

Logique & Analyse 207 (2009), 219–238

USING ILLOCUTIONARY LOGIC TO UNDERSTAND VAGUENESS

JOHN T KEARNS


1. *Two-Level Logical Systems*

Illocutionary logic is the logic of speech acts, or language acts. The study of illocutionary logic was introduced by John Searle and Daniel Vanderveken (in *Vanderveken and Searle 1985*), but I have a quite different “take” on this subject than they do, and my work in illocutionary logic is quite different from theirs. (Some of my recent papers are listed in the references.) In this paper, I am presenting my version of illocutionary logic.

A speech act, or language act, is a meaningful act performed by using expressions of a language. A person can perform a language act by speaking or writing, she can perform one by signing or by thinking with words. A person who reads or who listens with understanding is also performing language acts, although we commonly focus on those acts performed by the person who produces the expressions that are used.

I understand language acts to be the primary bearers of semantic features. Expressions, whether spoken or written (or signed or thought) are the bearers of syntactic features, and can themselves be regarded as syntactic objects. The meaning of a language act is the meaning that the language user intends. Expressions are conventionally used to perform acts with certain meanings, and it is common for a language user to intend meanings conventionally associated with the expressions that she uses. But a person can by mis-speaking produce the wrong expression for the act she performs — her act still has the meaning she intends, although her expression may mislead her addressees. She can also misunderstand the meaning conventionally associated with an expression, and use that expression to perform acts with the meaning that she mistakenly thinks is conventionally associated with the expression.

A *sentential act* is a language act performed by using a sentence. A *statement* is a sentential act that is true or false. This is a stipulated meaning for (my use of) the word, because statements are often understood to be





something like assertions. A sentential act can be performed with a certain *illocutionary force*, like the force of a promise or a request, of an apology or a threat. A sentential act performed with a certain illocutionary force constitutes an *illocutionary act*. Statements themselves can be performed with a variety of illocutionary forces. For example, a single statement can be asserted or denied, it can be supposed true or supposed false. A statement performed with a certain illocutionary force constitutes an illocutionary act, but in talking about *statements*, we abstract away from whatever forces they might have.

A *speech-act argument*, which is the kind of argument that people actually make, contains illocutionary acts as premisses and conclusions. A simple argument consists of one or more premiss acts and a conclusion inferred from them, and a complex argument contains (other) arguments as components. Our considering simple and complex arguments makes it inappropriate to characterize deductive arguments as valid or invalid, for the customary understanding of validity would apply only to simple arguments. For us, arguments are either *deductively correct* or not.

Let us consider how a focus on language acts impacts the study of logic, and how illocutionary logic is related to standard logic. It is now customary to carry out research in logic by developing *logical systems*, or *logical theories*. Such a theory consists of (1) a formal language, usually artificial, (2) a semantic account for that language, and (3) a deductive system for establishing one or another kind of result involving expressions of the language. Since people don't normally use expressions of formal languages to write, say, or think things, there is a sense in which these are not genuine languages. (Expressions of genuine languages are “scripts” for performing language acts.) I will continue to speak of *logical languages*, but I think it most appropriate to regard sentences of these languages as *representations* of language acts that people either do perform or might perform. The semantic account is then for the language acts that are represented, and the deductive system codifies certain “logically distinguished” expressions of the formal language.

In a standard system of logic, the sentences of the logical language represent statements, and the semantic account gives truth conditions of these statements. The deductive systems focus on truth-conditional logical consequence in one way or another. Standard systems make no provision for illocutionary force. To incorporate the illocutionary dimension, we add things to a standard system of logic. To obtain a theory/system of illocutionary logic, (1) illocutionary-force indicating expressions, or *illocutionary operators*, are added to the formal language; (2) the account of truth conditions



of statements is supplemented with an account of semantic features of illocutionary acts; (3) the deductive system is modified to accommodate illocutionary operators and illocutionary force.

In a system of illocutionary logic, the formal language and semantic account have two levels: the first level is ontological, or *ontic*, and the second level is *epistemic*. The ontic level considers the formal language without illocutionary operators, and provides a (familiar) account of the truth conditions of statements represented by sentences of this language. A standard system of logic is a first-level system of illocutionary logic.

The second-level (or full) system deals with illocutionary acts and *rational commitment*. In the first-level system, there is no special interest in who it is that makes the statements being considered. Different people can make essentially the same statement (or essentially similar statements), and the logically important features of a statement are independent of whoever it is that makes the statement. A first-level logical system is a "third person" system. It is different with illocutionary acts. Two people cannot, for example, make essentially the same assertion. For Jones' assertion commits Jones to make further assertions and denials, but make no claims on Smith, while Smith's assertion commits Smith but not Jones. Her commitments underlie the deductively correct (speech-act) arguments that a person makes. It is essential to an assertion, and is also logically important, just whose assertion it is.

An epistemic-level system of illocutionary logic is a first-person system developed for, and from the perspective of, a particular person. I generally consider an idealized person whom I call the *designated subject*, for whom I usually use feminine pronouns. The illocutionary acts represented by expressions in the epistemic-level logical language are the designated subject's acts. They could also be our own acts, if we use the language to represent them. The epistemic-level semantic account is concerned with the *rational commitments* of the designated subject.

