

Delimited continuations in natural language

Quantification and polarity sensitivity

Chung-chieh Shan

Harvard University
33 Oxford Street

Cambridge, MA 02138 USA

ccshan@post.harvard.edu

ABSTRACT

Making a linguistic theory is like making a programming language: one typically devises a type system to delineate the acceptable utterances and a denotational semantics to explain observations on their behavior. Via this connection, the programming language concept of delimited continuations can help analyze natural language phenomena such as quantification and polarity sensitivity. Using a logical meta-language whose syntax includes control operators and whose semantics involves evaluation order, these analyses can be expressed in direct style rather than continuation-passing style, and these phenomena can be thought of as computational side effects.

Categories and Subject Descriptors

D.3.3 [Programming languages]: Language constructs and features—*control structures*; J.5 [Linguistics]

General Terms

Languages, Theory

Keywords

Delimited continuations, control effects, natural language semantics, quantification, polarity sensitivity

1. INTRODUCTION

This paper is about computational linguistics, in the sense of applying insights from computer science to linguistics. Linguistics strives to scientifically explain empirical observations of natural language. Semantics, in particular, is concerned with phenomena such as the following. In (1) below, some sentences to the left *entail* their counterparts to the right, but others do not.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Continuation Workshop 2004 Venice, Italy

Copyright 200X ACM X-XXXXX-XX-X/XX/XX ...\$5.00.

- (1) Every student passed \vdash Every diligent student passed
No student passed \vdash No diligent student passed
A student passed $\not\vdash$ A diligent student passed
Most students passed $\not\vdash$ Most diligent students passed

The sentence in (2) is *ambiguous* between at least two readings. On one reading, the speaker must decline to run any spot that fails to substantiate any claims whatsoever. On another reading, there exist certain claims (anti-war ones, say) such that the speaker must decline to run any spot that fails to substantiate them.

- (2) We must decline to run any spot that fails to substantiate certain claims.¹

Finally, among the four sentences in (3), only (3a) is *acceptable*. That is, only it can be used in idealized conversation. The unacceptability of the rest is notated with asterisks.

- (3) a. No student liked any course.
b. *Every student liked any course.
c. *A student liked any course.
d. *Most students liked any course.

The linguistic entailments and non-entailments in (1) are facts about English, in that only a speaker of English can make these judgments. Nevertheless, they presumably have to do with corresponding logical entailments and non-entailments: both the English speaker who judges that *Every student passed* entails *Every diligent student passed* and the Mandarin speaker who judges that *Meige xuesheng dou jige-le* entails *Meige yonggong-de xuesheng dou jige-le* rely on knowing that, if every student passed, then every diligent student passed. Thus the typical linguistic theory specifies a semantics for natural language by translating declarative sentences into logical statements with truth conditions. The linguistic entailments in (1) hold, goes the theory, because the meanings—truth conditions—of the two sentences are such that any model that verifies the former also verifies the latter. Much work in natural language semantics aims in this way, as depicted in Figure 1, to explain the horizontal by positing the vertical. This approach is reminiscent of programming language research where an ill-understood language (perhaps one with a complicating feature like ex-

¹This sentence is part of a statement made by the cable television company Comcast after its CNN channel rejected an anti-war commercial hours before it was scheduled to air on January 28, 2003.




Figure 1: The translation/denotation approach to natural language semantics

ceptions) is studied by translation into a simpler language (without exceptions) that is better understood.

The translation target posited in natural language semantics is often some combination of the λ -calculus and predicate logic. For example, the verb *passed* might be translated as $\lambda x. \text{passed}(x)$. This paper argues by example that the translation target should be a logical metalanguage with first-class delimited continuations. The examples are two natural language phenomena: *quantification* by words like *every* and *most* in (1), and *polarity sensitivity* on the part of words like *any* in (3).

Quantification was first analyzed explicitly using continuations by Barker (2002). Building on that insight, this paper makes the following two contributions. First, I analyze natural language in direct style rather than in continuation-passing style. In other words, the logical metalanguage used here is one that includes control operators for delimited continuations, rather than a pure λ -calculus in which denotations need to handle continuations explicitly. Natural language is thus endowed with an operational semantics from computer science that is richer than just $\beta\eta$ -reduction.

Second, I propose a new analysis of polarity sensitivity that improves upon prior theories in explaining why *No student liked any course* is acceptable but **Any student liked no course* is not. This analysis crucially relies on the notion of evaluation order from programming languages, thus elucidating the role of control effects in natural language and supporting the broader claim that linguistic phenomena can be fruitfully thought of as computational side effects.

The rest of this paper is organized as follows. In §2, I introduce a simple grammatical formalism. In §3, I describe the linguistic phenomenon of quantification and show a straw man analysis that deals with some cases but not others. I then introduce a programming language with delimited continuations and use it to improve the straw man analysis: quantification in non-subject position is treated in §4, and inverse scope is covered in §5. In §6, I turn to the linguistic phenomenon of polarity sensitivity and show how a computationally motivated notion of evaluation order improves upon previous analyses. In §7, I place these examples in a broader context and conclude.

2. A GRAMMATICAL FORMALISM

In this section, I introduce a simple grammatical formalism for use in the rest of the paper. It is a notational variant of categorial grammar (as introduced by Carpenter (1997; chapter 4), for instance).

The verb *like* usually requires an object to its right and a subject to its left.

- (4) a. Alice liked CS187.
 b. *Alice liked.
 c. *Alice liked Bob CS187.

Intuitively, *like* is a function that takes two arguments, and the sentences (4b–c) are unacceptable due to type mismatch. We can model this formally by assigning types to the denotations of *Alice*, *CS187*, and *liked*, which we take to be atomic expressions.

- (5) $\llbracket \text{Alice} \rrbracket = \text{alice} : \text{Thing}$
 (6) $\llbracket \text{CS187} \rrbracket = \text{cs187} : \text{Thing}$
 (7) $\llbracket \text{liked} \rrbracket = \text{liked} : \text{Thing} \rightarrow \text{Thing} \rightarrow \text{Bool}$

Here **Thing** is the type of individual objects, and **Bool** is the type of truth values or propositions. Following (justifiable) standard practice in linguistics, we let *liked* take its object as the first argument and its subject as the second argument. For example, in (4a), the first argument to *liked* is *cs187*, and the second argument is *alice*.

As (4a) shows, there are two ways to combine expressions. A function can take its argument either to its right (combining *liked* with *CS187*) or to its left (combining *Alice* with *liked CS187*). We denote these two cases with two infix operators: “/” for forward combination and “\” for backward combination. (The tick marks depict the direction in which a function “leans on” an argument.)

- (8) $f / x = f(x) : \beta$ where $f : \alpha \rightarrow \beta, x : \alpha$
 (9) $x \backslash f = f(x) : \beta$ where $f : \alpha \rightarrow \beta, x : \alpha$

We can now derive the sentence (4a)—that is, prove it to have type **Bool**. The derivation can be written as a tree (10) or a term (11).



- (11) $\llbracket \text{Alice} \rrbracket \backslash (\llbracket \text{liked} \rrbracket / \llbracket \text{CS187} \rrbracket) = \text{liked cs187 alice} : \text{Bool}$

By convention, the infix operators / and \ associate to the right, so parentheses such as those in (11) above are optional.

Unfortunately, the system set up so far derives not only the acceptable sentence (4a) but also the unacceptable sentence (12), with the same meaning.

- (12) *Alice CS187 liked.

The reason the system derives (12) is that the direction of function application is unconstrained: in the derivation below, *liked* takes its first (object) argument to the left, which is usually disallowed in English.



