Final version is in H. Bunt and R. Muskens (eds.), *Computing Meaning, vol. 1*, Kluwer Acad. Publ., Dortrecht, 1999, 219-253.

Dynamic Skolemization

Lenhart Schubert
University of Rochester

1 Introduction

The phenomenon of 'donkey anaphora' has been much discussed in the linguistic semantics literature (e.g., see Carlson & Pelletier, 1995; Chierchia, 1995), and has given rise to new theories of logical form based on various kinds of 'dynamic semantics'.[1] Although these theories appear to account for donkey anaphora in a natural way, they have not had much impact so far on knowledge representation in AI. This is surprising since – as I will argue – much of the 'common-sense' knowledge needed for ordinary reasoning and language understanding characteristically involves donkey anaphora, when this knowledge is expressed in words. The reason may be that the development of knowledge representations has been influenced more by the need for efficient inference than the need for capturing linguistically conveyed information.

This chapter explores the extent to which the phenomena addressed by dynamic semantics theories can be systematically addressed within a more standard representational framework. The main proposal is a mechanism called *dynamic Skolemization* for conversion of logical forms containing anaphoric terms to ones containing constants and functional terms with a standard, non-dynamic semantics. The resultant formulas seem to be equivalent (in a sense) to corresponding formulas in logics using dynamic semantics, in cases where the latter are able to represent the anaphoric relations involved (esp. donkey anaphora). A fortuitous advantage of dynamic Skolemization is that it can be recast as automatic creation of script-like or frame-like structures (as understood in AI), and enables referring into these structures in ways beyond the capabilities of dynamic semantics. This is not to say that dynamic Skolemization rivals dynamic semantics accounts of donkey anaphora on a theoretical level, but it does appear to have practical potential, and to be connected in interesting ways to the latter.

The following sentences illustrate the problems encountered when we attempt to represent sentences involving donkey anaphora in standard first-order logic.

1. If *a farmer* buys *a donkey*, *he* pays for *it* in cash.
2. If John bought *a donkey* at the fair, he paid for *it* in cash.

Note that (1) has a generic flavor, in that it quantifies over cases where a farmer buys a donkey. (2) admits both a generic and a specific reading. On the generic reading, it quantifies over past instances of John buying a donkey (perhaps with different instantiations of "*the fair*", or perhaps with a single instantiation of a very long fair). On the specific reading, (2) seems to refer to just a single visit by John to the fair, and a possible donkey-purchase during that visit.

Both sentences and both readings of (2) illustrate the problem at issue, but the specific reading of (2) does so most simply and starkly. If we interpret indefinites in terms of existential quantification

[1] I would prefer to use Chierchia's term 'dynamic binding', but do not wish to imply particular theoretical commitments here.

(as has been customary in philosophical logic and much of AI), then the logical form of the specific reading of (2) appears to be

$$[(\exists x)D(x) \wedge B(J,x)] \;\rightarrow\; P(it, M),$$

ignoring implicit events. But how are we to disambiguate the anaphoric pronoun? We cannot set it to x, since this would lead to a free occurrence of x outside the scope of its quantifier, where this free occurrence is semantically disconnected from the prior bound occurrence. We also cannot appeal to a 'wide-scope reading' of the indefinite, giving the existential quantifier wider scope than the conditional and thus bringing the pronoun within its scope. This would give a reading to the effect that something x is such that if it is a donkey and John bought it at the fair, then he paid for it in cash. But such a statement is true virtually independently of John's donkey purchases and manner of payment, as long as there is *some*thing that is not a donkey, or *some*thing that John didn't buy, thus rendering the antecedent (protasis) of the conditional false and the conditional as a whole true.

Finally, we cannot appeal to some principle of quantifier scope widening analogous to conversion to prenex form, since this would convert the negatively embedded existential quantifier to a wide-scope universal. If it happens that John bought several donkeys at the fair, truth of (2) would then require John to have paid for *all* of them in cash. But this is wrong (though it is sometimes an implicature), as can be seen by augmenting (2) as follows:

3. John took enough money to the fair to buy a donkey. So, if he bought a donkey at the fair, he paid for it in cash.

Clearly we do not want to infer that if John was so impressed with the livestock offerings that he went on a donkey-buying spree, he paid for *all* his donkeys in cash. Yet (2) remains true in this context. The point is even clearer for the following variant of (2):

4. If John bought a donkey at yesterday's fair, he rode it home.

This seems perfectly true even in a case where John bought two donkeys, but rode only one home while leading the other. Another argument against the scope-widening solution is that it plainly leads to incoherence for variants of (1) involving explicit generic quantification, such as

5. If a farmer buys a donkey, he {*usually/ sometimes/ rarely*} pays for it in cash.

Why, then, should the KR community care about this problem? I claim that much of the knowledge needed for common-sense reasoning is of this form. For instance, the following are some general facts that everyone presumably knows:

6. If a hungry predator encounters a smaller creature, it may well try to attack and eat it.
7. If a conspicuous action occurs within plain sight of a person (in a normal state of alertness), the person will notice it.
8. If an adult attacks a child, that is an extremely wicked action.
9. If someone does something extremely wicked, and others find out about it, they will want to punish him severely.
10. If one (physically) enters something, one will then be inside it (and will have been outside it immediately before).
11. If you eat at a restaurant, you enter (it), get a table to sit at, select a meal from a menu, place your order with a waiter, etc.

The first five items are taken loosely from the encodings in (Schubert & Hwang, 1989, Hwang & Schubert 1993a) of the knowledge needed to understand some small fragments of a fairy tale (Little Red Riding Hood). We found that dozens of axioms of this sort, encoded in EL (Episodic Logic, an NL-like representation) were typically required to enable the EPILOG system to make sense of a sentence in a story (also represented in EL), i.e., to infer the 'obvious' causal, temporal, part-of, and other coherence relations. Of course, the claim that story understanding is a knowledge-intensive activity is a truism in AI nowadays (thanks to the work of Roger Schank and his co-workers, and many others in the 70's and 80's), but the present point concerns the particular form of the required facts when stated in ordinary language. Evidently, all five sentences involve at least one instance of donkey anaphora. (In the third sentence, the deictic *that* is in effect a reference to an action whose existence is implicit in the antecedent if-clause.)

The last sentence is intended to echo the kind of knowledge often encoded in script-like representations. Setting aside the role of *you*, the consequent clauses clearly involve tacit or explicit references to the restaurant in question, and these references are again of the 'donkey' variety.

The need to convey huge numbers of mundane generalities to machines remains a frustrating bottleneck in the effort to endow machines with ordinary understanding and common sense. There are several possible ways of attacking the problem, the most direct being the systematic hand-coding of the requisite facts (Lenat, 1995). However, this approach is very labor-intensive, and the task would be greatly lightened if the bulk of the necessary knowledge could be conveyed directly through language.

The problem of donkey anaphora presents a formidable obstacle in this endeavor. The next section includes a brief review of some proposed solutions based on dynamic semantics. These solutions deal very plausibly with examples like (1-4), but have some drawbacks from an AI perspective. One is that their implementation in a reasoning system would require a significant reworking of the usual inference mechanisms so as to allow for the context dependence of the meanings of dynamically interpreted formulas. In addition, I will show that the dynamic semantics approach has difficulty representing reference to dependent entities (e.g., ones that play a role in a certain type of event), where the dependency was established by prior discourse or hinges on general world knowledge. Such dependent entities would be frequently encountered in attempting to 'tell machines about the world'.

In Sect. 3, I will then develop the proposal I call *dynamic Skolemization*, showing how it deals both with standard donkey anaphora and the more challenging cases described in Sect. 2. Dynamic Skolemization can be augmented so as to introduce new concepts (predicates) that 'summarize' the content of the formulas comprising the scope of the Skolemized variables. When this is done, the representations becomes quite concise, and strongly reminiscent of scripts and frames in AI.

In Sect. 4 I discuss some further issues in the application of dynamic Skolemization, including limitations of the approach when it is used to represent reference relations across clauses whose connective lies in a negative polarity environment; as well, I will briefly discuss 'generic passages' in the sense of (Carlson & Spejewski, 1997), the issue of 'strong readings', and the 'proportion problem'. In the final section I reiterate my conclusions about the status and scope of dynamic Skolemization, and its potential utility in knowledge encoding.

2 Dynamic Semantics

The difficulties posed by donkey anaphora have spurred new developments in the logical analysis and representation of ordinary language. Sentence (1), repeated here as (12), will serve to illustrate two well-known approaches very briefly:

12. (a) If a farmer buys a donkey, he pays for it in cash

(b)
$$\boxed{\begin{array}{|l|}\hline \quad \mathrm{x}\ \ \mathrm{y} \\ \hline \mathrm{F(x)} \\ \mathrm{D(y)} \\ \mathrm{B(x,y)} \\ \hline \end{array} \Rightarrow \begin{array}{|l|}\hline \quad \mathrm{z}\ \ \mathrm{w} \\ \hline \mathrm{P(z,w)} \\ \mathrm{z = x} \\ \mathrm{w = y} \\ \hline \end{array}}$$

(c) $(x, y : F(x), D(y), B(x, y))$
$\Rightarrow (z, w : P(z, w), z = x, w = y)$

13. $[(\exists x, y)F(x) \wedge D(y) \wedge B(x, y)] \ \rightarrow \ P(x, y)$

Perhaps the best-known framework for dealing with anaphora is Discourse Representation Theory (DRT) (Kamp, 1981; Heim, 1982). The representation of (2a), repeated here as (12a), would be a Discourse Representation Structure (DRS) as shown in (12b). This DRS contains a conditional whose antecedent (protasis) and consequent (apodosis) are again DRS's. The indefinites "*a farmer*" and "*a donkey*" are represented as variables called *discourse referents*, comprising the syntactically separate *universe* of the antecedent DRS. Corresponding *conditions* $F(x)$ and $D(y)$ appropriately constrain them in the antecedent DRS. The treatment of the definites in the consequent DRS is analogous, and the anaphoric references are resolved by the equations in the consequent DRS. An alternative, non-pictorial syntax could be as illustrated in (12c).

