

Performatives in a Rationally Based Speech Act Theory*

Philip R. Cohen
Artificial Intelligence Center
and
Center for the Study of Language and Information
SRI International
333 Ravenswood Ave.
Menlo Park, CA 94025
and
Hector J. Levesque[†]
Department of Computer Science
University of Toronto

Abstract

A crucially important adequacy test of any theory of speech acts is its ability to handle performatives. This paper provides a theory of performatives as a test case for our rationally based theory of illocutionary acts. We show why "I request you..." is a request, and "I lie to you that p" is self-defeating. The analysis supports and extends earlier work of theorists such as Bach and Harnish [1] and takes issue with recent claims by Searle [10] that such performative-as-declarative analyses are doomed to failure.

*This paper was made possible by a contract from ATR International to SRI International, by a gift from the Systems Development Foundation, and by a grant from the Natural Sciences and Engineering Research Council of Canada. The views and conclusions contained in this document are those of the authors and should not be interpreted as representative of the official policies, either expressed or implied, of ATR International, the Systems Development Foundation, or the Canadian government.

[†]Fellow of the Canadian Institute for Advanced Research.

1 Introduction

There is something special about performative sentences, sentences such as "I promise to return": uttering them makes them true. How and when is this possible? Not all verbs can be uttered in the first-person present tense and thereby make the sentence true. In general, the successful verbs seem to correspond to those naming illocutionary acts, but not to perlocutionary ones such as "frighten." But, even some illocutionary verbs cannot be used performatively: e.g., "I lie to you that I didn't steal your watch" is self-defeating [12]. So, which verbs can be used performatively, and in Searle's words [10], "how do performatives work?"

Any theory of illocutionary acts needs to provide a solution to questions such as these. But, such questions are not merely of theoretical interest. Natural language database question-answering systems have been known to receive performative utterances [14], dialogue systems that recognize illocutionary acts (e.g., [6]) will need to infer the correct illocutionary force to function properly, dialogue translation systems [5] will have to cope with markers of illocutionary

force that function performatively (e.g., sentence final particles in Japanese), and proposals for “agent-oriented programming languages” [7, 13], as well as Winograd and Flores’ [15] COORDINATOR system, are based on performative communication. For all these systems, it is important to understand the semantics and pragmatics of such communicative acts, especially their intended effects. To do so, one needs a full theory of illocutionary acts, and a formal theory that predicts how utterances can be made true by uttering them.

The currently accepted theory of performatives is that they are in fact assertions, hence true or false, and additionally constitute the performance of the named illocutionary act, in the same way as an indirect reading of an illocutionary act is obtained from the direct illocutionary act. That is, the named illocutionary act is derived from the assertion as an indirect speech act. The most compelling defense of this performative-as-assertion analysis that we are aware of is that of Bach and Harnish [1], who address many of the linguistic phenomena discussed by Sadock [9], but who, we believe, have misanalyzed indirect speech acts. However, in a recent paper, Searle [10] forcefully criticizes the performative-as-assertion approach on the following grounds:

- Assertions commit the speaker to the truth of what is asserted
- Performative statements are self-referential
- “An essential feature of any illocutionary act is the intention to perform that act”

Searle claims that accounts based on self-referential assertions are “doomed to failure” because one cannot show that being committed to having the intention to be performing the named illocutionary act entails that one in fact has that intention. Moreover, he questions that one should derive the named illocutionary act from an assertion, rather than vice-versa. However, Searle has imparted into Bach and Harnish’s theory his notion of assertions as commitments to the truth without providing a precise analysis of commitment. What may be doomed to failure is any at-

tempt to base an analysis of performatives on such a theory of assertions.

This paper provides a formal analysis of performatives that treats them as declarative utterances, not initially as assertions, does not succumb to Searle’s criticisms, and does not require an entirely new class of illocutionary acts (the “declarations”) as Searle and Vanderveken [12] have proposed. The analysis is offered as another adequacy criterion for our theory of illocutionary acts. That theory, more fully explicated in [3], is based on an analysis of the individual rational balance agents maintain among their beliefs, goals, intentions, commitments, and actions [2].

As desiderata for the theory of performatives, we demonstrate that the analysis meets two properties:

- A sincere utterance of “I request you to open the door” is both a request and an assertion, yet neither illocutionary act characterization is derived from the other.
- “I lie that the door is open” is self-defeating.

Briefly, the ability to capture performatives is met almost entirely because such utterances are treated as indicative mood utterances, and because illocutionary acts are defined as attempts. Since attempts depend on the speaker’s beliefs and goals, and these mental states are introspectable in our theory if a speaker sincerely says, for example, “I request you to open the door,” he must believe he did the act with the requisite beliefs and goals. Hence, the utterance is a request.

