of classifier selection as a problem about agreement, this sort of analogy might be a natural one. But to draw this conclusion, we should show that classifier-induced domain restriction exhibits the standard behavior of presuppositions. One such behavior is the effect of inducing truth-valuelessness, or undefinedness, etc., when the presupposition is not satisfied (see e.g. Levinson 1983, Beaver 2001 for options); this seems a reasonable enough way of thinking about what happens when the classifier used is not an appropriate one. We also expect that the putative presupposition will escape from presupposition 'holes' such as negation and the scope of modal operators; this is also known as presuppositional invariance:

(10) a. It is not the case that John's daughter is beautiful.
$\dashrightarrow$ John has a daughter

b. John's daughter might be beautiful.
$\dashrightarrow$ John has a daughter

Both of the above require that John have a daughter to be interpretable, which is explained by the presuppositional nature of possessives. Does the content of classifiers follow this pattern? Examining examples yields a positive answer:

(11) # otoko-ga 2-satsu [haittekita to iu koto-wa nai | haittekite-nai]
man-NOM 2-CL entered C say thing-TOP COP.NEG come.in-NEG

‘It is not the case that 2 men (books) entered.’ (not metalinguistic)

The content of classifiers does project from negation, and similarly for the other standard presupposition tests. So far, the data is all consistent with the domain restriction being presupposed.

Another test for presuppositionality is so-called ‘binding’ behavior (cf. van der Sandt 1992). Binding of presuppositions can take place in conditionals and similar universal constructions. In essence, if a sentence S carries a presupposition P, and $S' \models P$, then in construction $S' \Rightarrow S[P]$, no presupposition is projected. Here the presupposition is, in some sense, bound or captured by the conditional antecedent. In the following example, for instance, the consequent presupposes that John has a daughter: Since this is precisely the content of the antecedent, no presupposition is projected:

(12) If John has a daughter, John’s daughter must be pretty.
$\dashrightarrow$ no presupposition

We can now ask: Can the putative presuppositions associated with classifiers be ‘bound’? It seems that they cannot. I have been unable to find any approximation of a reasonable sounding example in which (9) could be coherent. Here is my best attempt:

(13) otoko-ga hon dattara otoko-ga 2-satsu haittekita
man-NOM book COP-COND man-NOM 2-CL entered

‘If (the) men were books, 2 men entered.’
$\dashrightarrow$ the objects under discussion are books

This sentence is basically nonsense. If it can be made coherent at all, it sounds metalinguistic, as if the meaning of *otoko* ‘man/men’ is being redefined, or perhaps as if the meaning of the classifier itself is. But this is certainly not a ‘bound’ reading of the putative presupposition. Since binding is impossible, and therefore classifier domain restriction does not exhibit all the behavioral characteristics of presuppositions, I conclude that presupposition is not the right category for the content of classifiers.

Some further evidence for this conclusion comes from the behavior of classifier content in environments that function as ‘plugs’ for presupposition, so that the presuppositions do not project as usual, but rather in a modified form. An example is the complements of certain attitude verbs, e.g. *believe*: Here presuppositions appear as beliefs of the attitude holder, rather than as ‘objective’ facts or content taken for granted by the speaker:

(14) Mary believes that the king of Gambia is bald.
$\dashrightarrow$ Mary believes that Gambia has a king.

How does classifier content act in belief contexts? On the assumption that *omou* ‘think’ is a reasonable Japanese translation for *believe* (I include the standard *shinjiru* ‘believe’ as well, though its semantics look somewhat unlike the English version), the following example shows that the domain restrictions associated with classifiers pass unmodified through presupposition plugs:

(15) Shintaro-wa Taro-ga hon-o ni-satsu katta to
Shintaro-TOP Taro-NOM book-ACC 2-CL[printed.matter] bought C
{omotteiru|shinjiteiru}
thinks/believes

'Shintaro thinks/believes that Taro bought two books.'
⇢ The things Taro bought were bound, printed objects.

This observation is further evidence against taking classifier content to be presupposed. Together with the binding facts, we have good reason to look elsewhere for a meaning category for classifier content.

