

Degree vs. manner *well*: A case study in selective binding

Louise McNally

Departament de Traducció i Filologia
Universitat Pompeu Fabra
La Rambla, 30-32, E-08002 Barcelona
mcnally@upf.es

Chris Kennedy

Department of Linguistics
Northwestern University
Evanston, IL 60208
kennedy@northwestern.edu

But ‘well’ is semantically complex. It combines the features of ‘approval’ and ‘fulfillment’ in ways that defy separation of the two.
Bolinger 1972, p. 29

1 Introduction

Among the various strategies natural languages employ for expressing degree modification is the use of expressions which are demonstrably manner adverbs, including *well* and its negative counterparts *poorly* and *badly*, as illustrated in (1).¹

- (1) a. well-/poorly-acquainted with the facts
(high/low degree of familiarity)
b. a well-/badly-paid position (high/low degree of pay)

Such expressions can be distinguished from categorical degree modifiers such as *very* among other ways in their ability to feed subsequent degree modification, including comparatives:

- (2) a. Spade was less well-acquainted with the facts than his assistant was.
b. *Spade was less very happy about the fact than I was.

However, a degree reading for *well*, etc., is not always available, and its distribution is not random. Kennedy and McNally 1999 (hereafter, KM) observe that *well* allows both

¹Combining forms such as *over-*, *under-* and *ill-*, as in *overloaded*, *underpaid* and *ill-nourished* show a similar distribution and pattern of behavior, and are thus likely to be amenable to a similar analysis, but since their syntax and semantics is slightly different from that of manner adverbs, we will not discuss them here.

a degree reading and a “quality” reading in examples like (3a), but only a quality reading in (3b). (For purposes of brevity, we focus in this paper on the specific case of *well-loaded*, though this phenomenon is quite general.)

- (3) a. a well-loaded truck (high degree of loadedness or loaded in a skilled/neat etc. way)
b. well-loaded hay (loaded in an organized/skilled/neat etc. way)

Although KM accounted for this contrast by treating *well* as ambiguous between a degree reading and a quality reading and by placing constraints on its use as a degree modifier which exclude that use in cases like (3b), positing a lexically ambiguous *well* is obviously ad hoc and computationally undesirable. In this paper, we formalize Bolinger’s intuition that the two senses of *well* are deeply related. Specifically, we show how the attested readings are in fact predicted when a simple and completely standard representation of participles in the Generative Lexicon framework (hereafter, GL; see Pustejovsky 1995) is combined with an equally simple and unambiguous analysis of *well* via selective binding as used in GL analyses of adjectival modification (e.g. Pustejovsky 1995, Bouillon 1999, Badia and Saurí 1999). Moreover, the specific semantics we adopt for adjectives entails that selective binding, understood specifically as the possibility of acting on a variable in the telic or agentive qualia of an adjective (as opposed to the formal quale) is not merely an attractive option for capturing polysemy; it is, in fact, the *only* option for *well* in the cases discussed here.

2 Constraints on degree ‘well’

Since our analysis of *well* ultimately involves deriving its polysemy from the interaction of its lexical semantics with the lexical semantics of the expressions it modifies, it is necessary first to describe briefly exactly when modification by *well* is possible, and when a degree reading is available. (As we saw above, a quality reading is always possible.) As shown by KM, *well*-modification in general is possible only with participles that meet two semantic requirements: 1) they must denote gradable properties, and so make use of scalar representations as part of their semantics, and 2) they must be associated with scales which are closed on both ends. Closed-scale adjectives can be distinguished from open-scale adjectives in that only the former permit modification by endpoint-oriented modifiers such as *partially* or *fully*:

- (4) a. The truck was partially/fully loaded.
b. The hay was half/completely loaded.

As shown by (5), open-scale adjectives such as *worried* do not permit modification by *partially*, etc, nor do they permit modification by *well*:²

- (5) a. ??Marge was completely worried when she saw the flying pig.
b. ??Marge is still well-worried.