To meet the desiderata we need first to present, albeit briefly, the theory of rational interaction, the treatment of declarative mood utterances, and then the illocutionary act definitions for requesting and asserting. Finally, we combine the various analyses natural language processor’s task by making explicit the intended word sense of the action, and by reducing the combinatorics inherent in determining the attachment of the prepositional phrases.

2 Abbreviated theory of rational action

Below, we give an abbreviated description of the theory of rational action upon which we erect a theory of intention. The theory is cast in a modal logic of belief, goal, action, and time. Further details of this logic can be found in [2].

2.1 Syntax

The language we use has the usual connectives of a first-order language with equality, as well as operators for the propositional attitudes and for talking about sequences of events: (BEL x p) and (GOAL x p) say that p follows from x 's beliefs or goals (a.k.a choices) respectively; (AGT x e) says that x is the only agent for the sequence of events e ; $e_1 \leq e_2$ says that e_1 is an initial subsequence of e_2 ; and finally, (HAPPENS a) and (DONE a) say that a sequence of events describable by an action expression a will happen next or has just happened, respectively. Versions of HAPPENS and DONE specifying the agent (x) are also defined.

An action expression here is built from variables ranging over sequences of events using the constructs of dynamic logic [4]: $a;b$ is action composition; $a|b$ is nondeterministic choice; $a||b$ is concurrent occurrence of a and b ; $p?$ is a test action; and finally, a^* is repetition. The usual programming constructs such as IF/THEN actions and WHILE loops, can easily be formed from these. Because test actions occur frequently in our analysis, yet create considerable confusion, read $p?;a$ as "action a occurring when p holds," and for $a;p?$, read "action a occurs after which p holds." We use e as a variable ranging over sequences of events, and a and b for action expressions.

We adopt the following abbreviations and domain predicates.

$$(\text{BEFORE } a \ p) \stackrel{\text{def}}{=} (\text{DONE } p?;a)^1$$

$$(\text{AFTER } a \ p) \stackrel{\text{def}}{=} (\text{HAPPENS } a;p?)$$

$$\diamond p \stackrel{\text{def}}{=} \exists e (\text{HAPPENS } e;p?).$$

$$(\text{LATER } p) \stackrel{\text{def}}{=} \neg p \wedge \diamond p.$$

¹This differs from the BEFORE relation described in [3], which is here labelled PRIOR.

$$\Box p \stackrel{\text{def}}{=} \neg \diamond \neg p.$$

$$(\text{PRIOR } p \ q) \stackrel{\text{def}}{=} \forall c (\text{HAPPENS } c;q?) \supset \exists a (a \leq c) \wedge (\text{HAPPENS } a;p?).$$

The proposition p will become true no later than q .

$$(\text{KNOW } x \ p) \stackrel{\text{def}}{=} p \wedge (\text{BEL } x \ p).$$

(IMPERATIVE s) means that sentence s is an imperative.

(DECLARATIVE s) means that sentence s , a string of words, is a declarative.

(MAIN-VERB $s \ v$), (TENSE $s \ \textit{tense}$), (COMPLEMENT $s \ s'$), (D-OBJECT $s \ np$), (SUBJECT $s \ np$), are all syntactic predicates intended to have the obvious meanings.²

(TRUE $s \ e$) means that sentence s is true with respect to some event sequence e (which we will say has just been done.)

(REFERS $np \ x \ e$) means that noun phrase np refers to thing x with respect to event e .

(FULFILL-CONDS $s \ e \ e'$) means that event e fulfills the satisfaction conditions, relative to event e' , that are imposed by sentence s .³ For example, if s is "wash the floor," e would be a floor-washing event.

2.2 Assumptions

The model we are developing embodies various assumptions constraining beliefs and choices (goals). First, BEL has a "weak S5" semantics, and GOAL has a "system K" semantics.⁴ Among the remaining assumptions, the following will be used in this paper.⁵

Beliefs imply choice:

$$\models (\text{BEL } x \ p) \supset (\text{GOAL } x \ p).$$

²Feel free to substitute your favorite syntactic predicates.

³TRUE REFERS, and FULFILL-CONDS are just placeholders for semantic theories of truth, reference, and the meanings of imperatives, respectively. Their last event arguments would be used only in the interpretation of indexicals.

⁴See other work of ours [2] for a full model theory.

⁵In other words, we only deal with semantic structures where these propositions come out true.

