From the autonomy of syntax to the autonomy of linguistic semantics

Notes on the correspondence between the transparency problem and the relationship problem

DANIEL DOR

Current research on the syntax-semantics interface demonstrates the dramatic extent to which syntactic structures constitute transparent reflections of well-defined semantic regularities. As this paper shows, the empirical results accumulated within this framework strongly suggest that a theoretical distinction should be made between two distinct levels of meaning representation: A level of conceptual meaning on the one hand, and a uniquely linguistic level of meaning — Linguistic Semantics — on the other. The semantic notions and regularities which turn out to determine major syntactic phenomena are best interpreted as belonging to the level of Linguistic Semantics, rather than to the level of conceptual meaning. This view helps characterize language as a unique and functional system — a cognitive system whose function is defined at the level of Linguistic Semantics. It explains the fact, most recently highlighted by Levinson (1997), that the expressive power of language, as a tool for the communication of meanings, is constrained in non-trivial ways.

1. Introduction

In the last decade or so, a fair amount of linguistic research has concentrated on two seemingly unrelated theoretical problems. The first problem — let me

Pragmatics & Cognition Vol. 8(2), 2000. 325–356

dub it the *transparency problem* — has to do with the interface between *syntactic* and *semantic* representations. As I will show below, the empirical results accumulated in this field of research consistently demonstrate the dramatic extent to which syntactic phenomena are determined by semantic regularities. Thus, they strongly indicate that Chomsky's long-standing hypothesis of the autonomy of syntactic structures from meaning regularities should be abandoned in favor of an explicit, semantically-based and empirically-oriented theory of *transparent* meaning–form relations. The second problem — dubbed the *relationship problem* by Pederson and Nuyts (1997) — has to do with the interface between *conceptual* representations and *linguistic semantic* representations. Recent research in this field seems to suggest that a theoretical distinction should be made between (at least) two distinct levels of meaning representation: A level of *conceptual* meaning on the one hand, and a uniquely linguistic level of meaning — *linguistic semantics* — on the other. This view is most comprehensively developed and defended by Bierwisch and Schreuder (1992), and, more recently, by Levinson (1997).

In this paper, I will make the claim that the *transparency problem* and the *relationship problem* are intimately connected to each other: They are, so to speak, two sides of the same coin. As I will show, the semantic notions and regularities which turn out to determine major syntactic phenomena are best interpreted as belonging to the level of linguistic semantics, rather than to the level of conceptual meaning. Thus, research on the transparency problem actually lends vital support to the view of linguistic semantics as an independent level of representation. As a matter of fact, determining linguistic structures seems to be one of the constitutive properties of linguistic semantic regularities — demarcating them from conceptual meanings.

In the next section, I will review some of Levinson's (1997) arguments for the autonomy of linguistic semantics from conceptual meaning. Section 3 will be devoted to the transparency problem: I will discuss Chomsky's hypothesis of the autonomy of syntactic structures, and review the results of three research projects — each of which deals with a set of phenomena which seem to be autonomously syntactic, but upon closer inspection prove to be a transparent reflection of semantic structure. In Section 4, I will explain why the constitutive semantic notions involved in these analyses belong to the level of linguistic semantics. Finally, I will show that the *specific* semantic notions and regularities which turn out to determine

syntactic phenomena in natural languages help *explain* what is probably the major insight of the theory of linguistic semantics: The idea that the level of linguistic semantics constrains the expressive power of language in non-trivial ways.

2. The relationship problem

The idea of linguistic semantics as an autonomous level of semantic representation has been argued for by different scholars, most comprehensively by Levinson (1997). Levinson's major arguments are philosophical and empirical. The philosophical arguments make it clear that there *must* be something right about the idea of linguistic semantics, but I am not sure that all the empirical arguments presented by Levinson provide the supporting evidence which is so crucial for the claim to be established. As I will try to show below, research on the transparency problem provides exactly the type of empirical evidence we need to support Levinson's claim.

Philosophically speaking, the idea of linguistic semantics captures the intuition that language is just one mode of expression and representation of thought among many others, with its own intrinsic properties and limitations:

> … a picture cannot be non-specific about metric properties and shapes, while a linguistic description can hardly avoid being so. In the same way, there are notable failures in the adequacy of linguistic expression: We are helpless when we need to describe faces, smells, contour curves, and so on. These failures of the medium are presumably related both to the existence of other 'languages of thought', other kinds of mental representation in which we think, some more analogue than digital, and to the fact that natural language vocabularies are necessarily crudely general over possible distinctions if they are to remain of learnable size (Levinson 1997: 17).

A very similar point is made, in a different context, by Pinker and Bloom (1990) and by Aitchison (1996). As Levinson rightly claims, this view stands in sharp contrast to Searle's (1969) naive formulation of his *principle of expressibility* — the idea that "whatever can be meant can be said". Searle's principle reflects the view that the functional envelope of language, as a medium of thought expression, is nothing more nor less than the functional envelope of some general *language of thought* which we use to formulate all our ideas. The existence of more than one language of thought (cf. Keller

and Dixon Keller 1996) and more than one medium of expression of thought (cf. McNeill 1997), coupled with the above acknowledgment of the intrinsic limitations of linguistic expression, clearly conveys the message that the functional envelope of linguistic expression *cannot* be equated with the functional envelope of some general language of thought. This, in turn, means that there is a uniquely linguistic level of semantic representation — that determines the functional envelope of language as a specific mode of expression — and that we should attempt to figure out what belongs to this representational level and what does not. This view is reinforced by the fact that language forces a "linearization of thought and the taking of perspective" (Levinson 1997: 20), and by the very existence of *pragmatic inference*. The fact that every utterance necessarily generates intended or unintended implicatures shows that we never get to say exactly what we think.

According to Levinson, this view is supported by the following empirical arguments. First, different languages have both *accidental lexical gaps* and more systematic *missing semantic fields*:

> Tzeltal lacks a word for blue; it has a word *yax*, glossable as 'grue', i.e., green or blue: When Tzeltal speakers say in effect 'The sky is grue' they express a definite thought using the semantically general means at their disposal, in the same way that when I say *Elizabeth is my aunt* in English this does not correspond to a vague thought indifferent to whether Elizabeth is my mother's or father's sister, or the wife of my mother's or my father's brother. English just happens to lack a word for female sibling of my mother… It is not only the 'general vocabulary' that is subject to the vagaries of uneven linguistic coding: many languages lack causal and logical connectives of the familiar sort. Guugu Yimithirr lacks a word for 'if'; you can only express a conditional by saying in effect 'perhaps A, perhaps B', which might mean a host of other things, including 'perhaps A and B'. There is a simple moral: generally speaking, our thought is specific, but its linguistic expression is often necessarily general, non-specific, even imprecise… A language without a conditional construction does not imply that its speakers cannot have thoughts of the kind 'if p, q'; but they cannot exactly express just such a thought in the semantic representations available to them (Levinson 1997: 16–17).

Closely related to this argument is another observation, namely that languages have *language-specific obligatory categories*:

> … I can think 'You are mistaken', and speak my mind directly. A French speaker might have a corresponding thought, but she'll have to add the

> distinction between *tu* and *vous* if she is to voice it. I can think and say *I wish I owned that house*, but a Kwakiutl speaker will have to add the distinction 'house visible/invisible to speaker', since the determiners require such distinctions. I can say *Put the two jars on the table* without worrying about the shape of the jars, but a Tzeltal speaker has different verbs 'put on' for vessels of different shapes, and numeral classifiers that force distinctions between shapes too… These facts allow two immediate conclusions: First, unless we are to maintain that e.g., a Kwakiutl speaker's non-linguistic thoughts are systematically different from our thoughts (so he can't even entertain the idea of a house neutral with respect to visibility), we will have to hold that there is some kind of regimentation, reorganization, and embellishment of the original thought into a fully specified form that matches the grammatical and semantic structure of language. So we have once again an argument for some kind of separation of levels. Secondly, by showing that languages differ, by virtue of their semantico-grammatical structure, in the required nature of that fully specified form, we seem to have an argument of the following form: If CR is universal, SRs various, then it can't be the case that CR = SR (Levinson 1997: 21).