What is crucial here is the semantics, in particular the dynamic way in which discourse referents are added to the domain of a partial assignment function in the course of semantically evaluating a DRS. Roughly speaking, a conditional DRS like (12b) is taken to be true relative to a model and a partial assignment if every way of making the antecedent true by adding its discourse referents (here, x and y) to the domain of the assignment can be extended to a way of making the consequent true by adding its discourse referents (here, z and w) to the domain of the assignment.

Clearly, then, the way DRT treats the anaphoric pronouns in the consequent of a conditional is both syntactically and semantically dependent on the presence of co-referring discourse referents in the antecedent. Syntactically, the translations of the pronouns need to be equated to 'accessible' referents (such as those in the antecedent, or in the universe of an embedding DRS), and semantically the equations succeed in making the desired connection because of the way partial assignment functions get extended in the antecedent, and the way this extension is carried forward to the consequent.

A disadvantage of DRT as a meaning representation is that it requires a mapping from surface linguistic form to logical form that is hard to reconcile with a compositional semantics.[2] Dynamic Predicate Logic (DPL) (Groenendijk & Stokhof, 1991) provides a compositional alternative. DPL would assign (12) the logical form shown in (13). Note that in DPL we can actually substitute the quantified variables x and y for the anaphoric pronouns, much as we were tempted to do – but couldn't make formal sense of – in the case of *it* in (2). The way DPL makes sense of (13) is by treating formulas much like computer programs, where variables are assigned values in earlier subformulas, and these values persist to later subformulas. In particular, an existential formula $(\exists x)\Phi$ is thought of as denoting a *nondeterministic assignment* to x, followed by a *test* Φ. Truth corresponds to 'executability' (for some initial values of the variables and some nondeterministic choices during execution). Technically, this is formalized in terms of state-change semantics, where a state is a (total) variable assignment; i.e., the semantic value of $(\exists x)\Phi$ relative to a model consists

[2]But see (Muskens, 1994).

of the pairs of states such that the formula can (nondeterministically) transform one into the other. For a conditional $\Phi \to \Psi$, successful execution (truth) requires that for every way of executing Φ, the resultant state allows execution of Ψ.

Again, the dynamic change in the context of evaluation induced by the prior material is apparent, as is the need for 'accessible' prior referents. Another significant point to note is that both approaches lead to a universal ('strong') interpretation of indefinites in the antecedent of a conditional. In DRT, this is because of the iteration over *all* ways of making the antecedent true (in the present example, all ways of picking a farmer and a donkey such that the farmer buys the donkey). Similarly, in DPL this is because of the iteration over successful executions of the antecedent. So this semantics is somewhat appropriate for (12) but hardly for the specific reading of (2). However, one could define a 'specific conditional' within DRT or DPL that does justice to intuitions about (2) (see Chierchia, 1995). EL, mentioned above, would assign much the same logical form as DPL to (12) (neglecting events), but would employ a distinctive 'generic conditional' to obtain a quasi-universal reading. For the material conditional, EL's particular brand of DRT-like dynamic semantics would give the desired specific reading for (2).

The reliance of these theories on the presence of accessible referents for interpreting anaphora can be a serious handicap. This is apparent for passages where 'dependent entities' are introduced in one sentence, and then referenced in another, as in

14. There were several instances at the fair where a farmer bought a donkey. In each instance, *the farmer* paid for *the donkey* in cash.

The dynamic semantics approaches provide no easy way to logically represent the referential definites in the second sentence so that they refer to the appropriate farmer and donkey in each of the instances under discussion. This is apparent if we write down a DPL-like or EL-like logical form for these sentences (enriching the basic syntax slightly so as to allow restricted quantifiers and some set-theoretic constructs):

15. (a) $(\exists S : \textit{set-of-instances}(S))$
$(\forall e : e \in S)$
$(\exists x : F(x))(\exists y : D(y))B(x,y,e)$.
(b) $(\forall e : e \in S)(\exists e' : e' \sqsubseteq e)$
$(\textit{The } z : F(z))(\textit{The } w : D(w))P(z,w,e')$.

This interprets the 'instances' in (14) as donkey-purchases by farmers, and accordingly introduces a Davidson-like event variable into B ("*buys*") and likewise into P ("*pays in cash for*"). While I do not regard this approach to event representation adequate (see Hwang & Schubert, 1993b) it will do for present purposes.

Note, first of all, that the referential connection between the bound occurrence of S at the beginning and the free occurrence later on will be made successfully by the dynamic semantics approaches mentioned. However, the same is not true of the definites $(\textit{The } z : F(z))$ and $(\textit{The } w : D(w))$. Substituting the farmer-variable x from the first sentence for z, and the donkey-variable y for w, simply does not lead to a coherent result in any of the theories, since $(\exists x : F(x))$ and $(\exists y : D(y))$ lie within the scope of a universal quantifier. Intuitively, the first sentence posits various values of x and y depending on the choice of $e \in S$, and in the second sentence z and w should refer to these values as a function of e, but no such function is available.

In other words, when the desired referents are dependent entities embedded within a prior quantified context, they cannot do the required context-setting for interpreting the anaphoric descriptions. In fact, the only way we appear to have of solving the reference problem in (14) is to reformulate the second sentence so that it *repeats* the content of the first sentence antecedently, making the required referents directly accessible:

16. (a) In each instance, <u>if *a farmer* bought *a donkey*</u>, *the farmer* paid for *the donkey* in cash.

    (b) $(\forall e : e \in S)\ \underline{[(\exists x : F(x))(\exists y : D(y))B(x, y, e)]} \rightarrow$
    $(\exists e' : e' \sqsubseteq e)P(x, y, e')$

However, besides being cumbersome and *ad hoc*, this importation of material ought to be redundant. After all, the discourse prior to (16a) already characterized the instances in question as instances in which a farmer buys a donkey, so why should we have re-specify this property?

This is not the worst case. There are familiar kinds of reference where the required antecedent material is not merely submerged in the prior discourse, but hidden in the tacit knowledge of the speaker and hearer. The following is an example:

17. (a) John *dined at Mario's*. He left a big tip for *the waiter*.

    (b) $(\exists e)D(J, M, e)$. $(\exists x : T(x))(\exists e' : e' \sqsubseteq e)$
    $(\mathit{The}\ y : W(y))L(J, x, y, e')$

Any of the dynamic semantics approaches would now require retrieving and instantiating a 'dining out' scenario, at least in part, and incorporating this antecedently into the representation of the second sentence. Only then would we be able to resolve the definite description $(\mathit{The}\ y : W(y))$ to the appropriate referent, i.e., the waiter who served John. This approach would be laborious and hard to systematize, keeping in mind that stored scenarios could be quite complex, and that multiple scenarios could be relevant to a single sentence (e.g., "*At the business dinner at Mario's, the waiter spilled coffee on the CEO just when he was about to put up some pie-charts*"). These are just the sorts of problems that hampered full exploitation of 'scripts' in story-understanding research in the 70's and 80's.

My goal in the following is to provide a straightforward and concise method of interpreting sentences involving anaphora as in (14) and (17), as well as the more innocuous cases of donkey anaphora. The idea is to make available *Skolem functions* (constants, in the simplest cases) for establishing referential connections. For instance, in the case of (17) we want to make available a 'waiter function' that picks out the waiter of any dining-out event, as a function of that event. In the following I develop such an approach, beginning with simple cases of (non-donkey) anaphora.

3 Dynamic Skolemization

3.1 Skolem constants

Let's begin with a very simple example of cross-sentential anaphora, the sort that invites the use of Skolem constants – and has been dealt with in that way in countless natural-language systems:

18. (a) John bought a donkey. He paid for it in cash.

    (b) $(\exists x : D(x))(\exists e)B(J, x, e)$. $(\exists e')P(J, \mathit{it}, e')$.

    (c) $D(A) \wedge B(J, A, E)$. $P(J, A, E')$.

(18b) shows the preliminary logical form of (18a) (again with Davidsonian event arguments). In (18c), the variable x for the newly introduced donkey has been Skolemized to A and the variable e for the newly introduced buying event has been Skolemized to E. The anaphoric pronoun has also been replaced by A, unproblematically making the cross-sentential connection. Note that (as in some previous examples) there is also an implicit relation between the paying event and the buying event, i.e, the paying is understood to be part of the buying. Various theories have been proposed for deriving this sort of connection systematically (e.g., Webber, 1988; Hwang & Schubert, 1992; Lascarides *et al.*, 1992). By whatever means the connection is derived, the Skolemization of event variables facilitates its representation, e.g., as $E \sqsubseteq E'$.