This means that agents choose amongst worlds that are compatible with their beliefs.

Goals are known:

$$\models (\text{GOAL } x \text{ } p) \equiv (\text{BEL } x (\text{GOAL } x \text{ } p)).$$

Memory:

$$\models (\text{DONE } x (\text{BEL } x \text{ } p)?; e) \equiv (\text{BEL } x (\text{DONE } x (\text{BEL } x \text{ } p)?; e)).$$

That is, agents remember what their beliefs were.

3 Individual Commitments and Intentions

To capture one grade of commitment that an agent might have toward his goals, we define a persistent goal, P-GOAL, to be one that the agent will not give up until he thinks certain conditions are satisfied. Specifically, we have

Definition 1 (P-GOAL $x \text{ } p \text{ } q$) $\stackrel{\text{def}}{=}$

- (1) $(\text{BEL } x \neg p) \wedge$
- (2) $(\text{GOAL } x (\text{LATER } p)) \wedge$
- (3) $[\text{KNOW } x$
 $(\text{PRIOR } [(\text{BEL } x \text{ } p) \vee (\text{BEL } x \Box \neg p) \vee (\text{BEL } x \neg q)]$
 $\neg [\text{GOAL } x (\text{LATER } p)])]$.

That is, the agent x believes p is currently false, chooses that it be true later, and knows that before abandoning that choice, he must either believe it is true, believe it never will be true, or believe q , an escape clause (used to model subgoals, reasons, etc.) is false.

Intention is a species of persistent goal. We analyze two kinds of intentions, those to do actions and those to achieve propositions. Accordingly, we define INTEND_1 and INTEND_2 to take action expressions and propositions as arguments, respectively.

Definition 2 *Intention:*

$$(\text{INTEND}_1 \text{ } x \text{ } a \text{ } q) \stackrel{\text{def}}{=} (\text{P-GOAL } x [\text{DONE } x (\text{BEL } x (\text{HAPPENS } a))?; a] \text{ } q).$$

$$(\text{INTEND}_2 \text{ } x \text{ } p \text{ } q) \stackrel{\text{def}}{=} (\text{P-GOAL } x \exists e [\text{HAPPENS } x (\text{BEL } x \exists e' (\text{HAPPENS } x \text{ } e'; p?))?; e; p?] \text{ } q)$$

Intending to do an action a or achieve a proposition p is a special kind of commitment (i.e., persistent goal) to having done the action a or having achieved p .⁶ However, it is not a simple commitment to having done a or $e; p?$ for that would allow the agent to be committed to doing something accidentally or unknowingly. Instead, we require that the agent be committed to arriving at a state in which he believes he is about to do the intended action next.

This completes a brief discussion of the foundational theory of intention and commitment. Next, we proceed to define the more specific concepts needed for analyzing communicative action.

4 Utterance Events

We begin the analysis of utterance events by adopting a Gricean correlation of an utterance's features (e.g., syntactic mood or sentence-final particles in Japanese) with the speaker's mental state, termed a "core attitude" in [3, 8]. Very roughly, a declarative utterance s will be correlated with the speaker's believing the uttered sentence is true, and an imperative utterance will be correlated with the speaker's wanting the addressee to do some action that fulfills the conditions imposed by the sentence. Let us notate these correlations as:

$$\text{DECLARATIVE} \Rightarrow (\text{BEL } x (\text{TRUE } s \text{ } e))$$

$$\text{IMPERATIVE} \Rightarrow (\text{GOAL } x \diamond \exists e' (\text{DONE } y \text{ } e') \wedge (\text{FULFILL-CONDS } s \text{ } e' \text{ } e))$$

We formalize this notation below.

Someone who thinks he is observing an utterance event will come to believe the speaker is in the correlated mental state, unless he has other beliefs to the contrary. For example, if the observer thinks the speaker is lying, he believes that the speaker does not believe the uttered sentence is true. But, because he may think the speaker takes himself to be especially convincing, the observer may still believe that the speaker thinks the observer is deceived. Hence, he would believe the

⁶For simplicity, we omit here one condition from the definition of INTEND_2 in [2].

speaker thinks that he thinks the speaker believes p .

This type of reasoning can continue to further levels. In general, if an utterance is produced when there are no countervailing observer beliefs at a certain level of nesting, then the result will be, at the given level of nesting, that the speaker is taken to be in the correlated mental state [8]. To be able to state such conditions, we need to be able to refer easily to what a person x believes about what y believes about what x believes etc., to arbitrary depths. To do so, we use the notion of ABEL.