A second pragmatic option is conversational implicature. But this clearly cannot be the right way to go: Classifier content has all the wrong features. Conversational implicatures follow from reasoning on the part of hearers about the utterances of their interlocutors, on the basis of assumptions about cooperativity in communication (Grice, 1975).[2] As a result, the speaker of the sentence that induces the implicature is not responsible for its content, meaning that the implicature can be cancelled. Classifier content is neither inference-based nor cancellable. This option is just a nonstarter.

We have one remaining option, however: We can think of classifier content as conventionally implicated (cf. Bach 1999, Potts 2005). Conventionally implicated content has a number of properties which are more or less common ground to researchers working on the topic. First, it is triggered by particular lexical items or constructions, such as the nominal appositive in the following example:

(16) John, a swimmer, is a nice guy.
⇢ John is a swimmer

Second, conventionally implicated content is scopeless, or, otherwise put, it exhibits projection behavior like that seen with respect to presuppositional invariance above. It is worth noting that this observation does not hold in every case, as discussed in detail by Wang et al. (2005) and Amaral et al. (2008). Still, scopelessness is found in many (or even most) situations, and I will make use of this test in the present investigation. In the following sentence, then, the content of the nominal appositive does not fall in the scope of negation:

(17) It is false that John, a swimmer, is a nice guy.
⇢ John is a swimmer

Third, it is unable to be bound like presuppositions. This is the main feature of conventionally implicated content that distinguishes it from presupposition, or, at least, the main feature that has not been called into question (to my knowledge) by researchers working on the topic. In the following example, the content of the nominal appositive, again, 'projects' from the conditional; this seems to be the reason for the infelicity of the example, as, for reasons of Gricean economy, it is misleading (or just unnecessary) to conditionalize over content which is already taken for granted in a context; but this is just the result of using the nominal appositive:

[2] See Chierchia (2004), among others, for a different view according to which implicatures are computed in the grammar.

(18) If John is a swimmer, then John, a swimmer, is a nice guy.
$\dashrightarrow$ John is a swimmer

How does the classifier content fare with respect to these criteria? It in fact has all the properties of conventional implicature that I have mentioned: First, classifier content is plainly tied to the classifiers themselves, all lexical items; second, it shows projection behavior, as we have seen; and, third, it cannot be bound. It thus seems reasonable to take the classifier content to be conventionally implicated—indeed, all the other options, as far as I can see, fail. If this is right, the picture of the semantic effect of numeral classifers becomes the following: Numerals quantify over individuals in more or less the usual way, but classifiers do not restrict the domain of application but instead provide conditions that help guide the interpretation of utterances (like, perhaps, other conventional implicatures).

Interestingly, these sorts of facts do not seem to be restricted to 'classifier languages,' or even to languages which use bare nominals. We can even find related facts in English, in at least two areas. The first is the so-called 'terms of venery,' which are words originally used by aristocracy for animal groups such as *gaggle* for geese, *pride* for lions, and so on. Using the wrong term of venery yields very similar weirdnesses to using the wrong classifier. We also find similar behavior with respect to projection and binding:

(19) ?! John saw a pride of geese.

(20) John might see a pride of lions.

Another related set of facts involves measure phrases and quantifiers such as those in (21). If I use the wrong measure phrase, it is again odd in precisely the same way as the Japanese classifiers are; we also get similar projection and binding behavior as well:

(21) John cut two slices of bread/drank two cups of coffee.

(22) a. ??John drank two slices of coffee.
b. ??John might drink two slices of coffee.

I have argued extensively that the domain selection associated with classifiers is conventionally implicated. It is now time to return to the issue of individuation. Can this be viewed as conventionally implicated as well? The answer is negative, for two reasons. The first is technical: Without the classifier operating in the truth-conditional domain, the input to the quantifier will be left in a non-individuated form which is unsuitable for quantifying over. The second reason is that the choice of classifier can have a direct effect on truth conditions. Here is an example. Japanese has two distinct classifiers that can be used with chopsticks: A rather general one, *hon*, which is used for stick-shaped objects, and a rather specific one, *zen*, which is used for pairs of chopsticks only. Now consider the following scenario:

(23) Situation: there are 2 chopsticks.

a. hashi-ga 2-hon aru
chopstick-Nom 2-CL[stick-shaped] exist

‘There are 2 chopsticks.’

b. hashi-ga 2-zen aru
chopstick-Nom 2-CL[pair.of.chopsticks] exist

‘There are 2 pairs of chopsticks.’