For those participles that accept *well*-modification, the possibility of a degree reading is conditioned by a third feature: the nature of the participle’s “standard value” —

²It should be observed that maximality modifiers like *completely* and *totally* have both an endpoint-oriented use and a use that is roughly synonymous with *very*; these two uses are distinguished by their entailments. An endpoint-oriented use entails that the end of a scale has been reached, thus the sentence *The line is completely/totally straight, though you can make it straighter* is a contradiction. A nonendpoint-oriented use carries no such entailment, thus the contingency of *I’m completely/totally uninterested in phrenology, and Bob is even less interested than I am*.

the value on the adjective’s scale which determines whether or not its denotation truthfully holds of an entity. Specifically, the standard value cannot be the maximum value on the scale. This condition follows from the central semantic effect of the degree reading of *well*: it “boosts” the standard value for the attribute with which it combines. For example, holding all potentially variable factors constant, the standard of loadedness which must be reached for a vehicle to be considered well-loaded in any given situation is considerably higher than that required for the vehicle to qualify as simply loaded. Crucially, it is possible for *well* to boost the standard value associated with *loaded* as applied to vehicles because this standard value is demonstrably the minimum degree on the “loaded” scale: a truck qualifies as loaded as long as cargo has been placed upon it; it need not be the case that the cargo fills the entire available space in the truck. In contrast, *well* cannot boost the standard value associated with *loaded* as applied to contents, because the standard value in this case turns out to correspond to the maximum value on the “loaded” scale: in order for some cargo to qualify as loaded, it must be entirely on the truck. If the standard is already a maximal value, then it clearly cannot be meaningfully boosted. Thus, we correctly expect the degree reading of *well* to be blocked in the latter case, and indeed any time the standard value is a maximum.

While this analysis gets the facts right, it raises a more general question: how do we know when a standard value for some gradable property corresponds to a minimum or a maximum value on the relevant scale? According to KM, this is determined by looking at the semantic role that the object measured by the participle has in the corresponding verbal form. First, KM show that the scale structures associated with participles can be homomorphically related to (and, ultimately, derived from) aspects of their event structures, and that a given participle may be associated with more than one scale. Continuing with the “loading” example, a (maximal) loading event involving a vehicle *x* and con-

tents y can be divided into temporally and incrementally ordered subevents of loading x with subamounts amounts of y . The temporal endpoints of each of these subevents can be mapped onto a degree on a scale associated with the participle *loaded*. The endpoint of the first subevent of loading of the smallest amount of x onto y corresponds to the minimum on a “loaded” scale for both x and y . However, what constitutes the maximum value on the scale depends on the participle’s argument structure, since argument structure arguably affects the nature of the event described by the participle (see e.g. Dowty 1991, Levin and Rappaport-Hovav 1999). The endpoint of the last subevent of loading of the last bit of x onto y corresponds to the maximum on a scale for x . In contrast, the endpoint of the subevent of loading the last bit of x that fits onto y corresponds to the maximum on a scale for y . The result is that the “loadedness” of both arguments involves measurement on a closed scale, but the nature of the measurement is different in the two cases.

This difference has consequences for the nature of the participle’s standard value: it may vary depending on the thematic role borne by the participle’s argument in the event related to that participle. If the argument is a classic incremental theme (see Dowty 1991), the sort that Ramchand 1997 calls “Pat₌” (as in the case of *hay* in *loaded hay*), the standard is the maximal value on the scale. This is so because the conditions for truthful application of the participle are not met unless all of the incremental theme has undergone the event in question: (6a) is not true of the hay unless 100% of it has undergone loading. Further evidence that the standard for the scale is the maximum value can be derived from considering sentences like (4b), whose truth is evaluated by mapping the hay onto the scale associated with *loaded* and then computing the distance of the value assigned to the hay from the standard—the upper endpoint of the scale. Note that even though the truth of (4b) depends on the hay having some value on the “loaded” scale, (4b) crucially does not entail (6a); note that (6b) is a contradiction.

- (6) a. The hay is loaded.
 b. The hay may be only half-loaded, but it’s loaded (so I’m going home).