Definition 3 $(ABEL\ n\ x\ y\ p) \stackrel{def}{=} \underbrace{(BEL\ x\ (BEL\ y\ (BEL\ x\ \dots (BEL\ x\ p) \dots)) \dots)}_n$

That is, ABEL characterizes the n th alternating belief between x and y that p , built up "from outside in," i.e., starting with x 's belief that p . On this basis, one can define unilateral mutual belief — what one agent believes is mutually believed — as follows.

Definition 4 $(BMB\ x\ y\ p) \stackrel{def}{=} \forall n (ABEL\ n\ x\ y\ p)$

In other words, $(BMB\ x\ y\ p)$ is the infinite conjunction $(BEL\ x\ p) \wedge (BEL\ x\ (BEL\ y\ p)) \wedge \dots$. Finally, we define mutual belief and mutual knowledge as follows.

Definition 5 $(MB\ x\ y\ p) \stackrel{def}{=} (BMB\ x\ y\ p) \wedge (BMB\ y\ x\ p)$
 $(MK\ x\ y\ p) \stackrel{def}{=} p \wedge (MB\ x\ y\ p)$.

Utterance events can produce effects at any (or no) level of alternating belief. For example, the speaker may not be trying to communicate anything to an intended observer. Illocutionary acts will be defined to require that the speaker intend to produce BMBs. In what follows, it is important to keep in mind the distinction between utterance events and full-blooded communicative acts.

4.1 Notation for Describing Utterance Events

We now provide a formal notation for this correlation of utterance form and the speaker's mental

state as a kind of default axiom (cf. [8]). First, we specify who is speaking ($spkr$), who is observing (obs , which includes the speaker and addressee, but also others), who is being addressed ($addr$), and what kind of sentence (s) has been spoken (indicated by Φ). We shall assume that everyone knows that a given utterance is of a given syntactic type (e.g., declarative), that speakers and addressees are observers, and that observers are known by all to be observing.⁷

Definition 6 $\Phi \Rightarrow \alpha \stackrel{def}{=} \forall\ spkr,\ obs,\ addr,\ e,\ s,\ n$
 $(KNOW\ obs\ (DONE\ spkr\ e) \wedge$
 $(UTTER\ spkr\ addr\ s\ e) \wedge (\Phi\ s)) \wedge$
 $\sim (ABEL\ n\ obs\ spkr$
 $(BEFORE\ e$
 $\sim (GOAL\ spkr$
 $[AFTER\ e$
 $(KNOW\ addr$
 $(BEFORE\ e\ \alpha)])) \supset$
 $(ABEL\ n\ obs\ spkr$
 $(BEFORE\ e$
 $\alpha \wedge (GOAL\ spkr$
 $[AFTER\ e$
 $(KNOW\ addr$
 $(BEFORE\ e\ \alpha)])) \supset$

That is, $\Phi \Rightarrow \alpha$ is an abbreviation for a quantified implication roughly to the effect that if an observer obs knows that e was just done, where e was an uttering to addressee $addr$ of a sentence s in syntactic mood Φ , and obs does not believe that e was done when the speaker did not want the addressee to come to know that the core speaker-attitude α associated with utterances of that type held, then obs believes that the speaker in fact wanted the addressee to know that α , and so he, the observer, believes that α held just prior to the utterance. The notation states that at each level of alternating belief for which the antecedent holds, so does the consequent. The symbol ' \Rightarrow ' can now be understood as a textual-replacement "macro" operator.

Since these correlations are of the form $\forall n P(n) \supset Q(n)$, they imply $\forall n P(n) \supset \forall n Q(n)$.

⁷The case of unseen observers is straightforward, but omitted here.

As we quantify over the positive integers indicating levels of alternative belief, we can derive the conclusion that under certain circumstances, *addr* thinks it is mutually believed (in our notation, BMB'ed) that the speaker *spkr* wants *addr* to know α was true.

Notice that right after the utterance, we are concerned with what mental state the observer thinks the speaker chose to bring about in the observer with that utterance. That is, the condition on utterance events involves the speaker's wanting to get the observer to know something. Without this temporal dimension, our performative analysis would fail. The analysis of performatives will say that after having uttered such a sentence, or while uttering it, the speaker believes he has just done or is doing the named illocutionary act. Typically, prior to uttering a performative, the speaker has not just performed that speech act, and so he would believe his having just done so is false. So, if the condition on utterance events in Domain Axiom 1A involved only what the speaker believed or wanted to be true prior to the utterance, rather than after, all performatives would fail to achieve the observer's coming to believe anything.

We can now state the correlation between utterance form and a speaker's mental state as a domain axiom.