We should note in passing that Skolemization is not a logically sound transformation. Rather, it amounts to a logical strengthening, since predicating constraints on a constant will filter out some possible models. But the strengthening is trivial in the sense that it is guaranteed to preserve satisfiability, and that no new sentences are entailed, other than some that involve the Skolem constant. (See the 'conservative extension' proposition below. This is why Skolemization is legitimate in refutation proofs, such as resolution proofs.)

If one now considers a conditional sentence like (2) (repeated here as (19a)), one feels intuitively that the anaphoric connection is made just as in (18a), the non-conditional case.

19. (a) If John bought a donkey at the fair, he paid for it in cash.
    (b) $[(\exists x)D(x) \wedge B(J,x)] \rightarrow P(J,it)$
    (c) $[D(A) \wedge B(J,A)] \rightarrow P(J,A)$.

One conceptualizes the situation of John buying a donkey, and in doing so, establishes some conceptual token for the donkey. This token is then available for reference, whether the context-setting sentence supplies a fact or a proviso.

Of course, the discourse referents of DRT and the existential variables of DPL provide just such tokens, but we are seeking a solution admitting a standard semantics, and one extendable to functionally dependent entities. At first sight, Skolem constants and functions seem to fail us here. If we assign to (19a) the preliminary logical form (19b) (neglecting events for simplicity) and then attempt to Skolemize and resolve the referent as we did in (18c), the result (19c) is patently defective. This is apparent from the fact that (19c) is the Skolemized form of $(\exists x)[D(x) \wedge B(J,x) \rightarrow P(J,x)]$, i.e., "There is a thing x such that if it is a donkey and John bought it, then he paid for it in cash"; but this is trivially true if we can find any x that is not a donkey or was not bought by John, rendering the antecedent false.

However, it is interesting to note that (19c) fails as a representation of (19a) by being too weak, rather than too strong. Given (19a) and given that John bought a donkey, it should follow that there is a donkey that John bought, and paid for in cash. But if we replace (19a) as a premise in this argument by (19c), the conclusion no longer follows. This is because the existence of a donkey that John bought fails to entail $D(A) \wedge B(J,A)$, the antecedent of (19c).

This suggests a simple remedy: we supplement the Skolemization in (19c) with the following stipulation:

20. $(\exists x : D(x))B(J,x) \;\rightarrow\; D(A) \wedge B(J,A)$.

I will use the term *dynamic Skolemization* for the two steps of (i) Skolemizing an existential variable, and (ii) stipulating a supplementary condition like (20) above, relating the unskolemized existential formula to its Skolemized form. (This is stated more formally, and generalized, below.) I will also refer to the supplementary condition itself as a 'Skolem conditional'.

Given (20), falsity of $D(A) \wedge B(J,A)$ implies that John didn't buy a donkey. So $D(A) \wedge B(J,A)$ can no longer be trivially false, and (19c) trivially true, as a result of A not being a donkey or John not having bought it. Some additional points are worth noting. First, the Skolem conditional can be turned into a biconditional, by existential generalization. This makes it natural to think of (20) as a partial *definition* of the newly introduced constant A. As such it provides a *conservative extension* of any theory framed in a vocabulary exclusive of A, as shown later.

Second, when we resolve the pronoun in (19a,b) to the constant A, introduced through dynamic Skolemization, the resultant reading is *non-universal* ('weak'). To obtain a past-habitual reading of (19a), a generic or habitual quantifier (arguably, over events) would have to be introduced into the logical form (19b). Examples of this sort will be seen later.

A third, related point concerns the logical consequences of readings based on dynamic Skolemization. Suppose we are told after (19a) that John *did* buy a donkey (within the relevant time frame), namely Archie (A'):

$D(A') \wedge B(J, A')$.

It does *not* follow that John paid for Archie in cash. This is because it does not follow that $A' = A$. Yet it seems intuitively natural to draw that conclusion. I will assume that this intuition arises from a scalar implicature of existential sentences that there is at most one verifying instance of that sentence.[3] For (19b) the implicature could be written as

$[(\exists x : D(x))B(J, x)] \;\rightarrow\; (\exists! x : D(x))B(J, x)$,

or, using the Skolem constant in (19c), as

$(\forall x : (D(x) \wedge B(J, x)))(x = A)$.

This clearly leads to the desired conclusion concerning Archie. But note that the uniqueness assumption is cancellable, for instance in a context like "*I don't know how many donkeys John bought, if any. But if he did buy a donkey, he rode it home*". Here it may well be that John bought a donkey that he rode home, and another that he didn't ride home. I leave further discussion of the uniqueness assumption to Sect. 5.3.

Abstracting from the above example, we can define dynamic Skolemization as follows, for an existential sentence that occurs outside the scope of all other quantifiers:

21. Given an occurrence of a (closed) sentence of form $(\exists x : \Phi)\Psi$ in the provisional logical form of an English sentence being interpreted,
    (a) assert the Skolem conditional
    $(\exists x : \Phi)\Psi \;\rightarrow\; \Phi_{C/x} \wedge \Psi_{C/x}$,
    where C is a new constant;
    (b) replace the original occurrence of $(\exists x : \Phi)\Psi$ by the consequent of the above implication.

The notation $\Phi_{C/x}$ stands for the result of substituting C for all free occurrences of x in Φ.[4] The 'provisional logical form' of a sentence refers to the logical form we would obtain *before* finalizing the representation of pronouns and referential definites, but *after* analyzing phrase structure, applying

[3] E.g., (Kadmon, 1990). Such a uniqueness implicature is present in many contexts, and appears to be the result of the hearer's assumption that if the speaker thought there might have been multiple instances of the indefinite, and such instances are relevant to the purposes of the exchange, then he would have made this possibility explicit. This is similar to Gricean scalar implicatures in sentences such as "*Some / Most* students passed the exam", or "If *some / most* students passed the exam, John will be pleased", where the speaker seems to assume that not all students passed the exam.

[4] I use the particular variable x and constant C here for readability, but these are of course arbitrary.

rules of logical-form composition corresponding to phrase structure rules (grounding these in a lexicon that supplies logical translations of lexical items), and scoping quantifiers.

What do the representations of pronouns and referential definites look like before they are finalized? I would like to leave this somewhat open. One possible assumption is that they are represented as unscoped restricted quantifiers, as in the quasi-logical forms of Schubert & Pelletier (1982),[5] *coindexed* with their referents (simple or functional terms) within the logical forms of the prior discourse or knowledge base. They are not assumed to be interpretable as parts of their embedding expressions, but only as transient forms still reflecting surface syntax and 'waiting' to be replaced by interpretable terms. In assuming coindexing of anaphoric terms, I am setting aside the problem of determining their referents, which is separate from that of appropriately representing them in the final logical form. Another possible provisional representation of anaphoric terms would be as DPL-variables, suitably chosen to refer to the desired referents. This would have the advantage of being fully interpretable – but only under a dynamic semantics, and only as long as we are dealing with the sorts of local reference that DPL can handle (as opposed to reference into quantificational contexts).

One implicit assumption in the above rule is that the embedded existential sentence $(\exists x : \Phi)\Psi$ is itself a final logical form, free of provisional (or dynamic) representations of anaphoric terms; otherwise the meaning of the Skolem conditional would not be well-defined. Thus dynamic Skolemization of an existential sentence must be preceded by full anaphora processing within its scope.

Let us note some properties of dynamic Skolemization, beginning with the following simple

Proposition. *The Skolem conditional (21a) provides a* conservative extension *of any theory framed in a vocabulary exclusive of the new Skolem constant C.*

What this means is that no new sentences follow from the theory together with the definition, other than ones involving the defined term C. The proposition is easily proved by considering any formula Φ not involving C that is entailed by a given theory $\mathcal{T}$ (not involving C) together with the Skolem conditional. For any model $\mathcal{M}$ of theory $\mathcal{T}$, let $\mathcal{M}' = \mathcal{M}$ if the antecedent of the Skolem conditional is false in $\mathcal{M}$; otherwise let $\mathcal{M}'$ differ from $\mathcal{M}$ only in assigning a value to C that makes the consequent of the Skolem conditional true. Then $\mathcal{M}'$ is a model of $\mathcal{T}$ together with the Skolem conditional, and so by assumption is a model of Φ. But $\mathcal{M}$ and $\mathcal{M}'$ agree on Φ, hence $\mathcal{M}$ is a model of Φ. Thus $\mathcal{T}$ entails Φ. I leave the generalization of the proposition (and proof) for the case of Skolem *functions* (below) to the reader.

Another property that (21) ought to satisfy is that it should give the expected result in simple declarative contexts – keeping in mind that it was motivated by existential sentences within conditional antecedents. It is easy to see that it does. For instance, consider (once again) (18a). First, note that the Skolemization specified by (21b) yields the sentence $D(A) \wedge B(J, A)$ (*modulo* the choice of arbitrary new constant) as previously specified in (18c). But since this sentence occurs at the top level of the text being interpreted, it is *asserted*. As such it entails the Skolem conditional (previously given as (20)), since it coincides with the consequent. Hence the Skolem conditional is redundant. In other words, in the case of simple declaratives (21) just describes ordinary Skolemization of top-level existentials.

Finally, before proceeding to the more general case, we should note that the above rules of dynamic Skolemization will need some modification in certain negative polarity environments. Examples are "*John didn't buy a donkey and pay for it in cash*"; and "*If John bought a donkey and paid for it in cash, then he rode it home*". In both cases, "*a donkey*" occurs within a negative

[5] or as generalized quantifiers

polarity environment, and is referred to anaphorically within the same environment. For the time being we ignore such cases, but take them up again in Sect. 5.2.