Here (23a) is true, but (23b) is false. But the only difference between the two sentences (23a) and (23b) is the choice of classifier. This means that classifiers must also make a truth-conditional contribution.

In this section, I have argued that classifiers make a dual contribution to meaning: A conventionally implicated domain restriction, and an individuation of objects for quantification that takes place in the truth-conditional domain. In the next section, I will indicate how this contribution can be implemented in a formal system.

3 A Formal Treatment

If the above argumentation is right, two things are needed for an analysis: An analysis of the truth-conditional component of classifiers, and an analysis of their conventionally implicated part. It is also necessary to ensure that both parts of the analysis can be employed simultaneously. In this section, I will give a formal analysis with the indicated features.

The truth-conditional part of the content is relatively simple in outline, as it involves well-studied issues in semantic theory. I will build the analysis on top of the theory of Chierchia (1998), according to whom bare nominals in Japanese (etc.) denote mass terms with lattice structure (cf. Link 1987, i.a.). As usual, these lattices contain both atomic and nonatomic individuals. Classifiers can, in this picture, be viewed as picking out a subset of these individuals for quantification. In general, these will be the atomic individuals; but in some cases, as that of *zen*, they will be nonatomic (here chopstick pairs).

I will not make use of the full complex structure of the Chierchia analysis, which involves a complex (though intuitive) array of type-shifting rules. Instead, I will suppose a denotation for classifiers as operators on quantifiers. Plainly, this is a simplification, but it will be enough to show the general contribution of classifiers to truth-conditions, and to give the flavor of the mixed content that is the main focus of the paper. Sample lexical entries for the classifiers *hon* and *zen* are given below.[3]

(24) a. $[[hon]] = \lambda\mathcal{Q}\lambda P\lambda Q[\mathcal{Q}(P \cap At)(Q)]$

b. $[[zen]] = \lambda\mathcal{Q}\lambda P\lambda Q[\mathcal{Q}(\{A : Pr(A) \wedge P^*(A) \wedge At^*(A) \wedge \forall z, A'[z \in A \wedge A \neq A' \rightarrow z \notin A']\})(Q)]$

These lexical entries serve to restrict the elements quantified over to individuals of the appropriate kind. The entry for *hon* forces quantification to take place over only atomic elements of the relevant lattice (that denoted by the nominal complement). The entry for *zen* forces quantification over pairs of atomic elements (which will always turn out to be chopsticks in normal situations). These entries, while obviously not the final word on the subject, should be sufficient for the purposes of this paper.

[3] It is also necessary to add a condition to the entry for *zen* to ensure that there is no overlap in the pairs of chopsticks returned; I omit this complication here.

(*R*1) $\dfrac{\alpha : \sigma}{\alpha : \sigma}$

(*R*2) $\dfrac{\alpha : \langle \sigma^a, \tau^a \rangle, \beta : \sigma^a}{\alpha(\beta) : \tau^a}$

(*R*3) $\dfrac{\alpha : \langle \sigma^a, \tau^a \rangle, \beta : \langle \sigma^a, \tau^a \rangle}{\lambda X.\alpha(X) \wedge \beta(X) : \langle \sigma^a, \tau^a \rangle}$

(*R*4) $\dfrac{\alpha : \langle \sigma^a, \tau^c \rangle, \beta : \sigma^a}{\alpha(\beta) : \tau^c \bullet \beta : \sigma^a}$

(*R*5) $\dfrac{\alpha : t^c \bullet \beta : \tau^a}{\beta : \tau^a}$

(*R*6) $\dfrac{\alpha : \sigma}{\beta(\alpha) : \tau}$ (where β is a designated feature term)

Figure 1: $\mathcal{L}_{\mathrm{CI}}$ combinatorics.

The above takes care of the truth-conditional part of the content. The remainder of the content was argued above to be conventionally implicated. For this part of the content, I will give an analysis in the system of Potts (2005), as the most developed analysis of the computation of conventional implicatures currently on the market. As we will see, though, the basic system is not sufficient, and it will become necessary to use the modified system proposed by McCready (2010).