In contrast to what happens with a classic incremental theme, if the argument of the participle corresponds to some other argument of the verb (including e.g. arguments bearing Ramchand’s “Pat_{+/-}” role), then the standard corresponds to the minimum value on the scale. This is what happens in the case of the loaded vehicle, where the vehicle as a whole is involved in each subevent of loading, but where one of its properties, namely the degree to which its volume is occupied, changes incrementally. (7a) can be true as soon as the truck has undergone a minimal loading event; it is not necessary for its entire volume to be occupied, as seen in the fact that (7b) is not a contradiction.

- (7) a. The truck is loaded.
 b. The truck may be only half-loaded, but it’s loaded (so I’m going home).

To summarize, the standard for a closed-scale participle is a maximum value when that participle applies to a true incremental theme argument, and a minimum when it applies to other arguments. We now present an analysis of *well*+participle expressions in GL which captures the observations made in this section.

3 A GL analysis

We assume basic familiarity with the GL framework and only mention here those formal details which are specific to our analysis. For explicitness, we integrate our GL representations into a version of Head-Driven Phrase Structure Grammar (HPSG, see e.g. Pollard and Sag 1994). (In particular, we follow Badia and Saurí 2000 in making sets of qualia the values of the HPSG REST(RICTION) feature, though we diverge from them on other details.) To the

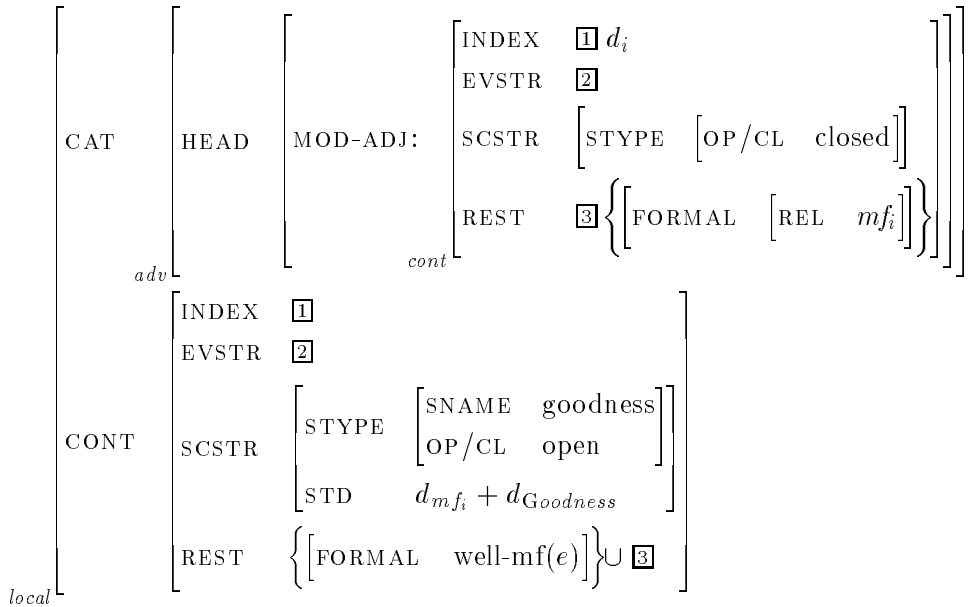


Figure 1: Lexical representation for *well*

three basic structures in a GL representation—argument structure (ARGSTR), event structure (EVSTR), and qualia structure (QSTR)—we add a fourth: *scale structure* (SCSTR), which identifies the scale (the SNAME feature, e.g. “loaded”) and specifies whether the scale is open or closed (the OP/CL feature) and what the standard value is (STD, whose value is a degree). The exact configuration we have given to scale structure is not crucial for the analysis, so for reasons of space we will not justify it here.

Following KM we treat *well* as a manner adverb which selects for adjectives.³ Our representation for *well* appears in Figure 1. We extend Pollard and Sag’s treatment of prenominal adjectives to preadjectival adverbs. *Well* is not type-changing; that is, it inherits the index and event structure of its adjective argument (as represented in *well*’s CONT feature). However, *well* does place conditions on the scale structure of the resulting *well*+participle

combination: it has an open scale (see KM for empirical support) whose standard value is boosted some context-dependent amount above that required by the participle alone (represented in (6) via the identity of the subscript on the left-hand component of *well*’s STD value and the description in the *formal* feature of the expression selected by *well*).

