A lexical marker of degrees of answerhood

Moria RONEN — *Bar Ilan University*
Yael GREENBERG — *Bar Ilan University*
Galit SASSOON — *Bar Ilan University*

Abstract. Questions and their answers have been discussed at length over the past few years. In this paper we present and analyze a Hebrew hedger, *be-gadol*, roughly translated as *basically*. We use the literature on questions, answers and the relation between them to suggest that *be-gadol* is an item which conveys a restriction on the context of utterance. This restriction concerns the relation between answers to the QUD on an answerhood scale, which is characterized as involving two notions, informativity (Roberts 1996) and resolution, defined using tools from decision theories (Ginzburg 1995; van Rooij 2003). This significantly supports the linguistic reality of these notions.

Keywords: formal semantics, decision theory, decision problems, utility, resolution, answerhood.

1. Introduction

This paper deals with the Hebrew particle ***be-gadol***, which literally means *'in-big'*, and can be roughly translated as *in principle, by-and-large, in theory,* or *basically*. *Be-gadol* is a focus sensitive hedger/approximator which has a variety of uses and effects. The abundance of the effects leads to an empirical challenge, as it is very difficult to find a unified lexical entry which captures the variety of meanings of *be-gadol.* Moreover, we will argue that only some of the uses we find with *be-gadol* can be analyzed using existing tools. This leaves us in even deeper waters. Given the difficulties mentioned, the question at the center of this paper is how to account for all the data and all of these readings in a unified way, and which tools can be used for doing that.

The structure of the paper is as follows: This section will discuss some data to demonstrate the different uses of *be-gadol* and the next section will present our intuition. In section three we present the proposal in detail as well as apply it to several cases. We also review the theoretical tools needed for our purpose, in particular tools which deal with question resolution, informativity and utility (e.g. Roberts 1996, Ginzburg 1995, van Rooij 2003). In section four we examine an additional use of *be-gadol* which seems to pose a problem for our theory. We then propose a solution based on theories of projective meanings and relationships between questions under discussions and their answers (e.g. Simons et al 2010). Section five lays out some questions for further research, summarizes and concludes.

1.1. Data: The empirical challenge

As discussed above, using *be-gadol* can yield a variety of effects. For example, some of the uses we find with *be-gadol* can be paraphrased using other particles, as in (1) which we refer to as the "approximative" use, and as in (2) which we refer to as the "temporal" use. Both

(1a) and (2a) paraphrase the specific use of *be-gadol* in (1) and (2), and (1b) and (2b) give alternative paraphrases using a correlating particle (e.g. *more or less* and *for the most part*).

(1) ha-xeder be-gadol naki
the.room in.big clean
"The room is be-gadol clean."
a. The room is not completely clean
b. The room is more or less clean

(2) Context: John and Mary are talking about a party, John is inquiring about the mutual friend Rina's actions at that party.

John: What did Rina do at the party?
Mary: Rina be-gadol rakda
Rina in.big danced
"Rina be-gadol danced."
a. Rina danced most of the time
b. Rina danced for the most part

Crucially, however, there are also uses which cannot be paraphrased using such expressions. We will argue that uses like the ones in (3) and (4) cannot be analyzed using the tools which have been proposed for *for the most part* or *more or less.* For example, (3) is similar to example (2) but is used differently. We refer to (3) as the "significance" use. Another example can be found in (4), which we refer to as the "not enough details" use. (3a) and (4a) paraphrase the meaning of (3) and (4), but (3b,c), and (4b,c) demonstrate that unlike (1b) and (2b), paraphrases with *more or less* and *for the most part* are not available for the examples below.[1]

(3) Context: Rina sang, drank beer and spoke to the barman as she usually does in parties, but her dancing was her most significant, e.g. enthusiastic activity, attracting most of the attention. Dancing was not necessarily the longest activity, or atypical.

John: What did Rina do at the party?
Mary: She be-gadol danced
a. The most significant thing Rina did in the party was to dance
b. ≠for the most part she danced
c. ≠she more or less danced

(4) A: What do you do? What's your occupation?
B: ani be-gadol orex sfarim
I in.big edit books
"I'm be-gadol a book editor."