The lexical facts described in the above passages are certainly important, but I am not sure that the arguments presented in these passages actually deliver the goods they are supposed to deliver. This is so, because they seem to be based on the presupposition that speakers of English, French, Tzeltal and Kwakiutl start out with *identical* conceptual representations, and are then forced to transform their identical thoughts into non-identical linguistic representations. Note that the identity of conceptual representation is a crucial component of the argument here, but as far as I can tell, it is far from being an established cognitive fact. On the contrary: It is a completely *open* question — the crux of the *linguistic relativity problem*. If it turns out, for example, that speakers of Tzeltal have a conceptual representation of *yax* which is *different* from English speakers' conceptual representations of either *blue* or *green*, then the empirical basis of Levinson's argument, as far as *color terms* are concerned, immediately collapses. The difference between the two lexicons corresponds to the difference between the two conceptual schemes, and the need for an independent level of linguistic semantics disappears. The same is true for all the other examples: To use the *aunt* example as evidence for the general argument, we need to *show*, not just assume, that speakers of English, and speakers of languages which *do* distinguish between mothers' sisters and fathers' sisters, actually think about

aunts in exactly the same way. To use the 'put on' example, we need to *show* that speakers of Tzeltal do *not* have different concepts of 'putting on' which are correlated with objects of different shapes. At least some of the people who are working on the relativity problem (e.g., Lucy 1996) insist that such conceptual variability *does* exist, in correlation with the facts of linguistic variability. As far as I can tell, Levinson's own work on spatial language and spatial conceptualization (cf. Levinson 1996, 1997) can be interpreted in exactly the same way. Note, moreover, that some amount of conceptual relativity does not immediately lead to radical relativity. Assuming that Tzeltal speakers have a conceptual representation of *yax* which is different from English speakers' representations of *blue* or *green* does *not* amount to the claim that Tzeltal speakers *cannot* in any way entertain the idea of *blue* or *green* — it just means that they do not, *in practice*, do that. In other words, speakers of English and Tzeltal may start out with non-identical thoughts, the one consisting of the concept *blue*, the other consisting of the concept *yax*, which means that their thoughts need not necessarily go through the process of "regimentation, reorganization, and embellishment of the original thought into a fully specified form that matches the grammatical and semantic structure of (their) language" (Levinson 1997: 21). Consequently, I do not see that lexical arguments of this type can be used as evidence for the autonomy of linguistic semantics from conceptual meaning.[1]

Levinson's third argument deals with *indexicals*:

> What is the thought corresponding to *Tomorrow I will leave here*? Could the meaning just be the thought expressed? If meanings were thoughts, then your saying *Tomorrow I will leave here* in New York on 1 September would express the same thought as my saying *Tomorrow I will leave here* in Paris on 2 August. Nor does the utterance *Tomorrow I will leave here* play the same psychic role in the individual as the corresponding thought: if days later I recollect the thought, what I think is not the meaning of *Tomorrow I will leave here*, which would have me ever itinerant. Indexical representations could not play a central role in memory; but that is just the role conceptual representations must play … So what do I retrieve when I retrieve from memory the thought corresponding to *Tomorrow I will leave here*? Could it be the semantical content — now identified as conceptual representation — together with the context of the utterance? If so, the 'language of thought' contains indexicals, which must be explicated by more thoughts, which might also contain more indexicals, and so on in potential infinite regress. The difficulty is that indexicals must be banished from the language of thought, while the substituting expressions

> must retain their cognitive immediacy … This, I take it, is a knock-down argument against identifying SR … with CR (Levinson 1997: 19–20).

If it turned out that indexicals must indeed be "banished from the language of thought", the above argument would certainly be a very strong one, because it would prove that a set of meanings, which *must* be a part of linguistic semantics, *cannot* be represented as such on the conceptual level. As Levinson claims, it would not only prove that linguistic semantics and conceptual structure are separate levels of semantic representation, but also that they are *non-isomorphic*.[2] The question is whether we can find valid demonstrations of this type, where we can show that some meaning component *cannot* be part of conceptual meaning. In principle, we should be looking for three types of cases:

a. cases where a language *disallows* the expression of a meaning, which corresponds to a thought which is conceptually *well-formed* for the speakers of the *same* language;
b. cases where a language is *blind* to a conceptual distinction which is *active* in the conceptualization patterns of the speakers of the same language;
c. cases where a semantic difference between two languages *cannot* be attributed to a parallel conceptual difference.

As I will show below, research on the transparency problem provides evidence of exactly the type we are looking for. Before delving into the empirical issues, however, I would like to discuss the transparency problem itself, especially the one view on this problem which has for the last 40 years been the most influential — Chomsky's hypothesis of the autonomy of syntactic structures. This digression will prove to be useful.

3. The problem of transparency and the autonomy hypothesis

The first thing we need to understand is that Chomsky's autonomy hypothesis is not a mere *existence claim*: It is not the claim that there are some syntactic phenomena in natural languages which are purely structural, and cannot be explained in terms of meaning. This would be a totally self-evident claim, and thus an uninteresting one, because structural phenomena of this type obviously exist: Linear order, for example, is such a phenomenon.[3] The autonomy hypothesis is a much stronger claim. It is the claim that

a significantly large set of core syntactic phenomena in natural languages *cannot* be theoretically *correlated* with a corresponding set of functions, which are formulated in terms of meaning. This, again, is not a mere existence claim. It is not the claim that such non-functional syntactic phenomena exist in natural languages. This, in itself, would also be a self-evident and uninteresting claim, because phenomena of this type are obviously in existence: Languages are full of historical accidents and irregularities. The crux of the autonomy hypothesis is that such phenomena comprise a *large*, *systematic*, *non-accidental* set of facts which lies at the *universal core* of our linguistic knowledge. The autonomy hypothesis takes the non-functionality of syntactic structures to be a *necessary* and *fundamental* property of human language, not a peripheral, accidental property of this or that construction in a specific language.

What this means is, first, that for anybody to be able to maintain, as a constitutive principle, that syntactic knowledge is in its very essence autonomous from meaning-related knowledge, much more should be done than merely point at structural phenomena which do not carry their functional significance on their sleeve. To use such structural pieces of knowledge as evidence for the autonomy hypothesis, at least two points should be established:

i. that an explanation of these phenomena in terms of meaning is not just unavailable at the moment, but somehow *impossible* in principle, and
ii. that the relevant structural phenomena constitute core domains of the linguistic system as a cognitive capacity, not accidental facts about a certain language at a certain point in time in its historical development.

These two points, to be sure, lie well beyond the realm of strict empirical research. They are theoretical issues *par excellence*. Again, the question is not whether there are pieces of grammatical knowledge to be found which cannot be accounted for in functional terms of meaning — there probably are many of those. It is the question of the theoretical *status* of these pieces of knowledge within the overall theory of language. Are they core phenomena or are they peripheral? Do they represent what is essential about language or are they nothing but accidental facts? Do they tell us something of importance about the nature of our genetic endowment for language, as is so often claimed, or do they constitute for the language-acquiring child an unwarranted level of complexity, to be reckoned with on an item-to-item basis?

With this perspective in mind, we may now look at what seems to be the most fundamental property of the autonomy hypothesis: As is so often

the case with such overarching, constitutive scientific principles, the autonomy hypothesis is perfectly *irrefutable*. This fact is rarely appreciated by generative linguists, but it is of the highest importance. The autonomy hypothesis is a *negative* claim. It is the claim that a significantly large set of core structural phenomena is immune to explanation in terms of a meaning-based theory. Any attempt to empirically *refute* a hypothesis of this type should, in principle, provide for *total* coverage of the entire set of structural-syntactic phenomena of, say, a whole language — preferably, of course, of all languages — and show that each and every syntactic phenomenon, which is not accidental, can be correlated with a well-formed functional theory of meaning. As no linguistic theory, regardless of ideological inclination, is anywhere close to a fully explicit description of a single language, let alone of the universal parameters of language as such, the autonomy hypothesis is in no danger of ever being refuted.