But how am I understanding rational commitment? Making a decision to do or not do something rationally commits the agent to do or not do it. Performing some intentional acts can rationally commit a person to do others. Some commitments are unconditional — deciding to get a haircut on the way to work establishes such a commitment for me. Other commitments are conditional, like my commitment to close the upstairs windows if it rains while I am at home. Accepting, or asserting, a certain statement as true commits me to accept other statements, but this is a conditional commitment. I am committed to act only if I have some interest in the matter and give it some

thought. It is irrational, for example, to accept $[A \vee B]$ and $\sim A$, but refuse to accept B ; however, I am not rationally required to consider whether B is true or not. (And I can lose the commitment if I give up one of the initial assertions.)

A group of statements either implies (or entails) a further statement or not. If there is implication, there is no distinction between immediate and mediate (or remote) implication. But some commitments are immediate while others aren't. If accepting A immediately commits me to accept B , and accepting B immediately commits me to accept C , accepting A may commit me only *mediately* to accept C . It is immediate commitment which motivates a person to act. In carrying out a complex deduction, the reasoner follows a chain of immediate commitment from the initial illocutionary acts to the concluding act. In an epistemic-level illocutionary theory, the deductive system is a first-person system for the designated subject to use in tracing the commitment consequences of her own illocutionary acts.

2. An Illustrative System Of Illocutionary Logic

The language L has indefinitely many atomic sentences, and compound sentences formed with the connectives ' \sim ,' ' \vee ,' and '&.' (The horseshoe of material implication is a defined symbol.) The atomic and compound sentences are the *plain sentences* of L . In addition, L has these illocutionary operators:

- \vdash – the sign of assertion \neg – the sign of denial
- \sqsubset – the sign of supposing true \neg – the sign of supposing false

An assertion is here understood to be an act of considering/making a statement and accepting it as being or representing what is the case, or an act of considering and reaffirming a statement if it has already been accepted. A denial is an act of considering and rejecting a statement, or of indicating one's continued rejection of the statement. So assertions and denials don't require an audience, and all assertions and denials are sincere.

A person can suppose a statement to be true, which is to temporarily accept the statement, or suppose a statement to be false, to reason from the supposition to further acts. When one infers a conclusion from one or more suppositions, that conclusion also has the status of a supposition, and will be called a supposition. (Even though, ordinarily, we say "suppose" to begin an argument, but don't say "suppose" for the conclusion.)

If A is a plain sentence of L , then the following: $\vdash A$, $\neg A$, $\ulcorner A$, $\neg A$ are *completed sentences* of L . There are no other completed sentences. So all completed sentences begin with an illocutionary operator. Illocutionary operators cannot be iterated, and a completed sentence cannot occur as a component of a larger sentence.

The ordinary connectives represent something a person says in making a statement, but the illocutionary operators represent what a person *does* in making a statement. It often happens that a person doesn't *say* anything to indicate what she is doing. (We don't usually say "I assert that" when making an assertion.) Completed sentences are not used by us to say what the designated subject does, they are for the designated subject to use to make assertions, denials, and suppositions. We can also use them to perform our own illocutionary acts. (Of course, sentences in the logical language, and proofs in the deductive system, are *representations* of acts performed by the designated subject, but they aren't representing statements *about* the designated subject.)

An epistemic-level deductive system should employ constructions which represent genuine (language act) arguments, whose premisses and conclusions are illocutionary acts. In a fuller presentation, I would develop a system S which is a natural deduction system employing tree proofs (or deductions). Proofs in such a system are especially perspicuous, since their deductive structure is apparent. Each step in one of these proofs is a completed sentence. An initial step in a tree proof can be an assertion, denial, or supposition. Only initial suppositions are hypotheses of the proof. Initial assertions and denials should be (or represent) knowledge or beliefs and disbeliefs of the designated subject (of whoever is making the argument).

Except for the illocutionary operators, S is a familiar sort of deductive system. Rules for constructing tree proofs (for making the inferences in tree proofs) take account of both truth conditions and illocutionary force. For example, the rule *& Introduction* allows one to infer the supposition of a conjunction when one premiss is a supposition:

$$\frac{\ulcorner A \quad \ulcorner B}{\ulcorner [A \ \& \ B]} \qquad \frac{\vdash A \quad \ulcorner B}{\ulcorner [A \ \& \ B]} \qquad \frac{\ulcorner A \quad \vdash B}{\ulcorner [A \ \& \ B]}$$

And to infer an assertion when both premisses are asserted:

$$\frac{\vdash A \quad \vdash B}{\vdash [A \ \& \ B]}$$

But it isn't correct to argue like this:

$$\frac{\perp A \quad \perp B}{\vdash [A \& B]}$$

even though the concluding statement is true whenever the premiss statements are true. This argument isn't invalid, but it is *deductively incorrect*.