[1]Other interpretations of the same sentence are possible (e.g. paraphrases with *more or less*). However, they do not generate what we give in (3a) and (4a). The other available interpretations with the other particles correlate with the uses mentioned in (1) and (2). For example, one could say that *For the most part Rina danced* is a felicitous sentence, but instead of the meaning in (3a) we get the meaning that what Rina did during the party was mostly to dance.

a. What I do is more specific (e.g. I edit books dealing with medieval history and religion)
b. ≠for the most part I edit books
c. ≠I more or less edit books

As the examples above show, *be-gadol* can lead to a wide range of hedging effects, and not all of them can be paraphrased using other particles. The main challenge is to unify all these examples despite their differences. Due to the diversity of uses, the hedging effect cannot be analyzed by resorting to quantificational tools, as done for example by Nakanishi and Romero (2004) for *for the most part*, or just by using coarser granularities/lower precision standards, as Sauerland and Stateva (2007) suggest for *more or less.* The characterization of the hedging effect induced by *be-gadol* should be thus more general and flexible. To proceed along that direction we will present our intuition regarding the common denominator between the different uses and propose a semantics for *be-gadol*. We then continue by presenting the tools we need to proceed with the analysis, and justify our proposed semantics for *be-gadol*. The last piece of the puzzle is what seems like a counterexample and its solution, followed by a summary and a few open questions.

2. Intuition and proposal

We propose that in all the uses presented in section 1.1 *be-gadol* is a marker of degrees of answerhood to the QUD, which indicates a hedging effect along two dimensions: the propositional dimension and the discourse dimension.

On the propositional dimension the use of *be-gadol* indicates that the prejacent *p* is not the most informative answer to the QUD (cf. Roberts 1996). There is a better answer to the QUD, p_{best}. This p_{best} entails *p* and rejects an implied enrichment of *p*, *q*, which is stronger than *p.* For example, in (3) *she be-gadol danced*, *p* (*she danced*), is not the most informative answer to the QUD (*what did she do in the party?*). There is a better answer, p_{best}, *she danced most of the time and sang a little bit*, which entails *p* and rejects an implied enrichment of *p, q,* namely *she danced all of the time*, which entails *p*. We will claim that we cannot be satisfied with hedging solely along the propositional level, and that the hedging along the discourse level is central to the semantics of *be-gadol*.

On the discourse dimension, the use of *be-gadol* indicates that *p* is not the most helpful answer relative to the roles and goals of the participants in the discourse. The concept of roles and goals can be represented more concretely by resorting to the speaker's decision problem. For example, in (3) the question *What did Rina do in the party?* can be asked to learn about Rina's preferences, e.g. in a context where John wants to throw Rina a party and has to choose between different activities for the party. In this context John can inquire about Rina's preferences to resolve a decision problem regarding which party activities she might enjoy.

The better answer p_{best} is required to be the most helpful answer to the QUD given this decision problem, but *p* is still close to being the most helpful answer given this decision problem. For example, in (3) (*she be-gadol danced)* we suggest that any possible p_{best} (e.g. *she danced most of the time and sang a little bit*) is not only more informative than *p*, but also the most helpful answer relative to the roles and goals of the discourse participants in any possi-

ble context. In contrast, *she danced* is not the most helpful answer, but is nonetheless close to being most helpful. For example, *be-gadol* will be useful when for contextual purposes it is significant whether Rina danced all the time or did something else some of the time, because this fact might have consequences for the participants' decisions. The rejection of the strong enrichment that Rina danced all the time is helpful. Moreover, the fact that *p* is not the best answer but close to it emphasizes that the missing details of p_{best} may be relevant as well. We will illustrate all these notions with concrete examples below.

If this intuition is on the right track then the theoretical challenge at hand is how to capture 'degrees of answerhood' to the QUD, what exactly a *best* answer is, what makes *p* less good than p_{best}, and in what sense is *p* close to the best answer.

To answer these questions we will use the literature on resolving answers (e.g. Ginzburg 1995), and tools from decision theory (e.g. van Rooij 2003). In the following section we lay out our proposal and explain and illustrate the above intuition in further detail.

2.1. Proposal

As mentioned above, we suggest that the hedging operation with *be-gadol* results from the interaction between hedging along two levels – in the discourse level and the propositional level. In all its uses *be-gadol p* indicates that *p* is not the best answer to the QUD, and that a better answer (p_{best}) to the QUD is true. This better answer p_{best} is related to *p* in two ways:

First, p_{best} is related to *p* since we take p_{best} to be more informative than *p*. In particular we require that p_{best} entails *p* and at the same time rejects a strong (but implied) enrichment of *p*. Defining p_{best} in this way is the first component in our proposed lexical semantics for *be-gadol* which leads to the hedging effect on the propositional level.