Take a moment to think about that. Empirical demonstrations of the functional nature of structural knowledge, as impressive as they may be, *cannot* refute the autonomy hypothesis — they can only bite at the borders of its empirical envelope. There are *always* going to be many more syntactic phenomena which, on the descriptive level, seem to be autonomous from meaning — simply because they have not yet been properly analyzed on the appropriate level of semantics.

Should we accept the autonomy hypothesis, then? Not at all. The crucial point we need to understand is that the fact that the autonomy hypothesis cannot be refuted means that its fate should be decided upon on the basis of empirical *demonstration*: Take some complex, syntactic phenomena, which everybody agrees are significant and non-accidental, and which certainly *seem* to be divorced from any considerations of meaning — and show that they *can* be given an explicit meaning-based explanation. Then, do it again. And again. To the extent that your analyses pass the test of empirical scrutiny, they demonstrate that an explanation of the relevant structural phenomena in terms of meaning is not just possible, but necessary. To the extent that the phenomena uncontroversially constitute core domains of linguistic knowledge, their meaning-based analyses serve to gradually weaken the value of the autonomy hypothesis as a *default assumption*. They strengthen the suspicion that the autonomy hypothesis reflects some stage in the history of the misunderstanding of the nature of meaning, rather than some deep insight into the nature of language. This is the challenge we face

when we work on the transparency problem, the challenge of empirical demonstration.[4]

Empirical demonstrations of this type have been accumulated in the last decade or so by semanticists and lexical semanticists, working on what is sometimes called the *syntax-semantics interface*, especially by those scholars working within the emerging domain of *argument structure*, or *event structure* (cf. Alsina, Bresnan, and Sells 1997; Butt and Geuder 1998; Dowty 1979, 1991; Frawley 1992; Goldberg 1995; Grimshaw 1990; Jackendoff 1983, 1990; Levin 1993; Parsons 1990; Van Valin and LaPolla 1997; and references therein). Some people, myself included, believe that enough results have been accumulated for us to be able to entertain the idea that grammars are *not* autonomous from meaning — and see where this idea can take us. Some other people, of course, do not accept this view. I will present three of these demonstrative cases, not only because I believe they are persuasive, but also because the specific semantic notions involved in the analyses will help us move on to the next step of my argument.

3.1. *Event structure and syntactic elasticity*

Rappaport-Hovav and Levin (1998) examine the way in which the lexical semantics of verbs determines their syntactic behavior. As Rappaport-Hovav and Levin indicate, the idea that a verb's meaning determines the number of arguments it takes, and their syntactic expressions, has generally been accepted as a working hypothesis in the field. A major challenge facing this hypothesis is the fact that some verbs appear in a bewildering range of syntactic contexts, whereas others seem to be extremely constrained in their syntactic behavior. To demonstrate, consider the contrasts between *sweep* and *break*, and between *run* and *go*. *Sweep* and *run* are of the 'elastic' type; *break* and *go* are of the rigidly constrained type:

(1) a. Terry swept the floor.
b. Terry swept.
c. Terry swept the crumbs into the corner.
d. Terry swept the leaves off the sidewalk.
e. Terry swept the floor clean.
f. Terry swept the leaves into a pile.

(2) a. Kelly broke the dishes.
b. *Kelly broke.

c. *Kelly broke the dishes off the table. (meaning: Kelly broke the dishes and as a result they fell off the table.)
d. *Kelly broke the dishes into a pile. (meaning: Kelly made a pile out of the dishes by breaking them.)

(3) a. Pat ran.
b. Pat ran to the beach.
c. Pat ran herself ragged.
d. Pat ran her shoes to shreds.
e. Pat ran clear of the falling rock.
f. The coach ran the athletes around the track.

(4) a. The students went.
b. The students went to the beach.
c. *The jetsetters went themselves ragged.
d. *The runner went his shoes to shreds.
e. *The pedestrian went clear of the oncoming car.
f. *The coach went the athletes around the track.

Note that this variation in syntactic behavior involves not only the number and syntactic type of the complements of the verbs, but also the allowable combinations of these complements. Thus, for example, *sweep* can take a single argument expressing the swept surface ('Terry swept the floor'), but it can only take the argument expressing the swept *object* if followed by a prepositional phrase. Thus, 'Terry swept the leaves off the sidewalk' is grammatical, but 'Terry swept the leaves' is ungrammatical.

How can we explain this diversity in syntactic behavior? Rappaport-Hovav and Levin suggest a semantic account. The first thing to note is the well-known observation that verbs belong to *verb classes*. The members of a verb class characteristically manifest similar properties, both in terms of their semantics and their syntactic behavior. *Sweep*, for example, belongs to the class of *surface contact* verbs, which also includes, among other verbs, *wipe* and *scrub*. These two verbs, which are obviously close in meaning to *sweep*, also appear in the same set of syntactic configurations that it appears in. *Break* is a *change of state* verb, just like *chip*, *crack* and *fracture* (Levin 1993). *Run* is a *manner of motion* verb, as opposed to *go*, which is a verb of *directed motion* (it does not specify the manner of motion, but rather the fact that it is goal oriented). As the names of the verb classes indicate, all the verbs in a specific verb class denote the same *type of event*. They differ

between themselves only with respect to the *idiosyncratic* properties of the events they denote: *run* and *walk*, for example, denote events of moving in a certain manner, and they only differ with respect to the idiosyncratic properties of the manner of movement itself.

A major insight, stemming from research in formal semantics, lexical semantics and the philosophy of language, is that all the different types of events, denoted by the different verb classes, are manifestations of a very constrained set of *basic* event types. Formal details aside, researchers agree that the inventory of basic event types includes no more than four event types: *activities, accomplishments, states* and *achievements*. *Manner of motion* verbs (*run, walk*) and *surface contact* verbs (*sweep, wipe*) denote *activities* of different types: *activities* happen in different ways (manners), and they do *not* necessarily cause any change of state. When we move a sweep back and forth over the surface of a perfectly clean floor, we do not change the state of the floor (we do not clean it by sweeping), but we nevertheless sweep the floor. When we run in place (on a treadmill, for example), we do not change our position, we do not 'go anywhere', but we run all the same. Verbs of *directed motion* (*go, come*) and verbs of *change of state* (*break, crack*) denote *accomplishments* — activities (with specified manners or without them) *resulting in some change of state*. If, for example, Kelly is reported to have broken the dishes, then she is reported to have done something to cause the dishes to break, and the dishes are reported to have broken as a result of that act. Verbs like *find out* and *discover* denote *achievements*, which involve an instantaneous change of state, and do not specify the cause of this change of state. Verbs like *sleep* and *know* denote *states* — situations that exist over time in a static fashion. Rappaport-Hovav and Levin represent this inventory of basic event types in the following fashion (the event type of *accomplishment* is represented in two versions, with and without a specification of the *manner* of the causing activity):

(5) a. Activity: [x ACT⟨manner⟩ (y)]
b. State: [x ⟨STATE⟩]
c. Achievement: [BECOME [x ⟨STATE⟩]]
d. Accomplishment: [[X ACT⟨manner⟩] CAUSE [BECOME [y ⟨STATE⟩]]]
e. Accomplishment: [x CAUSE [BECOME [y ⟨STATE⟩]]]

So, it turns out that those verbs which manifest syntactic elasticity — *surface contact* verbs, verbs of *manner of motion* — are *activity* verbs; the verbs which manifest the more rigid syntactic behavior — verbs of *directed motion*, verbs of *change of state* — are *accomplishment* verbs. The crucial thing to note, is that the extended usages of *sweep* and *run* (in (1c–f) and in (3c–f)), where non-subcategorized objects appear on the surface, all involve the *expansion* of the basic event type of the verbs (an activity) to yield various kinds of *accomplishments*. If Terry swept the crumbs into the corner, then the crumbs changed their state as a result of Terry's sweeping; if Pat ran her shoes to shreds, then Pat's shoes were destroyed as a result of her running. Rappaport-Hovav and Levin claim that this semantic observation is the key to the distinction in syntactic behavior. Informally, they suggest that (a) verbs can be used to mean *more* than they basically do (i.e., verbs can be used with extended meanings) *only* if their basic event type can be augmented to a more complex one, but that on the other hand, (b) verbs cannot be used to mean *less* than they basically do.