- (14) $\llbracket \text{Alice} \rrbracket \backslash \llbracket \text{CS187} \rrbracket \backslash \llbracket \text{liked} \rrbracket = \text{liked cs187 alice} : \text{Bool}$

To rule out this derivation of (12) in our type system, we split the function type constructor “ \rightarrow ” into two type con-

constructors “ $\dot{\rightarrow}$ ” and “ $\dot{\leftarrow}$ ”, one for each direction of application. Using these new type constructors, we change the denotation of *liked* to specify that its first argument is to its right and its second argument is to its left.

$$(15) \quad \llbracket \text{liked} \rrbracket = \text{liked} : \text{Thing} \dot{\rightarrow} \text{Thing} \dot{\leftarrow} \text{Bool}$$

We also revise the combination rules (8) and (9) to require different function type constructors.

$$(16) \quad f \dot{\leftarrow} x : \beta \quad \text{where } f : \alpha \dot{\leftarrow} \beta, x : \alpha$$

$$(17) \quad x \dot{\leftarrow} f : \beta \quad \text{where } f : \alpha \dot{\leftarrow} \beta, x : \alpha$$

The system now rejects (12) while continuing to accept (4a), as desired.

3. QUANTIFICATION

The linguistic phenomenon of quantification is illustrated by the following sentences.

- (18) a. Every student liked CS187.
 b. Some student liked every course.
 c. Alice consulted Bob before most meetings.

As with the previously encountered sentences, the natural language semanticist wants to translate English into logical formulas that account for entailment and other properties. More precisely, the problem is to posit translation rules that map these sentences thus. For instance, we would like to map (18a) to a formula like

$$(19) \quad \forall x. \text{student}(x) \Rightarrow x \dot{\leftarrow} \text{liked} \dot{\leftarrow} \text{cs187} : \text{Bool},$$

where the constants

$$(20) \quad \forall : (\text{Thing} \rightarrow \text{Bool}) \rightarrow \text{Bool}, \quad \Rightarrow : \text{Bool} \rightarrow \text{Bool} \rightarrow \text{Bool}$$

are drawn from the (higher-order) abstract syntax of predicate logic. To this end, what should the subject noun phrase *every student* denote? Unlike with *Alice*, there is nothing of type **Thing** that the quantificational noun phrase *every student* can denote and still allow the desired translation (19) to be generated. At the same time, we would like to retain the denotation that we previously computed for the verb phrase *liked CS187*, namely $\text{liked} \dot{\leftarrow} \text{cs187}$. Taking these considerations into account, one way to translate (18a) to (19) is for the determiner *every* to denote

$$(21) \quad \llbracket \text{every} \rrbracket = \lambda r. \lambda s. \forall x. r(x) \Rightarrow x \dot{\leftarrow} s \\ : (\text{Thing} \rightarrow \text{Bool}) \dot{\leftarrow} (\text{Thing} \dot{\leftarrow} \text{Bool}) \dot{\leftarrow} \text{Bool}.$$

Here the *restrictor* r and the *scope* s are λ -bound variables intended to receive, respectively, the denotations of the noun *student* (of type $\text{Thing} \rightarrow \text{Bool}$) and the verb phrase *liked CS187* (of type $\text{Thing} \dot{\leftarrow} \text{Bool}$). (More precisely, r and s are λ -bound variables; the tick mark again signifies the direction of function application.) In a non-quantificational sentence like (4a), the verb phrase takes the subject as its argument; by contrast, in the quantificational sentence (18a), the subject takes the verb phrase as its argument.

Extended with the lexical entry (21) for *every*, and assuming that *student* denotes

$$(22) \quad \llbracket \text{student} \rrbracket = \text{student} : \text{Thing} \rightarrow \text{Bool},$$

the grammar can derive the sentence (18a).

$$(23) \quad \begin{array}{c} \diagup \quad \diagdown \\ \text{every} \quad \text{student} \quad \text{liked} \quad \text{CS187} \end{array}$$

$$(24) \quad (\llbracket \text{every} \rrbracket \dot{\leftarrow} \llbracket \text{student} \rrbracket) \dot{\leftarrow} \llbracket \text{liked} \rrbracket \dot{\leftarrow} \llbracket \text{CS187} \rrbracket = (19)$$

The existential determiner *some* can be analyzed similarly: let *some* denote

$$(25) \quad \llbracket \text{some} \rrbracket = \lambda r. \lambda s. \exists x. r(x) \wedge x \dot{\leftarrow} s \\ : (\text{Thing} \rightarrow \text{Bool}) \dot{\leftarrow} (\text{Thing} \dot{\leftarrow} \text{Bool}) \dot{\leftarrow} \text{Bool}$$

to derive the sentence *Some student liked CS187*.

$$(26) \quad \begin{array}{c} \diagup \quad \diagdown \\ \text{some} \quad \text{student} \quad \text{liked} \quad \text{CS187} \end{array}$$

$$(27) \quad (\llbracket \text{some} \rrbracket \dot{\leftarrow} \llbracket \text{student} \rrbracket) \dot{\leftarrow} \llbracket \text{liked} \rrbracket \dot{\leftarrow} \llbracket \text{CS187} \rrbracket \\ = \exists x. \text{student}(x) \wedge x \dot{\leftarrow} \text{liked} \dot{\leftarrow} \text{cs187} : \text{Bool}$$

To summarize, we treat determiners like *every* and *some* as functions of two arguments: the restrictor and the scope of a quantifier, both functions from **Thing** to **Bool**. Such higher-order functions are a popular analysis of natural language determiners, and have been known to semanticists since Montague (1974) as *generalized quantifiers*. However, the simplistic account presented above only handles quantificational noun phrases in subject position, as in (18a) but not (18b) or (18c). For example, in (18b), neither forward nor backward combination can apply to join the verb *liked*, of type $\text{Thing} \dot{\leftarrow} \text{Thing} \dot{\leftarrow} \text{Bool}$, to its object *every course*, of type $(\text{Thing} \dot{\leftarrow} \text{Bool}) \dot{\leftarrow} \text{Bool}$. Yet, empirically speaking, the sentence (18b) is not only acceptable but in fact ambiguous between two available readings. This problem has prompted a great variety of supplementary proposals in the linguistics literature (Barwise and Cooper 1981; Hendriks 1993; May 1985; inter alia). The next section presents a solution using *delimited continuations*.

4. DELIMITED CONTINUATIONS

First-class continuations represent “the entire (default) future for the computation” (Kelsey, Clinger, Rees et al. 1998). Refining this concept, Felleisen (1988) introduced *delimited continuations*, which encapsulate only a prefix of that future. This paper uses *shift* and *reset* (Danvy and Filinski 1989, 1990), a popular choice of control operators for delimited continuations.

To review briefly, the shift operator (notated ξ) captures the current context of computation, removing it and making it available to the program as a function. For example, when evaluating the term

$$(28) \quad 10 \times (\xi f. 1 + f(2)),$$

the variable f is bound to the function that multiplies every number by 10. Thus the above expression evaluates to 21 via the following sequence of reductions. (The reduced subexpression at each step is underlined.)

$$(29) \quad [10 \times (\xi f. 1 + f(2))] \\ \triangleright [1 + (\lambda v. [10 \times v])(2)] \\ \triangleright [1 + [10 \times 2]] \triangleright [1 + [20]] \triangleright [1 + 20] \triangleright [21] \triangleright 21$$

Term reductions are performed deterministically in applicative order: call-by-value and left-to-right.

The reset operator (notated with square brackets $[]$) delineates how far shift can reach: shift captures the current context of computation up to the closest dynamically enclosing reset. Hence “ $3 \times$ ” below is out of reach.

$$(30) \quad [3 \times [10 \times (\xi f. 1 + f(2))]] \\ \triangleright [3 \times [1 + (\lambda v. [10 \times v])(2)]] \\ \triangleright [3 \times [1 + [10 \times 2]]] \triangleright [3 \times [1 + [20]]] \triangleright \dots \triangleright 63$$

Shift and reset come with an operational semantics (illustrated by the reductions above) as well as a denotational semantics (via the CPS transform). That both kinds of semantics are available is important to linguistics, because the meanings of natural language expressions (studied in semantics) need to be related to how humans process them (studied in psycholinguistics).