Now let us turn to the conditions that *well* imposes on its participial argument. The MOD-ADJ feature specifies the one condition explicitly imposed: the adjective’s scale must be closed. The other information represented in MOD-ADJ is characteristic of adjectives in general. Note that we follow Kennedy 1999, who argues extensively that adjectives denote measure functions which are converted to properties of individuals only after combining with (possibly phonetically null) degree morphology; thus, the adjective’s INDEX and FORMAL qualia features have degree values.⁴ But in addition to the explicit condition that

³Obviously, a fully general analysis of *well* should try to unify its adjective-modifying use with its verb-modifying use. In the interest of space we do not pursue such a unification here. Our analysis will similarly ignore certain other such generalizations which do not bear directly on the issue of how both the degree and non-degree readings of *well* can be derived from an unambiguous *well*.

⁴Evidence that *well* combines with the adjective before the latter takes on degree morphology comes from e.g. the incompatibility of *well* with comparative or superlative adjectives such as **well more acquainted with the facts*. Note also that we have suppressed here for lack of space an exposition of way in which the adjective or *well*+adjective combination combines with degree morphology to form a property of individuals.

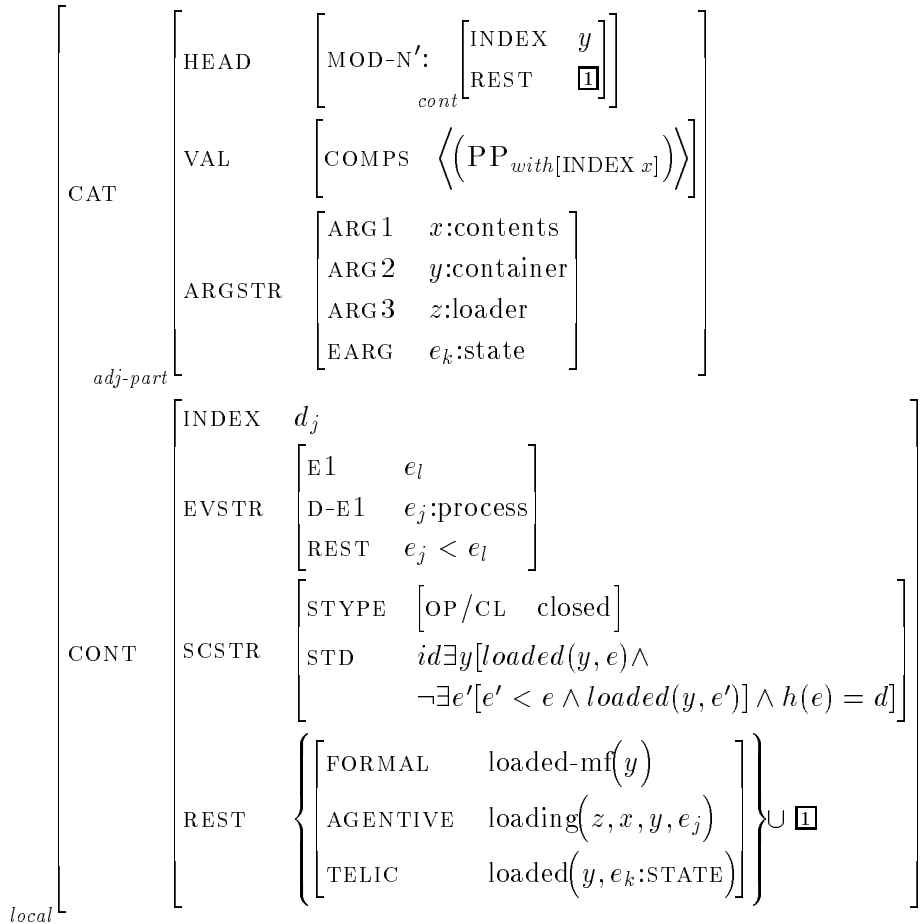


Figure 2: Lexical representation for *loaded-with*

the adjective’s scale be closed, the value of *well*’s CONT feature places a crucial implicit condition on the semantics of the adjective argument: Since *well* denotes a measure function on events, the semantic representation for the adjective must make available an event argument for the measure function to operate on, though it does not specify anything else about that event argument. This implicit condition is the key to understanding *well*’s polysemy.