The basic picture of the Potts system is as follows. Analysis takes place in a language called $\mathcal{L}_{\mathrm{CI}}$, which is a language consisting of a type specification and a set of combinatoric rules for the typed objects. The basic rules and types are the standard ones familiar from as far back as Montague 1974. The new part comes in a complication of the type system: The $\mathcal{L}_{\mathrm{CI}}$ type specification is split into two parts, a subdefinition for 'at-issue' types, which are for standard content (superscripted 'a'), and a subdefinition for 'CI' types, those introducing conventionally implicated content (superscripted 'c'). With the CI types are introduced special combinatoric rules for conventionally implicated meanings, as well as a special type definition: Here, an essential property of the system is that there are only functional types of the forms $\langle \sigma^a, \tau^a \rangle$ and $\langle \sigma^a, \tau^c \rangle$. No types take CI objects as input, so conventionally implicated content does not interact with at-issue content. This feature, together with the combinatoric rules in Figure 1 (of which rules 4 and 5 are especially relevant for the computation of conventional implicatures) results in the universal projection feature of conventional implicatures.[4] Further, all types are either CI types or at-issue types: There are no mixed objects, which will obviously be problematic for the analysis of classifiers under the present view.

[4]I follow the practice of McCready (2010) in putting the combinatorics in terms of derivation rules rather than conditions on trees. I do not think that this move has a large impact on their import.

The obvious implementation of the ideas above is to take the classifier to put a condition on the domain of the quantifier itself, restricting it to quantify only over objects of the relevant sort. This operation clearly corresponds to the intuition about classifier functions expressed above. It can be viewed as a modification of quantifier-typed meanings, roughly in a similar way to what was done with the truth-conditional part of the content. This is a local operation and thus unproblematic for $\mathcal{L}_{\text{CI}}$, which, unlike some categorial grammars, lacks operations that allow nonlocal modification (e.g. categorial grammars which use continuations, such as that proposed by Shan (2005)). The novelty of the proposal comes only in that this restriction takes place in the CI domain rather than in an at-issue setting.

Classifiers will then get lexical entries derived from the following schema for the conventionally implicated content. Here, classifiers are modeled as taking quantifiers as input and yielding a conventional implicature to the effect that only objects of the classifier-specified type are quantified over. If this proves to be incorrect (as in the examples of wrong classifier choices cited above), infelicity will result:

(25) $[[Cl]] = \lambda \mathcal{Q}[\mathbf{Dom}(\mathcal{Q}) \subseteq P] : \langle\langle\langle e,t\rangle^{\text{a}},\langle\langle e,t\rangle^{\text{a}},t\rangle\rangle^{\text{a}},t^{\text{c}}\rangle$, where P is the relevant sort of objects for the classifier.

Note that these domain restriction effects are only temporary; they hold only for the duration of the sentence. This makes sense in terms of the interpretation of conventionally implicated content offered by Potts: Conventional implicatures introduced by a sentence are interpreted in tandem with the truth-conditional content introduced by that sentence. It must be ensured that these domain restriction effects do not persist over discourse in the way pointed out by many authors (e.g. Stanley and Szabo (2000), Peters and Westerstahl (2006)). I will leave this as an implicit assumption here. In any case, it is not well understood exactly how domain restrictions persist and dissipate in discourse, so leaving things underspecified does not seem too pernicious at this point.

The following is an example of how the derivation of the conventionally implicated content of classifiers will go on the present proposal. We use the example in (26).

(26) hon-ga 2-satsu oitearu
book-NOM 2-CL exist

'There are 2 books there.'

The derivation is slightly simplified in that I allow conventionally implicated content (the content of type t^{c} to the right of the '$\bullet$') to 'percolate' up the tree undisturbed. The attentive reader will note that the rules in Table 1 do not allow for this to occur in that no provision is made for the existence of such conventional implicatures; they should in fact be removed by R5. Potts himself provides rules that allow for the presence of conventionally implicated propositions alongside the main derivation. I will provide a different approach in a moment, but for the present will let the derivation remain slightly inexact in order to show a full picture of the content:

$$\dfrac{\dfrac{\dfrac{2 : \langle\langle\langle e,t\rangle^{\text{a}},\langle\langle e,t\rangle^{\text{a}},t\rangle\rangle^{\text{a}} \qquad Cl_{\text{satsu}} : \langle\langle\langle e,t\rangle^{\text{a}},\langle\langle e,t\rangle^{\text{a}},t\rangle\rangle^{\text{a}},t^{\text{c}}\rangle}{2 : \langle\langle\langle e,t\rangle^{\text{a}},\langle\langle e,t\rangle^{\text{a}},t\rangle\rangle^{\text{a}} \bullet Cl_{\text{satsu}}(2) : t^{\text{c}}} \qquad hon : \langle e,t\rangle^{\text{a}}}{2(hon) : \langle\langle e,t\rangle^{\text{a}},t^{\text{a}}\rangle \bullet Cl_{\text{satsu}}(2) : t^{\text{c}}} \qquad oitearu : \langle e,t\rangle^{\text{a}}}{2(hon)(oitearu) : t^{\text{a}} \bullet Cl_{\text{satsu}}(2) : t^{\text{c}}}$$

This is precisely as desired: The classifier content is conventionally implicated, and therefore will have all the properties of such constructions, and the remaining content is truth-conditional. But there is a problem. As the reader will have noted, the lexical entry given does not include the truth-conditional part of the content, which is the part that enables quantification. This is obviously problematic. However, there is a reason for this omission. There is no provision in the $\mathcal{L}_{\text{CI}}$ type specification for terms of this type: There are no types for objects that simultaneously introduce at-issue and conventionally implicated content. It is necessary to use an extended system. Such a system is available, as already mentioned: The extension from $\mathcal{L}_{\text{CI}}$ to $\mathcal{L}_{\text{CI}}^{+}$ proposed by McCready (2010). I will now sketch this system and use it to provide a full semantics for classifier content. The reader interested in the full details should consult McCready (2010).

The main innovation in $\mathcal{L}_{\text{CI}}^{+}$ is the introduction of types for mixed content, and, as a necessary precondition, the introduction of resource-sensitive types for conventionally implicated content. In $\mathcal{L}_{\text{CI}}$, combining objects carrying conventionally implicated meanings with other semantic objects via R4 in Table 1 yields the result of the combination together with the original input, and so the logic is not (fully) resource-sensitive; but this property creates problems when coupled with mixed types. The new resource-sensitive types are called *shunting types*.

To handle objects which bear mixed content, we extend the type specification to include the following clause, for mixed type objects:

- If σ, ρ and τ are at-issue types for $\mathcal{L}_{\text{CI}}^{+}$ and υ is a shunting type for $\mathcal{L}_{\text{CI}}^{+}$, then $\langle\sigma,\tau\rangle \times \langle\rho,\upsilon\rangle$ is a mixed type for $\mathcal{L}_{\text{CI}}^{+}$.

Mixed type objects are paired with λ-terms of the form $\alpha\blacklozenge\beta$, so we get typed terms of the form $\alpha\blacklozenge\beta : \langle\sigma,\tau\rangle \times \langle\rho,\upsilon\rangle$. We thus add types operating simultaneously in at-issue and CI dimensions. A new rule for mixed type application is also necessary; this is given in R7.

$$\text{(R7)}\quad \frac{\alpha\blacklozenge\beta : \langle\sigma^{a},\tau^{a}\rangle \times \langle\sigma^{a},\upsilon^{s}\rangle, \gamma : \sigma^{a}}{\alpha(\gamma) : \tau^{a} \bullet \beta(\gamma) : \upsilon^{s}}$$

Intuitively: 'Applying an object of mixed type to the correct sort of input yields an output of at-issue and CI types, conjoined by a $\bullet$.'[5]

With these tools we can provide a full analysis of the CL case. We need only conjoin the two elements proposed cross-dimensionally, so that the individuating condition will live in the at-issue dimension and the domain condition will live in the dimension associated with conventional implicatures. As an example, the classifier *hon* will get the denotation

$$\lambda\mathcal{Q}\lambda P\lambda Q[\mathcal{Q}(P\cap At)(Q)]\blacklozenge\lambda\mathcal{Q}[\mathbf{Dom}(\mathcal{Q}) \subseteq stick.like]$$

which is typed as

$$\langle\langle\langle e,t\rangle,\langle\langle e,t\rangle,t\rangle\rangle,\langle\langle e,t\rangle,\langle\langle e,t\rangle,t\rangle\rangle\rangle^{a} \times \langle\langle\langle e,t\rangle^{a},\langle\langle e,t\rangle^{a},t\rangle\rangle^{a},t^{s}\rangle.$$