Domain Axiom 1 *Declaratives and Imperatives:*

A. $\models \text{DECLARATIVE} \Rightarrow (\text{BEL } \text{spkr} (\text{TRUE } s \ e))$

B. $\models \text{IMPERATIVE} \Rightarrow (\text{GOAL } x \ \diamond \exists e' (\text{DONE } y \ e') \wedge (\text{FULFILL-CONDS } s \ e' \ e))$

Below, we present our definitions of illocutionary acts. Further justification can be found in [3].

5 Illocutionary Acts as Attempts

Searle [11] points out that an essential condition for requesting is that the speaker be attempting to get the addressee to perform the requested action. We take this observation one step further and define all illocutionary acts as attempts, hence de-

fining in terms of the speaker's mental states. Attempts involve both types of goal states, GOAL (merely chosen) and INTEND (chosen with commitment), as noted below.

Definition 7 $\{\text{ATTEMPT } x \ e \ p \ q \ t_1\} \stackrel{\text{def}}{=} t_1?; \{(\text{BEL } x \ \sim p \wedge \sim q) \wedge (\text{INTEND}_1 \ x \ t_1?; e; p? \ (\text{GOAL } x \ \diamond q)) \wedge (\text{GOAL } x \ \diamond q)\}?: e$

That is, an attempt to achieve *q* via *p* is a complex action expression in which *x* is the agent of event *e* at time *t*₁, and prior to *e* the agent believes *p* and *q* are both false, chooses that *q* should eventually be true, and intends, relative to that choice, that *e* should produce *p*. So, *q* represents some ultimate goal that may or may not be achieved by the attempt, while *p* represents what it takes to make an honest effort.

5.1 Definitions of Request and Assert

To characterize a request or, for that matter, any illocutionary action, we must decide on the appropriate formulas to substitute for *p* and *q* in the definition of an attempt. We constrain illocutionary acts to be those in which the speaker is committed to understanding, that is, to achieving a state of BMB that he is in a certain mental state. Below is a definition of a speaker's requesting an addressee to achieve *p*.

Definition 8 $\{\text{REQUEST } \text{spkr } \text{addr } e \ p \ t_1\} \stackrel{\text{def}}{=} \{ \text{ATTEMPT } \text{spkr } e \ [\text{BMB } \text{addr } \text{spkr} \ (\text{BEFORE } e \ (\text{GOAL } \text{spkr} \ \diamond p \wedge \ [\text{AFTER } e \ (\text{INTEND}_2 \ \text{addr } p \ [[(\text{GOAL } \text{spkr} \ \diamond p) \wedge (\text{HELPFUL } \text{addr } \text{spkr})]]]]]] \exists e' (\text{DONE } \text{addr } e'; p?) \ t_1 \}$

That is, event *e* is a request at time *t*₁ if it is an attempt at that time to get the addressee to

achieve some condition p while being committed to making public that the speaker wanted: first, that p eventually be achieved; and second, that the addressed party should intend to achieve it relative to the speaker's wanting it achieved and relative to the addressee's being helpfully disposed towards the speaker.

The illocutionary act of asserting will be defined as an attempt to make the speaker's believing the propositional content mutually believed.

Definition 9 $\{\text{ASSERT } \text{spkr } \text{addr } \text{e } \text{p } \text{t}_1\} \stackrel{\text{def}}{=} \{ \text{ATTEMPT } \text{spkr } \text{addr } \text{e} \text{ [BMB } \text{addr } \text{spkr} \text{ (BEFORE } \text{e} \text{ [GOAL } \text{spkr} \text{ (AFTER } \text{e} \text{ [KNOW } \text{addr} \text{ (BEFORE } \text{e} \text{ (BEL } \text{spkr } \text{p}))])])]) \text{ (BMB } \text{addr } \text{spkr} \text{ (BEFORE } \text{e} \text{ (BEL } \text{spkr } \text{p})) \text{ t}_1\} \}$

More precisely, assertions at time t_1 are defined as attempts in which to make an "honest effort," the speaker is committed to getting the addressee to believe that it is mutually believed that the speaker wanted prior to the utterance that the addressee would come to know that the speaker believed p held then. That is, just like a request, an assertion makes public that the speaker wants the addressee to know what mental state he was in. Although he is committed to that, what the speaker has chosen to achieve is not merely to make public his *goal* that the addressee know what mental state he was in, but to make public that he was in fact in that state of believing p . For an INFORM, the speaker would choose to achieve (KNOW *addr* p).

