Inverse scope as metalinguistic quotation in operational semantics

Chung-chieh Shan

Rutgers University
`ccshan@rutgers.edu`

Abstract. We model semantic interpretation *operationally*: constituents interact as their combination in discourse evolves from state to state. The states are recursive data structures and evolve from step to step by context-sensitive rewriting. These notions of *context* and *order* let us explain inverse-scope quantifiers and their polarity sensitivity as metalinguistic quotation of the wider scope.

1 Introduction

An utterance is both an action and a product [1]. On one hand, when we use an utterance, it acts on the world: its parts affect us and our conversation. For example, a referring expression reminds us of its referent and adjusts the discourse context so that a pronoun later may refer to it. On the other hand, linguistics describes the syntax and semantics of a sentence as a mathematical object: a product, perhaps a tree or a function. But trees and functions do not remind or adjust; at their most active, they merely contain or map. How then does the meaning of an utterance determine its effects on the discourse and its participants? In particular, how do utterances affect their *context*, and in the *order* that they do?

We approach this "mind-body problem" the way many programming-language semanticists approach an analogous issue. On one hand, a program such as

$$(1) \qquad\qquad 2 \to n;\ n \cdot 3$$

expresses a sequence of actions: store the number 2 in the memory location n, then retrieve the contents of n and multiply it by 3. On the other hand, computer science describes the syntax and semantics of the program as a tree or a function. To view the same program as both a dynamic action and a static product, programming-language theory uses *operational semantics* alongside denotational semantics. A denotational semantics assigns a denotation to each expression—compositionally, by specifying the denotation of a complex expression in terms of the denotations of its parts. In contrast, an operational semantics defines a *transition relation* on states—for each state, what it can become after one step of computation [2]. Technical conditions relate denotational and operational semantics. For example, *adequacy* requires that two utterances with the same denotation must have the same operational outcome.

To analyze the example (1), we could define a state as an ordered pair of an integer n and a current program, and the transition relation $\leadsto$ as a set that includes the three transitions

$$(2) \qquad \langle 0,\ (2 \to n;\ n \cdot 3) \rangle \leadsto \langle 2,\ n \cdot 3 \rangle \leadsto \langle 2,\ 2 \cdot 3 \rangle \leadsto \langle 2,\ 6 \rangle.$$

These transitions together say that the program (1) starting with the memory location n initialized to 0 yields the final result 6 with the contents of n changed to 2. Unlike in denotational semantics, we specify the outcome of a program without recursively specifying what each part of the program denotes. Rather, we specify rewriting rules such as

$$(3) \qquad \langle x,\ (y \to n;\ P) \rangle \leadsto \langle y,\ P \rangle \quad \text{for any numbers } x, y \text{ and any program } P$$

to be elements of the transition relation, so that the first computation step in (2) goes through. In other words, we include

$$(4) \quad \big\{ \, \big\langle \langle x,\ (y \to n;\ P) \rangle,\ \langle y,\ P \rangle \big\rangle \mid x \text{ and } y \text{ are two numbers; } P \text{ is a program} \, \big\}$$

in the transition relation as a subset.

Following a longstanding analogy between natural and formal languages, we model natural-language semantics operationally. In Section 2, we review a standard operational semantics for a λ-calculus with *delimited control*. In Section 3, we use this operational semantics to explain quantification and polarity sensitivity. We focus on these applications because they go beyond phenomena such as anaphora and presupposition, where dynamic semantics has long been fruitful. In particular, metalinguistic quotation (or *multistage programming*) extends our account to inverse scope. We conclude in Section 4.

2 Context and order

We use the λ-calculus to introduce the notions of *context* and *order*, then let an expression interact actively with its context. These computational notions recur in various linguistic guises, especially in analyses of quantification, as Section 3 makes concrete.

2.1 Expressions and contexts

The set of λ-calculus expressions is recursively defined by the following grammar. In words, an expression is either a variable x, an abstraction $\lambda x.\,E$, or an application EE.

$$(5) \qquad E \to x \qquad\qquad E \to \lambda x.\,E \qquad\qquad E \to EE$$

The derivation below shows that $\lambda x.\,xx$ is an expression, even though no function x can apply to itself in standard set theory.

$$(6) \qquad E \to \lambda x.\,E \to \lambda x.\,EE \to \lambda x.\,Ex \to \lambda x.\,xx$$

In programming practice, constants such as 2 and multiplication · are also expressions. Not shown in the grammar is the convention that we equate expressions that differ only by variable renaming, so $\lambda x. \lambda x. x = \lambda x. \lambda y. y = \lambda y. \lambda x. x \neq \lambda x. \lambda y. x$.