The polysemy associated with *well* is highly reminiscent of that associated with adjectives. Badia and Saurí 1999, extending Pustejovsky’s 1995 treatment of verb polysemy and suggestions of his concerning the treatment of adjectives, propose that the different interpretations of e.g. *fast* in *fast car* (drives fast) vs. *fast cake* (made/baked fast) are a consequence of the adjective’s ability to act on

an event variable in either the telic quale of a noun (in the case of *car*, where the telic quale specifies a driving event) or its agentive quale (in the case of *cake*, where the agentive quale specifies a making/baking event; see also Bouillon 1999 for a similar treatment of *vieux* ‘old’ in French). Indeed, if adverbs are like adjectives in being able to act on different event variables in their complements’ representations via selective binding (see Pustejovsky 1995:129 for a definition), and if adjectives and participles are like nouns in being potentially specified for telic and agentive qualia, the polysemy that *well* exhibits is exactly what we would expect.

Consider the representation of *loaded-with* in Figure 2.⁵ First, observe that *loaded-with*

⁵ Although we assign *loaded* two different lexical entries corresponding to its two argument structures, these two entries are highly redundant and could

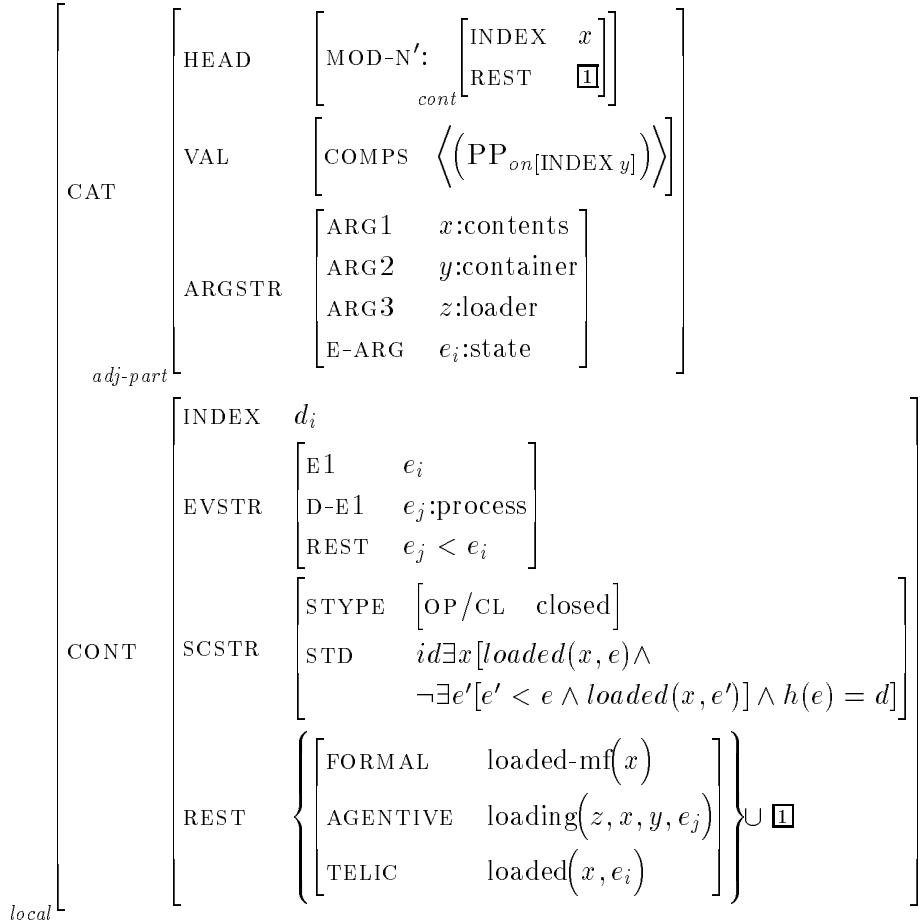


Figure 3: Lexical representation for *loaded-on*

satisfies the explicit conditions imposed by *well* (see KM for arguments that these participles can be adjectival). It satisfies the implicit condition as well: since the state of being loaded with some contents is achieved via a loading process, we can specify an agentive quale whose value is a description of a loading event. Since its being loaded with some contents is a result of a loading process, we assign the loaded state as the value of the telic quale.⁶

The representation of *loaded-on* in Figure 3 is identical, except that the arguments are crucially rearranged. Note that in both representations, the standard value for the partici-

ple is that degree which corresponds homomorphically to the smallest event of the type described by the participle involving (all of) the participant described by the participle's argument.

What happens when *well* combines with these representations? Neither the formal quale nor the index feature provide any event variable for *well* to act on. Thus, there remain only the event variables in the agentive and telic qualia. Modification of the event variable in the agentive quale corresponds to a manner/quality reading of *well*: the loading process is assigned a value on a scale of goodness which intuitively involves approval of objective aspects of the event—neatness, rapidity, skill, and, doubtless, degree of completion. All participles give rise to this reading. However, note that degree of completion will not be an interesting dimension of the

be partially unified in a hierarchical lexicon (Koenig 1999). To save space we have represented the values of the qualia as linear formulae, rather than as feature-structures.