Second, we take *p* to be a proposition which is not the best answer to the QUD, but which is still 'close' to such a best answer, p_{best}. The notion of a 'best' answer, and of being 'close' to a best answer, will be defined relative to a decision problem. This second component leads to the hedging on the discourse level.

Given these two components our proposal is found in (5).

(5) $||\textit{be-gadol}||_{w,dp} = \lambda p \in QUD.\ \lambda w\colon \exists q\ [q \in QUD \wedge p \sim\!\!> q \wedge q \subset p \wedge \exists BEST \neq \emptyset$, where
$BEST = \{p_{best} \in QUD\colon p_{best} \subset [p \wedge \neg q] \wedge best_{dp}(p_{best}) \wedge Close_{dp}(p_{best}, p)\}]$.
$\exists p_{best} \in BEST\ [w \in p_{best}]$

In words, *be-gadol* (relative to w and a decision problem dp) combines with its prejacent, i.e. a proposition *p* which is a member of the QUD, and with a world w. It is defined if and only if there is a proposition *q* which is also a member of the QUD, which is stronger than (i.e. asymmetrically entails) *p* and is implied by it (i.e. it is a strong enrichment of *p*). In addition, there is a nonempty set, BEST, of propositions p_{best} which are also members of the QUD, such that for any p_{best} in this set BEST: p_{best} entails *p* and the negation of *q* (i.e. entails *p* and

rejects its stronger enrichment), p_{best} is 'best' relative to the decision problem and p is 'close' to p_{best} relative to the decision problem.

Clearly, the notion of being a 'best' answer, and being 'close' to such a best answer, relative to the decision problem are the main novel parts of the definition. To explain these notions, and their necessary presence in the lexical entry of *be-gadol*, we will first consider in the following section the hedging effect of *be-gadol* along the propositional domain, i.e. the fact that some true answer, p_{best}, is required to entail p and reject a strong (implied) enrichment of it. We then show that this hedging effect is not enough to account for the full range of facts concerning *be-gadol*. To get a fuller picture we characterize precisely the 'best' and 'close' components in the lexical entry in (5), which capture the hedging along the discourse domain, and demonstrate how adding these components account for the remaining data.

2.2. Illustrations and support

To try and account for the data by assuming hedging along the propositional dimension alone, consider the following examples. In every example (a) gives the prejacent, p, (b) gives the strong enrichment of p, q, and (c) gives an example of a possible p_{best}. Example (6) includes the gradable adjective *clean*. We follow von Stechow (1984) in assuming that gradable adjectives denote relations between degrees and objects, and in the positive form as in (6a) a null morpheme *pos* introduces a standard degree, s.t. (6) is true iff the room is at least as clean as the standard (formally, $\exists d\ [d \geq stand_{clean} \wedge clean(the\ room,d)]$). Kennedy and McNally (2005) assume that the standard in adjectives like *clean* is maximal, although pragmatical shifts to lower standards are often used in actual context. Rotstein and Winter (2004), however, assume an interval of degrees in the maximal end of the scale, and hence semantically the standard is not always maximal. We follow Rotstein and Winter (2004) as well as McNally (2011), and Sassoon and Toledo (2011) in assuming that the standard for *cleanliness* can be identified with a point which is not exactly at the maximum of the scale.

(6) The room is be-gadol clean
   a. p: The room is clean
   b. q: The room is maximally clean
   c. p_{best}: The room is clean except for the windows

(7) Rina be-gadol danced
   a. p: Rina danced
   b. q: Rina danced all of the time (e.g. of the party)[2]
   c. p_{best}: Rina danced most of the time and sang a little bit.

(8) Mary is be-gadol a book editor
   a. p: Mary is a book editor
   b. q: Mary edits books in general.
   c. p_{best}: Mary is a book editor specializing only in medieval history and religion

[2]We think this implication is derived as a scalar implicature, namely that Rina's dancing eventuality was the only activity in the party. We don't go now into how precisely this is derived (e.g. how exhaustification scopes wrt the existential over events).

In all of these cases *be-gadol p* indeed indicates that a more informative answer, p_{best} (as in the (c) part of (6-8)) is true. This true answer indeed entails *p* (as in (a)) and rejects a stronger enrichment of it (as in (b)).