We specify certain expressions to be *values*, namely those generated by the following grammar. In words, a value is any variable or abstraction, but not an application.

$$(7) \qquad\qquad V \to x \qquad\qquad\qquad V \to \lambda x. E$$

For example, the identity expression $\lambda x. x$ (henceforth I) is a value. In practice, constants such as 2 and multiplication · are also values. Intuitively, a value is an expression that is done computing, so the expression $2 \cdot 3$ is not a value.

An operational semantics is a relation on states. (More precisely, we consider *small-step* operational semantics.) Our states are just closed expressions (that is, expressions with no unbound variables), and the transition relation is in fact a partial function that maps each closed expression to what it becomes after one step of computation. We define the transition relation as follows. A *context* is an expression with a gap, which we write as [ ]. For example, $\lambda x. x(y[\,])$ is a context. The set of *evaluation contexts* $C[\,]$ is defined by the grammar below, in the style of [3]. The notation $C[\dots]$ means to replace the gap [ ] in the context $C[\,]$ by an expression or another context.

$$(8) \qquad C[\,] \to [\,] \qquad\qquad C[\,] \to C\big[[\,]E\big] \qquad\qquad C[\,] \to C\big[V[\,]\big]$$

The first production above says that the null context [ ]—the context with nothing but a gap—is an evaluation context. The second production says that, whenever $C[\,]$ is an evaluation context, replacing its gap by $[\,]E$ (that is, the application of a gap to any expression E) gives another evaluation context. The third production says that, whenever $C[\,]$ is an evaluation context, replacing its gap by $V[\,]$ (that is, the application of any value V to a gap) gives another evaluation context. For example, the derivation below shows that $I\big([\,](II)\big)$ is an evaluation context. (Recall from above that I stands for $\lambda x. x$.)

$$(9) \qquad C[\,] \to C\big[[\,](II)\big] \to C\big[I([\,](II))\big] \to I\big([\,](II)\big)$$

We now define the transition relation $\rightsquigarrow$ to be the set of expression-pairs

(10) $\{\,\langle C[(\lambda x. E)V], C[E']\rangle \mid C[\,]$ is an evaluation context; E' substitutes the value V for the variable x in the expression $E\,\}$,

or for short,

$$(11) \qquad\qquad C[(\lambda x. E)V] \rightsquigarrow C[E\,\{x \mapsto V\}].$$

In words, a transition is a λ-conversion in an evaluation context where the argument is a value expression. Letting $C[\,] = I\big([\,](II)\big)$, $V = I$, and $E = x$ gives

$$(12) \qquad\qquad I\big((II)(II)\big) \rightsquigarrow I\big(I(II)\big).$$

In fact, this is the only step that a computation starting at $I\big((II)(II)\big)$ can take. It takes three more steps to reach a value, namely I:

$$(13) \qquad\qquad I\big(I(II)\big) \rightsquigarrow I(II) \rightsquigarrow II \rightsquigarrow I.$$

In practice, λ-conversions are joined by other transitions such as $1 + 2 \cdot 3 \rightsquigarrow 1 + 6$.

This operational semantics is not the only reasonable one for the λ-calculus. Instead of or in addition to the transition (12), $I\big((II)(II)\big)$ could transition to $(II)(II)$ or $I\big((II)I\big)$. Different operational semantics regulate the *order* in which to run parts of a program differently, by constraining the set of evaluation contexts and the rewriting that takes place inside. Our transition relation in (10)–(11) implements a *call-by-value, left-to-right* evaluation order: it only performs λ-conversion when the argument is a value (the V in (10)–(11) rules out a transition to $(II)(II)$) and everything to the left is also a value (the V in (8) rules out a transition to $I\big((II)I\big)$).

2.2 Delimited control

The simple expression $I\big((II)(II)\big)$ above is not sensitive to the evaluation order: pretty much any reasonable order will bring it to the final value I after a few transitions. However, as we add functionality to the λ-calculus as a programming language, different operational semantics result in different program outcomes. A particularly powerful and linguistically relevant addition is *delimited control* [3, 4], which lets an expression manipulate its context. To illustrate, we add 's delimited-control operators *shift* and *reset* [5, 6, 7] to the λ-calculus.