3.2 Skolem functions, concept definitions

The generalization of dynamic Skolemization (21) to existential formulas containing free variables (i.e., embedded within the scope of other quantifiers) is straightforward:

22. Given an occurrence of a formula $(\exists y : \Phi)\Psi$, containing free variables $x_1, ..., x_m$, in the provisional logical form of an English sentence being interpreted,
    (a) assert the Skolem conditional
    $(\forall x_1)...(\forall x_m)[(\exists y : \Phi)\Psi \rightarrow \Phi_{f(x_1,...,x_m)/y} \wedge \Psi_{f(x_1,...,x_m)/y}]$,
    where f is a new m-place function symbol;
    (b) replace the original occurrence of $(\exists y : \Phi)\Psi$ by the consequent of the above implication.

As before, we assume that the embedded formula $(\exists y : \Phi)\Psi$ is free of unresolved referring expressions. A simple illustration is provided by the following type of quantified donkey sentence, often used as central example in the literature; (we neglect events):

23. (a) Every farmer who bought a donkey paid for it in cash.
    (b) $(\forall x : F(x) \wedge (\exists y : D(y))B(x, y))\, P(x, it)$
    (c) Assert:
    $(\forall x)[(\exists y : D(y))B(x, y) \rightarrow D(f(x)) \wedge B(x, f(x))]$
    Subst.: $(\forall x : F(x) \wedge D(f(x)) \wedge B(x, f(x)))\, P(x, f(x))$

Note the resolution of the anaphoric pronoun in (23c). The reading this encodes is *non-universal* for donkeys, i.e., any one farmer purchasing multiple donkeys need only have paid for one in cash. I believe this is correct, based on examples analogous to (3) and (4), although we often generalize to a universal interpretation (see Sect. 5.3). For example, an analog of (4) is

24. Every farmer who bought a donkey rode it home,

which seems true even if the farmers involved rode only one donkey each.

It is clear that repeated application of (22) to a formula with multiple existential quantifiers, $(\exists y_1 : \Phi_1)...(\exists y_n : \Phi_n)\Psi$ will be equivalent to simultaneous Skolemization of the y_i to $f_1, ..., f_n$. I will henceforth freely apply it in this way. The following is an interesting variant of (1), illustrating the case $n = 2$; (25b) is the Skolem conditional, and (25c) the logical form of (25a) after substitution:

25. (a) Whenever a farmer bought a donkey, he paid for it in cash.
    (b) $(\forall e)[(\exists x : F(x))(\exists y : D(y))B(x, y, e) \rightarrow F(f(e)) \wedge D(g(e)) \wedge B(f(e), g(e), e)]$

(c) $(\forall e)[F(f(e)) \wedge D(g(e)) \wedge B(f(e), g(e), e) \rightarrow$
$(\exists e')\, P(f(e), g(e), e')]$

The *whenever*-clause calls for universal event quantification, and this leads to the introduction of a farmer-function f and donkey-function g on events. In this case, we obtain a reading that effectively quantifies universally over farmers and donkeys, assuming that in distinct action predications such as $B(x, y, e)$, $B(x', y', e')$, e and e' cannot be the same buying-event unless x and y are the same as x' and y' respectively. This seems intuitively correct.

There remains a certain inelegance in these examples of dynamic Skolemization from a practical computational perspective, in that we are repeating variants of the original existential formula, $(\exists y_1 : \Phi_1)...(\exists y_n : \Phi_n)\Psi$, three times: twice in the Skolem conditional (in the antecedent and consequent) and once in the logical form of the given text. I will therefore introduce a space-saving notation, called *concept definition* that avoids redundancy.

Before defining the notation formally, the idea is just this: when we encounter a wff $(\exists y_1)...(\exists y_n)\Phi$ that is to be Skolemized, we first define a new predicate that is simultaneously equivalent to (i) this wff, for all values of the free variables (those not bound by the $(\exists y_i)$), and (ii) the Skolemized version of the wff. We then substitute the newly defined predicate into the source text. Formally, the concept definition schema is this:

26. For π a new m-place predicate constant, $f_1, ..., f_n$ new m-place function constants, and Φ a formula possibly containing free occurrences of variables $x_1, ..., x_m$ and $y_1, ..., y_n$ (and no others),

    $(\mathbf{Def}\ \pi\ \ (x_1, ..., x_m)\ (f_1, ..., f_n)\ \Phi_{\underline{f}/\underline{y}})$

    abbreviates (with underlining indicating vectors)

    $(\forall x_1)...(\forall x_m)[\pi(x_1, ..., x_m)$
    $\leftrightarrow (\exists y_1)...(\exists y_n)\Phi \leftrightarrow \Phi_{\underline{f}(x_1,...,x_m)/\underline{y}}]$

$x_1, ..., x_m$ are the *variables* of the Def-schema, $f_1, ..., f_n$ are the *roles* (role functions, Skolem functions), and $\Phi_{\underline{f}/\underline{y}}$ is the *body* of the definition. Note that in this body we are using function symbols (without arguments) in place of variables, so that the result will not be a well-formed formula. This avoids having to repeat the variables $x_1, ..., x_m$ on which each role depends.

Applying the Def-schema to (25) as described above, we get a more concise and transparent representation of the dynamic Skolemization (now formulated as a concept definition, (27a)) and the resultant logical form, (27b); *FBD* is the newly defined concept (mnemonic for "farmer buys donkey"):

27. (a) $(\mathbf{Def}\ \mathit{FBD}\ (e)\ (f, g)\ F(f) \wedge D(g) \wedge B(f, g, e))$
    (b) $(\forall e)\ \mathit{FBD}\ (e) \rightarrow (\exists e')P(f(e), g(e), e')$

We are now well-equipped to return to the farmers at the fair, and Mario's restaurant. Looking back at (15a), note that we will first Skolemize S, the set of instances mentioned, to some new constant (but I will retain S here). The next existential sentence encountered is $(\exists x : F(x))(\exists y :$

$D(y))B(x,y,e)$. Hence we get the definition

(**Def** $FBD(e)$ (f,g) $F(f) \wedge D(g) \wedge B(f,g,e)$).

Substituting the defined concept in the provisional logical form (15a), we obtain (28a); we can now easily resolve the references in (15b) to $f(e)$ and $g(e)$, as shown in (28b):

28. (a) $(\forall e : e \in S) FBD(e)$.
    (b) $(\forall e : e \in S)(\exists e' : e' \sqsubseteq e) P(f(e), g(e), e')$.

To deal with the example of Mario's restaurant, (17), we naturally need to presuppose some prior knowledge of 'what happens when' a person dines out. While people are unlikely to first acquire this information in the form of a tidy verbal package, it would help with knowledge bootstrapping if this were possible for computers. So suppose we tell our NLP system something like the following:

> Generally, when a person dines at a restaurant, s/he enters the restaurant, gets seated at a table, selects a meal from a menu, tells the order to a waiter, etc.

I will assume that *generally* quantifies over dining-at-restaurant episodes here, binding an episodic variable:

$$
\begin{array}{l}
(\mathit{Gen}\ e : (\exists x : P(x))(\exists y : R(y)) D(x,y,e)) \\
\quad (\exists e_1 : e_1 \sqsubset e \wedge \mathit{starts}(e_1, e))\ \mathit{Enter}(\hat{x}, \hat{y}, e_1) \wedge \\
\quad (\exists e_2 : e_2 \sqsubset e \wedge \hat{e} \prec e_2)\ \mathit{Get\text{-}seated}(\hat{x}, e_2) \wedge \\
\quad (\exists e_3 : e_3 \sqsubset e \wedge \hat{e}_2 \prec e_3)\ \mathit{Select}(\hat{x}, e_3) \wedge \\
\quad (\exists w : W(w))(\exists e_4 : e_4 \sqsubset e \wedge \hat{e}_3 \prec e_4) \\
\qquad\quad \mathit{Tell\text{-}order}(\hat{x}, w, e_4) \wedge \\
\qquad\quad \ldots \qquad)
\end{array}
$$

Note that the first line of this expression comprises the *restrictor* of the *Gen* quantifier. The symbols with carets are an *ad hoc* abbreviation for anaphoric terms, used for readability. For instance, $\mathit{Enter}(\hat{x}, \hat{y}, e_1)$ abbreviates

$(\mathit{The}\ z : R(z)) \mathit{Enter}(\mathit{s/he}, z, e_1)$,

i.e., s/he enters the restaurant; the particular symbols used, such as $\hat{x}$ and $\hat{y}$, are intended to suggest to the reader the expected referents of these terms.

The first existential sentence encountered is the restrictor of the *Gen* quantifier. Thus we apply dynamic Skolemization and concept definition to it:

(**Def** PDR (e) (f,g) $P(f) \wedge R(g) \wedge D(f,g,e)$).