In complex arguments, a sub-argument, as a whole, often serves as a premiss. We see this in the following:

$$\frac{\frac{x}{\perp A} \quad \vdash B}{\perp [A \& B]} \quad \& \text{ Introduction}$$

$$\frac{\perp [A \& B]}{\perp B} \quad \& \text{ Elimination}$$

$$\frac{\perp B}{\vdash [A \supset B]} \quad \supset \text{ Introduction, discharge } \perp A$$

The entire sub-proof ending with ' $\perp B$ ' is a premiss for the final conclusion. The hypothesis ' $\perp A$ ' of this sub-proof is discharged by the application of \supset *Introduction* so that it is not a hypothesis of the completed proof. (An ' x ' is placed above hypotheses that are discharged.) The conclusion is an assertion because it depends on no undischarged hypotheses. If the initial assertion ' $\vdash B$ ' were replaced by a supposition ' $\perp B$,' then the conclusion would be the supposition ' $\perp [A \supset B]$.'

It is important to understand that proofs in the deductive system are to be made by the designated subject — we are merely onlookers. Although it is actually us who are constructing the proofs, we are doing this “on her behalf.” (On the other hand, we can place ourselves in the role of the designated subject, and construct arguments for ourselves.)

3. Two Semantic Levels

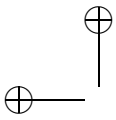
It is to some extent arbitrary where we draw the line around the field to be called *semantics*. Semantics might, for example, be limited to the truth conditions of statements, or it might be the study of the meanings conventionally associated with expressions. I conceive semantics more broadly, although as I conceive it, semantics includes the study of the truth conditions of statements, and also of meanings conventionally associated with expressions.

My goal is to understand language and the use of language, so that I can understand what topics and issues belong together, and further understand how they fit together. To this end, I am trying to develop and deploy concepts that "carve nature at its joints," and to use these concepts in investigating illocutionary logic. With topics that fit together, we can use similar techniques to explore them. The two semantic levels associated with an illocutionary system of logic are both concerned with meanings intended by language users, and in studying these levels we employ functions that assign things to expressions in the logical language.

The first semantic level of the language L is entirely familiar. *Interpreting functions* assign truth and falsity to atomic plain sentences of the language, and each of these functions determines a valuation of the plain sentences of L in a standard way. A set X of plain sentences of L *implies* a plain sentence A iff there is no interpreting function of L for which all sentences in X are true but A is false. A set X of plain sentences is (*semantically*) *consistent* iff there is an interpreting function of L for which all sentences in X are true.

The second-level semantic account is for completed sentences of L , and should deal with semantic features of these sentences in a way analogous to that in which the first-level account deals with truth and truth conditions. An assertion or a denial or a supposition isn't true or false, except in some derivative sense. But the designated subject, or any of us, will be *committed* at a time to accept or reject certain statements, or to suppose them true or suppose them false. We can develop the second-level account for the designated subject at some particular time. It is convenient to focus on either knowledge or (justified) belief. At the time in question, there are some statements that the designated subject has actually considered and accepted, which she continues to accept, and some statements she has considered and rejected, which she continues to reject. If she has accepted the statements with the force of knowledge claims and rejected them with an analogous force, her acceptings and rejectings constitute her *explicit knowledge*; otherwise they are her *explicit beliefs* and *disbeliefs*. These acts commit her at that time to accept some further statements and to reject still others.

If the designated subject has already accepted A , and hasn't forgotten or changed her mind, she is committed to perform: $\vdash A$, and if she hasn't yet accepted A but is committed to do so, then she is also committed to perform ' $\vdash A$.' If we use the symbol '+' for the value of assertions and denials (completed sentences ' $\vdash A$ ' and ' $\neg B$ ') which the designated subject is committed at that time to "perform," then a *commitment valuation* will be a function which assigns + to some assertions and denials of L .



Commitment and truth have important relations to one another. A commitment valuation V is *based on* an interpreting function f iff the function makes true those statements whose assertions have value $+$, and makes false those statements whose denials have value $+$. That is, (i) if $V(\vdash A) = +$, then $f(A) = T$, and (ii) if $V(\neg A) = +$, then $f(A) = F$. If V registers the knowledge (or the belief) of the designated subject at a time, and V is based on f , then f is one way the world might be, given what the designated subject knows, or believes.



A commitment valuation is *coherent* iff it is based on an interpreting function. If commitment valuation V is based on interpreting function f , then $\langle f, V \rangle$ is a *coherent pair*.

If sentences of L represent certain statements, it may happen that the meanings of those statements, or some aspects of their meanings, determine some interpreting functions not to be appropriate. For example if distinct sentences A, B represent statements with the same meaning, then an interpreting function which assigns T to A and F to B isn't appropriate. We often limit our attention to a class W of *admissible* interpreting functions.

If a coherent valuation V_0 registers the designated subject's explicit knowledge or belief at a time, there will be some sentences that are true for every (admissible) interpreting function on which V_0 is based, and others which are false for every such function. We anticipate that the designated subject will be committed to accept all those statements that always come out true, and to reject the ones that must be false. The valuation V which assigns $+$ to all the assertions and denials which the designated subject is committed to perform is the *completion* of V_0 .