As the representations of the basic event types in (5) indicate, activities are *obligatory components* of accomplishments; activities can thus be augmented to accomplishments, and activity verbs can be used as accomplishment verbs (with an additional object indicating the resulting change of state). Accomplishments, on the other hand, *cannot* be augmented to a more complex event type, because they are the *maximally* complex event type in the inventory. This is why accomplishment verbs cannot be used to describe an event of change of state resulting in another change of state (e.g., *'Kelly broke the dishes off the table'). Moreover, accomplishments cannot be used as activity verbs (e.g., *'Kelly broke'), because verbs cannot be used to mean less than they basically do, and the meaning component of the resulting change of state is an obligatory component of the basic meaning of accomplishments. What makes this analysis so interesting is the fact that the meanings of the verbs, as they are captured within the overall framework of event semantics, are not only used to derive the specific distributions of the different verbs across different syntactic configurations, but also to derive such abstract structural properties as *syntactic elasticity* and *syntactic rigidity* — properties which, at least superficially, seem to be purely syntactic.

3.2. *Denominal verbs and syntactic optionality*

Following Hale and Keyser (1992), Kiparsky (1997) discusses the behavioral patterns of two sets of denominal verbs. Verbs of the first type, *locatum* verbs, include *butter, carpet, coat, glove, label, saddle, salt, soap* and many others; they denote acts of *causing things to have something on*: carpeting the floors, salting the food, and so on. Verbs of the second type, *location* verbs, include, among others, *barrel, bench, bottle, cage, corral, garage, imprison* and *pocket;* they denote acts of *causing things to be put in something*: bottling the wine, caging the lion. As Kiparsky demonstrates, speakers of English know much more about the behavioral patterns of these two types of verbs than they could have possibly learned in any direct fashion from their linguistic environment. Thus, for example, the sentences in (6)

(6) a. Bill saddled the horse.
b. Bill corralled the horse.

do not only mean that Bill caused the horse to have a saddle on (6a), and that Bill put the horse into a corral (6b), but also that these activities are done in a *certain* way. (6a) cannot be used to describe "putting a saddle into a basket on the horse's back, or putting it on the wrong part of the horse's anatomy". (6b) implies that the horse is alive; this is why (7a) is a perfectly regular sentence, whereas (7b) is infelicitous:

(7) a. We put the dead horse in the corral.
b. *We corralled the dead horse.

As Kiparsky shows, the constraints on the properties of the events described by denominal verbs are surprisingly specific. For example, it is entirely natural to speak of *flagging a table* or *imprisoning a politician*, but the following usages of the same verbs are ruled out:

(8) a. *The table was filthy, but we flagged it clean and shiny.
b. *The motels were full, but the authorities managed to imprison all the victims of the flood.

Moreover, the process of coining new denominal verbs, which is extremely productive in English, is nevertheless severely constrained. Sentences like 'I bushed some fertilizer' (which is supposed to mean something like 'I put some fertilizer on the bush'), or 'I housed a coat of paint' (which is supposed

to mean something like 'I put a coat of paint on the house') are ruled out as ungrammatical.

More relevant to the topic of our paper, denominal verbs do not participate in the syntactic *unaccusativity alternation*, where the so-called 'semantic object' of a verb, which usually surfaces as the object of the verb (in the transitive constellation), is also capable of surfacing as its subject (in the intransitive constellation). As the following examples demonstrate, this is possible with the deadjectival verbs *redden* and *thin*, but not with the denominal verbs *shelve* and *paint*:

(9) a. John reddened the solution.
b. The solution reddened.

(10) a. John thinned the solution.
b. The solution thinned.

(11) a. John painted the room.
b. *The room painted.

(12) a. John shelved the book.
b. *The book shelved.

Hale and Keyser (1992) set out to provide these phenomena with a purely syntactic explanation. Kiparsky shows that their analysis does not get anywhere close to capturing the relevant facts, and suggests in its stead a semantic analysis, based on a careful examination of the meanings of the predicates. He shows that the process of denominalization is governed by the following semantic constraint: "If an action is named after a thing, it involves a canonical use of the thing" (Kiparsky 1997: 482). According to this semantic constraint, information about the canonical usage of the thing denoted by the noun, becomes part of the meaning of the action denoted by the denominal verb. Thus, a corral is canonically used for housing *live* domesticated animals, so dead horses cannot be corralled. A saddle is canonically put on the horse's back in a *specific fashion*, intended to facilitate riding, so a horse is not saddled if the saddle is put in a basket on its back. People are canonically put in prison as a means of *punishment*, not as a temporary housing solution; and having coats of paint on is *not* a canonical use of a house.

Moreover, Kiparsky shows that the participation of a verb in the unaccusativity alternation is determined by "the nature of the Agent's involvement in the event (denoted by the verb)". Verbs like *paint* and *shelve*

"denote processes requiring the direct initiation and continuous participation of a causing Agent. When John stops painting the wall, the painting stops" (Kiparsky 1997: 495). Thus, the Agent in these cases is what Kiparsky calls a *constitutive argument*. Constitutive arguments, says Kiparsky, are not omissible; they must appear on the surface, which means that verbs like *paint* and *shelve* can only appear in the transitive forms, (11a) and (12a). Verbs like *redden* and *thin*, on the other hand, "denote processes which can be initiated without the participation of a causing Agent (e.g., The sky is reddening, Fred's hair is thinning)… and which, once initiated, can continue without it" (p. 495). Thus, the Agent in these cases is *not* constitutive, which means it can be omitted. Consequently, verbs like *redden* and *thin* can appear in the intransitive forms, (9) and (10). This analysis seems to hold for all verbs, not just for the denominal and deadjectival ones. Thus, for example, it captures the truly curious contrast between *roll* (which participates in the unaccusativity alternation) and *push* (which does not):

(13) a. Mary pushed the cart.
b. *The cart pushed.

(14) a. Mary rolled the cart.
b. The cart rolled.

Pushing a cart involves a specific activity of the agent, which must continue for the entire duration of the event. When John stops pushing the cart, the pushing *stops*, even though the cart may continue to move. Thus, the Agent of *push* is constitutive; it cannot be omitted. If John *rolls* a cart, on the other hand, and then lets go of it, he stops rolling the cart, but the rolling event continues as long as the cart is in motion. Thus, the Agent of *roll* is omissible, and *roll* can appear in the intransitive form.

Like the Rappaport-Hovav and Levin analysis, Kiparsky's theory does not only capture the behavioral facts at hand, but also suggests a principled account of a complex and abstract syntactic phenomenon — syntactic optionality, the fact that some semantic arguments are omissible, whereas some others must always appear on the surface.

3.3. *Concealed questions and the accessibility of meanings*

Considering the importance of the autonomy hypothesis as a constitutive presupposition in the field, it is striking that there are very few explicit

attempts in the literature to prove its validity on empirical grounds. The distributional pattern of *concealed questions* (definite NPs which are interpreted as interrogative wh-complements) has provided the basis for what has probably been the most explicit attempt to do that — Grimshaw's (1979) famous argument for the autonomy of *syntactic subcategorization* from semantic selection. This is Grimshaw's argument, in skeletal form.

Some verbs in English, like *know* and *ask*, semantically select for wh-complements. Other verbs, like *believe* and *think*, do not:

(15) a. John *knew* what the time of the meeting was.
b. John *asked* what the time of the meeting was.