6 Performatives

To illustrate how performatives work, we show when both assertions and requests can be derived from the utterance of the performative "I request you to <do act>." The important point to notice here is that we have not had to add to our machinery; performative utterances will be treated

exactly as declarative utterances, with the exception that the content of the utterance will make reference to an utterance event.

6.1 Request Reports

Let us characterize the truth conditions of the family of declarative sentences " x requests y to (imperative sentence s')." Let s be such a sentence. Let α be $\exists e_1(\text{DONE } y \text{ e}_1) \wedge (\text{FULFILL-CONDS } s' \text{ e}_1 \text{ e})$. We ignore most syntactic considerations and indexicality for reasons of space.

Domain Axiom 2 Present tense requests

$$\models \forall x, y, e, t_1, (\text{DONE } t_1?; e) \wedge (\text{SUBJECT } s \text{ } x) \wedge (\text{D-OBJECT } s \text{ } y) \wedge (\text{REFERS } x \text{ } x \text{ } e) \wedge (\text{REFERS } y \text{ } y \text{ } e) \supset (\text{TRUE } s \text{ } e) \equiv (\text{DONE } x \text{ } \{\text{REQUEST } x \text{ } y \text{ } e \alpha \text{ } t_1\})$$

That is, if event e is happening and the sentence s is a present tense declarative sentence whose main verb is "request," whose subject x refers to person x , and whose direct object y refers to person y , then the sentence is true iff x is requesting the addressee y to fulfill the conditions of imperative sentence s' . A bare present (or present progressive) tense sentence is true when the event being described is contemporaneous with the event of uttering it.⁸ This definition applies equally well to "John requests Mary to ..." as it does when I utter "I request you to ..." For the former, such sentences are likely to be narrations of ongoing events.⁹ For the latter, the event that is happening that makes the utterance true is the speaker's uttering of the sentence.

By our definition of request, for x to request y to achieve p , x has to attempt to get y to do some action intentionally to fulfill the sentence s' , by making that goal mutually believed between them. Thus, to say x requested y to do something is only to say that x had the right beliefs, goals, and intentions.

⁸Searle [10] correctly points out that performatives can be uttered in the passive, and in the first-person plural.

⁹We are ignoring the habitual reading of bare present tense sentences because we do not have a semantics for them.

6.2 Performatives Used as Requests

Next, we treat performative sentences as declaratives. This means that the effects of uttering them are described by Domain Axiom 1A. We sketch below a proof of a general theorem regarding performative requests, with s being the declarative sentence “I request you to (imperative sentence s_1), and α being $\exists e_1(\text{DONE addr } e_1) \wedge (\text{FULFILL-CONDS } s_1 e_1 e)$. We take the uttering of a sentence to be a unitary utterance event.

Theorem 1 *A Performative Request*

$$\begin{aligned} & \models \forall \text{ spkr, addr, e, n, } t_1, \\ & (\text{MK spkr addr } (\text{DONE spkr } t_1?;e) \wedge \\ & \quad (\text{UTTER spkr addr e s})) \wedge \\ & (\text{BEFORE } t_1?;e \\ & \quad (\text{GOAL spkr} \\ & \quad \quad [\text{AFTER } t_1?;e \\ & \quad \quad \quad (\text{KNOW addr} \\ & \quad \quad \quad \quad [\text{BEFORE } t_1?;e \\ & \quad \quad \quad \quad \quad (\text{BEL spkr } (\text{TRUE s e}))]])) \supset \\ & (\text{DONE } \{\text{REQUEST spkr addr e } \alpha t_1\}) \end{aligned}$$

That is, we need to show that if the sentence “I request you to <imperative sentence>” has just been uttered at time t_1 sincerely, i.e., when the speaker wanted the addressee to know that he believed the sentence was true, then a direct request has taken place at t_1 .

Proof sketch: Essentially, one invokes the domain axiom for declaratives at the first level of ABEL, entailing that the speaker believes that he believes that he has just done a REQUEST. Then, one expands the definition of REQUEST into an ATTEMPT, and then into its parts. The definition of ATTEMPT is based on BEL, GOAL and INTEND, the first two of which are obviously introspectable. That is, if one believes one has them one does, and vice-versa. Hence, by the memory assumption, the speaker actually had them prior to the utterance. More critically, intending to act at time t_1 is also introspectable at time t_1 because agents know what they are doing at the next instant and because there is no time to drop their commitment [2]. Thus, one can repackage these mental states up into an ATTEMPT and then a REQUEST.

6.3 Performatives Used as Assertions

We have shown that the speaker of a sincere performative utterance containing an illocutionary verb has performed the illocutionary act named by that verb. Under somewhat stronger conditions, we can also prove that the speaker has made an assertion. As before, let s be “I request you to <imperative sentence>.”