Commitment valuations do not award values to suppositions, because suppositions come and go too quickly. We often introduce a supposition in the course of an argument and discharge that supposition before the argument is finished. Assertions and denials, at least for an idealized person, have more permanence. To explain second-level counterparts to implication and (semantic) inconsistency, we introduce a concept of *satisfaction*. Satisfaction involves both truth and commitment, although commitment itself respects truth conditions.

A coherent commitment valuation *satisfies* those assertions and denials which have value $+$ for its completion. An interpreting function f *satisfies* a positive supposition $\sqcup A$ iff A is true for the valuation determined by f , and it *satisfies* a negative supposition $\sqcap A$ iff A is false for that valuation. Finally,



a coherent pair $\langle f, V \rangle$ satisfies a completed sentence B iff (i) B is an assertion or denial, and V satisfies B , or (ii) B is a supposition, and f satisfies B .

With this conceptual apparatus, we can introduce illocutionary counterparts to implication and (semantic) consistency:

Let X be a set of completed sentences of L and A be a completed sentence of L . Then X logically requires A and A is a commitment consequence of X iff (i) A is an assertion or denial and every coherent pair that satisfies the assertions and denials in X also satisfies A , or (ii) A is a supposition and every coherent pair that satisfies all the sentences in X also satisfies A .

A set X of completed sentences is coherent iff there is a coherent pair which satisfies every sentence in X .

The definition of logical requiring has two clauses, because suppositions make no demands on assertions and denials. The set $X = \{\lrcorner A, \lrcorner \lrcorner A, \lrcorner B\}$ logically requires the positive and the negative supposition of every plain sentence, but it is only the denial of B which leads to further assertions and denials.

A second-level deductive system in an illocutionary logical theory enables us to derive the commitment consequences of initial assertions, denials, and suppositions. It is a straightforward matter to develop these systems, and to establish that they are sound and complete in suitable senses. That such a system is sound and complete shows, for the system in question, that the commitment associated with logical form adequately tracks or traces the truth-conditional consequences of knowledge, or of coherent beliefs and disbeliefs.

Systems of illocutionary logic are useful for representing, and understanding, what people are doing when they say things, and when they construct arguments or proofs. These systems provide the resources for solving, or resolving, a number of puzzles concerning language. For example, if 'B' is a belief operator for the designated subject, so that 'B(A)' is true iff the designated subject (explicitly) believes A at this moment, then a sentence '[A & ~B(A)]' will be consistent, but its assertion ' $\lrcorner [A \& \sim B(A)]$ ' will be incoherent for the designated subject to perform. This both explains, and dissolves, Moore's Paradox. And if 'K' is a knowledge operator, so that 'K(A)' is true iff the designated subject knows A (at this moment), then a sentence '[A & ~K(A)]' can be true, but the designated subject can't know it (at this moment). The assertion ' $\lrcorner [A \& \sim K(A)]$ ' is incoherent for the designated subject to perform with the force of a knowledge claim. This is sometimes regarded as a paradox, or puzzle, but there is nothing paradoxical about it.

4. *Reconceiving Studies Of Language*

A system of illocutionary logic is developed in order to help us understand and explain our practice of using language to perform a variety of illocutionary acts, and our practice of constructing deductive arguments. An adequate account of these practices must accommodate both truth and commitment, and systems of illocutionary logic are equipped to do this. These systems faithfully represent realistic arguments, provide the resources for distinguishing assertions and denials from suppositions, and help us better understand indirect arguments that introduce and discharge suppositions.

The study of illocutionary acts and illocutionary force is often thought to belong to pragmatics rather than semantics. In fact, our distinction between the ontic and the epistemic levels of a theory of illocutionary logic is sometimes taken to demarcate the semantic and pragmatic dimensions of language. But in systems of illocutionary logic, the treatment of commitment and completed sentences parallels the more familiar account of truth conditions and plain sentences. There are second-level counterparts of implication and consistency which can be explored by formal techniques similar to those used to investigate truth conditional ideas.

I think it is common to regard a language as a kind of "free standing" entity composed of expressions possessing syntactic and semantic features, where the semantic features are concerned with what might be called "descriptive content." The language user simply employs items in this structure, taking advantage of their semantic features, and sometimes contributing extra features to those that are already there. For example, the language user supplies illocutionary force, she exploits the meanings she finds to achieve new meanings in cases of irony or metaphor, and on occasion manages to do other things to achieve various conversational implicatures.

From our speech act perspective, matters of meaning and force which are intentionally supplied by a language user, especially by the language user who produces the expressions she uses, fall within the area of semantics. We can distinguish customary, conventional meanings from other sorts, and semantic studies commonly focus on conventional meanings. But there is no mystery about how a language producer manages to mean what she does mean, or about how she knows what she means — it is what she intends. What needs explaining is how her addressees are able to determine what she means. I think a third study of language (in addition to syntax and semantics), which might as well be called *pragmatics*, is appropriately concerned with how meaning (including illocutionary force) is communicated, with the

cues and clues that addressees use to determine what the speaker/writer intends. Actually, pragmatics should be conceived more broadly as investigating how what one says and doesn't say (as in "damning with faint praise") can serve to communicate what a speaker intends.