⁶Here we differ from Pustejovsky 1995 in the treatment of result states.

loading event to be affected by *well* in the case of *loaded-on*, because truthful use of the phrase *loaded hay* entails that the entire referent of *hay* has undergone the loading process in question; in other words, the loading will have been completed to the maximal degree, as discussed in Section 2. Thus, no degree reading for *well* can be derived when it modifies the event variable in the agentive quale of *loaded-on*. In contrast, when *well* modifies the event variable in the agentive quale of *loaded-with*, it will derive a quality reading which is not incompatible with the implication that the loading was carried out to a high degree (though not necessarily completely—recall that truthful use of the phrase *loaded truck* doesn't entail that the truck is 100% loaded).

Be this as it may, we suspect that *well*'s degree reading arises principally when it modifies an event variable in the telic quale of its argument. Modification of the event variable in this quale corresponds to the assignment of a value on the goodness scale to the result state (here, the state of the contents or the vehicle being loaded). If the adverb is restricted, as seems to be empirically the case, to describing objective aspects of this state as opposed to e.g. the speaker's opinion as to its utility or appropriateness, then there would appear to be little to be evaluated as “good” other than the degree to which the state holds—such parameters as rapidity or skill cannot be evaluated, and neatness is clearly a reflection of the manner in which the loading process was carried out. But, as noted, any interesting evaluation of the degree to which some contents is loaded on a vehicle is impossible because the standard for its being considered loaded at all is the maximum value on the scale. Thus, *well* will not have a degree reading when it modifies the event variable in *loaded-on*'s telic quale, either. In contrast, such a reading is possible when *well* modifies the corresponding variable in *loaded-with*'s telic quale because the standard is a minimal value on the scale.

4 Conclusion

We have sketched an account of *well*'s polysemy in combination with participles which avoids the undesirable ambiguity posited in KM. While we maintain two lexical entries for the participle exemplified, this ambiguity is independently motivated on argument structure grounds and has been claimed to have semantic consequences (by e.g. Dowty 1991). Moreover, maintaining an ambiguity in the participle turns out to correctly predict that the phenomenon found in (3) occurs only with participles for which multiple semantic representations can be independently motivated.

At the theoretical level, we want to make two observations. First, the fact that existing analyses of adjectival modification can be adapted so straightforwardly to this case of adverbial polysemy lends further support to the selective binding approach in general.