Before specifying shift and reset formally, let us first examine some example transition sequences among arithmetic expressions. Reset by itself does not perform any computation, so the programs $1 + 2 \cdot 3$ and $1 + \text{reset}(2 \cdot 3)$ yield the same final value:

$$(14) \qquad\qquad 1 + 2 \cdot 3 \rightsquigarrow 1 + 6 \rightsquigarrow 7,$$

$$(15) \qquad\qquad 1 + \text{reset}(2 \cdot 3) \rightsquigarrow 1 + \text{reset}\, 6 \rightsquigarrow 1 + 6 \rightsquigarrow 7.$$

Shift means to *remove* the surrounding context—up to the nearest enclosing reset—into a variable. This functionality lets an expression manipulate its context. For example, the variable f below receives the context that multiplies by 2, so the program computes $1 + 3 \cdot 2 \cdot 2 \cdot 5$.

$$
\begin{aligned}
(16) \quad & 1 + \text{reset}\big(2 \cdot \text{shift}\, f.\, (3 \cdot f(f(5)))\big) \\
& \rightsquigarrow 1 + \text{reset}\big(3 \cdot (\lambda x.\, \text{reset}(2 \cdot x))((\lambda x.\, \text{reset}(2 \cdot x))(5))\big) \\
& \rightsquigarrow 1 + \text{reset}\big(3 \cdot (\lambda x.\, \text{reset}(2 \cdot x))(\text{reset}(2 \cdot 5))\big) \\
& \rightsquigarrow 1 + \text{reset}\big(3 \cdot (\lambda x.\, \text{reset}(2 \cdot x))(\text{reset}\, 10)\big) \\
& \rightsquigarrow 1 + \text{reset}\big(3 \cdot (\lambda x.\, \text{reset}(2 \cdot x))10\big) \\
& \rightsquigarrow 1 + \text{reset}\big(3 \cdot \text{reset}(2 \cdot 10)\big) \rightsquigarrow 1 + \text{reset}(3 \cdot \text{reset}\, 20) \\
& \rightsquigarrow 1 + \text{reset}(3 \cdot 20) \rightsquigarrow 1 + \text{reset}\, 60 \rightsquigarrow 1 + 60 \rightsquigarrow 61
\end{aligned}
$$

To take another example, the shift expression below does not use the variable f and so discards its surrounding context and supplies the reset with the result 4 right away.

(17) $\qquad 1 + \text{reset}\,(2 \cdot 3 \cdot (\text{shift}\,f.\,4) \cdot 5) \rightsquigarrow 1 + \text{reset}\,4 \rightsquigarrow 1 + 4 \rightsquigarrow 5$

We call reset a *control delimiter* because it delimits how much context an enclosed shift expression manipulates.

For concreteness, the rest of this subsection formalizes the delimited-control operators shift and reset; it can be skipped if the examples above suffice. We add two productions for expressions.

(18) $\qquad\qquad\qquad E \rightarrow \text{reset}\,E \qquad\qquad\qquad E \rightarrow \text{shift}\,f.\,E$

The expression "shift $f.\,E$" binds the variable f in the body E, so for instance the expressions "shift $f.\,f$" and "shift $x.\,x$" are equal because they differ only by variable renaming.

We add one production for evaluation contexts.

(19) $\qquad\qquad\qquad\qquad\qquad C[\,] \rightarrow C\big[\text{reset}[\,]\big]$

We call $D[\,]$ a *subcontext* if it is an evaluation context built without this new production. Finally, we add two new kinds of transitions to our transition relation, one for reset and one for shift:

(20) $\{\langle C\big[\text{reset}\,V\big],\,C\big[V\big]\rangle \mid C[\,]$ is an evaluation context and V is a value $\}$,

(21) $\{\langle C\big[\text{reset}\,D[\text{shift}\,f.\,E]\big],\,C\big[\text{reset}\,E'\big]\rangle$

$\qquad\qquad\qquad\qquad \mid C[\,]$ is an evaluation context; $D[\,]$ is a subcontext;
$\qquad\qquad\qquad\qquad E'$ substitutes $\lambda x.\,\text{reset}\,D[x]$ for the variable f
$\qquad\qquad\qquad\qquad$ in the expression E, where x is a fresh variable $\}$,

or for short,

(22) $\qquad\qquad\qquad C\big[\text{reset}\,V\big] \rightsquigarrow C\big[V\big],$

(23) $\qquad\qquad C\big[\text{reset}\,D[\text{shift}\,f.\,E]\big] \rightsquigarrow C\big[\text{reset}\,E\,\{f \mapsto \lambda x.\,\text{reset}\,D[x]\}\big].$

3 Linguistic applications

Linguistic theory traditionally views syntax, semantics, and pragmatics as a pipeline, in which semantics maps expressions to denotations. We envision a semantic theory for natural language that is operational in the sense that it specifies transitions among representations rather than mappings to denotations. That is, whereas denotational semantics interprets a syntactic constituent (such as a verb phrase) by mapping it to a separate semantic domain (such as of functions), operational semantics interprets a constituent by rewriting it to other constituents. The rewriting makes sense insofar as the constituents represent real objects. Thus we may specify a fragment of operational semantics as a set of states, a transition relation over the states, and an ideally trivial translation from utterances to states.