(16) a. *John *believed* what the time of the meeting was.
b. *John *thought* what the time of the meeting was.

Concealed questions can only appear on the surface with verbs that can take sentential interrogative complements. Thus, *know* and *ask* can take concealed questions, whereas *believe* and *think* cannot:

(17) a. John *knew* the time of the meeting.
b. John *asked* the time of the meeting.

(18) a. *John *believed* the time of the meeting.
b. *John *thought* the time of the meeting.

The above distributional facts are explained by semantic selection, because sentential interrogative complements and concealed questions are supposed to have the same meaning. If a verb selects for this meaning, it will be able to take both types of complements. The picture is complicated, however, by the fact that some verbs, which take sentential interrogatives, do *not* take concealed questions:

(19) a. John wondered what the time of the meeting was.
b. *John wondered the time of the meeting.

(20) a. John inquired what the time of the meeting was.
b. *John inquired the time of the meeting.

How can these facts be explained? Grimshaw claims that they can only be explained by the independent constraints imposed by syntactic subcategorization. Verbs like *ask* subcategorize for NPs, on the syntactic level, regardless of the semantic nature of the NP. Verbs like *inquire* do not subcategorize for

NPs, and cannot take them as their complements — again, regardless of their meanings. Consequently,

> the lexicon must contain both subcategorization and selection frames, operating independently to explain the combinatorial properties of predicates and complements… The semantic types which appear in selection frames and the syntactic categories which appear in subcategorization frames are not in one-to-one correspondence. Thus, it is impossible to reduce the syntactic categories to the semantic types or the semantic types to the syntactic categories (Grimshaw 1979: 316).

Grimshaw's argument is actually much stronger than it may seem at first sight. To explain the above facts on semantic grounds, we should find a way to correlate the behavioral difference between, say, *ask* and *inquire*, with a difference between their meanings. This may sound rather trivial — obviously there *are* meaning differences between the two verbs — but the problems begin when we start looking for an explicit meaning difference which can do the job. Remember: We need an explicit and stable meaning difference, which lumps together verbs like *ask* and *know*, together with a very large set of other verbs which accept the concealed question construction, and leaves out *wonder* and *inquire*, together with some others, all of which reject concealed questions. When asked about the meanings of these verbs, different speakers characteristically suggest different definitions. It seems that the verbs mean different things for different speakers. Even to the extent that speakers agree about the definitions of the verbs, their definitions fall short of explicating the line of demarcation between *know*, *ask* and their likes on the one hand, and *wonder* and *inquire* and their likes on the other. And remember: The same speakers who struggle with their semantic intuitions have a perfectly clear and stable pattern of syntactic intuitions. Everybody agrees that the sentences in (17) are good, whereas the sentences in (19b) and (20b) are bad. In light of these facts, Grimshaw's solution seems to be right on target.

But is it really? In Dor (1996), I have shown that the distribution of concealed questions is perfectly predictable from a set of meaning considerations, having to do with the meanings of the embedding verbs and the meanings of the two types of complements. This solution is presented here in a simplified and informal fashion. The distribution of concealed questions is determined by the cross-section of two constitutive semantic distinctions — the distinction between *facts* and *non-facts* and the distinction between

knowledge and *ignorance*. Complement-embedding verbs can either be *factual*, i.e., have *facts* as their objects, or *non-factual*, i.e., have *non-facts* as their objects.[5] *Know, ask, wonder* and *inquire* are factual verbs; *think* and *believe* are non-factual verbs. The set of factual verbs can be further divided into verbs that are compatible with the notion of *knowledge* and verbs that are compatible with the notion of *ignorance.* Verbs can take concealed questions only if they are semantically compatible with *factual knowledge*. Verbs cannot take concealed questions if they are either *non-factual* verbs or verbs of *factual ignorance*.

To get a better sense of what is involved in this analysis, let us take another look at *ask* and *inquire*. As we saw before, speakers find it extremely difficult to consciously figure out the meaning distinction between these two verbs. However, when asked to use the verbs in different contexts, carefully designed to highlight the meaning differences between them, speakers do *that* easily, naturally and with a very high rate of inter-speaker agreement. The reader is encouraged to try the following test.

Assume a murder occurred somewhere in town, in a big hotel. You, the reader, were at the murder site at the time a police officer started an investigation. You saw the officer approaching witnesses and asking questions. Telling a friend about these events, you will probably be able to use both verbs, as in (21):

(21) a. The police officer asked who was in the room at the time of the murder.
b. The police officer inquired who was in the room at the time of the murder.

Now, let me add one crucial fact to the story. When the police officer approached his witnesses with his question, he had already established the answer to his query. The only motivation for the questioning session (so he told you) was his hope that one of the witnesses would lie, and the discovery of the liar would get him closer to the next step: Finding the murderer. With this new bit of knowledge in your mind, you will no doubt still be able to use (21a) to describe the event, but will you be able to use (21b) for the same purpose? The native speakers I have worked with had a very clear intuition here: (21b) is no longer a viable option. To describe the activity of asking questions as *inquiring*, one has to assume something like a sincere intention to get the right answer, which in turn presupposes *ignorance* on

behalf of the inquirer as to what this answer actually is. To describe the activity of asking questions as *asking*, on the other hand, one does not have to make this assumption. I can ask you a question if I know the answer in advance — I can even ask a rhetorical question, and answer it myself. Thus, *ask* is compatible with *factual knowledge*; *inquire* is not. All the verbs that can take concealed questions are compatible with factual knowledge; the verbs that cannot take concealed questions are not compatible with it. Verbs like *think* are non-factual, and verbs like *inquire* are verbs of ignorance.

The formal details of the analysis — which explain *why* concealed questions are sensitive to the notion of factual knowledge — should not interest us here. The point to be made here is that meaning-bearing elements in language do not wear their meanings on their sleeve. The native speakers I worked with were genuinely surprised to find out that their usage of the relevant verbs reflected a deeper level of lexical knowledge than their introspective intuitions allowed them to discover. As native speakers of the language, they knew more about the meanings of the verbs than they thought they knew. This psycholinguistic phenomenon raises some interesting questions regarding the psychological accessibility of word meanings via introspection. Generative linguistics has always assumed, as a part of its set of metatheoretical presuppositions, that the really important components of our linguistic knowledge are untouchable by introspection — we do not have intuitions about filters or assimilation processes in phonology, for example. We only have intuitions about the *outputs* of such mechanisms. As the above example demonstrates, the very same working hypothesis should be adopted with respect to word meanings.

What the example should tell us (at least as a preliminary working hypothesis) is that we simply do not have access to word meanings via introspection. What we do have access to are the *outputs* of word meanings — which enable us to make introspective judgments about the proper usage of words within specific semantic-pragmatic contexts. This working hypothesis, in turn, means that no decision can be made about what a word means prior to a serious investigation of the proper usage of the relevant word in a set of different semantic-pragmatic contexts. More significantly for our current discussion, no decision on the autonomy of some grammatical generalization from meaning considerations can be made on the basis of semantic introspection — the fact that some grammatical generalization *seems* to be divorced from meaning simply cannot be taken as evidence for anything.

Rather than prove the autonomy of syntactic subcategorization, the distribution of concealed questions proves the dramatic extent to which syntactic phenomena are correlated with a level of semantic representation which is so fundamental and basic to language that it is virtually inaccessible to introspection.

4. Back to the relationship problem

In each of the above analyses, a specific set of semantic notions turned out to determine a certain subset of syntactic phenomena. In Rappaport-Hovav and Levin's analysis, the constitutive notions were the five *event types* and the notion of *semantic expansion*. Kiparsky makes use of such notions as *constitutive argument* and *involvement* of the Agent in the event. In my own analysis, the constitutive notions were *factuality*, *knowledge* and *ignorance*. A natural question arises: What is the theoretical status of these semantic notions? Where do they belong in the general cognitive scheme of linguistic knowledge? This is a crucial question. The majority of scholars working on the syntax-semantics interface assume, as if by default, that the relevant semantic notions belong to the level of conceptual structure. The most important proponent of this view is Jackendoff (1987, 1990). As I will try to show below, we have very good reasons to reject this idea, and assume instead that the relevant semantic notions belong to the representational level of *linguistic semantics*.