Derivations proceed in the obvious way. The first step is the crucial one. For *ni-hon* '2-Cl', we get, via R7,

[5]This rule is a minor simplification of the official system in McCready 2010; in brief, as it stands, it does not allow for cases where mixed types take multiple arguments in the conventionally implicated side. This capacity is not needed for the present application and so the simplification is harmless for the purposes of this paper.

$$\lambda P \lambda Q[2(P \cap At)(Q)] : \langle\langle e,t\rangle, \langle\langle e,t\rangle, t\rangle\rangle^{a}$$
$$\bullet$$
$$\mathbf{Dom}(2) \subseteq stick.like : t^{s}$$

The CI term on the right hand side of the • can be stripped off via (R5), and then the derivation goes on as usual in the at-issue dimension for quantified sentences. This looks satisfactory. The proposal, then, is able to handle both aspects of the meaning of classifiers.

4 Conclusion and Prospects

To briefly sum up, I have argued that classifiers introduce both at-issue and conventionally implicated content, that the at-issue content has the function of individuating the mass term-like denotations of bare nominals for quantification, and that the conventionally implicated content serves to temporarily restrict the domain of quantification to objects of the proper kind. I provided a formal implementation of these ideas in $\mathcal{L}_{\mathrm{CI}}{}^{+}$, an extension of Potts's (2005) logic for conventional implicatures which allows for terms carrying mixed content.

Mixed content of this sort is not hard to find, so $\mathcal{L}_{\mathrm{CI}}{}^{+}$ is useful in a variety of settings. McCready (2010) examines pejoratives, the Japanese benefactive, formal pronouns and honorifics, and evidentials, among other phenomena; details of the results of these investigations can be found in McCready 2010. With respect to the data in the present paper, the analysis I have given here should apply more or less equally to the English facts about terms of venery (19) and measure terms (21), though I will leave the details for another occasion.

Some readers may find the conclusions I have reached in this paper surprising: Classifiers may not be the first linguistic items that leap to mind when looking for examples of conventionally implicating terms. Still, it is not really obvious exactly how to think about the function of the domain restricting aspect of classifier meaning in the first place. This aspect of the classifier content is always entailed by the 'host' NP (if one is present): For instance, the content of *book* entails *printed matter*. In general, $[[Cl]] \subseteq [[NP]]$ for any classifier and host (excluding those like *zen* that have effects on number). If this content was truth-conditional, it would be absolutely useless, given the (even covert) presence of a host NP. In this sense, it seems reasonable that CL content should have its function outside truth conditions.

Are there other expressions that have a character similar to that I have claimed for classifiers, other than terms of venery and measure phrases? The answer seems positive. We are interested in expressions that seem, in some sense, presuppositional (and may even have been analyzed so), but that do not pass all the tests for presuppositionality. In particular, classifiers failed to pass the 'binding' test, as shown by (13). They also project from presupposition plugs. It appears that there are many other such expressions. For example, gender features on pronouns, which have been analyzed as presuppositional by e.g. Sauerland 2008, fail to be bindable in the general case. It seems at least plausible that they could, or even should, be analyzable as conventionally implicating as well. Other kinds of 'grammatical' phenomena might be thought of in the same way—in fact, conventional implicature in the present sense could be viewed as a general mechanism for analyzing the meaning of purely grammatical content, such as agreement and failures thereof. For another kind of example, we could consider selectional restrictions, which are also standardly viewed as presuppositional (McCawley 1968 is the classical reference).

(27) a. #The beer can is dead.

b. # The beer can is not alive.

But selectional restrictions are obviously not bindable either.

(28) # If beer cans could be alive, then the beer can {is|would be} dead.

Selectional restrictions then are also prime candidates for conventional implicature on the present view. Generalizing, it seems to me that many cases where presupposition is invoked as a tool to handle restrictions that are, in a broad sense, grammatical are ripe for reevaluation using conventional implicature.