Theorem 2 *Performatives Used as Assertions*

$$\begin{aligned} & \models \forall \text{ spkr, addr, e, n, } t_1, \\ & (\text{MK spkr addr } (\text{DONE spkr } t_1?;e) \wedge \\ & \quad (\text{UTTER spkr addr e s})) \wedge \\ & [\text{BEFORE } e \\ & \quad (\text{BEL spkr} \\ & \quad \quad [\text{AFTER } e \\ & \quad \quad \quad \forall n \sim (\text{ABEL n addr spkr} \\ & \quad \quad \quad \quad (\text{BEFORE } e \\ & \quad \quad \quad \quad \quad \sim (\text{GOAL spkr} \\ & \quad \quad \quad \quad \quad \quad [\text{AFTER } e \\ & \quad \quad \quad \quad \quad \quad \quad (\text{KNOW addr} \\ & \quad \quad \quad \quad \quad \quad \quad \quad [\text{BEFORE } e \\ & \quad \quad \quad \quad \quad \quad \quad \quad \quad (\text{BEL spkr } (\text{TRUE s e}))]]))]]] \end{aligned}$$

This default condition says that before the utterance, the speaker believed there would be no addressee belief after the utterance event (at any level n) to the effect that prior to that event the speaker did not want the addressee to come to know that the speaker believed (TRUE $s e$). Given Domain Axiom 1A, and the fact that BEL entails GOAL, this suffices to entail the definition of assertion. Notice that whereas requesting was derived in virtue of the content of the utterance, an assertion was derived by default assumptions regarding lack of belief in the speaker’s insincerity.

7 ‘Lie’ is not a performative

Some illocutionary verbs such as “lie,” “hint,” “insinuate,” cannot be achieved performatively. The following analysis shows a general model for why such verbs naming covert acts cannot be performatively achieved.

A reasonable definition of lying is the following complex action:

Definition 10 $\{\text{LIE } \text{spkr } \text{addr } \text{e } \text{p}\} \stackrel{\text{def}}{=} (\text{BEL } \text{spkr } \sim \text{p})?; \{\text{ASSERT } \text{spkr } \text{addr } \text{e } \text{p } \text{t}_1\}$

That is, a lie is an assertion performed when the speaker believes the propositional content is false.

For “I lie to you that the door is open” to be a successful performative utterance, it would have to be true that the utterance is a lie. We would have to show that the uttering of that declarative sentence results in a lie’s having been done. More generally, we provide a putative statement of the truth conditions of “ x lies to y that <declarative sentence s' >.” Call the main sentence s .

Domain Axiom 3 *Supposed Truth Conditions for Performative Lying*

$$\models \forall e, x, y, t_1. (\text{DONE } t_1?; e) \wedge (\text{REFERS } x \ x \ e) \wedge (\text{REFERS } y \ y \ e) \supset$$

$$(\text{TRUE } s \ e) \equiv$$

$$(\text{DONE } \{\text{LIE } x \ y \ e \ (\text{TRUE } s' \ e) \ t_1\})$$

That is, if s and s' are declarative sentences of the appropriate syntactic form, x refers to x and y refers to y , then s is true iff in performing it at time t_1 , x was lying that sentence s' is true.

So we can prove the following. Let the sentence s be “I lie to you that <declarative sentence s' >.”

Theorem 3 *Lies are not performative*

$$\models \forall \text{spkr}, \text{addr}, e, n$$

$$(\text{MK } \text{spkr } \text{addr } [(\text{DONE } \text{spkr } t_1?; e) \wedge (\text{UTTER } \text{spkr } \text{addr } e \ s)]) \supset$$

$$\sim (\text{DONE } \{\text{LIE } \text{spkr } \text{addr } e \ (\text{TRUE } s \ e) \ t_1\})$$

In other words, you cannot perform a lie by saying “I lie that ...”

Proof Sketch: Assume that it is mutually believed that the speaker has uttered declarative sentence s . Now, apply Domain Axiom 1A. By assumption, the first conjunct of the antecedent holds. There are then two cases to consider. First, assume (**) the second conjunct holds (say, at level $n = 1$), i.e., the addressee does not believe the speaker did not want him to know that he believed s' was true. In virtue of the supposed truth conditions on lying, spkr would have to have been lying. By expanding its definition, and using the memory and introspectability properties of BEL,

GOAL, and INTEND the addressee can conclude that, before the utterance, the speaker wanted him not to know that the speaker believes that in uttering S , he was lying. But, this contradicts the assumption (**). Since the speaker in fact uttered the sentence, that assumption is false, and the addressee believes the speaker did not in fact want him to know that he believed the sentence was true. This renders impossible the intentions to be achieved in asserting, which are constitutive of lying as well.