Having substituted the defined predicate in the restrictor, and knowing that x is now $f(e)$ and y is $g(e)$, we can also resolve $\hat{x}$ and $\hat{y}$ to $f(e)$ and $g(e)$ respectively in the matrix of the *Gen* construct. Next we apply dynamic Skolemization and concept definition to the matrix. Here I allow myself a slight liberty, Skolemizing all of the conjoined existential sentences at once (processing them in succession would yield multiple definitions rather than a single one, but would otherwise be equivalent):

$$
\begin{array}{l}
(\mathbf{Def}\ \mathit{PDRsteps}\ (e)\ (f_1, f_2, f_3, h, f_4, \ldots) \\
\quad f_1 \sqsubset e \wedge \mathit{starts}(f_1, e) \wedge \mathit{Enter}(f, g, f_1) \wedge \\
\qquad \ldots \\
\quad W(h) \wedge f_4 \sqsubset e \wedge f_3 \prec f_4 \wedge \mathit{Tell\text{-}order}(f, h, f_4) \wedge \\
\qquad \ldots \qquad)
\end{array}
$$

Substituting in the matrix of *Gen*, we then obtain

$(\mathit{Gen}\ e : \mathit{PDR}(e)) \mathit{PDRsteps}(e)$.

Having accommodated the background knowledge in this way, we are now ready to interpret the definite description *the waiter* in (17), repeated here as (29a). The result is (29b), and after further Skolemization, (29c):

29. (a) John dined at Mario's. He left a big tip for *the waiter.*
    (b) $D(J, M, E)$.
    $(\exists x : T(x))(\exists e : e \sqsubset E) L(J, x, h(E), e)$.
    (c) $D(J, M, E)$. $T(C) \wedge E' \sqsubset E \wedge L(J, C, h(E), E')$.

There is still a hidden ambiguity, however. As far as the stored background knowledge is concerned, the agent of E is $f(E)$ and the restaurant is $g(E)$. To link these to *John* and *Mario's* respectively, we again need to invoke a 'uniqueness of roles' assumption concerning dining-out events:

$$(\forall e)(\forall x, y, x', y') D(x, y, e) \wedge D(x', y', e) \rightarrow$$
$$x = x' \wedge y = y'.$$

Again, this is an intuitively reasonable assumption.

4 Scripts, Frames and Knowledge Bootstrapping

As previously hinted, there is a close resemblance between the above encodings of the farmers-at-the-fair and dining-at-Mario's examples on the one hand, and scripts and frames as understood in AI (e.g., Schank & Abelson, 1977; Minsky, 1975) on the other. For instance, Minsky suggests that a person entering an office retrieves and instantiates an office frame featuring *slots* for the expected parts and furnishings (floor, ceiling, walls, door, desk, chair, etc.). These slots quickly become bound to perceived parts and furnishings, and the knowledge associated with them helps to guide further perception and action. Now this is just the sort of structured representation we would obtain by Skolemizing a description of the form, "*An office generally has a floor, a ceiling, four walls, a door, a desk, etc., with such-and-such properties and related in such-and-such ways...*". Assuming a provisional logical form like

$$(Gen\ x : Office(x))\ [(\exists y_1 : Floor(y_1)\ ...\ ],$$

we would define of a concept such as $Office\text{-}interior(x)$ corresponding to the body of the *Gen*-formula, with roles (Skolem functions) that pick out the parts and furnishings of the office. It would then be straightforward to interpret the definites in general and particular sentences like

30. An office belonging to a professor often has piles of books and papers on *the desk.*

31. Mary went into her office and sat down at *the desk.*

Note that the *Gen*-formula for an office does not involve any existential quantifiers in the restrictor, and so does not give rise to a concept definition for the restrictor (in contrast with the examples of the farmers at the fair and Mario's restaurant). Thus the formula supplies information about the pre-existing concept of an office, rather than about newly introduced episodic notion such as that of dining at a restaurant. This is typically how frames have been employed in knowledge-based systems, i.e., they are used to supply information about familiar, usually non-episodic concepts such as that of an office, a wage earner, a particular product, etc., expressible with a lexical or compound noun.

In the case of general episodic knowledge, the tie-in with scripts is by now clear. The dining-out example was of course intended as a parallel to Schank and Abelson's (1977) restaurant script. While their script is not formally interpreted, the name of the script can be viewed as a defined concept similar to *PDR* above; further, it involves roles for the participating entities that are clearly analogous to some of the Skolemized roles in the definitions of *PDR* and *PDRsteps*, and it involves steps that are analogous to the Skolemized subepisodes in *PDRsteps*. An interesting difference is that it is not possible to access the roles in a script 'from the outside', as seems to be required for representing general or particular sentences such as

31. When a person dines at a *fancy* restaurant, s/he pays *the waiter* directly, not at a cash register.

32. When Mary dined at Mario's, *the waiter* turned out to be an an ex-classmate of hers.

For (31), a traditional script-based approach would have to create a variant script for fancy dining, as a context in which to place the new fact (assuming that the pre-existing restaurant script is noncommittal about how the bill is paid). This could be done in some *ad hoc* way by modifying a copy of the regular script, or directly inserting an alternative track into that script. In the approach based on dynamic Skolemization, we would automatically create a separate concept of dining at a fancy restaurant, and, recognizing that dining at a fancy restaurant entails dining at a restaurant, make use of the role functions in the *PDR* and *PDRsteps* definitions to pick out the waiter (and the implicitly referenced bill). For (32), the traditional approach would expand out an instance of the restaurant script, thus gaining access to the waiter in that script instance. In the approach based on dynamic Skolemization, it is sufficient to recognize Mary's dining at Mario's as an instance of a person dining at a restaurant (hence instantiating *PDR*); this immediately makes available the role functions in *PDRsteps* (as was illustrated for 29a).

It has been known for a long time that frame slots involve Skolem-like functions, logically speaking (Hayes, 1979). Furthermore, role-functions have been proposed in the planning literature for picking out steps in goal-oriented courses of action (e.g., Kautz 1991). However, the representational devices employed in scripts, frames and plans were conceived on purely computational and introspective grounds, with no formal links to the syntax/semantics interface. So it seems auspicious that considerations arising from problems in linguistic anaphora lead directly to script/frame-like representations, and in doing so also provide solutions to some traditionally difficult reference problems.

Schank and his collaborators were very interested in the problem of *learning* scripts by generalization from particular narratives containing repetitive patterns of events (Schank, 1982). While this type of generalization learning is undoubtedly of crucial importance for AI, the present work clears some obstacles from a more direct path to the acquisition of such knowledge, namely explicit description of the general patterns in ordinary language. Lenat (1995), as mentioned earlier, has already mounted a large-scale effort to encode general knowledge directly in a frame-like representation. But this has depended so far on familiarizing the contributors to this effort with the internal representation, and entrusting to their intuitions the correct formulation of common-sense facts in this representation. Such an effort would surely benefit from an effective linguistic input channel. Language is the shared medium for the exchange of common-sense human knowledge, and as such would provide a much faster, much more natural means for knowledge bootstrapping; as well, if the mapping from the internally stored information to language were straightforward and systematic, browsing and verifying the stored knowledge, and any inferences drawn from it, would become much easier.

5 Further Issues

A variety of issues remain concerning the scope and details of dynamic Skolemization. Most require further research, but I want to at least provide a preliminary discussion.

5.1 Generic passages

In the discussion of the restaurant scenario so far, I have oversimplified matters. In particular, I assumed just one tacit all-encompassing generic quantifier and restrictor (roughly, "*For most situations e in which there is a person x and a restaurant y such that x dines at y, ...*"). More realistic examples are considered in (Carlson & Spejewski, 1997) under the heading of *generic passages.* These may (tacitly or explicitly) involve a succession of different quantifiers and restrictors. For instance, the restaurant scenario might contain a subpassage such as

> If the restaurant is informal, the menus are *often* already on the table. If it is formal the waiter *usually* brings the menus. When the waiter arrives at the table, he *usually* begins by requesting orders for drinks. He *always* asks the ladies first...

It can be seen that successive sentences can add their own restrictions to one implied by the preceding context, and can supply their own generic quantifier. (Arguably the quantifier is always there, explicitly or tacitly.) The quantification implicitly refers to the overall event type (i.e., dining-out) and may in addition be restricted by the subordinate situation types specified in previous sentences. For instance, the last sentence above (about asking the ladies first) quantifies implicitly over events of requesting orders for drinks (introduced in the *previous* sentence) that are part of a dining-out event.

Thus the logical form of a generic passage will be something like

$(Ge : \phi(e))$

$\quad (G_1 e_1 : \phi_1(e_1))\psi_1(e_1) \wedge ... \wedge (G_n e_n : \phi_n(e_n))\psi_n(e_n),$

where $G, G_1, ..., G_n$ are various generic event quantifiers (such as "generally", "usually", "sometimes", etc.). The ϕ_i and ψ_i can of course contain arbitrary additional quantifiers, and in particular the ψ_i will typically introduce new existentially quantified events/situations (e.g., the waiter's arriving at the table, the waiter's requesting orders for drinks, etc.) Now there is a difficulty here very similar to the one encountered in the 'farmers at the fair' example. We have a series of quantificationally closed sentences, yet we need to relate events introduced in the ψ_i to each other temporally (and in other ways); and in addition, as noted the restrictors (i.e., the ϕ_i) need to refer to event-types introduced in previous sentences, as well as to the overall event type.

In a DRT-based approach such as is developed in (Carlson & Spejewski, 1997) these phenomena require a great deal of forward propagation of material from both restrictor clauses and main clauses, in order to restrict each successive quantifier appropriately, and make temporal and other connections. (This is similar to techniques used in the DRT literature for modal subordination.) Something quite analogous would have to be done in a DPL representation.