5. *Ideal Completions*

If a coherent commitment valuation V registers the explicit knowledge, or the explicit beliefs and disbeliefs, of the designated subject at a certain time, and V is based on interpreting function f , we have thought of f as one way the world might be, given what is currently known, or believed and disbelieved. But to think of f in this way is to regard the language L , or the language whose statements are represented by (sentences of) L , as *completely developed*. Every sentence of L would then represent a statement which is determinately true or determinately false.

The languages we speak are never like this. We are making them up as we go along. With respect to the meanings with which expressions are conventionally used, many expressions are vague: when used conventionally, they have borderline cases. A predicate ' F ' has a criterion associated with it, either conventionally, or idiosyncratically by a given language user. The predicate can be truly applied to an object if the object satisfies the criterion, and falsely applied if the object fails to satisfy that criterion. A borderline case α for predicate ' F ' is not a case that fails to satisfy the criterion. If α is a borderline case, then there are good reasons for applying ' F ' to α , and good reasons to withhold ' F ' from α . But, with apologies to paraconsistency logicians, an object cannot both be F and not be F .

A vague expression can certainly be predicated of a borderline case for that expression. The resulting language act is significant: it represents (say) α as satisfying ' F 's criterion. But in seriously predicating an expression of an object, a language user intends to make a statement — to say what is either true because it fits the world or false because it fails to fit. Someone who predicates a vague expression of a borderline case fails to realize this intention, for what she says doesn't clearly fit the world, and it doesn't clearly fail to fit. Her language act fails to be a statement at all. A statement that isn't true is false (and vice versa); but while a failed attempt to make a statement isn't true, it isn't false either. Only genuine statements are either true or false.

When a vague expression is predicated of a borderline case for that expression, the resulting speech act fails to be a statement, but there are different *ideal completions* of the language in which that predicate yields different

values when predicated of the object in question. In these ideal completions, it is the same object which is either F or *not* F , but in these completions, that object will be one or the other. It is the same predicate as well, with the criterion for applying it "tightened up" in one way or another.

If we think of a logical language as being, or representing, a natural language, it isn't interpreting functions which determine the conventional meanings of expressions. In actual speech, it is the meanings of what we say that determine whether our statements are true or false. Interpreting functions just reflect the way statements can turn out, given their meanings and the world that they describe. But the meanings of our expressions, together with the world as it is, don't determine for everything that we can say, that it is clearly true or clearly false. An interpreting function that is faithful to our actual situation will assign T to some atomic sentences, and F to others, but will leave many sentences with no values.

Given a function that is faithful to our actual situation, there will be many *classical* interpreting functions which extend our faithful function, and assign either T or F to each atomic sentence. These classical functions represent *ideal completions* of our language — they represent the different ways things can turn out in ideal completions of our language.

However, given the meanings of expressions (and language acts) at present, not all interpreting functions can represent ideal completions. Even if we begin with a set of admissible interpreting functions, it may be that not all classical interpreting functions can represent ideal completions. For example, imagine that ' $F(x)$ ' means *x is red*, and that α and β are distinct borderline cases for red, but these two borderline cases are of exactly the same color. The faithful interpreting function will not assign a value to either ' $F(\alpha)$ ' or ' $F(\beta)$.' But we may not think an admissible classical extension of our faithful function is appropriate if it assigns T to ' $F(\alpha)$ ' and F to ' $F(\beta)$,' or the reverse. It won't be appropriate if we intend for ' $F(x)$ ' to continue to mean *x is red*. However, we might be willing to slightly change the meaning so that ' $F(x)$ ' then means *x is red and φ ,* where φ is some further condition satisfied by all currently red objects, and by α , but not satisfied by β . Which classical extensions of our faithful interpreting function are *allowable (ideal) completions* will depend on our intentions for our language and the expressions it contains.


The designated subject, who is an idealized language user, has complete and perfect knowledge of the language she speaks, and the statements she accepts are true, while those she rejects are false. With respect to a context in which we deal with sentential acts which fail on account of vagueness

to be statements, we can understand the designated subject to have an intention for the meaning of every expression in the language. Let V be a commitment valuation which reflects the designated subject's knowledge at a given time. We can attend to allowable (classical) interpreting functions f which extend the faithful interpreting function. Since the designated subject's alleged knowledge is genuine knowledge, the commitment valuation V will be based on all of these functions, and these functions will respect the designated subject's intentions for expressions in the language. However, the valuation V has little importance for the allowable completions of the language. They depend on how the world actually is, and not on what the designated subject knows or believes.

To deal with vagueness in a second-level system of illocutionary logic, we should consider only interpreting functions which are admissible apart from the designated subject's current knowledge. An admissible (rather than allowable) classical interpreting function f on which V is based shows one way the world might be in an ideal completion of our language, given what the designated subject now knows. We aren't concerned with the way things actually are, but with how they might be, given what the designated subject now knows. And the designated subject's intentions for the meanings of different expressions may change as she acquires additional knowledge, so that her current intentions are not decisive for the ideal completions that will eventually turn out to be allowable. Our treatment should reflect the way things may (turn out to) be (in an ideal completion), given what the designated subject now knows.