3.1 Quantification

We now use delimited control to model quantification. We assume that the sentence

(24) Somebody saw everybody

translates to the program (state)

$$(25) \qquad \mathrm{reset}\big((\overbrace{\mathrm{shift}\,f.\,\exists x.\,fx}^{somebody})^{<}(\overbrace{\mathrm{saw}}^{saw}\,{}^{>}(\overbrace{\mathrm{shift}\,g.\,\forall y.\,gy}^{everybody}))\big),$$

where $^{<}$ and $^{>}$ indicate backward and forward function application. The occurrences of shift in this translation arise from the lexical entries for the quantifiers *somebody* and *everybody*, whereas the occurrence of reset is freely available at every clause boundary. This program then makes the following transitions to yield the surface-scope reading of (25).

$$\rightsquigarrow \mathrm{reset}\big(\exists x.\,(\lambda x.\,\mathrm{reset}(x^{<}(\mathrm{saw}^{>}\,\mathrm{shift}\,g.\,\forall y.\,gy)))x\big)$$
$$\rightsquigarrow \mathrm{reset}\big(\exists x.\,\mathrm{reset}(x^{<}(\mathrm{saw}^{>}\,\mathrm{shift}\,g.\,\forall y.\,gy))\big)$$
$$\rightsquigarrow \mathrm{reset}\big(\exists x.\,\mathrm{reset}(\forall y.\,(\lambda y.\,\mathrm{reset}(x^{<}(\mathrm{saw}^{>}y)))y)\big)$$
$$\rightsquigarrow \mathrm{reset}\big(\exists x.\,\mathrm{reset}(\forall y.\,\mathrm{reset}(x^{<}(\mathrm{saw}^{>}y)))\big)$$
$$\rightsquigarrow \mathrm{reset}\big(\exists x.\,\mathrm{reset}(\forall y.\,x^{<}(\mathrm{saw}^{>}y))\big)$$
$$\rightsquigarrow \mathrm{reset}\big(\exists x.\,\forall y.\,x^{<}(\mathrm{saw}^{>}y)\big) \rightsquigarrow \exists x.\,\forall y.\,x^{<}(\mathrm{saw}^{>}y).$$

The last three transitions assume that the final formula and its subformulas are values.

This example illustrates that the context of a shift operator is the scope of a quantifier, and the order in which expressions are evaluated is that in which they take scope. This analogy is appealing because it extends to other apparently noncompositional phenomena [8–11] and tantalizing because it links meaning to processing.

3.2 Polarity sensitivity

We now extend the model above to capture basic polarity sensitivity: a polarity item such as *anybody* needs to take scope (right) under a polarity licensor such as existential *nobody*.

(26) Nobody saw anybody.

(27) Nobody saw everybody.

(28)*Somebody saw anybody.

Following [12, 13], we treat a polarity license ℓ as a dynamic resource [14] that is produced by *nobody*, required by *anybody*, and disposed of implicitly. To this end, we translate *nobody* to

$$(29) \qquad\qquad\qquad \mathrm{shift}\,f.\,\neg\exists x.\,fx\ell,$$

translate *anybody* to

(30) $$\text{shift } g.\, \lambda\ell.\, \exists y.\, gy\ell,$$

and add license-disposal transitions of the form

(31) $$C[V\ell] \leadsto C[V].$$

In (30), $\lambda\ell$ denotes a function that must take the constant ℓ as argument.

The sentence (26) translates and computes as follows. The penultimate transition disposes of the used license using (31).

(32) $$\text{reset}\big((\overbrace{\text{shift } f.\,\neg\exists x.\, fx\ell}^{nobody})^{<}(\overbrace{\text{saw}}^{saw}\,{}^{>}(\overbrace{\text{shift } g.\,\lambda\ell.\,\exists y.\, gy\ell}^{anybody})))\big)$$

$$\leadsto \text{reset}\big(\neg\exists x.\,(\lambda x.\,\text{reset}(x^{<}(\text{saw}^{>}\text{shift } g.\,\lambda\ell.\,\forall y.\, gy\ell)))x\ell\big)$$

$$\leadsto \text{reset}\big(\neg\exists x.\,\text{reset}(x^{<}(\text{saw}^{>}\text{shift } g.\,\lambda\ell.\,\forall y.\, gy\ell))\ell\big)$$

$$\leadsto \text{reset}\big(\neg\exists x.\,\text{reset}(\lambda\ell.\,\forall y.\,(\lambda y.\,\text{reset}(x^{<}(\text{saw}^{>}y)))y\ell)\ell\big)$$