The first important point we should discuss is the following: The semantic notions which turn out to play a constitutive role in determining syntactic generalizations seem to belong to a very constrained *subset* of all the semantic notions which we can use to think and conceptualize about the world. Even more significantly, these semantic notions cut across conceptual structures in ways that seem, from the conceptual point of view, quite arbitrary. In the literature on the syntax-semantics interface, this point has mostly been made as a methodological principle: The fact that some syntactic pattern cannot be explained as a reflection of some candidate meaning component does not necessarily mean that it is determined by purely syntactic considerations. It may very well be the case that the meaning component examined was simply the wrong one, which does not play a role in determining syntactic generalizations. But the distinction between linguistically relevant and

linguistically irrelevant meanings goes a long way beyond the methodological issue. It is a fundamental fact about the relationship between meaning and form in language. Some semantic notions turn out to be linguistically relevant, again and again, whereas others consistently do not participate in the linguistic game.

The very fact that verbs belong to *semantic classes* is probably the most fundamental piece of evidence for this claim. The reader may remember that Rappaport-Hovav and Levin distinguished between *change-of-state* and *manner-of-motion* verbs. The verb *sweep*, which appeared in examples (1a–f), belongs to one subclass of manner-of-motion verbs — the class of *surface-contact* verbs. Some other members of this class are *wipe*, *scrape*, *scratch*, and *scrub* (Levin 1993). These verbs denote different versions of the same event type — movement in contact with a surface. The important point is that, as far as we can tell, the meaning distinctions between the different members of the same verb class — the meaning distinctions between, say, *sweep* and *wipe* — are completely irrelevant as far as the grammar of English is concerned. There does not seem to be a syntactic generalization that is sensitive to these meaning distinctions. The same is true, for example, for the subclass of change-of-state verbs which includes *break*, *smash*, *crash*, *fracture*, *shatter* and many others. It is as if, from the point of view of the grammar, all these verbs are the same one. Needless to say, we *do* distinguish between *sweep* and *wipe*, and between *break* and *smash*, on the conceptual level, but as far as language is concerned, *sweep* and *wipe* are indistinguishable, and so are *break* and *smash*. It is only their *event type* that is somehow singled out and isolated by language, and is thus a constitutive determinant of its structure.

Examples of this phenomenon abound. In Dor (1996), I looked at sentence embedding verbs, which we have already discussed: *think*, *believe*, *know*, *ask* and *inquire*. Now, some of the relevant verbs have an additional, *emotional* meaning component. Verbs like *complain*, *hope* or *be afraid* tell us something about the emotional attitude of their subjects towards their propositional object. As it turns out, however, this emotional meaning component does *not* seem to determine any of the syntactic behavioral patterns of these verbs. As far as language is concerned, *complain*, *hope* and *be afraid* are just like *believe* and *think* — propositional attitude verbs of the *non-factual* type.

To take a third example, Frawley (1992) discusses and compares two important and robust distinctions of meaning: The distinction between *natural* and *nominal kinds* on the one hand, and the distinction between *animate* and *inanimate objects* on the other. The distinction between natural and nominal kinds, thoroughly discussed in the philosophical and psychological literature, has to do with the difference between common nouns that denote *compositionally* and common nouns that denote *inherently* (Pulman 1983; Cruse 1986; Kripke 1980). The distinction between natural and nominal kinds is visible throughout the lexicon: *tiger, gold, hepatitis, heat, pain,* and *red* are natural kinds; *car, wheel, coat, wedding, divorce* and *president* and nominal kinds. As Frawley (1992: 9) notes, however, the distinction between these two robust sets of common nouns "has little bearing on the grammatical structure of English". Natural kind terms have no unique structural properties:

> … both natural and nominal kinds take articles (*the gold, the wedding*) and quantifiers (*some gold, many weddings*). English does not differentiate the two in terms of pluralization (*two tigers, two divorces*) or verb choice (*the tiger fell down, the coat fell down, pain annoys me, the president annoys me*)… In fact, no language appears to differentiate natural from nominal kinds by any sort of formal device… (Frawley 1992: 9–10).

The distinction between animate and inanimate objects, on the other hand, is extremely relevant for linguistic structure. As Frawley says, "the linguistic evidence shows that in every language there appears to be some grammatical reflex of the difference between animate and inanimate objects" (p. 10). Frawley concludes that the fundamental question of philosophical semantics — what kinds of meaning are possible — contributes to the identification of a variety of potential meanings that language may encode. But only some of the results of an inquiry driven by this question are relevant. Not all possible meanings are grammaticalized; not all have empirical status. There may be an ontological distinction between natural kinds and nominal kinds, but that does not affect semantic analysis because languages do not signal the difference structurally (p. 10).

Needless to say, the distinctions between *sweep* and *wipe*, between *complain* and *believe*, and between *natural* and *nominal* kinds are extremely important for us as human beings. They definitely play an active role in our *conceptualization* processes. They just do not seem to play a role as determinants of grammatical generalizations.

All this is tightly connected to the next point that needs to be made with respect to the empirical analyses we have looked at. The semantic notions that demarcate the grammatical sentences from the non-grammatical ones, in each of the analyses, do *not* constrain our ability to freely conceptualize about the meanings of the non-grammatical sentences. Think about Kiparsky's analysis. We do not really have a conceptual problem with the idea of putting dead horses in corrals, or with the idea of putting the victims of a flood in a prison. What we have a problem with is describing these events as "corralling dead horses" and "imprisoning the victims of the flood". The problem is linguistic, not conceptual. It has nothing to do with the thought, and everything to do with the linguistic expression of the thought. By the same token, there is nothing inherently false, or conceptually impossible, about the ungrammatical sentences in Rappaport-Hovav and Levin's analysis: We may cause dishes to fall off a table by breaking them (imagine the dishes glued to the surface of the table, close to the edge, such that breaking them will cause their fragments to fall off the table). More obviously, a coach may make his athletes go around the track. What *is* impossible is *describing* these event as "breaking the dishes off the table" and "going the athletes around the track", and this is a fact about the *categories of language*, not about the way we think about the world.

Finally, the semantic notions we have looked at may vary from language to language, just like the notions discussed by Levinson. This time, however, we may be pretty sure that these patterns of variation are *not* correlated with differences in the way speakers of the different languages think about the corresponding concepts. Thus, for example, some of the sentences in (1) – (4), which *are* grammatical in English, are nevertheless *ungrammatical* in Modern Hebrew. The Modern Hebrew equivalents of "Pat ran her shoes to shreds" and "Pat ran herself ragged" are ruled out. Explaining this fact is far from trivial, but one thing is certain: English speakers and Modern Hebrew speakers do *not* have different concepts of *running*. It is *not* the case that Modern Hebrew speakers have a concept of running which is limited to the activity, whereas English speakers have a concept of running which somehow includes the accomplishment of ruining a good pair of shoes. Where speakers of English and Modern Hebrew differ is in the scope of the principle of semantic expansion — the extent to which verbs in the two languages can be used with extended meanings. As a first approximation, it seems that the scope of the principle of semantic expansion is much

narrower in Modern Hebrew. The set of possible meaning extensions in Modern Hebrew is a subset of the possible meaning extensions in English. Thus, the semantic distinction between the two languages is linguistic in nature, not conceptual.