Quantificational expressions in natural language can be thought of as phrases that manipulate their context. In a sentence like *Alice liked CS187* (4a), the context of *CS187* is the function mapping each thing x to the proposition that Alice liked x . Compared to the proper noun *CS187*, what is special about a quantificational expression like *every course* is that it captures its surrounding context when used.

$$(31) \quad \text{Alice liked [every course].}$$

Thus, loosely speaking, the meaning of the sentence (31) no longer has the overall shape $\text{alice} \setminus \text{liked} / \dots$ once the occurrence of *every course* is considered, much as the meaning of the program (28) no longer has the overall shape $10 \times \dots$ once the shift expression is evaluated. Let us add shift and reset to the target language of our translation from English. We can then translate *every course* as

$$(32) \quad \llbracket \text{every course} \rrbracket = \xi s. \forall x. \text{course}(x) \Rightarrow s(x) : \text{Thing}_{\text{Bool}}^{\text{Bool}}.$$

The type notation α_γ^δ here indicates an α with a control effect; the CPS transform maps it to $(\alpha \rightarrow \gamma) \rightarrow \delta$. The denotation of *every course* behaves locally as a **Thing**, but requires the current context to have the answer type **Bool** and maintains that answer type.

To see the new denotation (32) in action, let us derive the sentence (31). The type of *every course* is $\text{Thing}_{\text{Bool}}^{\text{Bool}}$, similar to the type **Thing** of *CS187*, so the derivation of (31) is analogous to (10–11).

$$(33) \quad \begin{array}{c} \text{Alice} \\ \diagdown \quad \diagup \\ \text{liked} \quad \text{every course} \end{array}$$

$$(34) \quad \llbracket \llbracket \text{Alice} \rrbracket \setminus \llbracket \text{liked} \rrbracket / \llbracket \text{every course} \rrbracket \rrbracket \\ = [\text{alice} \setminus \text{liked} / \xi s. \forall x. \text{course}(x) \Rightarrow s(x)] \\ \triangleright [\forall x. \text{course}(x) \Rightarrow (\lambda v. [\text{alice} \setminus \text{liked} / v])(x)] \triangleright \dots \\ \triangleright \forall x. \text{course}(x) \Rightarrow \text{alice} \setminus \text{liked} / x : \text{Bool}$$

Like the straw man analysis in §3, the denotation in (32) generalizes to determiners other than *every*: we can abstract the noun *course* out of *every course*, and deal with *some student* similarly.

$$(35) \quad \llbracket \text{every} \rrbracket = \lambda r. \xi s. \forall x. r(x) \Rightarrow s(x),$$

$$(36) \quad \llbracket \text{some} \rrbracket = \lambda r. \xi s. \exists x. r(x) \wedge s(x) \\ : (\text{Thing} \rightarrow \text{Bool}) \dot{\rightarrow} \text{Thing}_{\text{Bool}}^{\text{Bool}}$$

(We require here that the restrictor r have the type $\text{Thing} \rightarrow \text{Bool}$, not a type of the form $\text{Thing} \rightarrow \text{Bool}_\gamma^\delta$, so r cannot incur control effects when applied to x . Any control effect in the restrictor, such as induced by the quantificational noun phrase *a company* in the sentence *Every representative of a company left*, must be contained within reset.)

More importantly, unlike the straw man analysis, the new analysis works uniformly for quantificational expressions in subject, object, and other positions, such as in (18a–c). Intuitively, this is because shift captures the context of an expression no matter how deeply it is embedded in the sentence.² By adding control operators for delimited continuations to our logical metalanguage, we arrive at an analysis of quantification with greater empirical coverage.

Figure 2 shows a logical metalanguage that formalizes the basic ideas presented above. It is in this language that denotations on this page are written and reduced. Refining Danvy and Filinski’s original shift-reset language, we distinguish between *pure* and *impure* expressions. An impure expression may incur control effects when evaluated, whereas a pure expression only incurs control effects contained within reset (Danvy and Hatcliff 1992, 1994; Nielsen 2001; Thielecke 2003). This distinction is reflected in the typing judgments: an impure judgment

$$(37) \quad \Gamma \vdash E : \alpha_\gamma^\delta$$

not only gives a type α for E itself but also specifies two answer types γ and δ . By contrast, a pure judgment

$$(38) \quad \Gamma \vdash E : \alpha$$

only gives a type α for E itself. As can be seen in the Lift rule, pure expressions are polymorphic in the answer type.

As mentioned in §2, the use of directionality in function types to control word order is not new in linguistics, but the use of delimited control operators to analyze quantification is. It turns out that we can tie the potential presence of control effects in function bodies to directionality. That is, only directional functions—those whose types are decorated with tick marks—are potentially impure; all non-directional functions we need to deal with, including contexts captured by shift, are pure. Another link between directionality and control effects is that the $\dot{\rightarrow}E$ and $\dot{\rightarrow}E$ rules for directional function application are not merely mirror images of each other: the answer types γ_0 through γ_3 are chained differently through the premises. This is due to left-to-right evaluation.

Having made the distinction between pure and impure expressions, we require in our Shift rule that the body of a shift expression be pure. This change from Danvy and Filinski’s original system simplifies the type system and the CPS transform, but a shift expression $\xi f. E$ in their language may need to be rewritten here to $\xi f. [E]$.

The CPS transform for the metalanguage follows from the typing rules and is standard; it supplies a denotational semantics. The operational semantics for the metalanguage specifies a computation relation between complete terms; it is also standard and shown in Figure 3.

The present analysis is almost, but not quite, the direct-style analogue of Barker’s (2002) CPS analysis. Put in direct-style terms, Barker’s function bodies are always pure, whereas function bodies here can harbor control effects. In

²No matter how deep, that is, up to the closest dynamically enclosing reset. Control delimiters correspond to *islands* in natural language (Barker 2002).

Directions	$\Delta ::= / \mid \backslash$
Types	$\alpha, \beta, \gamma, \delta ::= \text{Thing} \mid \text{Bool} \mid \alpha \rightarrow \beta \mid \alpha \xrightarrow{\Delta} \beta_{\gamma}^{\delta}$
Antecedents	$\Gamma ::= x_1 : \alpha_1, \dots, x_n : \alpha_n$
Terms	$E, F ::= c \mid x \mid \lambda x. E \mid \lambda^{\Delta} x. E \mid FE \mid F \upharpoonright E \mid E \setminus F \mid [E] \mid \xi f. E$
Constants $c : \alpha$	

$$\begin{array}{c}
\overline{\forall : (\text{Thing} \rightarrow \text{Bool}) \rightarrow \text{Bool}} \quad \overline{\exists : (\text{Thing} \rightarrow \text{Bool}) \rightarrow \text{Bool}} \\
\overline{\Rightarrow : \text{Bool} \rightarrow \text{Bool} \rightarrow \text{Bool}} \quad \overline{\wedge : \text{Bool} \rightarrow \text{Bool} \rightarrow \text{Bool}} \\
\overline{\text{student} : \text{Thing} \rightarrow \text{Bool}} \quad \overline{\text{liked} : \text{Thing} \xrightarrow{\Delta} (\text{Thing} \xrightarrow{\Delta} \text{Bool}_{\delta}^{\gamma})_{\gamma}} \quad \dots
\end{array}$$