$$\leadsto \text{reset}\big(\neg\exists x.\,(\lambda\ell.\,\forall y.\,(\lambda y.\,\text{reset}(x^{<}(\text{saw}^{>}y)))y\ell)\ell\big)$$

$$\leadsto \text{reset}\big(\neg\exists x.\,\forall y.\,(\lambda y.\,\text{reset}(x^{<}(\text{saw}^{>}y)))y\ell\big)$$

$$\leadsto \text{reset}\big(\neg\exists x.\,\forall y.\,\text{reset}(x^{<}(\text{saw}^{>}y))\ell\big)$$

$$\leadsto \text{reset}\big(\neg\exists x.\,\forall y.\, x^{<}(\text{saw}^{>}y)\ell\big)$$

$$\leadsto \text{reset}\big(\neg\exists x.\,\forall y.\, x^{<}(\text{saw}^{>}y)\big) \leadsto \neg\exists x.\,\forall y.\, x^{<}(\text{saw}^{>}y)$$

Similarly for (27), but not for (28): the transitions for (28) get stuck, because $\lambda\ell$ in (30) needs a license and does not get it.

As in Fry's analyses, our translations of *nobody* and *anybody* integrate the scope-taking and polarity-licensing aspects of their meanings. Consequently, *nobody* must take scope *immediately* over one or more occurrences of *anybody* in order to license them. Essentially, (29)–(31) specify a finite-state machine that accepts strings of scope-taking elements in properly licensed order. It is easy to incorporate into the same system other scope-taking elements, such as the positive polarity item *somebody* and its surprising interaction with *everybody* [15–17]: surface scope is possible in (33) but not the simpler sentence (34).

(33) Nobody introduced everybody to somebody.

(34) Nobody introduced Alice to somebody.

3.3 Quotation

A deterministic transition relation predicts wrongly that quantifier scope can never be ambiguous. In particular, the transition relation above enforces call-by-value, left-to-right evaluation and thus forces quantifiers that do not contain each other to take surface scope with respect to each other. We refine our empirical predictions using the notion of metalinguistic quotation, as expressed in operational semantics by *multistage programming* [18, inter alia].

Briefly, multistage programming makes three new constructs available in programs: *quotation*, *splice*, and *run*. Quoting an expression such as $1+2$, notated $\lceil 1+2 \rceil$, turns it into a static value, a piece of code. If x is a piece of code, then it can be spliced into a quotation, notated $\lfloor x \rfloor$. For example, if x is $\lceil 1+2 \rceil$, then $\lceil \lfloor x \rfloor \times \lfloor x \rfloor \rceil$ is equivalent to $\lceil (1+2) \times (1+2) \rceil$. Finally, $!x$ notates running a piece of code. The transitions below illustrate.

(35)
$$\big(\lambda x.\,!\lceil \lfloor x \rfloor \times \lfloor x \rfloor \rceil\big)\big(\lceil 1+2 \rceil\big)$$
$$\rightsquigarrow\ !\lceil \lfloor \lceil 1+2 \rceil \rfloor \times \lfloor \lceil 1+2 \rceil \rfloor \rceil$$
$$\rightsquigarrow\ !\lceil (1+2) \times \lfloor \lceil 1+2 \rceil \rfloor \rceil \ \rightsquigarrow\ !\lceil (1+2) \times (1+2) \rceil$$
$$\rightsquigarrow\ (1+2) \times (1+2) \rightsquigarrow 3 \times (1+2) \rightsquigarrow 3 \times 3 \rightsquigarrow 9$$

We need to add to our transition relation the general cases of the second, third, and fourth transitions above. We omit these formal definitions, which involve augmenting evaluation contexts too.

We contend that inverse scope is an instance of multistage programming. Specifically, we propose that the sentence (24) has an inverse-scope reading because it translates not just to the program (25) but also to the multistage program

(36)
$$\text{reset}!\lceil \text{reset}\big((\overbrace{\text{shift } f.\,\exists x.\, fx}^{somebody})^{<}(\overbrace{\text{saw}}^{saw}\, ^{>}\lfloor \overbrace{\text{shift } g.\,\forall y.\, g\lceil y \rceil}^{everybody}\rfloor)\big)\rceil.$$

This latter program makes the following transitions, even under left-to-right evaluation.