So far, everything seems to be working. However, as we said above, we require also that p_{best} be a 'best' answer to the QUD and that *p* be 'close' to it. To justify the two extra conditions on p_{best}, we turn to look at a case where hedging in the propositional level isn't enough, and using it alone yields over-generalization.

(9) Context: John is interested in arranging a birthday party for Rina and wants to have activities that she enjoys. He knows that Sarah and Rina went to a party last week. To be able to choose the activities that Rina will most enjoy in her birthday party, John wants to learn about her preferred party activities. For that purpose John is asking Sarah about Rina's behavior during the party. Sarah is aware of John's purpose. She saw that Rina danced most of the time of the party last week, sang for a little while, and spoke with the barman.

John: What did Rina do in last week's party?
a. Sarah: She be-gadol danced
b. Sarah: #She be-gadol danced and sang

(9a) and (9b) have correlating *p*'s, *q*'s and p_{best}*'s*, given in (10) and (11).

(10) Sarah: She be-gadol danced
a. p: she danced
b. q: she danced all the time
c. p_{best}: Rina danced and sang.

(11) Sarah: #She be-gadol danced and sang
a. p: Rina danced and sang
b. q: The only things she did in the party was to sing and dance.
c. p_{best}: Rina danced, sang, spoke with the barman.

In both cases there is a p_{best} which entails *p* and rejects a strong enrichment of it. Nonetheless, given the details of the context in (9), *be-gadol* is felicitous in (9a), and infelicitous in (9b).

To account for the infelicity of (9b), we propose the following intuition: *She be-gadol danced and sang* is odd because the information that Rina spoke with the barman does not seem to add anything relevant to what John is interested in. Thus, although a more informative answer is possible, e.g. the proposed p_{best} in (11c), it cannot count as a real appropriate 'best' answer. In fact, given John's goals, it is not better than the prejacent *She be-gadol danced and sang* – it does not add any relevant information.

Moreover, changing the goals of the participants in the discourse can affect the felicity of *be-gadol* in these examples. The answer in (12) is judged better in the following situation:

(12) Context: John suspects that Rina is a spy. He sends Sarah, a private detective, to follow Rina and report back to him. John and Sarah are talking as Sarah recounts Rina's actions during the party. Rina danced most of the time of the party, sang for a little while and spoke with the barman.

John: What did Rina do in last week's party?
Sarah: She be-gadol danced and sang

Intuitively, Sarah's answer is felicitous because the information that Rina spoke with the barman might be relevant under this context. That is, unlike what we saw in (9b), given these roles and goals of John and Sarah, *p* in (12) is really not the best answer anymore, so *be-gadol* achieves the hedging effect in the discourse domain. The conclusion from this is that characterizing the hedging effects of *be-gadol* using only informativity tools is not enough, and we need an additional type of hedging which is sensitive to the roles and goals of the participants in the discourse.

The observations we have made regarding *be-gadol* seem very similar to observations made in e.g. Ginzburg (1995) and van Rooij (2003) regarding answers to questions. Similarly to what we observed with *be-gadol*, these theories noted that for answers to be considered good answers, it is not enough that they are fully informative. Rather, they need to be also 'helpful' relative to the roles and goals of the participants in the discourse. Hence, we now turn to relevant details of these theories, so we can use the tools they developed to formally capture the complex hedging effect of *be-gadol*.

2.3. Hedging along the discourse domain

According to Ginzburg (1995), even if an answer is exhaustive (i.e. maximally informative), it is not always helpful (enough). In addition to that, a helpful answer can be non-exhaustive. Consider (13), inspired by example (15), on p.469 in Ginzburg (1995).

(13) Context: A scientist is invited to give a lecture at a university he's unfamiliar with. The scientist asks who will attend the talk, with the goal of giving a good and clear talk.

Question: Who is going to attend the talk?
a. Answer 1: (Provides an exhaustive list of names).
b. Answer 2: A number of cognitive phoneticians and Willshaw-net experts.

While (13a) is exhaustive, it is not helpful given the scientist's goal and information state. In contrast, (13b) is not exhaustive, but is helpful given A's and information state.

Van Rooij (2003) formalizes sensitivity to goals and information states using notions from probabilistic decision theory. A decision problem of an agent can be modeled as a triple <P,U,A>, where *P* is the agent's probability function which represents the agent's beliefs about the world, *U* is the utility function, which assigns each action a value, representing the agent's goals, and *A* is the set of alternative actions the agent considers. The expected utility (EU) of an action *a* in *A* is defined as in (14):

(14) $EU(a) = \Sigma_w P(w) \times U(a,w)$

In words, the expected utility of an action *a* is the result of the summing operation over worlds of the probability of being in each specific world *w*, times the utility value of the action *a* in that world.