Systems of illocutionary logic don't enable us to "solve" the problem of vagueness, because there is no problem that calls for a solution. But reflecting on systems of illocutionary logic helps us better understand the phenomenon of vagueness. One puzzle that arises in connection with vagueness is this: If we have a heap of grains of sand, or a man with hair, one who clearly isn't bald, then taking away one grain of sand or one hair from the man's head doesn't change things very much. We still have a heap, or a man with hair. Yet everyone realizes that taking away one grain of sand at a time over and over again will eventually result in there no longer being a heap, and taking away one hair at a time (ouch!) will eventually produce a man who is bald.

We seem to either have conflicting intuitions, or else have an intuition which is contradicted by what we know better. But let us reflect on the designated subject who is an ideal, or an idealized, language user. She has perfect knowledge of our language, so far as it has been developed. And she also has some knowledge of the world. She knows of collections of





sand particles which are genuine heaps, which she has judged to be heaps, and knows lots of men who are definitely not bald, some of whom she has judged to be not bald. For all of the collections which she has judged to be heaps, and all of the men she has judged to be not bald, it is the case that taking away one grain of sand will still leave a heap, and taking away one hair will still leave a man who isn't bald. Our language is sufficiently developed that for those collections she has judged to be heaps, it is true that taking away a grain will leave a heap. This is the source of her intuition, and ours, about taking away one relevant item from a heap or head of hair. But if we are going to keep taking items away, and for some reason want to be able to make yes or no judgments about whether we still have a heap or head of hair, then we will need to amend our language in the “direction” of an ideal completion.

6. *The Logic Of Ideal Completions*

Given that a language is never completely developed, we might adopt the project of describing and investigating interpreting functions which reflect the development of a language at a time. Some sentences are true or false, others lack a value. If we are trying to reflect ordinary usage, some issues may be controversial, or simply unclear. I myself am unsure whether a disjunction with one true disjunct and one disjunct which fails to be a statement is a true disjunction or is one which fails to be a statement. However that turns out, we can devise (first-level) deductive systems to track truth-conditional consequence. Another project for dealing with vagueness would involve assigning rational or real numbers ≤ 1 to reflect the “degree” to which an object satisfies the criterion associated with a predicate. It is probably an interesting task to carry out such investigations, or to develop systems of fuzzy logic. But such systems don't shed much light on our ordinary ways of speaking, writing, and thinking. However, one of my goals in developing systems of illocutionary logic is to illuminate our actual practice, or practices.

The fact that our actual language is incompletely developed doesn't slow us down much. We don't reason in different ways when we do and don't have vague expressions and potential borderline cases. I think we can use illocutionary logic to understand why no unusual or unorthodox logical principles are required. My proposal involves the epistemic level of a system of illocutionary logic, and might ordinarily be considered to be pragmatic as opposed to semantic. But it is semantic in my broad sense. To understand this proposal, we need to recognize that there are different *modes*, or *manners*, of accepting and rejecting statements. The simplest and most straightforward



manner is *factual*: we accept a statement as being, or representing, the way things are, and we reject one for being at odds with the way things are. Even with the factual mode, we can distinguish between accepting a statement with the force of a knowledge claim, and accepting one with the force of justified belief. As I understand it, the factual mode of accepting and rejecting statements is all or nothing — we don't partially accept a statement as being the case.

It is also possible to accept a statement *as a basis for action*, without accepting that the statement indicates the way things really are. Accepting a statement as a basis for action is a matter of more and less, and is characteristically indicated by speaking of what is or isn't probable. If I accept as a basis for action that it will (probably) rain today, that is a good reason to take an umbrella. But the weather report is often wrong, and I do not simply accept that it will rain. Similarly, I judge that I probably won't win if I buy a lottery ticket, but I might buy one anyway. It doesn't cost so much, and there is always that chance. Accepting or rejecting statements as a basis for action, either strongly or weakly, will commit a person to accept or reject other statements as a basis for action, and logical systems can be devised to investigate this commitment. This is how I understand studies of subjective probability.

Studies of subjective probability are concerned with *probabilistic* assertion and denial, although they are not usually conceived of in this way. Another mode of accepting and rejecting, and of supposing as well, is *fictional*. John Searle has argued that there is no special illocutionary force associated with stories and story telling; instead, he says, authors are *pretending* to perform acts of factual assertion and denial. I don't disagree that we can regard authors as pretending, but *we* aren't pretending when we discuss fictional characters and events in other people's stories.

The author of a story, when she is writing the story, is something like an actor in a play. Each is pretending to make assertions, denials, etc. And neither is committed by what they say to perform further acts in the way they would be committed if their acts were genuine instead of pretended. Pretending to assert, deny, etc. is not performing an act with a special illocutionary force. However, when we talk about a fictional world, we *do* perform genuine illocutionary acts. In talking about the world of Sherlock Holmes, say, we can accept and reject statements as being, or representing, the way the story world is. I accept with respect to the stories that Sherlock Holmes was a brilliant detective, for example, and I deny that he was married. This is a distinct mode, or manner, of accepting and rejecting statements — more precisely, of accepting and rejecting sentential acts. We are committed by

our assertions and denials of fiction to perform further acts with the same manner or mode.