Now, assume (**) is false, so the addressee in fact believes the speaker did not want him to know that s' was true. Again, this immediately makes the speaker’s intentions in asserting, and hence lying, impossible to achieve. So, in neither case is the utterance a lie. If the addressee believes the speaker is a competent speaker of the language, the speaker must have intended some other interpretation.

8 Conclusion

Requesting works well as a performative verb because requesting requires only that the agent has made an attempt, and need not have succeeded in getting the hearer to do the requested action, or even to form the right beliefs. Some verbs cannot be used performatively, such as “frighten,” because they require something beyond a mere attempt. Hence, such verbs would name action expressions that required a particular proposition p be true after the utterance event. When the utterance event does not guarantee such a p , the use of the performative verb will not be possible.

On the other hand, certain utterances (performative or not), when performed by the right people in the right circumstances, make certain institutional facts hold. So, when a clergyman, judge, or ship captain says “I now pronounce you husband and wife,” the man and woman in question are married. In our framework, there would be a domain axiom whose antecedent characterizes the circumstances, participants, and nature of the utterance event, and whose consequent asserts that an institutional fact is true. The axiom is justified not by the nature of rational action, but by the existence of an institution. Such utterances could be

made with a performative prefix provided such attempts are made into successes by the institution.

This paper has shown that treating performative utterances as declarative sentences is a viable analysis, in spite of Searle's criticisms. The performative use of an illocutionary verb is self-guaranteeing when the named illocutionary act consists in the speaker's making an attempt to make public his mental state. In such cases, if the speaker thinks he has done so, then he has. However, we do not derive the named illocutionary act from the assertion, nor vice-versa. Instead, both derivations may be made from the utterance event, but the assertive one is in fact harder to obtain as it has extra conditions that need to be satisfied.

References

- [1] K. Bach and R. Harnish. *Linguistic Communication and Speech Acts*. M. I. T. Press, Cambridge, Massachusetts, 1979.
- [2] P. R. Cohen and H. J. Levesque. Intention is choice with commitment. *Artificial Intelligence*, 42(3), 1990.
- [3] P. R. Cohen and H. J. Levesque. Rational interaction as the basis for communication. In P. R. Cohen, J. Morgan, and M. E. Pollack, editors, *Intentions in Communication*. M.I.T. Press, Cambridge, Massachusetts, in press.
- [4] D. Harel. *First-Order Dynamic Logic*. Springer-Verlag, New York City, New York, 1979.
- [5] K. Kogure, H. Iida, K. Yoshimoto, H. Maeda, M. Kume, and S. Kato. A method of analyzing Japanese speech act types. In *Second International Conference on Theoretical and Methodological Issues in Machine Translation of Natural Languages*, 1986.
- [6] D. J. Litman and J. F. Allen. A plan recognition model for subdialogues in conversation. Technical report, Department of Computer Science, Univ. of Rochester, Rochester, New York, November 1984.
- [7] J. McCarthy. ELEPHANT: a programming language based on speech acts. Unpublished ms., Dept. of Computer Science, Stanford University, 1989.
- [8] C. R. Perrault. An application of default logic to speech act theory. In P. R. Cohen, J. Morgan, and M. E. Pollack, editors, *Intentions in Communication*. M.I.T. Press, Cambridge, Massachusetts, in press.
- [9] J. Sadock. *Toward a Linguistic Theory of Speech Acts*. Academic Press, New York, 1984.
- [10] J. Searle. How performatives work. *Linguistics and Philosophy*, 12:535-558, 1989.
- [11] J. R. Searle. *Speech acts: An essay in the philosophy of language*. Cambridge University Press, Cambridge, 1969.
- [12] J. R. Searle and D. Vanderveken. *Foundations of Illocutionary Logic*. Cambridge Univ. Press, New York City, New York, 1985.
- [13] Y. Shoham. Agent oriented programming. Unpublished ms., Dept. of Computer Science, Stanford University, October 1989.
- [14] H. Tennant. Evaluation of natural language processors. Technical Report T-103, Coordinated Science Laboratory, University of Illinois, Urbana, Illinois, November 1980. Ph. D. Thesis.
- [15] T. Winograd and F. Flores. *Understanding Computers and Cognition: A New Foundation for Design*. Ablex Publishing Co., Norwood, New Jersey, 1986.