Pure expressions $\Gamma \vdash E : \alpha$

$$\begin{array}{c}
\frac{c : \alpha}{\Gamma \vdash c : \alpha} \text{Const} \quad \frac{}{\Gamma, x : \alpha \vdash x : \alpha} \text{Var} \quad \frac{\Gamma \vdash E : \alpha_{\alpha}^{\beta}}{\Gamma \vdash [E] : \beta} \text{Reset} \\
\frac{\Gamma, x : \alpha \vdash E : \beta}{\Gamma \vdash \lambda x. E : \alpha \rightarrow \beta} \rightarrow \text{I} \quad \frac{\Gamma, x : \alpha \vdash E : \beta_{\gamma}^{\delta}}{\Gamma \vdash \lambda^{\Delta} x. E : \alpha \xrightarrow{\Delta} \beta_{\gamma}^{\delta}} \Delta \text{I} \quad \frac{\Gamma \vdash F : \alpha \rightarrow \beta \quad \Gamma \vdash E : \alpha}{\Gamma \vdash FE : \beta} \rightarrow \text{E}
\end{array}$$

Impure expressions $\Gamma \vdash E : \alpha_{\gamma}^{\delta}$

$$\begin{array}{c}
\frac{\Gamma \vdash E : \alpha}{\Gamma \vdash E : \alpha_{\gamma}^{\delta}} \text{Lift} \quad \frac{\Gamma, f : \alpha \rightarrow \gamma \vdash E : \delta}{\Gamma \vdash \xi f. E : \alpha_{\gamma}^{\delta}} \text{Shift} \\
\frac{\Gamma \vdash F : (\alpha \xrightarrow{\Delta} \beta_{\gamma_3}^{\gamma_2})_{\gamma_1}^{\gamma_0} \quad \Gamma \vdash E : \alpha_{\gamma_2}^{\gamma_1}}{\Gamma \vdash F \upharpoonright E : \beta_{\gamma_3}^{\gamma_0}} \xrightarrow{\Delta} \text{E} \quad \frac{\Gamma \vdash E : \alpha_{\gamma_1}^{\gamma_0} \quad \Gamma \vdash F : (\alpha \xrightarrow{\Delta} \beta_{\gamma_3}^{\gamma_2})_{\gamma_2}^{\gamma_1}}{\Gamma \vdash E \setminus F : \beta_{\gamma_3}^{\gamma_0}} \xrightarrow{\Delta} \text{E}
\end{array}$$

Figure 2: A logical metalanguage with directionality and delimited control operators

Values	$V ::= U \mid \lambda x. E \mid \lambda^{\Delta} x. E$
Unknowns	$U ::= c \mid UV \mid U \upharpoonright V \mid V \setminus U$
Contexts	$C \langle \rangle ::= \langle \rangle \mid (C \langle \rangle)E \mid C \langle \rangle \upharpoonright E \mid C \langle \rangle \setminus F$ $\quad \mid V(C \langle \rangle) \mid V \upharpoonright C \langle \rangle \mid V \setminus C \langle \rangle$
Metacontexts	$D \langle \rangle ::= \langle \rangle \mid C \langle [D \langle \rangle] \rangle$
Computations $E \triangleright E'$	

$$\begin{array}{l}
D \langle C \langle (\lambda x. E) V \rangle \rangle \triangleright D \langle C \langle E \{x \mapsto V\} \rangle \rangle \\
D \langle C \langle (\lambda x. E) \upharpoonright V \rangle \rangle \triangleright D \langle C \langle E \{x \mapsto V\} \rangle \rangle \\
D \langle C \langle V \setminus (\lambda x. E) \rangle \rangle \triangleright D \langle C \langle E \{x \mapsto V\} \rangle \rangle \\
D \langle C \langle [V] \rangle \rangle \triangleright D \langle C \langle V \rangle \rangle \\
D \langle C \langle \xi f. E \rangle \rangle \triangleright D \langle E \{f \mapsto \lambda x. [C \langle x \rangle]\} \rangle
\end{array}$$

Figure 3: Reductions for the logical metalanguage

other words, function bodies are allowed to shift, as in the determiner denotations in (35) and (36). By contrast, Barker uses *choice functions* to assign meanings to determiners.

5. QUANTIFIER SCOPE AMBIGUITY

Of course, natural language phenomena are never as simple as a couple of programming language control operators. Quantification is no exception, so to speak. For example, the sentence *Some student liked every course* (18b) is ambiguous between the following two readings.

$$(39) \quad \exists x. \text{student}(x) \wedge \forall y. \text{course}(y) \Rightarrow x \setminus \text{liked} \upharpoonright y$$

$$(40) \quad \forall y. \text{course}(y) \Rightarrow \exists x. \text{student}(x) \wedge x \setminus \text{liked} \upharpoonright y$$

In the *surface scope* reading (39), *some* takes scope over *every*. In the *inverse scope* reading (40), *every* takes scope over *some*. Given that evaluation takes place from left to right, the shift for *some student* is evaluated before the shift for *every course*. Our grammar thus predicts the surface scope reading but not the inverse scope reading. This prediction can be seen in the first few reductions of the (unique) derivation for (18b):

$$\begin{aligned}
(41) \quad & [([\text{some}] \upharpoonright [\text{student}]) \setminus [[\text{liked}] \upharpoonright [[\text{every}] \upharpoonright [\text{course}]]]] \\
& = [((\lambda r. \xi s. \exists x. r(x) \wedge s(x)) \upharpoonright \text{student}) \\
& \quad \setminus \text{liked} \upharpoonright ((\lambda r. \xi s. \forall y. r(y) \Rightarrow s(y)) \upharpoonright \text{course})] \\
& \triangleright [(\xi s. \exists x. \text{student}(x) \wedge s(x)) \\
& \quad \setminus \text{liked} \upharpoonright ((\lambda r. \xi s. \forall y. r(y) \Rightarrow s(y)) \upharpoonright \text{course})] \\
& \triangleright [\exists x. \text{student}(x) \wedge (\lambda v. [v \\
& \quad \setminus \text{liked} \upharpoonright ((\lambda r. \xi s. \forall y. r(y) \Rightarrow s(y)) \upharpoonright \text{course})])](x)
\end{aligned}$$

Regardless of what evaluation order we specify, as long as our rules for semantic translation remain deterministic, they will only generate one reading for the sentence. Hence our theory fails to predict the ambiguity of the sentence (18b).

To better account for the data, we need to introduce some sort of nondeterminism into our theory. There are two natural ways to proceed. First, we can allow arbitrary evaluation order, not just left-to-right. This change would render our term calculus nonconfluent, a result unwelcome for most programming language researchers but welcome for us in light of the ambiguous natural language sentence (18b). This route has been pursued with some success by Barker (2002) and de Groote (2001). However, there are empirical reasons to

maintain left-to-right evaluation, one of which appears in §6.

A second way to introduce nondeterminism is to maintain left-to-right evaluation but generalize shift and reset to a *hierarchy* of control operators (Barker 2000; Danvy and Filinski 1990; Shan and Barker 2003), leaving it unspecified at which level on the hierarchy each quantificational phrase shifts. Following Danvy and Filinski, we extend our logical metalanguage by superscripting every shift expression and pair of reset brackets with a nonnegative integer to indicate a level on the control hierarchy. Level 0 is the highest level (not the lowest). When a shift expression at level n is evaluated, it captures the current context of computation up to the closest dynamically enclosing reset at level n or higher (smaller). For example, whereas the expression

$$(42) \quad [3 \times [10 \times (\xi^5 f. 1 + f(2))]^5]^0$$

evaluates to 63 as in (30), the expression

$$(43) \quad [3 \times [10 \times (\xi^3 f. 1 + f(2))]^5]^0$$

evaluates to $1 + 3 \times 10 \times 2$, or 61. The superscripts can be thought of “strength levels” for shifts and resets.

Danvy and Filinski (1990) give a denotational semantics for multiple levels of delimited control using continuations of higher-order type. We can take advantage of that work in our quantificational denotations (35–36) by letting them shift at any level. The ambiguity of (18b) is then predicted as follows. Suppose that *some student* shifts at level m and *every course* shifts at level n .

$$(44) \quad \text{Some}^m \text{ student liked every}^n \text{ course.}$$

If $m \leq n$, the surface scope reading (39) results. If $m > n$, the inverse scope reading (40) results. In general, a quantifier that shifts at a higher level always scopes over another that shifts on a lower level, regardless of which one is evaluated first. This way, evaluation order does not determine scoping possibilities among quantifiers in a sentence unless two quantifiers happen to shift at the same level.