$$\rightsquigarrow \text{reset}\big(\forall y.\,\big(\lambda y.\,\text{reset}!\lceil \text{reset}((\text{shift } f.\,\exists x.\, fx)^{<}(\text{saw}^{>}\lfloor y \rfloor))\rceil\big)\lceil y \rceil\big)$$
$$\rightsquigarrow \text{reset}\big(\forall y.\,\text{reset}!\lceil \text{reset}((\text{shift } f.\,\exists x.\, fx)^{<}(\text{saw}^{>}\lfloor \lceil y \rceil \rfloor))\rceil\big)$$
$$\rightsquigarrow \text{reset}\big(\forall y.\,\text{reset}!\lceil \text{reset}((\text{shift } f.\,\exists x.\, fx)^{<}(\text{saw}^{>}y))\rceil\big)$$
$$\rightsquigarrow \text{reset}\big(\forall y.\,\text{reset reset}((\text{shift } f.\,\exists x.\, fx)^{<}(\text{saw}^{>}y))\big)$$
$$\rightsquigarrow \text{reset}\big(\forall y.\,\text{reset reset}(\exists x.\,(\lambda x.\,\text{reset}(x^{<}(\text{saw}^{>}y)))x)\big)$$
$$\rightsquigarrow \text{reset}\big(\forall y.\,\text{reset reset}(\exists x.\,\text{reset}(x^{<}(\text{saw}^{>}y)))\big)$$
$$\rightsquigarrow \text{reset}\big(\forall y.\,\text{reset reset}(\exists x.\, x^{<}(\text{saw}^{>}y))\big)$$
$$\rightsquigarrow \text{reset}\big(\forall y.\,\text{reset}(\exists x.\, x^{<}(\text{saw}^{>}y))\big)$$
$$\rightsquigarrow \text{reset}\big(\forall y.\,\exists x.\, x^{<}(\text{saw}^{>}y)\big) \rightsquigarrow \forall y.\,\exists x.\, x^{<}(\text{saw}^{>}y)$$

The intuition behind this translation is that the scope of *everybody* in the inverse-scope reading of (24) is metalinguistically quoted. The program (36) may be glossed as "Everybody y is such that the sentence *Somebody saw y* is true", except it makes no sense to splice a person into a sentence, but it does make sense to splice the quotation $\lceil y \rceil$ in (36) into a sentence [19].

Several empirical advantages of this account of inverse scope, as opposed to just allowing non-left-to-right evaluation, lie in apparently noncompositional

phenomena other than (but closely related to) quantification. In particular, we explain why a polarity licensor cannot take inverse scope over a polarity item [13, 20].

(37) *Anybody saw nobody.

Surface scope is unavailable for (37) simply because *anybody* must take scope (right) under its licensor. All current accounts of polarity sensitivity capture this generalization, including that sketched in Section 3.2. A more enduring puzzle is why inverse scope is also unavailable. Intuitively, our analysis of inverse scope rules out (37) because it would gloss it as "Nobody y is such that the sentence *Anybody saw y* is true" but *Anybody saw y* is not a well-formed sentence.

Formally, we hypothesize that quotation only proceeds by enclosing the translation of clauses in reset!$\lceil I_t \text{ reset} \ldots \rceil$, where I_t is an identity function restricted to take proposition arguments only. In other words, the only transition from a program of the form $C[I_t V]$, where $C[\ ]$ is any evaluation context, is to $C[V]$ when V is a proposition (rather than a function such as $\lambda \ell. \ldots$). We hypothesize I_t not because the operational semantics forces us to, but to express the intuition (some might say stipulation) that quotation applies to propositions only.

Replacing reset!$\lceil \text{reset} \ldots \rceil$ in (36) by reset!$\lceil I_t \text{ reset} \ldots \rceil$ does not hamper the transitions there, because $\exists x. x^<(\text{saw}^> y)$ is a proposition. In contrast, even though (37) translates successfully to the program

$$
(38) \quad \text{reset!} \left\lceil I_t \text{ reset} \big((\overbrace{\text{shift } f. \lambda \ell. \exists x. f x \ell}^{anybody})^< (\overbrace{\text{saw}}^{saw}{}^> \lfloor \overbrace{\text{shift } g. \neg \exists y. g \lceil y \rceil \ell}^{nobody} \rfloor) \big) \right\rceil,
$$

it then gets stuck after the following transitions.