I am not understanding a story world to be a possible world in some set W . A story world depends on a story teller, and on her representing activity. What is true in a story world depends on what the story says, together with the default assumption that in most ordinary respects, the story world is like the real world — unless the story says otherwise. Story worlds are never complete. For example, Sherlock Holmes had a mother, and his mother had a mother, and that grandmother had a mother, because the stories don't say that Sherlock Holmes was peculiar in this respect. But there is no answer to the question of whether Holmes' mother's mother's mother had blue eyes. She either did or she didn't, because excluded middle is one of the default assumptions. But it isn't true in the story world that she had blue eyes, and it isn't true in that world that she had eyes of another color.

With respect to a given story or collection of stories, we can understand the designated subject to be an ideal reader, or listener. She knows all the stories, and she knows the default assumptions for the stories. She accepts a number of statements as true of the story world, and rejects others; she is committed to accept and reject still more. We can identify truth in the story world with the statements she accepts or is committed to accept, and falsity with the statements she rejects or is committed to reject. So if V is the completion of a coherent commitment valuation which records the designated subject's knowledge of a given story, or story world, then a statement A is true if the assertion $V(\vdash A) = +$, and B is false if $V(\vdash B) = +$. The commitment valuation determined by the designated subject's knowledge of the story determines what is true or false about the story world.

If A is the statement that Sherlock Holmes' mother's mother's mother had blue eyes, then A isn't true or false in the story world. Neither ' $\vdash A$ ' nor ' $\vdash \neg A$ ' has value $+$. But ' $\vdash [A \vee \sim A]$ ' is true in the story world, for ' $\vdash [A \vee \sim A]$ ' has value $+$. The (epistemic-level) logic for the story world is the *same* as the logic for the real world, in the sense that the same deductive system is correct for making (speech-act) arguments about the respective worlds.

With respect to a system of illocutionary logic, truth in a story world is determined by a coherent commitment valuation awarding $+$ to some assertions and denials. But these commitment valuations are still based on classical interpreting functions which assign T and F to sentences. But for a story world, T and F are not truth and falsity. So what are they? If commitment valuation V determines truth and falsity in a given story world, and V is based on interpreting function f , then f presents an *ideal completion*


of the story world. The story will never be completed — no story teller can narrate a complete story. But the story as we have it is open to infinitely many ideal completions. And it turns out that the sentences that are true of the story world are those that are true for every ideal completion of the story, while the false sentences are false for every ideal completion of the story.

We can accept a statement as true of a given story world — this amounts to accepting it as true in every ideal completion of the story world, and we can reject a statement as false of a story world. To be false of a story world is different from not being true of the story world, for many sentences (sentential speech acts) are neither true nor false. We can also suppose a statement to be true of a story world, which amounts to considering an ideal completion in which the statement is true.

7. *Remedying Vagueness*

In dealing with the logic of fiction, we understand truth and falsity in the story world to be determined by a commitment valuation, and an interpreting function which is (or which reflects) an *ideal completion* of the story is an admissible function on which the commitment valuation is based. But in considering vagueness and speech acts which fail on account of vagueness to be statements, we understand the truth, falsity, and failure of sentences (speech acts) to be determined by the way the world is. An admissible classical interpreting function on which the commitment valuation is based which reflects the designated subject's current knowledge (or belief) is an *ideal completion* of the *current language*, a completion in which no sentences fail on account of vagueness. However, both sorts of ideal completion are classical interpreting functions which determine a value for every sentence in the logical language. Our understanding of the logic of fiction can help us understand and develop the logic of a language in which some speech acts fail on account of vagueness to be statements.

Ordinarily, in dealing with factual assertion, denial, and supposition, to suppose a sentence true is to consider how things are (not how they would be) if the sentence is true. Let us be operating in a context where assertions and denials have the status of knowledge claims. Let V be a commitment valuation which records the designated subject's current explicit knowledge. Then to suppose that A is true is, in effect, to consider an (admissible) interpreting function on which V is based, which function also makes A to be true. Now imagine that we have a logical language (and the language acts which this represents) in which some sentences (some attempted statements) fail on account of vagueness. And let A be a sentence, or speech act, about





whose status we are ignorant. For all we know, A could be true, or false, or fail on account of vagueness. If we limit our attention to the factual mode of accepting, rejecting, and supposing sentences (statements), then we cannot successfully accept or suppose true a failing sentence. We can, however, reject or negatively suppose such sentences, even though we can't accept or suppose true the negations of these sentences. (This is analogous to the treatment of paradoxical sentences in my paper on the Liar Paradox: *Kearns 2007*). It is then a straightforward exercise to provide a second-level semantic account and to design a (second-level) deductive system for constructing arguments.

Even though we can come up with a suitable semantics and deductive system for representing how vagueness “works,” there seems something a little idle about doing this. My goal in developing systems of illocutionary logic is to represent our actual practice of using language, and to better understand or even resolve certain puzzles or perplexities about our use of language. But in actual practice, we are not very often troubled by vagueness, and we don't carry out reasoning that “makes provision” for vagueness. A deductive system that accommodates speech acts that fail to be statements on account of vagueness doesn't represent or reflect our actual practice, and it doesn't remedy a shortcoming of our actual practice.