This view may be questioned. Is it really right to place grammatical meanings of this kind on a par with content like that found in nominal appositives, or even with more expressive meanings such as that of *damn* or the German particle *ja*, which have been analyzed as conventionally implicating by some authors? It seems obvious that all these meanings are not of an identical type. But there is no reason to assume that all content which exhibits the general properties of conventional implicature is of a single kind. Rather, there are a number of subdivisions to be made: Expressive vs nonexpressive content (cf. Potts 2007), supplementary conventional implicatures vs main conventional implicatures (McCready, 2010), among other distinctions yet to be clarified. I do not wish to claim that conventional implicature is necessarily the correct diagnosis of these meanings, or even necessarily of the meaning of classifiers. But what is clear is that the content of classifiers, gender features, and selectional restrictions (to name a few) is of a different kind than e.g. the presuppositions associated with definite descriptions or factive verbs. At our present state of knowledge, conventional implicature appears to be the best way to think of this kind of content. Future progress in understanding the nature of conventionally implicated meanings, and of other kinds of meaning that might exist in natural language, will determine whether conventional implicature is indeed the right category, or not. The present paper is intended as a first step toward an answer to these questions.

References

Amaral, P., C. Roberts, and E. Smith. 2008. Review of ‘the Logic of Conventional Implicatures. *Linguistics and Philosophy* 30:707–749.

Bach, K. 1999. The myth of conventional implicature. *Linguistics and Philosophy* 22:327–366.

Beaver, D. 2001. *Presupposition and assertion in dynamic semantics*. Stanford, CA: CLSI.

Chierchia, G. 1998. Reference to kinds across language. *Natural Language Semantics* 339–405.

Chierchia, G. 2004. Scalar implicatures, polarity phenomena, and the syntax/pragmatics interface. In *Structure and beyond*, ed. A. Belletti, 39–103. Oxford: Oxford UP.

von Fintel, Kai. 1994. Restrictions on quantifier domains. Doctoral Dissertation, University of Massachusetts at Amherst, Amherst, MA.

Gauker, C. 1998. What is a context of utterance? *Philosophical Studies* 91:149–172.

Grice, H. 1975. Logic and conversation. In *Syntax and semantics iii: Speech acts*, ed. P. Cole and J. Morgan, 41–58. New York: Academic Press.

Levinson, S. 1983. *Pragmatics*. Cambridge: Cambridge UP.

Link, G. 1987. Algebraic semantics for event structures. In *Proceedings of the sixth amsterdam colloquium*, ed. J. Goenendijk, M. Stokhof, and F. Veltman. ITLI.

McCawley, J. 1968. Concerning the base component of a transformational grammar. *Foundations of Language* 243–269.

McCready, E. 2010. Varieties of conventional implicature. To appear in *Semantics and Pragmatics*.

Montague, R. 1974. The proper treatment of quantification in ordinary English. In *Formal philosophy: Selected papers of richard montague*, ed. R. H. Thomason, 188–221. New Haven, CT: Yale UP.

Peters, S., and D. Westerstahl. 2006. *Quantifiers in language and logic*. Oxford: Oxford UP.

Potts, C. 2003. Expressive content as conventional implicature. In *Proceedings of NELS 33*, ed. M. Kadowaki and S. Kawahara. Amherst, MA: GSLA.

Potts, C. 2007. The expressive dimension. *Theoretical Linguistics* 33:165–198.

Potts, Christopher. 2005. *The Logic of Conventional Implicatures*. Oxford: Oxford UP.

van der Sandt, R. 1992. Presupposition projection as anaphora resolution. *Journal of Semantics* 9:333–377.

Sauerland, U. 2008. On the semantic markedness of phi-features. In *Phi theory*, ed. D. Harbour, D. Adger, and Béjar. Oxford: Oxford UP.

Schlenker, P. 2007. Expressive presuppositions. *Theoretical Linguistics* 33:237–245.

Shan, C.-C. 2005. Linguistic side effects. Doctoral Dissertation, Harvard University, Cambridge, MA.

Stanley, J., and Z. Szabo. 2000. On quantifier domain restriction. *Mind and Langauge* 15:219–261.

Wang, L., B. Reese, and E. McCready. 2005. The projection problem of nominal appositives. *Snippets* 11:13–14.

Williamson, T. 2009. Reference, inference and the semantics of pejoratives. In *The philosophy of David Kaplan*, ed. J. Almog and P. Leonardi, 137–159. Oxford: Oxford UP.