To make the best choice between these actions the agent can ask a question. The answer to this question can help resolve the decision problem, i.e. help the agent calculate which action is the most likely to be most helpful. Given this calculation, the agent knows which is the action with the highest expected utility. Given an answer C, the maximal expected utility of the decision problem equals the maximal value of summing the conditional probability of each world given C, times the utility of each action *a*:

(15) $Max(\{\Sigma_w P_C(w) \times U(a,w): a \in A\})$

Given this, the agent can know which is the action with the highest expected utility (EU) given the answer C.[3]

Finally, for van Rooij, an answer which resolves the decision problem, a 'resolving answer', is one which leaves the agent with exactly one action with the highest EU.

With these theoretical tools at hand, we can now analyze the second component from our proposed lexical entry, the one where *be-gadol* hedges along the discourse domain.

Consider again our proposed lexical entry, repeated here as (16).

(16) $||\textit{be-gadol}||_{w,dp} = \lambda p \in QUD.\ \lambda w: \exists q\ [q \in QUD \wedge p \rightsquigarrow q \wedge q \subset p \wedge \exists BEST \neq \emptyset$, where
$BEST = \{p_{best} \in QUD: p_{best} \subset [p \wedge \neg q] \wedge best_{dp}(p_{best}) \wedge Close_{dp}(p_{best},p)\}]$.
$\exists p_{best} \in BEST\ [w \in p_{best}]$

Above we suggested that the prejacent of *be-gadol, p,* is neither the most informative answer to the QUD, nor the most helpful, but still 'close' to being most helpful. This is captured in (16) by taking another answer, which we called p_{best}, to be the true answer to the QUD which entails *p* and rejects a stronger enrichment of *p*. In addition, p_{best} must be also best relative to the decision problem, and close to *p* relative to the decision problem.

We now define the notion of a best answer relative to the decision problem, and being 'close' to such a best answer more precisely, as in (17) and (18) respectively:

(17) p_{best} is the best answer relative to the decision problem iff
a. It is resolving, i.e. leaves us with exactly one action, which has a maximal EU[4]

[3]Benz (2006) and Benz and van Rooij (2007) have argued that a decision problem must also take into account information about the speaker in order to define "goodness" of answers.

[4]Formally, following ideas in van Rooij (2003), we require that $|\mathbf{p_{best,\ A^*}}| = \mathbf{1}$, where $p_{best,\ A^*}$ is the set of propositions a*, of the form "you should choose action *a in A*", which are consistent with p_{best} ($p_{best} \cap a^* = \emptyset$). If the cardinality of this set is 1, then p_{best} is said to be 'resolving' – i.e. it leaves the agent with exactly one action: The one with the highest expected utility.

b. And any other distinct answer is either non-resolving, or leads to an equal or lower maximal EU:
$\forall s \in QUD\ [s \neq p_{best} \rightarrow [[\neg \text{resolving}(s)] \vee \text{maxEU}(P_{pbest},U,A) \geq \text{maxEU}(P_s,U,A)]]$, i.e. no other answer is more helpful, and leads to a higher maximal EU.

(18) p is 'close' to p_{best} iff
p is also resolving i.e. it also leaves the agent with single action, one with the highest EU, and $\text{Small}(\text{maxEU}(P_{pbest},U,A) - \text{maxEU}(P_p,U,A))$

To explain (18) in words, *p* is close to p_{best} iff the maximal EU of the action with the maximal EU that we are left with after learning *p,* is lower than the one we are left with after learning p_{best} but only slightly so. This boils down to a situation where the agent is a bit more certain (less at risk) that the action chosen is the right one after learning p_{best} than after learning *p*.[5] In this case, p_{best} leaves us with one action, with the highest EU, and *p* leaves us with one action, but with a slightly lower EU.

3. Illustration

Consider the example from (3), repeated here in more detail as (19). John asks Sarah a question, and Sarah, who is aware of the decision problem, answers with *be-gadol*.

(19) Context: John wants to throw a party for Rina, and plans an activity that she likes, out of the three potential activities in the following set of actions *A: {prepare a dancing activity, prepare a karaoke competition, prepare a jumping activity}*. John knows that Sarah and Rina participated in a party last week. To learn about what Rina likes to do in parties he asks Sarah, thinking that if Sarah tells him what Rina like to do he would know better which action to choose.