If A is a sentence (attempted statement) which fails on account of vagueness, then if we suppose A to be true, it is most reasonable to think we are interested in exploring how things are if A is true. Even though A fails on account of vagueness, we don't want for this to “spoil” our supposition, turning it into a failed attempt at supposing. In this case, we aren't engaged in *factual* supposing. We are, in effect, considering an ideal completion of our language in which A is true. This is a different mode, or manner, of supposing than factual supposing.

As well as supposing that a language act is true in an ideal completion of our current language, we can also accept a language act as *true in every ideal completion of our language*, or reject a language act for being *false in every ideal completion of our language*. (If a language act is *already* a true statement or a false one, it will be true, or false, in every ideal completion.) When we are operating in circumstances where vagueness “intrudes” or might intrude, and we have no particular interest in exploring vagueness, it would make sense for us to employ the *indifferent to vagueness* manner of performing illocutionary acts — this is an *ideal-completion* manner of accepting, rejecting, and supposing speech acts. The illocutionary logic, the deductively correct speech-act arguments, are the same for ideal-completion manners of performing illocutionary acts as they are for factual asserting, denying, and



supposing — when there are no failed attempts at making statements. Since people don't, in general, pay attention to vagueness, or provide for failure due to vagueness, in carrying out reasoning and making arguments, we can imagine that they are already employing the indifferent to vagueness manner of performing illocutionary acts.

The suggestion that there is, or might be, an *indifferent to vagueness* manner of performing illocutionary acts, to accommodate language acts in which vague expressions are predicated of borderline cases for those expressions, *rescues*, or *redeems*, Kit Fine's proposal (in *Fine 1996*) that it is supervvaluations that determine the values of statements when vague expressions are predicated of what may or may not be borderline cases for those expressions. For my suggestion provides a *motivation* for Fine's proposal. If f is a function which assigns T or F to some atomic sentences of L_1 , but which may assign no value to others, then the *supervaluation determined by f* awards T to every sentence that is true for every ideal completion of f , and awards F to those sentences that are false for every ideal completion.

With respect to systems of illocutionary logic, Fine has proposed an ontic-level solution to a problem that requires an epistemic-level solution. (In terms of our earlier discussion, Fine's proposal would have us focus on allowable ideal completions rather than admissible ones.) His proposal attributes meanings to expressions (or acts of using them) which they don't possess. The values of our ordinary statements, or language acts, are not determined by supervvaluations. For a disjunction to be true, one disjunct (anyway) must be true. So that if both A and $\sim A$ are statements that fail on account of vagueness, then $[A \vee \sim A]$ also fails (to represent a statement). We don't need to attribute to sentences (and speech acts) meanings which they don't have in order to recognize different manners of accepting disjunctions (and other speech acts). If we factually accept a disjunctive statement, our attempt to do this fails if both disjuncts fail to be statements. But we can accept a disjunctive language act *as true in every ideal completion of our language* — we can do this even for a disjunctive language act which fails to be a statement.

Although probably no one actually thinks this matter through, and carefully distinguishes between accepting statements as being the case and accepting sentential speech acts as being true in every ideal completion of our language, looking at things this way accommodates our ordinary practice. We don't alter, and don't need to alter, the inference principles we ordinarily employ when vague expressions are predicated of borderline cases for those expressions. But these are principles for reasoning with illocutionary acts

(with completed sentences of the artificial language), not principles for exploring (truth-conditional) logical consequence.

My proposal for "handling" language acts that fail on account of vagueness to be statements requires us to employ a two-stage system of illocutionary logic. At the first level, or the ontic level, statements or attempted statements have values determined by the fit (or absence of fit) between their meanings and the world. It is at the epistemic level that we are in a position to recognize different modes of performing illocutionary acts, and to investigate which speech-act arguments are correct for a given mode. It is deductively correct arguments for factual illocutionary acts which are the same as the correct arguments for ideal-completion illocutionary acts.

Department of Philosophy
and Center for Cognitive Science University at Buffalo, SUNY
E-mail: kearns@buffalo.edu

REFERENCES

- Kit Fine. 1996. "Vagueness, Truth and Logic," *Vagueness, A Reader*, R. Keefe and P. Smith (eds.). Cambridge, Massachusetts: The MIT Press, 119–150.
- John T. Kearns. 2006. "Conditional Assertion, Denial, and Supposition as Illocutionary Acts," *Linguistics and Philosophy* 29, 455–485.
- John T. Kearns. 2007. "An Illocutionary Logical Explanation of the Liar Paradox," *History and Philosophy of Logic* 28, 31–66.
- John T. Kearns. 2008. "Illocutionary Origins of Familiar Logical Operators," *The Logica Yearbook 2007*. Michal Peliš ed., Filosofia, Prague, 55–66.
- Daniel Vanderveken and John Searle. 1985. *Foundations of Illocutionary Logic*. Cambridge: Cambridge University Press.