To summarize the discussion so far, whether we introduce nondeterministic evaluation order or a hierarchy of delimited control operators, we can account for the ambiguity of the sentence (18b), as well as more complicated cases of quantification in English and Mandarin (Shan 2003). For example, both the nondeterministic evaluation order approach and the control hierarchy approach predict correctly that the sentence below, with three quantifiers, is 5-way ambiguous.

$$(45) \quad \text{Every representative of a company saw most samples.}$$

Despite the fact that there are three quantifiers in this sentence and $3! = 6$, this sentence has only 5 readings. Because *a company* occurs within the restrictor of *every representative of a company*, it is incoherent for *every* to scope over *most* and *most* over *a*. The reason neither approach generates such a reading can be seen in the denotation of *every* in (35): “ $\forall x$ ” is located immediately above “ \Rightarrow ” in the abstract syntax, with no intervening control delimiter, so no control operator can insert any material (such as *most*-quantification over samples) in between.

There exist in the computational linguistics literature algorithms for computing the possible quantifier scopings of a sentence like (45) (Hobbs and Shieber 1987; followed by Lewin 1990; Moran 1988). Having related quantifier scoping to control operators, we gain a denotational understanding

of these algorithms that accords with our theoretical intuitions and empirical observations.

An extended logical metalanguage with an infinite hierarchy of control operators is shown in Figure 4. This system is more complex than the one in Figure 2 in two ways. First, instead of making a binary distinction between pure and impure expressions, we use a number to measure “how pure” each expression is. An expression is pure up to level n if it only incurs control effects at levels above n when evaluated. Pure expressions are the special case when $n = 0$. The purity level of an expression is reflected in its typing judgment: a judgment

$$(46) \quad \Gamma \vdash E : \alpha!n$$

states that the expression E is pure up to level n . Here $\alpha!n$ is a *computation type* with n levels: as defined in the figure, it consists of $2^{n+1} - 1$ value types that together specify how a computation that is pure up to level n affects answer types between levels 0 and $n - 1$. In the special case where $n = 1$, the computation type $\alpha!1$ is of the familiar form $\alpha!_1$.

In the previous system in Figure 2, directional functions are always impure (that is, pure up to level 1) while non-directional functions are always pure (that is, pure up to level 0). In the current system, both kinds of functions declare in their types up to what level their bodies are pure. For example, the determiners *every* and *some*, now allowed to shift at any level, both have not just the type

$$(47) \quad (\text{Thing} \rightarrow \text{Bool}) \dot{\rightarrow} \text{Thing}_{\text{Bool}_{\gamma!n}^{\delta!n}}$$

but also the type

$$(48) \quad (\text{Thing} \rightarrow \text{Bool}_{\gamma!n}^{\delta!n}) \dot{\rightarrow} \text{Thing}_{\text{Bool}_{\gamma!n}^{\delta!n}}$$

(but see the second technical complication below). As the argument type $\text{Thing} \rightarrow \text{Bool}_{\gamma!n}^{\delta!n}$ above shows, the first argument to these determiners, the restrictor, is non-directional yet can be impure (that is, pure up to level $n + 1$).

To traverse the control hierarchy, we add a new Reset rule, which makes an expression more pure, and a new Lift rule, which makes an expression less pure. (Consecutive nested resets like $[[E]^{n+1}]^n$ can be abbreviated to $[E]^n$ without loss of coherence.)

A second complication in this system, in contrast to Figure 2, is that we can no longer encode logical quantification using a higher-order constant like $\forall : (\text{Thing} \rightarrow \text{Bool}) \rightarrow \text{Bool}$, because such a constant requires its argument—the logical formula to be quantified—to be a pure function. This requirement is problematic because it is exactly the impurity of quantified logical formulas that underlies this account of quantifier scope ambiguity. On one hand, we want to quantify logical formulas that are impure; on the other hand, we want to rule out expressions like

$$(49) \quad \forall x. \xi f. x,$$

where the logical variable x “leaks” illicitly into the surrounding context. This issue is precisely the problem of classifying open and closed terms in staged programming (see Taha and Nielsen 2003 and references therein): the types Thing and Bool really represent not individuals or truth values but staged programs that compute individuals and truth values. For this paper, we adopt the simplistic solution of adjoining to these types a set of free logical variables for

Directions	$\Delta ::= / \mid \setminus$
Value types	$\alpha, \beta, \gamma, \delta ::= \mathbf{Thing}\{\bar{p}\} \mid \mathbf{Bool}\{\bar{p}\} \mid \alpha \rightarrow \beta!n \mid \alpha \xrightarrow{\Delta} \beta!n$
Computation types	$\alpha!0 ::= \alpha, \quad \alpha!(n+1) ::= \alpha_{\gamma!n}^{\delta!n}$
Antecedents	$\Gamma ::= x_1 : \alpha_1, \dots, x_n : \alpha_n$
Terms	$E, F ::= c \mid x \mid \lambda^n x. E \mid \lambda^{\Delta^n} x. E \mid FE \mid F / E \mid E \setminus F \mid [E]^n \mid \xi^n f. E$
Constants $c : \alpha$	

$\overline{p : \mathbf{Thing}\{\bar{p}\}}$	$\overline{\forall p : \mathbf{Bool}\{\bar{p}, \bar{q}\} \rightarrow \mathbf{Bool}\{\bar{q}\}}$	$\overline{\exists p : \mathbf{Bool}\{\bar{p}, \bar{q}\} \rightarrow \mathbf{Bool}\{\bar{q}\}}$
$\Rightarrow : \mathbf{Bool}\{\bar{p}\} \rightarrow \mathbf{Bool}\{\bar{q}\} \rightarrow \mathbf{Bool}\{\bar{p} \cup \bar{q}\}$	$\wedge : \mathbf{Bool}\{\bar{p}\} \rightarrow \mathbf{Bool}\{\bar{q}\} \rightarrow \mathbf{Bool}\{\bar{p} \cup \bar{q}\}$	
$\overline{\text{student} : \mathbf{Thing}\{\bar{p}\} \rightarrow \mathbf{Bool}\{\bar{p}\}}$	$\overline{\text{liked} : \mathbf{Thing}\{\bar{p}\} \xrightarrow{\Delta} (\mathbf{Thing}\{\bar{q}\} \xrightarrow{\Delta} \mathbf{Bool}\{\bar{p} \cup \bar{q}\}_{\delta!m}^{\gamma!n})} \dots$	
Expressions $\Gamma \vdash E : \alpha!n$		
$\frac{c : \alpha}{\Gamma \vdash c : \alpha} \text{Const}$	$\frac{}{\Gamma, x : \alpha \vdash x : \alpha} \text{Var}$	$\frac{\Gamma \vdash E : \alpha_{\alpha}^{\beta}}{\Gamma \vdash [E]^0 : \beta} \text{Reset}$
		$\frac{\Gamma \vdash E : \alpha_{\alpha_{\gamma!n}^{\beta!n}}^{\beta!n}}{\Gamma \vdash [E]^{n+1} : \beta!(n+1)} \text{Reset}$
$\frac{\Gamma, x : \alpha \vdash E : \beta!n}{\Gamma \vdash \lambda^{\Delta} x. E : \alpha_{\beta!n}^{\Delta}} (\Delta)I$	$\frac{\Gamma \vdash F : \alpha \rightarrow \beta!n \quad \Gamma \vdash E : \alpha}{\Gamma \vdash FE : \beta!n} \rightarrow E$	
$\frac{\Gamma \vdash E : \alpha}{\Gamma \vdash E : \alpha_{\gamma}^{\gamma}} \text{Lift}$	$\frac{\Gamma \vdash E : \alpha_{\gamma!n}^{\gamma!n}}{\Gamma \vdash E : \alpha_{\beta_{\gamma!n}^{\gamma!n}}^{\beta_{\gamma!n}^{\gamma!n}}} \text{Lift}$	$\frac{\Gamma, f : \alpha \rightarrow \gamma!n \vdash E : \delta!n}{\Gamma \vdash \xi^n f. E : \alpha_{\gamma!n}^{\delta!n}} \text{Shift}$
$\frac{\Gamma \vdash F : (\alpha \xrightarrow{\Delta} \beta_{\gamma!n}^{\gamma!n})_{\gamma!n}^{\gamma!n} \quad \Gamma \vdash E : \alpha_{\gamma!n}^{\gamma!n}}{\Gamma \vdash F / E : \beta_{\gamma!n}^{\gamma!n}} \xrightarrow{\Delta} E$	$\frac{\Gamma \vdash E : \alpha_{\gamma!n}^{\gamma!n} \quad \Gamma \vdash F : (\alpha \xrightarrow{\Delta} \beta_{\gamma!n}^{\gamma!n})_{\gamma!n}^{\gamma!n}}{\Gamma \vdash E \setminus F : \beta_{\gamma!n}^{\gamma!n}} \xrightarrow{\Delta} E$	
$\frac{\Gamma \vdash F : (\alpha \xrightarrow{\Delta} \beta) \quad \Gamma \vdash E : \alpha}{\Gamma \vdash F / E : \beta} \xrightarrow{\Delta} E$	$\frac{\Gamma \vdash E : \alpha \quad \Gamma \vdash F : (\alpha \xrightarrow{\Delta} \beta)}{\Gamma \vdash E \setminus F : \beta} \xrightarrow{\Delta} E$	