The promise of dynamic concept definition lies in the fact that it would automatically supply both Skolem functions for the events introduced in the ψ_i, and new event predicates corresponding to ϕ (the main event-type being described) and those of the ϕ_i and ψ_i that involve existential quantification. These are precisely what we need to relate the events in the generic passage in time (etc.), and to succinctly express the restrictions ϕ_i, *without* carrying forward any of the prior material. Thus it appears the earlier remarks about the advantages of adding dynamic concept definition to the process of logical-form calculation gain further support when we move from simplified scenarios to genuine generic passages. The technique should enable us to obtain direct, succinct representations of such passages, in a form that admits a non-dynamic semantics and allows the application of familiar inference techniques.

5.2 Anaphora within negative environments

In all of the examples of anaphora seen so far, the indefinite E_i and the coindexed pronoun or definite D_i were 'positively linked' in the following sense: they occurred within a sentential complex of form

$\Phi(E_i)\ op\ \Psi(D_i)$,

where the connecting operator *op* lies in a positive polarity (upward entailing) environment. For instance, this is true for anaphoric links across the two components of a (top-level) conjunction or a (top-level) if-then sentence. This bias in our examples (and elsewhere in the literature) in favor of positive environments is a consequence of the fact that extended negative environments are rare in natural language.

Nonetheless, it is of interest how dynamic Skolemization fares in negative environments. My tentative conclusion is that dynamic Skolemization needs to be modified in such environments. In particular, the Skolem conditionals need to take account of all of the material within such environments – i.e., in terms of the above schema, both the Φ-material and the Ψ-material.

The following are some relevant examples, with corresponding provisional logical forms shown in (b) (using free variables, as in DPL or EL, to provisionally represent any anaphoric pronouns). The (c)-items show the corresponding dynamically Skolemized representations that would be obtained if we applied the same rewriting rules to negatively linked indefinite-definite pairs as to positively linked pairs. In each case the unmodified Skolem conditional would be

$(\exists x : D(x))B(J, x) \ \rightarrow\ D(A) \wedge B(J, A)$.

33. (a) John did not buy a donkey
    (b) $\neg(\exists x : D(x))B(J, x))$
    (c) $\neg(D(A) \wedge B(J, A))$

34. (a) John didn't buy a donkey and pay for it in cash.
    (b) $\neg[((\exists x : D(x))B(J, x)) \wedge P(J, x)]$
    (c) $\neg(D(A) \wedge B(J, A) \wedge P(J, A))$

35. (a) If John bought a donkey and paid for it in cash, then he rode it home.
    (b) $[((\exists x : D(x))B(J, x)) \wedge P(J, x)] \rightarrow R(J, x)$
    (c) $(D(A) \wedge B(J, A) \wedge P(J, A)) \rightarrow R(J, A)$

36. (a) It's not the case that if John bought a donkey, he paid for it in cash.
    (b) $\neg[(\exists x : D(x))B(J, x) \rightarrow P(J, x)]$
    (c) $\neg[(D(A) \wedge B(J, A)) \rightarrow P(J, A)]$

The first example, (33), involves an explicitly negative context for Skolemization, but no anaphora. Despite the flagrant violation of the conditions under which ordinary Skolemization would be licensed, the Skolemized representation, (33c), evidently has the desired truth conditions, taken together with the Skolem conditional.

(34) is our first example of a negative linkage between an indefinite and an anaphoric pronoun. Does (34c) (along with the Skolem conditional) capture the intuitive meaning of (34a) (which corresponds to the DPL meaning of (34b))? Well, it does as long as we let a uniqueness assumption stand. For, if John bought just one donkey, then by the Skolem conditional, $D(A) \wedge B(J, A)$ holds, and hence from $\neg(D(A) \wedge B(J, A) \wedge P(J, A))$, $\neg P(J, A)$ holds, i.e., John did not pay in cash for the one donkey he bought. However, we can preempt the uniqueness implicature by prefacing (34a)

with the disclaimer, "*I don't know how many donkeys, if any, John bought, but...*". Even in this context (34a) says that John did not pay in cash for *any* donkey he bought, no matter how many he bought; whereas (34c) incorrectly allows for the possibility that John bought two donkeys and paid for one (but not the other) in cash.

As already indicated, I believe that the correct Skolem conditional for (34) would take account of all of the material in the negatively embedded conjunction; i.e., we need to define A as any donkey that John bought *and paid for in cash*, if there is such a donkey. It is as if we tacitly widened the scope of the existential quantifier in (34b) so that its scope is the entire conjunction, allowing it to bind the pronoun (non-dynamically). Here it is worth citing an analogous example involving a universal quantifier, apparently requiring expansion of the quantifier binding scope:[6]

37. John didn't go to each donkey and give it a sugar cube.

Example (35) further illustrates anaphoric binding within a negatively embedded conjunction, since the conjunction comprises the antecedent of a (positively occurring) conditional. Again (c) fails to provide correct truth conditions for (a) if John bought more than one donkey. For if A picks out any donkey John bought, and he happened to pay for that donkey by credit card, then (35c) will be true (in virtue of the falsity of the conjunctive antecedent) – even if John bought a second donkey, paid for it in cash, and didn't ride any donkey home. Once again the correct Skolem conditional would be one that amalgamates the content of both clauses of the conjunction.

Example (36) is a little harder to get an intuitive hold on, but I think it illustrates the same point. In this case the anaphoric linkage bridges the clauses of a negatively embedded conditional, rather than a negatively embedded conjunction. (36a) implies that it's not the case that John either didn't buy a donkey, or bought a donkey that he paid for in cash. Hence John did buy a donkey, but didn't buy a donkey that he paid for in cash. But again (36c) merely entails that John bought a donkey A and didn't pay for A in cash – incorrectly leaving open the possibility that he bought another donkey A' and paid for it in cash. The correct Skolem conditional should again express that A is a donkey that John paid for in cash, if there is such a donkey, i.e., it should amalgamate the material from the antecedent and consequent of the negatively embedded conditional.[7] Then (36c) would entail (*via* the Skolem conditional) that there is no donkey that John bought and paid for in cash, while also entailing $D(A)$ and $B(A)$ (he did buy a donkey).

The details of the above modification of dynamic Skolemization for negative anaphoric linkages, when the coindexed terms lie within arbitrarily complex sentences, are a matter for further research. The need to change the rules of dynamic Skolemization for negative anaphoric linkages would diminish its appeal as an *alternative* to dynamic semantics, if it were offered as such. However, that is not the intention; the very fact that dynamic Skolemization involves the stipulation of new definitions disqualifies it as a *semantic* theory (especially if we are partial to compositional semantics). Its promise lies elsewhere, namely in the way it links up dynamic semantic representations with 'standard' knowledge representations, and in the way it implements functional anaphoric connections.

Finally, let me mention a very tangentially related point concerning apparent cases of external reference into negated contexts, as in

38. John didn't buy a DONKEY. It was a mule.

[6]However, I am not suggesting quantifier scope widening as the actual mechanism to be be employed in such cases, but rather a direct expansion of the Skolem conditional to incorporate the extra material.

[7]If we wish to view this amalgamation in terms of scope widening, we would first have to rewrite the conditional $\Phi \rightarrow \Psi$ as $\neg\Phi \vee (\Phi \wedge \Psi)$, confining the scope-widening to the conjunction. But this would have the undesirable consequence of generating two Skolem conditionals, one for each disjunct.

The result of dynamically Skolemizing the indefinite in the negated context would again be unsatisfactory. In particular, it would fail to capture the entailment that John bought a mule. But this is to be expected, since the interpretation of (38) depends upon a presupposition, mediated by the indicated stress on A DONKEY, that John did buy SOMEthing. It appears, then, that the correct referent for the pronoun comes from that presupposition, rather than from the negative sentence.

5.3 Uniqueness and strong readings

I invoked uniqueness implicatures in connection with arguments like the following.

> If John bought a donkey, he paid for it in cash. John bought a donkey named Archie.
> Therefore, John paid for Archie in cash.

Now while the existence of uniqueness implicatures is widely acknowledged, their role in donkey anaphora remains controversial. One point that is frequently made is that uniqueness implicatures cannot be used to justify a 'lazy definite' account of anaphoric pronouns (such as suggested by Geach, 1962; Partee, 1975; Evans, 1980). On such an account, we could reformulate "*If John bought a donkey, he paid for it in cash*" as "*If John bought a donkey, he paid for the donkey he bought in cash*". Here successful reference by the phrase "*the donkey he bought*" depends on the uniqueness of that donkey, if John bought one at all. But since uniqueness is merely implicated, the implicature may be cancelled, as in the following sentence due to Kratzer (1995:158) (a variant of Heim's 'sageplant' sentences):

39. When a house has a barn, it often has a second one right next to it (the first barn).

Clearly the final "*it*" cannot be rephrased as "*the house's barn*". However, this is not a problem for the Skolem-function approach. Even when the uniqueness implicature is cancelled, the Skolem function still refers uniquely. In other words, the Skolem-function approach is more akin to a 'lazy indefinite' account of anaphoric pronouns in this respect.

However, there is a general objection to any sort of pronoun-of-laziness account, which would also carry over to dynamic Skolemization: we are left with no account of strong readings of donkey sentences, which appear to quantify universally over all verifying instances of an existential sentence. For instance, "*If John bought a donkey, he paid for it in cash*" is claimed to have a strong reading, even on a nongeneric construal. According to this reading, John paid for *every* donkey he bought in cash. Note that this would provide an alternative justification for the above argument concerning Archie the donkey.