It seems, then, that the semantically-based analyses we have looked at provide us with the exact type of material we need to demonstrate the autonomy of linguistic semantics from conceptual structure. As we have seen, language sometimes disallows the expression of a meaning which corresponds to a thought which is conceptually well-formed for the speakers of the same language; language is sometimes blind to conceptual distinctions which are active in the conceptualization patterns of the speakers of the same language; and languages sometimes manifest semantic differences which cannot be attributed to a parallel conceptual difference. All this strongly suggests that the set of semantic categories which determine grammatical regularities in natural languages (and probably some pragmatic notions, such as *topic* and *focus*, which sometimes play a similar role) constitute a *system*, whose primitive terms occupy an autonomous representational level, distinct from conceptual structure.[6]

This theoretical state-of-affairs brings us back to Levinson's major philosophical insight, the idea that the representational level of linguistic semantics severely *constrains* the functional capacities of language as a tool for the communication of meanings. Language is *not* a general-purpose tool for the communication of all possible meanings, but a *specific* tool for the communication of a constrained *subset* of meanings — those meanings which are represented on the level of linguistic semantics. One of Levinson's demonstrations of this state-of-affairs deserves special attention here, because it directly connects his framework with the specific insights accumulated by researchers working on the syntax-semantics interface. As Levinson shows, some extremely simple messages sometimes force complex linguistic representation:

> Entertain the idea of a capital T. Now describe this shape to an illiterate friend. You could say: *Draw a horizontal line. Find the mid-point, draw a vertical line of equal length downwards.* Or you could say *Draw a vertical line. Now draw a horizontal line that touches the top, with equal halves on either side of the vertical line.* There are still other ways that do not depend on notions like 'vertical': *Draw a line. Now from the end point go left 90 degrees by half the length of the first line. Go back to that point. Go right 90 degrees for the same distance.* Which system of coordinates we use and in which order

> we trace the shape are matters that in the spoken expression of the thought have to be determinate. But our original idea of the T might be a single gestalt, or it could embody multiple coordinate systems. This linearization of information is part of the preparation for speech production and can be empirically investigated… it is particularly evident in route directions, descriptions of visual stimuli, and the like (Levinson 1997: 20–21).

The same point is made, in a different context, by Aitchison:

> Consider the instructions for doing one of the simplest knots, a figure of eight:
>
> 1. Pass the end of the rope over the standing part.
> 2. Take the end under the standing part away from the loop.
> 3. Bring the end of the rope back up over itself towards the loop.
> 4. Pass the end down through the loop.
> 5. Pull tight.
>
> Without its accompanying diagram, this description is difficult to follow, although accurate. In this case, a picture is truly 'worth a thousand words'.
>
> Or take the following information from a widely used guidebook to Brasil:
>
> The hike begins… where a paved jogging track runs for 1,200 metres… At the end of the track pick up the trail on the other side of the cement tank in the tall grass. Follow this trail (always taking the uphill forks) for 100 metres. At the old foundations, some 30 metres above the water, the trail ascends steeply for 60 metres until levelling off on the narrow ridge… the trail to follow is up the far left-hand side ridge. At the base of the rock the trail deviates slightly to the right.
>
> These instructions are possibly as clear as language allows — but a map would have made things clearer (Aitchison 1996: 19).

What do these demonstrations tell us? According to Aitchison, they mean that the primary function of language is not conveying information. We will get back to this general question later on. According to Levinson, the demonstrations are a result of the fact that speech forces a linearization on the expression of thoughts. The linear nature of speech, according to Levinson, is another reason why conceptual and linguistic meanings are different. But notice: Visual information can be linearized in *many* ways. I can tell you how to draw a capital T, for example, by giving you the coordinates of the relevant points — the endpoints of the two lines — on imaginary x and y axes, and I can do *that* in a linear fashion, giving you the values of the coordinates one by one. What happens when language is used to convey spatial and visual information — how to draw a T, how to tie a

knot, how to follow a trail — is *not* just the linearization of the information. Much more importantly, the information is *reframed*: As I have shown (Dor 1998), the information is transformed into the type of information which is expressible in terms of linguistic semantic notions. Take another look at Aitchison's instructions for tying a knot. The single event of tying the knot is broken down into a series of smaller events, each of which *separately* meets the conditions of the linguistic semantic set of possible event types (the set of event types used by Rappaport-Hovav and Levin to explain the grammatical phenomena they worked on). Each of the smaller events is an *accomplishment* event, where the person holding the rope causes *the end of the rope* to move in space, and reach a new position in relation to *the standing part* and *the loop*. Note, and this is a crucial point, that if you learn to tie the knot in non-verbal ways — from a picture, or by looking at someone doing it — you may never have to conceptualize about such entities as *the end of the rope*, *the standing part* or *the loop*. The above linguistic instructions force us to bring these entities to life exactly because the different event types allowed by language set specific conditions on the types of entities participating in them.

Just as the number of event types is severely restricted, so is the number of possible types of participants — this is the essence of the theory of *thematic roles* (cf. Jackendoff 1990; Frawley 1992). *The end of the rope* participates in each of these events as the *theme*, the moving participant. *The standing part* and *the loop* participate in the different events in the spatial roles of *path*, *goal* and *location*. Note that the rope itself — which is, conceptually speaking, a single physical unit, the object which is tied into a knot — does not appear anywhere in the instructions. This is so for a very good reason: It cannot be accommodated within the linguistic framework I just detailed. In each of the subevents described in the instructions, the rope functions as the theme, the path, the goal and the location, all at the same time. As the theory of thematic roles teaches us, a participant cannot have more than one spatial role per event. Needless to say, this is a uniquely linguistic semantic constraint, not a conceptual one.

5. Conclusion

Let us take stock. In the above discussion, the empirical results we have looked at served as arguments for two claims which are usually treated in the literature as unrelated to each other. First, we discussed them in the context of the autonomy hypothesis. We saw why results of this type help weaken the status of the autonomy hypothesis as a default presupposition, and strengthen the hypothesis that grammars of natural languages actually constitute structural reflections of built-in complexities of meaning. Then, we interpreted the *same* empirical results as arguments for the autonomy of linguistic semantics — a uniquely linguistic level of semantic representation — from conceptual meaning. I take these two interpretive results to be two sides of the same coin. Think about the following point. For a few decades now, we have made remarkable progress in our understanding of *general communication* — we have a much better understanding, for example, of pragmatic inference — but this progress has *not* made a difference to our understanding of natural language grammars. This is exactly the point made, for example, by Sperber and Wilson (1986), for whom there is a clear division of labor between the study of communication (linguistic or otherwise) and the study of linguistic form.[7] However, whenever we make progress on the syntax-semantics front — each time we manage to reduce the level of arbitrary complexity of our syntactic theory — we add another explanatory component to our theory of linguistic semantics. Every additional insight into the structure of linguistic semantics means one syntactic stipulation less. As we develop our understanding of the syntax-semantics interface, we seem to be gradually replacing the autonomy of syntax with the autonomy of linguistic semantics.

This fact about the correspondence between the transparency problem and the relationship problem has far-reaching consequences for the overall theory of language. Throughout the 20th century, linguistic discourse has been characterized by an insistence on black-and-white, all-or-nothing binary distinctions. It was either *nature* or *nurture*, *universalism* or *relativism*, *form* or *meaning*, *empiricism* or *rationalism*, *structuralism* or *functionalism*. This state-of-affairs kept alive a major theoretical tension — the tension between the *functional* nature of language on the one hand, and its *uniqueness* as a cognitive domain on the other. Chomsky's conception of grammar as an autonomous system captured the fundamental fact — reflected in the process

of language acquisition, the structures of our brains, the details of our physiology and the discrepancy between human and primate communication — that language is an independent and unique cognitive system, irreducible to some general cognitive capacity. The price, however, was the total rejection of another fundamental truth about language — its functionality. Functionalists, on the other hand, described a functional system of communication which, in this way or the other, was nothing but another instantiation of man's general intelligence. The correspondence between the transparency problem and the relationship problem shows the way out of this dilemma. It allows for a new, synthetic conception of language — neither as an autonomously-formal generative system, nor as a functional reflection of conceptual structure, but as a transparent mapping-system, mapping a constrained subset of uniquely linguistic notions of meaning onto the linear channel of speech. It suggests a characterization of language as a cognitive system whose function *is defined on the representational level of linguistic semantics*, a system that is at the very same time uniquely linguistic *and* functionally based.[8]

Tel Aviv University

Acknowledgments

I would like to thank the two anonymous referees for their comments and suggestions.