$$
\leadsto \text{reset}\big(\neg \exists y. \big(\lambda y. \text{reset!} \lceil I_t \text{ reset}((\text{shift } f. \lambda \ell. \exists x. f x \ell)^< (\text{saw}^> \lfloor y \rfloor)) \rceil \big) \lceil y \rceil \ell \big)
$$

$$
\leadsto \text{reset}\big(\neg \exists y. \big(\text{reset!} \lceil I_t \text{ reset}((\text{shift } f. \lambda \ell. \exists x. f x \ell)^< (\text{saw}^> \lfloor \lceil y \rceil \rfloor)) \rceil \big) \ell \big)
$$

$$
\leadsto \text{reset}\big(\neg \exists y. \big(\text{reset!} \lceil I_t \text{ reset}((\text{shift } f. \lambda \ell. \exists x. f x \ell)^< (\text{saw}^> y)) \rceil \big) \ell \big)
$$

$$
\leadsto \text{reset}\big(\neg \exists y. \big(\text{reset}(I_t \text{ reset}((\text{shift } f. \lambda \ell. \exists x. f x \ell)^< (\text{saw}^> y))) \big) \ell \big)
$$

$$
\leadsto \text{reset}\big(\neg \exists y. \big(\text{reset}(I_t \text{ reset}(\lambda \ell. \exists x. (\lambda x. \text{reset}(x^< (\text{saw}^> y))) x \ell)) \big) \ell \big)
$$

$$
\leadsto \text{reset}\big(\neg \exists y. \big(\text{reset}(I_t(\lambda \ell. \exists x. (\lambda x. \text{reset}(x^< (\text{saw}^> y))) x \ell)) \big) \ell \big)
$$

The scope of *nobody*, namely *anybody saw* __, is a function rather than a proposition, so the intervening I_t blocks licensing. (The I_t would block licensing even if we quote the license ℓ—that is, even if we replace the last ℓ in (38) by $\lceil \ell \rceil$.) In general, a polarity item must be evaluated after its licensor, because a quantifier can take inverse scope only over a proposition.

Our operational semantics of metalinguistic quotation, like Barker and Shan's analyses of polarity sensitivity [8, 21], thus joins a syntactic notion of order to a semantic notion of scope in an account of polarity—as desired [13, 20]. The general strategy is to proliferate clause types: a clause in the analysis above may denote not a proposition but a function from ℓ to propositions. We can

generalize this strategy to more clause types in order to account for additional quantifiers and polarity items, such as in English [15], Dutch, Greek, Italian [22], and Hungarian [23].

3.4 Other linguistic phenomena

Quantification and polarity sensitivity are just two out of many apparently non-compositional phenomena in natural language, which we term *linguistic side effects* [11]. Two other linguistic side effects are anaphora and interrogation. As the term suggests, each effect finds a natural treatment in operational semantics. For example, it is an old idea to treat anaphora as mediated by a record of referents introduced so far in the discourse [24–26]. We can express this idea in an operational semantics either by adding the record to the state as a separate component, as sketched in Section 1, or by integrating the record into the evolving program as it rewrites [27]. Another old idea is that wh-words take scope to circumscribe how an asker and an answerer may interact [28]. Our use of delimited control extends to this instance of scope taking.

The payoff of recasting these old ideas in a general operational framework goes beyond conceptual clarity and notational simplicity. Our notions of *context* and *order* apply uniformly to all linguistic phenomena and make borne-out predictions. For example, left-to-right evaluation explains not just the interaction between quantification and polarity sensitivity but also crossover in binding and superiority in questions [8, 9].

4 Conclusion

We have shown how an operational semantics for delimited control and multi-stage programming in natural language helps explain inverse scope, polarity sensitivity, and their interaction. We are actively investigating the foundations and applications of our metalanguage, where many open issues remain. In particular, there is currently no type system or denotational semantics on the market that soundly combines delimited control and quotation, so we have no way to understand type-logically or statically what it means for a program such as (38) to get stuck, much as we would like to.

Our approach extends dynamics semantics from the intersentential level to the intrasentential level, where side effects are incurred not only by sentences (*A man walks in the park*) but also by other phrases (*nobody*). Thus discourse context is not a sequence of utterances in linear order but tree of constituents in evaluation order. This view unifies many aspects of human language—syntactic derivation, semantic evaluation, and pragmatic update—as compatible transitions among a single set of states. A tight link between operational and denotational semantics [5] promises to strengthen the connection between the views of language as product and language as action [1].

References

[1] Trueswell, J.C., Tanenhaus, M.K., eds.: Approaches to Studying World-Situated Language Use: Bridging the Language-as-Product and Language-as-Action Traditions. MIT Press, Cambridge (2005)

[2] Plotkin, G.D.: A structural approach to operational semantics. Technical Report DAIMI FN-19, Department of Computer Science, University of Aarhus (1981) Revised version submitted to *Journal of Logic and Algebraic Programming*.

[3] Felleisen, M.: The Calculi of λ_v-CS Conversion: A Syntactic Theory of Control and State in Imperative Higher-Order Programming Languages. PhD thesis, Computer Science Department, Indiana University (1987) Also as Tech. Rep. 226.