Figure 4: Extending the logical metalanguage to a hierarchy of control operators

tracking purposes, denoted p, q, \dots . Unfortunately, we also need to stipulate that these logical variables be freshly α -renamed (“created by gensym”) for each occurrence of a quantifier in a sentence.

As before, the CPS transform and reductions for this metalanguage are standard; the latter appears in Figure 5.

The present analysis is almost, but not quite, the direct-style analogue of Shan and Barker’s (2003) CPS analysis, even though both use a control hierarchy. Each level in Shan and Barker’s hierarchy is intuitively a staged computation produced at one level higher. More concretely, a computation type with n levels in that system has the shape

$$(50) \quad (\alpha_{\gamma!1}^{\gamma!0})_{\delta!1}^{\delta!0}$$

rather than the shape

$$(51) \quad \alpha_{\gamma!1}^{\gamma!0} \delta!1 \delta!2 \delta!3$$

The issue above of how to encode logical quantification over impure formulas receives a more satisfactory treatment in Shan and Barker’s system: no stipulation of α -renaming is necessary, because there is no analogue of (49) to prohibit. The relation between that system and staged programming with effects has yet to be explored.

6. POLARITY SENSITIVITY

Because the analysis so far focuses on the truth-conditional meaning of quantifiers, it equates the determiners a

and *some*—both are existential quantifiers with the type and denotation in (36). Furthermore, sentences like *Has anyone arrived?* suggest that the determiner *any* also means the same thing as a and *some*. To the contrary, though, the determiners a , *some*, and *any* are not always interchangeable in their existential usage. The sentences and readings in (52) show that they take scope differently relative to negation (in these cases the quantifier *no*).

- (52) a. No student liked some course. (unambiguous $\exists \neg$)
b. No student liked a course. (ambiguous $\neg \exists, \exists \neg$)
c. No student liked any course. (unambiguous $\neg \exists$)
d. Some student liked no course. (unambiguous $\exists \neg$)
e. A student liked no course. (ambiguous $\neg \exists, \exists \neg$)
f. *Any student liked no course. (unacceptable)

The determiner *any* is a *negative polarity item*: to a first approximation, it can occur only in *downward-entailing* contexts, such as under the scope of a *monotonically decreasing* quantifier (Ladusaw 1979). A quantifier q , of type $(\mathbf{Thing} \rightarrow \mathbf{Bool}) \rightarrow \mathbf{Bool}$, is *monotonically decreasing* just in case

$$(53) \quad \forall s_1. \forall s_2. (\forall x. s_2(x) \Rightarrow s_1(x)) \Rightarrow q(s_1) \Rightarrow q(s_2).$$

The quantificational noun phrases *no student* and *no course* are *monotonically decreasing* since, for instance, if no student liked any course in general, then no student liked any computer science course in particular.

Whereas *any* is a *negative polarity item*, *some* is a *positive polarity item*. Roughly speaking, *some* is allergic to down-

Values	$V ::= U \mid \lambda x. E \mid \lambda^{\Delta} x. E$
Unknowns	$U ::= c \mid UV \mid U \uparrow V \mid V \setminus U$
Contexts	$C\langle \rangle ::= \langle \rangle \mid (C\langle \rangle)E \mid C\langle \rangle \uparrow E \mid C\langle \rangle \setminus F \mid V(C\langle \rangle) \mid V \uparrow C\langle \rangle \mid V \setminus C\langle \rangle$
Metacontexts at level n	$D^n\langle \rangle ::= \langle \rangle \mid D^{n+1}\langle D^{n+2}\langle \dots \langle C\langle [D^n\langle \rangle]^n \rangle \dots \rangle \rangle$
Computations $E \triangleright E'$	

$$\begin{aligned}
D^0\langle D^1\langle \dots \langle C\langle (\lambda x. E)V \rangle \dots \rangle \rangle &\triangleright D^0\langle D^1\langle \dots \langle C\langle E\{x \mapsto V\} \rangle \dots \rangle \rangle \\
D^0\langle D^1\langle \dots \langle C\langle (\lambda x. E) \uparrow V \rangle \dots \rangle \rangle &\triangleright D^0\langle D^1\langle \dots \langle C\langle E\{x \mapsto V\} \rangle \dots \rangle \rangle \\
D^0\langle D^1\langle \dots \langle C\langle V \setminus (\lambda x. E) \rangle \dots \rangle \rangle &\triangleright D^0\langle D^1\langle \dots \langle C\langle E\{x \mapsto V\} \rangle \dots \rangle \rangle \\
D^0\langle D^1\langle \dots \langle C\langle [V] \rangle \dots \rangle \rangle &\triangleright D^0\langle D^1\langle \dots \langle C\langle V \rangle \dots \rangle \rangle \\
D^0\langle D^1\langle \dots \langle D^n\langle D^{n+1}\langle \dots \langle C\langle \xi^n f. E \rangle \dots \rangle \rangle \dots \rangle \rangle &\triangleright D^0\langle D^1\langle \dots \langle D^n\langle E\{f \mapsto \lambda x. [D^{n+1}\langle \dots \langle C\langle x \rangle \dots \rangle]^n \rangle \dots \rangle \rangle \dots \rangle \rangle
\end{aligned}$$

Figure 5: Reductions for the extended logical metalanguage

ward-entailing contexts (especially those with an overtly negative word like *no*). These generalizations regarding polarity items cover the data in (52a–e): in principle, goes the theory, all these sentences are ambiguous between two scopings, but the polarity sensitivity of *some* and *any* rule out one scoping each in (52a), (52c), (52d), and (52f). These four sentences are thus predicted to be unambiguous, but it remains unclear why (52f) is downright unacceptable.