Notes

1. All this does *not* mean that linguistic semantics does not exist, or that the semantic notions discussed by Levinson in the above paragraphs do not belong to linguistic semantics; it only means that the *arguments* are not as conclusive as they should be.
2. I am not sure, however, that notions like *today*, *tomorrow*, *here* and *there* are not represented as such on the conceptual level. Susan Carey (p.c.), for example, claims that notions such as *here* and *now* are very clearly represented conceptually not just by humans, but also by primates.
3. The assertion that linear order is a purely structural phenomenon does *not* deny the obvious fact that it is often used to mark semantic or pragmatic relations. The point here is that the structural phenomenon of linear order *as such*, i.e., the fact that words and constituents appear before and after each other on the speech channel, has to be described in terms of the structural properties of the speech channel, regardless of its possible usages.

4. One of the anonymous referees of this paper suggested that "empirical proof of any single syntactic phenomenon which is directly semantically imposed is enough to break the absoluteness of the (autonomy) claim, making it futile since it only makes sense as an absolute claim". I disagree: Generative grammar has never suggested that the autonomy hypothesis is absolute in the suggested sense, i.e., that there are no syntactic phenomena to be found which are "semantically imposed". In fact, generative grammar explicitly *accepts* the claim that *some* syntactic facts, e.g., some subcategorization phenomena, are derivable from semantic generalizations (S-selection). The generative claim has always been that a certain *set* of syntactic phenomena must be explained by a system "whose primitive terms are nonsemantic and nondiscourse-derived syntactic elements and whose principles of combination make no reference to system-external factors" (Newmeyer 1998: 23).
5. The notion of *factuality* is an extension of Kiparsky and Kiparsky's (1970) notion of *factivity*. As a first approximation, factive verbs are factual, and non-factive verbs are non-factual. The notion of factuality, however, is applicable not just to propositional-attitude verbs, but also to *information-seeking* verbs, such as *ask* and *inquire.* It captures the intuition that the object of interest of the inquirer, the thing that the inquirer is seeking, is a *fact*, not a hypothesis, a thought or a guess.
6. Note that this claim does *not* deny the obvious *interrelationship* between linguistic semantics and conceptual structure. As one of the anonymous referees of this paper has indicated, some of the principles of linguistic semantics may be "motivated by conceptual semantic elements, even if to some extent they go their own way with them, due to dimensions which are not applicable at the conceptual level, such as the 'logic' of communication". The crucial question is exactly the extent to which the principles of linguistic semantics *do* set themselves apart from the principles of conceptual structure. I hope to have shown that they indeed set themselves apart in significant and systematic ways.
7. Here, of course, I find myself in agreement with the traditional view of generative grammar (cf. Newmeyer 1998) and in disagreement with some proponents of linguistic functionalism.
8. Dor and Jablonka (forthcoming) suggest that this characterization of language as a unique functional system allows for a novel answer to the question of the biological-cultural *evolution* of language. In Dor (1999), I show that this view is surprisingly similar to Ernst Cassirer's view of language as a *symbolic form*.

References

Aitchison, J. 1996. *The Seeds of Speech: Language Origin and Evolution.* Cambridge: Cambridge University Press.

Alsina, A., Bresnan, J., and Sells, P. (eds) 1997. *Complex Predicates.* Stanford, CA: CSLI Publications.

Bierwisch, M. and Schreuder, R. 1992. "From conceptual to lexical items". *Cognition* 42: 23–60.

Butt, M. and Geuder, W. (eds) 1998. *The Projection of Arguments: Lexical and Compositional Factors*. Stanford, CA: CSLI Publications.

Cruse, D. A. 1986. *Lexical Semantics*. Cambridge: Cambridge University Press.

Dor, D. 1996. *Representations, Attitudes and Factivity Evaluations: An Epistemically Based Analysis of Lexical Selection*. Stanford Dissertation.

Dor, D. 1998. "On linguistics as a study of culture". Lecture at the Cohen Institute for the History and Philosophy of Science, Tel Aviv University.

Dor, D. 1999. "From lexical semantics to symbolic forms: Where modern linguistics and Cassirer's philosophy start to converge". *Science in Context* 12(4): 493–511.

Dor, D. and Jablonka, E. Forthcoming. "How language changed the genes: Towards an explicit account of the evolution of language". In J. Trabant (ed), *Essays on the Origins of Language*. Berlin: Mouton de Gruyter.

Dowty, D. 1979. *Word Meaning and Montague Grammar*. Dordrecht: Reidel.

Dowty, D. 1991. "Thematic proto-roles and argument selection". *Language* 67: 547–619.

Frawley, W. 1992. *Linguistic Semantics*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Goldberg, A. E. 1995. *Constructions: A Construction Grammar Approach to Argument Structure*. Chicago: The University of Chicago Press.

Grimshaw, J. 1979. "Complement selection and the lexicon". *Linguistic Inquiry* 10(2): 279–326.

Grimshaw, J. 1990. *Argument Structure*. Cambridge, MA: The MIT Press.

Gumperz, J. J. and Levinson S. C. (eds). 1996. *Rethinking Linguistic Relativity*. Cambridge: Cambridge University Press.

Hale, K. and Keyser, J. 1992. "The syntactic character of thematic structure". In I. M. Roca (ed.), *Thematic Structure: Its Role in Grammar*. Berlin: Foris, 107–143.

Jackendoff, R. 1983. *Semantics and Cognition*. Cambridge, MA: The MIT Press.

Jackendoff, R. 1990. *Semantic Structures*. Cambridge, MA: The MIT Press.

Keller, C. M. and Dixon Keller, J. 1996. "Imaging in iron, or thought is not inner speech". In Gumperz and Levinson, 115–129.

Kiparsky, P. 1997. "Remarks on denominal verbs". In Alsina et al., 473–499.

Kiparsky, P. and Kiparsky, C. 1970. "Fact". In M. Bierwisch and K. Heidolph (eds), *Progress in Linguistics*. The Hague: Mouton, 345–369.

Kripke, S. 1980. *Naming and Necessity*. Oxford: Oxford University Press.

Levin, B. 1993. *English Verb Classes and Alternations: A Preliminary Investigation*. Chicago: The University of Chicago Press.

Levinson, S. C. 1996. "Relativity in spatial conception and description". In Gumperz and Levinson, 177–202.

Levinson, S. C. 1997. “From outer to inner space: Linguistic categories and non-linguistic thinking”. In Nuyts and Pederson, 13–45.

Lucy, J. A. 1996. “The scope of linguistic relativity: An analysis and review of empirical research”. In Gumperz and Levinson, 37–69.

McNeill, D. 1997. “Growth points cross-linguistically”. In Nuyts and Pederson, 90–112.

Newmeyer, F. J. 1998. *Language Form and Language Function*. Cambridge, MA: The MIT Press.

Nuyts, J. and Pedersen, (eds). E. 1997. *Language and Conceptualization.* Cambridge: Cambridge University Press.

Pederson, E. and Nuyts, J. 1997. “Overview: On the relationship between language and conceptualization”. In Nuyts and Pederson, 1–12.

Pinker, S. and Bloom P. 1990. “Natural language and natural selection”. *Behavioral and Brain Sciences* 13: 707–84.

Pulman, S. G. 1983. *Word Meaning and Belief*. London: Croom Helm.

Rappaport-Hovav, M. and Levin, B. 1998. *Building Verb Meanings*. In Butt and Geuder, 97–134.

Searle, J. 1969. *Speech Acts*. Cambridge: Cambridge University Press.

Sperber, D. and Wilson, D. 1986. *Relevance*. Oxford; Blackwell.

Van Valin, R. D. Jr. and LaPolla, R. J. 1997. *Syntax: Structure, Meaning and Function*. Cambridge: Cambridge University Press.