Kamp (1981) and Heim (1982) and most researchers after them took it as their central task to account for strong readings, and formulated their dynamic semantics accordingly. Schubert and Pelletier (1989) pointed out the importance of weak (nonuniversal) readings, taking these as basic while suggesting that strong readings arise from generic construals. Chierchia (1995) considers these matters in detail and concludes that weak readings are basic, in the sense that dynamic binding yields weak rather than strong readings. Nonetheless, he maintains (with Kanazawa, 1994) that many donkey sentences, even nongeneric ones, have strong readings, and offers a novel account of these. An example he attributes to Heim is

40. Every man who owned a slave owned his offspring

(Chierchia, 1995:64). Certainly one feels that this sentence tends to rule out cases where a slave owner owns the offspring of some, but not all, of his slaves. Chierchia suggests that in addition to a dynamically bound reading, donkey pronouns can also have an E-type reading (Evans, 1980), specifically one which uses a contextually determined function whose value may be a plural entity. On this account, there is a reading of "*his offspring*" in (40) tantamount to "the offspring of *the slaves owned by* the given slave-owner"; evidently this yields the strong reading. So we are invoking a plural-valued (or at least number-neutral) function "*the slaves owned by*", somehow supplied by the linguistic context (where the subject noun phrase presumably plays the key role).

From the present perspective such a suggestion is rather attractive insofar as we already have in hand a simple mechanism for extracting functions from context; we could implement Chierchia's suggestion by allowing for definition of set- or collection-valued Skolem functions.[8] However, my own intuitions about examples like (40) are uncertain. It seems possible to me that this sentence has no direct entailments concerning a slave owner's 'extra slaves', if any. Rather, it may make a definite claim only about one slave per slave owner – but may invite inference of a stronger statement to cover the cases left open by the uniqueness implicature.[9] Such inferences may derive from our disposition to make sense of what we are told in terms of more general principles and trends, and to reconcile the phrasing and content of the claims made with our knowledge of the world and of conversational maxims. In the case of (40), we might make sense of the weak reading (which has no entailments for a slave owner's 'extra slaves') in terms of some presumed law of ownership likely to be in effect during times of slavery, and conclude that there is no reason why this law would not extend to the 'extra slaves', thus arriving at the strong 'reading'.

I would point out that such strengthened interpretations seem to arise even in the absence of explicit indefinites, for instance in

41. Everyone who burped said "excuse me",

where we tend to conclude that repeat-burpers made repeated apologies. It would be a stretch to invoke a function from people to sets of burps here, as a way of capturing this inference. Rather (as in the case of the slave-ownership sentence), we seem to generalize over the event type at issue (burping) in a way that goes beyond the explicit quantification over people who burped at least once (cf. the appeal to implicit quantification over times in Evans, 1980). Another example where it is hard to see how context could supply a suitable function for a strong reading is the following:

42. If there is an extraterrestrial civilization, it hasn't contacted humanity.

This appears to rule out contact with all extraterrestrial civilizations, however many there might be. We would need something like a function from sectors of the universe to the set of civilizations in that sector (where the contextually relevant sector is the universe exclusive of Earth), to implement Chierchia's suggestion here. My inclination is again to regard the strong interpretation as an explanatory inference: why would anyone make a claim like (42), if not as a comment on the lack of evidence for an extraterrestrial civilization? And if there is no evidence for one, there is no evidence for any.

More familiar sentences of this form, such as (43), appear susceptible to the same kind of account:

---

[8] In the case where $(\exists x : \Phi)\Psi$ contains no free variables, we would stipulate
$\forall x.(x \in C) \leftrightarrow (\Phi \wedge \Psi)$.
When there are free variables $x_1, ..., x_m$ in the existential formula, we would use Skolem function $f(x_1, ..., x_m)$ in place of C and add external $\forall$-quantifiers for $x_1, ..., x_m$.

[9] Some unconditional sentences where a uniqueness presupposition is likely to be undermined by general knowledge also appear to have a 'strong reading': "*I have no money in my pocket*"; "*John has a strong arm and a sharp eye*"; "*John is scarcely aware of his next-door neighbor*"; etc.

43. (a) If John bought a donkey he didn't pay for it in cash.

    (b) $[((\exists x : D(x))B(J,x)] \rightarrow \neg P(J,x)]$

    (c) $\neg(D(A) \wedge B(J,A) \wedge P(J,A))$

Again a strong reading appears to be dominant,[10] and I would again surmise that this derives from a generalization from the weak reading. John's not paying in cash for one donkey suggests he was unable or unwilling to come up with the cash, and this makes it very unlikely that he would pay cash for any *additional* donkeys. For comparison note the following variant of a sentence due to Kanazawa (1994), where the weak reading does prevail:

44. If John has an umbrella, he surely didn't leave it at home today.

5.4 The proportion problem

A further question of interest is how dynamically Skolemized donkey sentences fare with respect to the 'proportion problem' (e.g., Kadmon, 1987). For example, consider the following sentence:

45. In most instances, if a bee stings a man, he survives.

The problem is that if we count every 'stinging event' e of the sort $Stings(x,y,e)$ as distinct, a single instance where a man died after being stung by a thousand bees will outweigh hundreds of nonfatal stingings by isolated bees, and lead to the incorrect prediction that the sentence is false.

One option here is to treat "*in most instances*" as a quantifier that selectively binds times and men, but not stinging events and bees. This corresponds to a paraphrase of (45) as "*For most pairs* (t,y) *of times and men such that there is a bee* x *and a stinging event* e *at time* t *where* x *stings* y *in* e, y *survives*". A disadvantage of this approach is its *ad hoc* introduction of a time variable and its nonuniform treatment of indefinites as sometimes bound by adverbial quantifiers and sometimes existentially quantified (with a dynamic semantics). Be that as it may, such an approach would be compatible with dynamic Skolemization. We would obtain the following Skolem conditional and Skolemized sentence:

$$
\begin{aligned}
&(\forall y)(\forall t)[[(\exists x : B(x))(\exists e : At\!-\!time(e,t))(M(y) \wedge Stings(x,y,e))] \\
&\qquad\quad \leftrightarrow [B(f(y,t)) \wedge At\!-\!time(g(y,t),t) \wedge \\
&\qquad\qquad M(y) \wedge Stings(f(y,t),y,g(y,t))] \\
&(Most\, y,\, t : B(f(y,t)) \wedge M(y) \wedge At\!-\!time(g(y,t),t) \wedge \\
&\qquad\qquad Stings(f(y,t),y,g(y,t)))\ Survives(y),
\end{aligned}
$$

ignoring the event variable of *Survives*, which should be related in some way to the time t. Though the Skolem functions pick out a particular bee and stinging by that bee for any given man and time, there is no presumption that the selected bee is acting alone.

To me a more attractive option is to treat the adverbial as quantifying only over events (more generally, situations/ eventualities/ episodes), but not Davidsonian ones. We regard events as being determined not just by atomic predications, but also by more complex formulas such as $(\exists x : B(x))Stings(x,y)$, as in Situation Semantics (Barwise & Perry, 1983) or Episodic Logic

[10] If we try to express the strong reading of (43) unambiguously, we have to be careful about phrasing. Since that reading is supposed to express iteration over all verifying instances of the indefinite, we might try to express it as "*If John bought a donkey, he didn't pay for every donkey he bought in cash*". However, this is essentially the weak reading, since the negation in the matrix clause tends to outscope the quantifier, making the matrix clause equivalent to "*for some donkey John bought, he didn't pay for that donkey in cash.*"

(Hwang & Schubert, 1993b). Then we can quantify over 'coarse-grained' stinging events that may involve multiple bees. Without going into details (which would require a digression into SS or EL), such an approach would treat both indefinites in (45) as existentially quantified. If we narrow-scope the bee-quantifier and wide-scope the man-quantifier with respect to the operator that connects event variables with sentences describing the events ("$\models$" in SS and "**" in EL), the preferred reading emerges. Dynamic Skolemization can be applied to existential quantifiers outside the scope of this operator, and in the case of (45) this leads to an event-dependent function supplying the victim of a bee-stinging (where any number of bees may be involved).

6 Concluding Remarks

I have shown how one can dynamically Skolemize indefinites encountered in natural language input, so that the resultant constants and functions are available for subsequent reference. This provides a new way of dealing with donkey anaphora and other cases of anaphora standardly handled through some version of dynamic semantics. The advantages are that the resulting representations are context-independent and admit a standard semantics, and that referential connections can be made for instances of functional anaphora that are problematic for dynamic semantics.

The kinds of anaphora at which dynamic Skolemization is aimed pervade ordinary language, particularly the sort of language used to express commonsense generalizations about the world. So the possibility of dealing straightforwardly with heretofore refractory cases should prove very helpful in the effort to impart commonsense knowledge to machines through language.

Moreover, I showed that the representations of general facts obtained through dynamic Skolemization, aided by concept definitions, are strikingly similar to the frame-like and script-like representations that are the stock-in-trade in knowledge representation and reasoning. This convergence between certain kinds of linguistically and nonlinguistically motivated representations provides some reason for optimism about the possibility of a unified approach to representation and inference.