In the type-theoretic tradition of linguistics, polarity sensitivity is typically implemented by splitting the answer type **Bool** into several types, each a different functor applied to **Bool**, that are related by subtyping (Bernardi 2002; Bernardi and Moot 2001; Fry 1999). For instance, to differentiate the determiners in (52) from each other in our formalism, we can add the types **BoolPos** and **BoolNeg** alongside **Bool**, such that both are supertypes of **Bool** (but not of each other).

$$(54) \quad \overline{\text{Bool} \leq \text{BoolPos}} \quad \overline{\text{Bool} \leq \text{BoolNeg}}$$

We also extend the subtyping relation between (value and computation) types with the usual closure rules, and allow implicit coercion from a subtype to a supertype.

$$(55) \quad \frac{\alpha' \leq \alpha \quad \beta!n \leq \beta'!n}{\alpha \leq \alpha \quad \alpha \langle \Delta \rangle \beta!n \leq \alpha' \langle \Delta \rangle \beta'!n}$$

$$(56) \quad \frac{\alpha' \leq \alpha \quad \gamma'!n \leq \gamma!n \quad \delta!n \leq \delta'!n}{\alpha_{\gamma'!n}^{\delta!n} \leq \alpha_{\gamma!n}^{\delta'!n}}$$

$$(57) \quad \frac{\Gamma \vdash E : \alpha!n \quad \alpha!n \leq \beta!n}{\Gamma \vdash E : \beta!n} \text{Sub}$$

We then add a side condition to **Reset**, requiring that the produced answer type be **Bool** or **BoolPos**, not **BoolNeg**.

$$(58) \quad \frac{\Gamma \vdash E : \alpha_{\alpha}^{\beta}}{\Gamma \vdash [E] : \beta} \text{Reset} \quad \text{where } \beta \leq \text{BoolPos}$$

$$(59) \quad \frac{\Gamma \vdash E : \alpha_{\gamma!n}^{\beta!(n+1)}}{\Gamma \vdash [E] : \beta!(n+1)} \text{Reset} \quad \text{where } \beta \leq \text{BoolPos}$$

Finally, we refine the types of our determiners from (47) to

$$(60) \quad \llbracket \text{a} \rrbracket : (\text{Thing} \rightarrow \text{Bool}) \dot{\rightarrow} \text{Thing}_{\text{BoolNeg}_{\gamma!n}^{\delta!n}, \text{Bool}_{\gamma!n}^{\delta!n}}$$

$$(61) \quad \llbracket \text{some} \rrbracket : (\text{Thing} \rightarrow \text{Bool}) \dot{\rightarrow} \text{Thing}_{\text{BoolPos}_{\gamma!n}^{\delta!n}, \text{BoolPos}_{\gamma!n}^{\delta!n}}$$

$$(62) \quad \llbracket \text{a} \rrbracket : (\text{Thing} \rightarrow \text{Bool}) \dot{\rightarrow} \text{Thing}_{\text{Bool}_{\gamma!n}^{\delta!n}, \text{Bool}_{\gamma!n}^{\delta!n}}$$

$$(63) \quad \llbracket \text{any} \rrbracket : (\text{Thing} \rightarrow \text{Bool}) \dot{\rightarrow} \text{Thing}_{\text{BoolNeg}_{\gamma!n}^{\delta!n}, \text{BoolNeg}_{\gamma!n}^{\delta!n}}$$

The chain of answer-type transitions from one quantificational expression to the next acts as a finite-state automaton, shown in Figure 6. The states of the automaton are the three supertypes of **Bool**; the ϵ -transitions are the two subtyping relations in (54); and the non- ϵ transitions are the determiners in (60–63).




Figure 6: An automaton of answer-type transitions

This three-state machine enforces polarity constraints as follows. Any valid derivation for a sentence assigns each of its quantifiers to shift at a certain level in the control hierarchy. For each level, the quantifiers at that level—in the order in which they are evaluated—must form

- either a path from **BoolPos** to **Bool** in the state machine, in other words a string of determiners matching the regular expression “some* (a|no any*)*”; or
- or a path from **Bool** to **Bool** in the state machine, in other words a string of determiners matching the regular expression “(a|no any*)*”.

Furthermore, the **BoolPos**-to-**Bool** levels in the hierarchy must all be higher than the **Bool**-to-**Bool** levels. Every assignment of quantifiers to levels that satisfies these conditions gives a reading for the sentence, in which quantifiers at higher levels scope wider, and, among quantifiers at the same level, ones evaluated earlier scope wider.

Consider now the two alternative ways to characterize scope ambiguity suggested in §5. The first approach is to allow arbitrary evaluation order (and use a degenerate control hierarchy of one level only). If we take this route, we can account for all of the acceptability and ambiguity

judgments in (52a–e), but we cannot distinguish the acceptable sentence (52c) from the unacceptable (52f). In other words, it would be a mystery how the acceptability of a sentence hinges on the linear order in which the quantifiers *no* and *any* appear. This mystery has been noted by Ladusaw (1979; §9.2) and Fry (1999; §8.2) as a defect in current accounts of polarity sensitivity.

The second approach, using a control hierarchy with multiple levels, fares better by comparison.³ We can stick to left-to-right evaluation, under which—as desired—an *any* must be preceded by a *no* that scopes over it with no intervening *a* or *some*. Indeed, the variations in ambiguity and acceptability among sentences in (52) are completely captured. For intuition, we can imagine that the hearer of a sentence must first process the trigger for a downward-entailing context, like *no*, before it makes sense to process a negative polarity item, like *any*.⁴ Intuition aside, the programming-language notion of evaluation order provides the syntactic hacker of formal types with a new tool with which to capture observed regularities in natural language.

7. LINGUISTIC SIDE EFFECTS

This paper outlines how quantification and polarity sensitivity in natural language can be modeled using delimited continuations. These two examples support my claim that the formal theory and computational intuition we have for continuations can help us construct, understand, and maintain linguistic theories. To be sure, this work is far from the first time insights from programming languages are applied to natural language:

- It has long been noted that the intensional logic in which Montague grammar is couched can be understood computationally (Hobbs and Rosenschein 1978; Hung and Zucker 1991).
- *Dynamic semantics* (Groenendijk and Stokhof 1991), which relates anaphora and discourse in natural languages to nondeterminism and mutable state in programming languages (van Eijck 1998), has been applied to a variety of natural language phenomena, such as verb-phrase ellipsis (van Eijck and Francez 1995; Gardent 1991; Hardt 1999).

However, the link between natural language and continuations has only recently been made explicit, and this paper’s use of control operators for a direct-style analysis is novel.

The analyses presented here are part of a larger project, that of relating computational side effects to *linguistic side effects*. The term “computational side effect” here covers all programming language features where either it is unclear what a denotational semantics should look like, or the “obvious” denotational semantics (such as making each arithmetic expression denote a number) turns out to break referential transparency. A computational side effect of the first

³Although this paper uses Danvy and Filinski’s control hierarchy, polarity sensitivity can be expressed equally well in Shan and Barker’s system.

⁴The syntactic distinction among the types `Bool`, `BoolPos`, and `BoolNeg` may even be semantically interpretable via the formulas-as-types correspondence, but the potential for such a connection has only been briefly explored (Bernardi and Nilsen 2001) and we do not examine it here. In this connection, Krifka (1995) and others have proposed on pragmatic grounds that determiners like *any* are negative polarity items because they indicate extreme points on a scale.

kind is jumps to labels; one of the second kind is mutable state. By analogy, I use the term “linguistic side effects” to refer to aspects of natural language where either it is unclear what a denotational semantics should look like, or the “obvious” denotational semantics (such as making each clause denote whether it is true) turns out to break referential transparency. Besides quantification and polarity sensitivity, some examples are:

- (64) a. Bob *thinks* Alice likes CS187. (Intensionality)
 b. *A man* walks. *He* whistles. (Variable binding)
 c. *Which* star did Alice see? (Interrogatives)
 d. Alice *only* saw VENUS. (Focus)
 e. *The king of France* whistles. (Presuppositions)

To study linguistic side effects, I propose to draw an analogy between them and computational side effects. Just as computer scientists want to express all computational side effects in a uniform and modular framework and study how control interacts with mutable state (Felleisen and Hieb 1992), linguists want to investigate properties common to all linguistic side effects and study how quantification interacts with variable binding. Furthermore, just as computer scientists want to relate operational notions like evaluation order and parameter passing to denotational models like continuations and monads, linguists want to relate the dynamics of information in language processing to the static definition of a language as a generative device. Whether this analogy yields a linguistic theory that is empirically adequate is an open scientific question that I find attractive to pursue.