[4] Felleisen, M.: The theory and practice of first-class prompts. In: POPL '88: Conference Record of the Annual ACM Symposium on Principles of Programming Languages, New York, ACM Press (1988) 180–190

[5] Danvy, O., Filinski, A.: A functional abstraction of typed contexts. Technical Report 89/12, DIKU, University of Copenhagen, Denmark (1989) http://www.daimi.au.dk/~danvy/Papers/fatc.ps.gz.

[6] Danvy, O., Filinski, A.: Abstracting control. In: Proceedings of the 1990 ACM Conference on Lisp and Functional Programming, New York, ACM Press (1990) 151–160

[7] Danvy, O., Filinski, A.: Representing control: A study of the CPS transformation. Mathematical Structures in Computer Science **2** (1992) 361–391

[8] Barker, C., Shan, C.c.: Types as graphs: Continuations in type logical grammar. Journal of Logic, Language and Information **15** (2006) 331–370

[9] Shan, C.c., Barker, C.: Explaining crossover and superiority as left-to-right evaluation. Linguistics and Philosophy **29** (2006) 91–134

[10] Shan, C.c.: Linguistic side effects. In Barker, C., Jacobson, P., eds.: Direct Compositionality, New York, Oxford University Press (2007) 132–163

[11] Shan, C.c.: Linguistic Side Effects. PhD thesis, Harvard University (2005)

[12] Fry, J.: Negative polarity licensing at the syntax-semantics interface. In Cohen, P.R., Wahlster, W., eds.: Proceedings of the 35th Annual Meeting of the Association for Computational Linguistics and 8th Conference of the European Chapter of the Association for Computational Linguistics, San Francisco, Morgan Kaufmann (1997) 144–150

[13] Fry, J.: Proof nets and negative polarity licensing. In Dalrymple, M., ed.: Semantics and Syntax in Lexical Functional Grammar: The Resource Logic Approach. MIT Press, Cambridge (1999) 91–116

[14] Kiselyov, O., Shan, C.c., Sabry, A.: Delimited dynamic binding. In: ICFP '06: Proceedings of the ACM International Conference on Functional Programming, New York, ACM Press (2006) 26–37

[15] Shan, C.c.: Polarity sensitivity and evaluation order in type-logical grammar. In Dumais, S., Marcu, D., Roukos, S., eds.: Proceedings of the 2004 Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics. Volume 2., Somerset, NJ, Association for Computational Linguistics (2004) 129–132

[16] Kroch, A.S.: The Semantics of Scope in English. PhD thesis, Massachusetts Institute of Technology (1974) Reprinted by New York: Garland, 1979.

[17] Szabolcsi, A.: Positive polarity—negative polarity. Natural Language and Linguistic Theory **22** (2004) 409–452

[18] Taha, W., Nielsen, M.F.: Environment classifiers. In: POPL '03: Conference Record of the Annual ACM Symposium on Principles of Programming Languages, New York, ACM Press (2003) 26–37

[19] Quine, W.V.O.: Word and Object. MIT Press, Cambridge (1960)

[20] Ladusaw, W.A.: Polarity Sensitivity as Inherent Scope Relations. PhD thesis, Department of Linguistics, University of Massachusetts (1979) Reprinted by New York: Garland, 1980.

[21] Shan, C.c.: Delimited continuations in natural language: Quantification and polarity sensitivity. In Thielecke, H., ed.: CW'04: Proceedings of the 4th ACM SIGPLAN Continuations Workshop. Number CSR-04-1 in Tech. Rep., School of Computer Science, University of Birmingham (2004) 55–64

[22] Bernardi, R.: Reasoning with Polarity in Categorial Type Logic. PhD thesis, Utrecht Institute of Linguistics (OTS), Utrecht University (2002)

[23] Bernardi, R., Szabolcsi, A.: Partially ordered categories: Optionality, scope, and licensing. `http://ling.auf.net/lingBuzz/000372` (2006)

[24] Kamp, H.: A theory of truth and semantic representation. In Groenendijk, J.A.G., Janssen, T.M.V., Stokhof, M.B.J., eds.: Formal Methods in the Study of Language: Proceedings of the 3rd Amsterdam Colloquium, Amsterdam, Mathematisch Centrum (1981) 277–322

[25] Groenendijk, J., Stokhof, M.: Dynamic predicate logic. Linguistics and Philosophy **14** (1991) 39–100

[26] Heim, I.: The Semantics of Definite and Indefinite Noun Phrases. PhD thesis, Department of Linguistics, University of Massachusetts (1982)

[27] Felleisen, M., Friedman, D.P.: A syntactic theory of sequential state. Theoretical Computer Science **69** (1989) 243–287

[28] Karttunen, L.: Syntax and semantics of questions. Linguistics and Philosophy **1** (1977) 3–44