I have also tried to explore some areas of potential difficulty for dynamic Skolemization, including anaphora within negative polarity environments, strong readings, and the proportion problem. This indicated the need for some refinement of the technique for certain kinds of embeddings of coindexed terms, but as far as I can tell there are no insurmountable obstacles to its broad application in natural language processing.

There appear to be some interesting links between my proposal and an approach to anaphora proposed by Meyer Viol (1995). (See also Meyer Viol *et al.*, this volume.) The latter is based on *instantial logic*, a development of Hilbert's ϵ-calculus. As noted by Meyer Viol, there is a close connection between Skolemization and *epsilon*-terms. In fact, the indefinite in a sentence such as

John owns a donkey

would be represented by an ϵ-term

$\epsilon x(D(x) \wedge O(J,x))$,

which can be thought of a picking out an arbitrary donkey owned by John, if there is such a donkey. However, Meyer Viol's goal is not to provide 'standard' logical translations of NL sentences, but rather to develop a logic containing ϵ-terms, and to employ that logic to represent quantification and reference (among other things) in NL. His account of anaphoric pronouns is, in fact, a lazy-indefinite account. The reconstructed content of a pronoun often depends not only on the material in the sentence that introduces the indefinite antecedent of the pronoun, but potentially also on subsequently added material. In particular, successive occurrences of a pronoun that we might naively take as having the same implicit content are taken to have different content; e.g., this is the case for the two occurrences of "*he*" in

A man walked in. He was wearing a hat. He smiled.

where the first is interpreted as having the content "*a man who walked in*", while the second has the content "*a man who walked in and was wearing a hat*". By contrast, dynamic Skolemization would resolve both pronouns to the same Skolem constant. Meyer Viol offers a novel account of strong readings of donkey sentences, based on incorporating the entire content of such sentences into the reconstructed pronoun. For example, the pronoun in

Every farmer who owns a donkey beats it,

on a strong reading, is rendered as (roughly) "*some donkey owned by the farmer which he doesn't beat*". Evidently this entails that each farmer beats all of his donkeys, for if he had any donkeys he doesn't beat, such a donkey would be picked out by the ϵ-term, rendering the matrix clause false (indeed logically false). It would be possible in principle to incorporate this idea into dynamic Skolemization, but I would hesitate to complicate the dynamic Skolemization rule to that extent.

As a final comment, I would also mention the possibility of transmuting dynamic Skolemization into a truly semantic theory. The idea would be incorporate the creation of functions, corresponding to indefinites in quantified contexts, into dynamic predicate logic. This may be possible by generalizing variable assignments so that variables can be dynamically bound to functions, as well as to individuals, and requiring all occurrences of existential quantifiers to bind distinct variables (which could then be used freely outside their binding context). The attraction of such a move would be that it would dispense with separately asserted Skolem conditionals, preserve a DPL-like compositional semantics for donkey sentences, and extend this semantics to functional anaphora. Script- and frame-like structures would be created even more directly than in the scheme proposed herein, since it would no longer be necessary to substitute Skolem terms for indefinites – the names of the corresponding existential variables would already denote the requisite individuals or functions.

Acknowledgements

This work was supported in part by NSF grant IRI-9503312 and by ARPA/SSTO grant F30602-95-1-0025. I am grateful to Greg Carlson for useful pointers to the literature (including the draft of his paper with Beverly Spejewski on generic passages), and to Wilfried Meyer Viol and Ruth Kempson for alerting me to the cited work on instantial logic, and providing helpful comments on its application to anaphora.

References

Allen, J. F. and L. K. Schubert (1991) The TRAINS project. Tech. Rep. 382, Dept. of Computer Science, Univ. of Rochester, Rochester, NY. Also slightly revised as "Language and discourse in the TRAINS project", in A. Ortony, J. Slack, and O. Stock (eds.), *Communication from an Artificial Intelligence Perspective.* Heidelberg: Theoretical Springer-Verlag, 91-120, 1993.

Barwise, J. and J. Perry (1983) *Situations and Attitudes.* Cambridge, MA: Bradford Books, MIT Press.

Carlson, G. N. and F. J. Pelletier (1995) *The Generic Book.* Chicago: Univ. of Chicago Press.

Carlson, G. N. and B. Spejewski (1997) Generic passages. *Natural Language Semantics 5*, 1-65 (to appear).

Chierchia, G. (1995) *Dynamics of Meaning.* Chicago: Univ. of Chicago Press.

Evans, G. (1980) Pronouns, quantifiers and relative clauses (I). In M. Platts (ed.), *Reference, Truth and Reality*, New York: Routeledge, Kegan Paul.

Geach, P. (1962) *Reference and Generality.* Ithaca, N.Y.: Cornell University Press.

Groenendijk, J. and M. Stokhof (1991) Dynamic predicate logic. *Linguistics and Philosophy 14*, 39-100.

Hayes, P. J. (1979) The logic of frames. In D. Metzing (ed.), *Frame Conceptions and Text Understanding*, New York: de Gruyter.

Heim, I. (1982) *The Semantics of Definite and Indefinite Noun Phrases.* Ph.D. Dissertation, U. Mass.

Hwang, C. H. and L.K. Schubert (1993a) Episodic Logic: a comprehensive, natural representation for language understanding. *Minds and Machines 3*(4): Special Issue on KR for NLP, 381-419.

Hwang, C. H. and L.K. Schubert (1993b) Episodic Logic: a situational logic for natural language processing. In P. Aczel, D. Israel, Y. Katagiri and S. Peters (eds.), *Situation Theory and its Applications, v.3 (STA-3)*, CSLI, Stanford, CA, 303-338.

Hwang, C. H. and L.K. Schubert (1992) Tense trees as the 'fine structure' of discourse. *Proc. of the 30th Ann. Meet. of the Assoc. for Computational Linguistics (ACL-92)*, U. Delaware, Newark, DE, June 28 - July 2, 232-240.

Kadmon, N. (1987) *On Unique and Non-Unique Reference and Asymmetric Quantification.* Ph.D. Dissertation, U. Mass., Amherst, MA.

Kadmon, N. (1990) Uniqueness. *Linguistics and Philosophy 13*, 273-324.

Kamp, H. (1981) A theory of truth and semantic representation. In J. Groenendijk, T. Janssen, and M. Stokhof (eds.), *Formal Methods in the Study of Language*, Mathematical Centre-tracts, U. Amsterdam, Amsterdam, Netherlands.

Kanazawa, M. (1994) Weak vs, strong readings of donkey sentences and monotonicity inferences in a dynamic setting. *Linguistics and Philosophy 17*(2), 109-158.

Kautz, H. A. (1991) A formal theory of plan recognition and its implementation. In J.F. Allen, H.A. Kautz, R.N. Pelavin, and J.D. Tenenberg, *Reasoning about Plans*, San Mateo, CA: Morgan Kaufmann, 69-125.

Kratzer, A. (1995) Stage-level and individual-level predicates. In (Carlson and Pelletier, 1995), 125-175.

Lascarides, A., N. Asher, and J. Oberlander (1992) Inferring discourse relations in context. In *Proc. of the 30th Ann. Meet. of the Assoc. for Computational Linguistics (ACL-92)*, U. Delaware, Newark, DE, June 28 - July 2, 1-8.

Lenat, D. (1995) CYC: A large-scale investment in knowledge infrastructure. *Comm. of the ACM 38*(11), 33-38.

Meyer Viol, W. P. M. (1995) *Instantial Logic.* IILC Dissertation Series 1995-11, Inst. for Logic, Language and Computation, Univ. of Amsterdam.

Minsky, M. (1975) A framework for representing knowledge. In P.H. Winston (ed.), *The Psychology of Computer Vision*, New York: McGraw-Hill, 211-277.

Muskens, R. (1994) A compositional discourse representation theory. In *Proc. of the 9th Amsterdam Colloquium*, ed. by P. Dekker and M. Stokhof, 467-486.

Partee, B. (1975) Bound variables and other anaphors. In D. L.,Waltz (ed.) *Theoretical Issues in Natural Language Processing-2 (TINLAP-2)*, July 25-27, Univ. of Illinois at Urbana-Champaign, 79-85.

Schank, R. C. (1982) *Dynamic Memory: A Theory of Reminding and Learning in Computers and People.* New York: Cambridge Univ. Press.

Schank, R. C. and R. P. Abelson (1977) *Scripts, Plans, Goals and Understanding.* Hillsdale, NJ: Lawrence Erlbaum Assoc.

Schubert, L. K. and C. H. Hwang (1989) An episodic knowledge representation for narrative texts. *Proc. of the 1st Int. Conf. on Principles of Knowledge Representation and Reasoning (KR-89)*, May 15-18, Toronto, Canada, 444-458.

Schubert, L. K. and F. J. Pelletier (1989) Generically speaking, or, using discourse representation theory to interpret generics. in G. Chierchia, B. Partee, and R. Turner (eds.), *Properties, Types, and Meanings II*, Dortrecht: Reidel, 193-268.

Schubert, L. K. and F. J. Pelletier (1982) From English to logic: Context-free computation of 'conventional' logical translations. *Am. J. of Computational Linguistics 8*, 27-44, 1982. Reprinted in B. J. Grosz, K. Sparck Jones, and B. L. Webber (eds.), *Readings in Natural Language Processing.* Los Altos, CA: Morgan Kaufmann, 293-311, 1986.

Webber, B. L. (1988) Tense as discourse anaphor. *Computational Linguistics 14*(2), 61-73.