Second, and more importantly, our analysis of adjectives as denoting measure functions raises issues concerning how to represent the three facets of adjectives (as measure functions, event descriptions, and properties of individuals) and their semantic compositional behavior. We have made the decision to use the INDEX and FORMAL features in parallel to reflect (what we take to be the adjective's) semantic type: the INDEX corresponds to the degree assigned as the output of the measure function denoted by the adjective on its nominal argument, and the value of the FORMAL feature constitutes a description of that degree, much in the same way that a verb in conjunction with its arguments can be assigned an event as index. However, it is by no means obvious what the relationship between these two features really should be, or how our proposal should be adjusted to better represent the eventive properties of adjectives (see Badia and Saurí 2000 for a different view on the role of the *index* feature in adjective representations).

In the case studied here, the option of using selective binding obviates the need to assign *well* to two semantic types: that of a measure function on events for its verb modifier use,

and that of a function from measure functions to measure functions for its adjective modifier use. However, it comes at a price: in the case of adjective-noun modification on the traditional analysis (or analyses such as that in Larson 1998), selective binding captured a simpler relationship between modifier and modified: the adjectival property could be said to apply to an entity in virtue of the entity's participation in an event. At this point, the distance between modifier and modified is rather greater in the case of *well*: a measure function on events applies to a measure function on individuals in virtue of the fact that the measure function on individuals is related to an event. The question is how complex can such indirect modification relations be? There are two paths to follow in pursuit of an answer to this question. One is to ask what kinds of constraints should be put on selective binding or its implementational equivalent if it is to tell us something interesting about natural language modification. The other is to begin integrating the long tradition on the scalar semantics of adjectives into the treatments of adjectives from the perspective of event semantics, in the hopes of eventually reducing or rationalizing the complexity of the modification relation instantiated by *well*. In making a first attempt to integrate scale structure into GL, we hoped to have taken a first step on this latter path.

References

- Badia, Toni and Roser Saurí. 1999. Semantic disambiguation of adjectives in local context: A generative approach. In Pierette Bouillon and Evelyne Viegas (eds.), *Proceedings of TALN 99*, 163–180.
- Badia, Toni and Roser Saurí. 2000. Enlarging HPSG with lexical semantics. In A. Gelbukh (ed.), *Proceedings of CICLing 2000*, 101–122, Mexico City, Mexico. Computer Research Center, National Polytechnic Institute.
- Bolinger, Dwight. 1972. *Degree Words*. The Hague: Mouton.
- Bouillon, Pierette. 1999. The adjective “vieux”: The point of view of “generative lexicon”. In Evelyne Viegas (ed.), *Breadth and Depth of Semantic Lexicons*, 148–166. Dordrecht: Kluwer.
- Dowty, David R. 1991. Thematic proto-roles and argument selection. *Language* 67:547–619.
- Kennedy, Christopher. 1999. *Projecting the Adjective: The Syntax and Semantics of Gradability and Comparison*. New York: Garland.
- Kennedy, Christopher and Louise McNally. 1999. From event structure to scale structure: Degree modification in deverbal adjectives. In Tanya Matthews and Devon Strolovitch (eds.), *Proceedings from SALT IX*, 163–180, Ithaca, NY. CLC Publications.
- Koenig, Jean-Pierre. 1999. *Lexical Relations*. Stanford, CA: CSLI Publications.
- Larson, Richard. 1998. Events and modification in nominals. In Devon Strolovitch and Aaron Lawson (eds.), *Proceedings from SALT VIII*, 145–168, Ithaca, NY. CLC Publications.
- Levin, Beth and Malka Rappaport-Hovav. 1999. Two structures for compositionally derived events. In Tanya Matthews and Devon Strolovitch (eds.), *Proceedings from SALT IX*, 199–223, Ithaca, NY. CLC Publications.
- Pollard, Carl and Ivan Sag. 1994. *Head Driven Phrase Structure Grammar*. Chicago: CSLI.
- Pustejovsky, James. 1995. *The Generative Lexicon*. Cambridge: MIT Press.
- Ramchand, Gillian C. 1997. *Aspect and Predication*. Oxford: Clarendon Press.