8. ACKNOWLEDGMENTS

Thanks to Stuart Shieber, Chris Barker, Raffaella Bernardi, Barbara Grosz, Pauline Jacobson, Aravind Joshi, William Ladusaw, Fernando Pereira, Avi Pfeffer, Chris Potts, Norman Ramsey, Dylan Thurston, Yoad Winter, and anonymous referees. This work is supported by the United States National Science Foundation Grants IRI-9712068 and BCS-0236592.

9. REFERENCES

- Barker, Chris. 2000. Notes on higher-order continuations. Manuscript, University of California, San Diego.
- . 2002. Continuations and the nature of quantification. *Natural Language Semantics* 10(3):211–242.
- Barwise, Jon, and Robin Cooper. 1981. Generalized quantifiers and natural language. *Linguistics and Philosophy* 4: 159–219.
- Bernardi, Raffaella. 2002. Reasoning with polarity in categorical type logic. Ph.D. thesis, Utrecht Institute of Linguistics (OTS), Utrecht University.
- Bernardi, Raffaella, and Richard Moot. 2001. Generalized quantifiers in declarative and interrogative sentences. *Journal of Language and Computation* 1(3):1–19.
- Bernardi, Raffaella, and Øystein Nilsen. 2001. Polarity items in type logical grammar: Connection with DMG. Slides for talk at Learning Logic and Grammar workshop, Amsterdam.
- Carpenter, Bob. 1997. *Type-logical semantics*. Cambridge: MIT Press.

- Danvy, Olivier, and Andrzej Filinski. 1989. A functional abstraction of typed contexts. Tech. Rep. 89/12, DIKU, University of Copenhagen, Denmark. <http://www.daimi.au.dk/~danvy/Papers/fatc.ps.gz>.
- . 1990. Abstracting control. In *Proceedings of the 1990 ACM conference on Lisp and functional programming*, 151–160. New York: ACM Press.
- Danvy, Olivier, and John Hatcliff. 1992. CPS-transformation after strictness analysis. *ACM Letters on Programming Languages and Systems* 1(3):195–212.
- . 1994. On the transformation between direct and continuation semantics. In *Mathematical foundations of programming semantics: 9th international conference (1993)*, ed. Stephen D. Brookes, Michael G. Main, Austin Melton, Michael W. Mislove, and David A. Schmidt, 627–648. Lecture Notes in Computer Science 802, Berlin: Springer-Verlag.
- van Eijck, Jan. 1998. Programming with dynamic predicate logic. Report INS-R9810, Centrum voor Wiskunde en Informatica, Amsterdam. Also as Research Report CT-1998-06, Institute for Logic, Language and Computation, Universiteit van Amsterdam.
- van Eijck, Jan, and Nissim Francez. 1995. Verb-phrase ellipsis in dynamic semantics. In *Applied logic: How, what, and why: Logical approaches to natural language*, ed. László Pólos and Michael Masuch. Dordrecht: Kluwer.
- Felleisen, Matthias. 1988. The theory and practice of first-class prompts. In *POPL '88: Conference record of the annual ACM symposium on principles of programming languages*, 180–190. New York: ACM Press.
- Felleisen, Matthias, and Robert Hieb. 1992. The revised report on the syntactic theories of sequential control and state. *Theoretical Computer Science* 103(2):235–271.
- Fry, John. 1999. Proof nets and negative polarity licensing. In *Semantics and syntax in lexical functional grammar: The resource logic approach*, ed. Mary Dalrymple, chap. 3, 91–116. Cambridge: MIT Press.
- Gardent, Claire. 1991. Dynamic semantics and VP-ellipsis. In *Logics in AI: European workshop JELIA '90*, ed. Jan van Eijck, 251–266. Lecture Notes in Artificial Intelligence 478, Berlin: Springer-Verlag.
- Groenendijk, Jeroen, and Martin Stokhof. 1991. Dynamic predicate logic. *Linguistics and Philosophy* 14(1):39–100.
- de Groote, Philippe. 2001. Type raising, continuations, and classical logic. In *Proceedings of the 13th Amsterdam Colloquium*, ed. Robert van Rooy and Martin Stokhof, 97–101. Institute for Logic, Language and Computation, Universiteit van Amsterdam.
- Hardt, Daniel. 1999. Dynamic interpretation of verb phrase ellipsis. *Linguistics and Philosophy* 22(2):185–219.
- Hendriks, Herman. 1993. Studied flexibility: Categories and types in syntax and semantics. Ph.D. thesis, Institute for Logic, Language and Computation, Universiteit van Amsterdam.
- Hobbs, Jerry R., and Stanley J. Rosenschein. 1978. Making computational sense of Montague's intensional logic. *Artificial Intelligence* 9:287–306.
- Hobbs, Jerry R., and Stuart M. Shieber. 1987. An algorithm for generating quantifier scopings. *Computational Linguistics* 13(1–2):47–63.
- Hung, Hing-Kai, and Jeffery I. Zucker. 1991. Semantics of pointers, referencing and dereferencing with intensional logic. In *LICS '91: Proceedings of the 6th symposium on logic in computer science*, 127–136. Washington, DC: IEEE Computer Society Press.
- Kelsey, Richard, William Clinger, Jonathan Rees, et al. 1998. Revised⁵ report on the algorithmic language Scheme. *Higher-Order and Symbolic Computation* 11(1):7–105. Also as *ACM SIGPLAN Notices* 33(9):26–76.
- Krifka, Manfred. 1995. The semantics and pragmatics of polarity items. *Linguistic Analysis* 25:209–257.
- Ladusaw, William A. 1979. Polarity sensitivity as inherent scope relations. Ph.D. thesis, Department of Linguistics, University of Massachusetts. Reprinted by New York: Garland, 1980.
- Lewin, Ian. 1990. A quantifier scoping algorithm without a free variable constraint. In *COLING '90: Proceedings of the 13th international conference on computational linguistics*, vol. 3, 190–194.
- May, Robert. 1985. *Logical form: Its structure and derivation*. Cambridge: MIT Press.
- Montague, Richard. 1974. The proper treatment of quantification in ordinary English. In *Formal philosophy: Selected papers of Richard Montague*, ed. Richmond Thomason, 247–270. New Haven: Yale University Press.
- Moran, Douglas B. 1988. Quantifier scoping in the SRI core language engine. In *Proceedings of the 26th annual meeting of the Association for Computational Linguistics*, 33–40. Somerset, NJ: Association for Computational Linguistics.
- Nielsen, Lasse R. 2001. A selective CPS transformation. In *Proceedings of MFPS 2001: 17th conference on the mathematical foundations of programming semantics*, ed. Stephen Brooks and Michael Mislove. Electronic Notes in Theoretical Computer Science 45, Amsterdam: Elsevier Science.
- Shan, Chung-chieh. 2003. Quantifier strengths predict scopal possibilities of Mandarin Chinese *wh*-indefinites. Draft manuscript, Harvard University; <http://www.eecs.harvard.edu/~ccshan/mandarin/>.
- Shan, Chung-chieh, and Chris Barker. 2003. Explaining crossover and superiority as left-to-right evaluation. Draft manuscript, Harvard University and University of California, San Diego; <http://semanticsarchive.net/Archive/TBjZDQ3Z/>.
- Taha, Walid, and Michael Florentin Nielsen. 2003. Environment classifiers. In *POPL '03: Conference record of the annual ACM symposium on principles of programming languages*, 26–37. New York: ACM Press.
- Thielecke, Hayo. 2003. From control effects to typed continuation passing. In *POPL '03: Conference record of the annual ACM symposium on principles of programming languages*, 139–149. New York: ACM Press.