John: What did Rina do in last week's party?
a. Sarah: She be-gadol danced
b. Sarah: #She be-gadol danced and sang

The felicity of (a) and the infelicity of (b) hold, for example, in the following two scenarios:

(20) The 'temporal' scenario: Rina danced most of the time of the party, sang for a short while, and spoke with the barman.

(21) The 'significance' scenario: Rina danced most enthusiastically in the party, sang, as she usually does, and spoke with the barman. Dancing didn't occupy most of the time.

In both scenarios, we can find an answer, p_{best}, which meets the conditions in our definition, namely (22a) and (22b):

[5]There can be additional ways in which *p* can be 'close' to p_{best} (cf. van Rooij 2003). E.g., if *p* leaves us with more actions in the decision problem than p_{best} does, but not with many more. Alternatively, if given the decision problem our goal is just to gain more information, but *p* leaves us with less information than p_{best} does, but not with much less. We do not further discuss these options in the present paper.

(22) a. p_{best} for the temporal scenario: Rina danced most of the time and sang for a while.
b. p_{best} for the 'significance' scenario: Rina danced extremely enthusiastically and sang.

In both cases p_{best} is a true answer which entails p (*she danced*) and at the same time rejects a strong implication of it (*she danced all of the time*). Moreover, it is indeed the best answer with respect to the decision problem because learning it we can rather safely conclude that the dancing activity is the activity she will enjoy most. More formally, learning one of these p_{best} answers, we are indeed left with one action with the highest expected utility (i.e. the one which is most likely to be most helpful).

Assuming that this is $p_{best,}$ Sarah's answer in (19a) (*She be-gadol danced*) is felicitous, since although the prejacent *She danced* is not the best answer, it is close to it. Despite the fact that like p_{best}, p also leads to choosing dancing as the activity for the party, the expected utility of this answer is a bit lower. Upon learning *She danced*, all John would know is that there was a dancing eventuality by Rina in the party. However, John will not know how long it lasted, whether it was significant or not, whether there were other relevant events which lasted more or were more significant, etc. Thus, learning that Rina danced in last week's party, John might choose the dancing competition for her party, but he will be less certain about whether this is the best choice to make. The EU of 'prepare a dancing activity' is still maximal compared to the other activities in the set, but lower given p than given p_{best}.

The infelicitous case in (19b) is discussed using the same context and decision problem as in (19). In this case, *She be-gadol danced and sang* seems odd since here p, namely *She danced and sang,* is not 'resolving': i.e. it leaves us with TWO possible actions (namely dancing and singing), and not one.

Moreover, if p does not end up less helpful than p_{best}, we get infelicity even when it IS resolving. To demonstrate this we will modify the context in (19) for the one in (23).

(23) Context: John wants to choose a pair of activities for the party, out of the following pairs: *{singing & dancing, dancing & jumping, singing & jumping}*. In reality, Rina danced for a long time, sang for a short while and spoke with the barman.

John: What did Rina do in last week's party?
Sarah: #She be-gadol danced and sang

Why is Sarah's answer odd in this case? After all, p (*She danced and sang*) leaves us with one action with the highest EU (namely choosing the pair 'dancing and singing' for the party). One way to think about this is that in this case there is no appropriate p_{best}. The more informative *She danced for a long time, sang for a while and spoke with the barman* cannot function here as p_{best} since, given John's goals it does not have a higher expected utility than *She danced and sang*. In other words, it is not more helpful, and hence, *be-gadol*'s required hedging in the discourse domain cannot be achieved.

To demonstrate the proposal further, we will consider two more cases, the 'not enough details' use in (24) and (25), and the 'approximative' use in (26) and (27).

(24) Context: John is having a sports event next week and wants to know whether to invite Danny (who is in general a good candidate for consideration). To decide on this matter John is asking about Danny. For example:

John: Is Danny healthy?
Mary: He is be-gadol healthy

In this case, we take *p, q* (the rejected strong enrichment of *p*), and p_{best} to be as in (25):

(25) a. p: Danny is healthy to a degree which is at least as high as the standard of health in the context.
b. q: Danny is healthy to a maximal degree
c. p_{best}: Danny is healthy except for a cold.

Learning *p* will encourage John to invite Danny. However, John will not be completely certain this is the right action to choose. Even given *p,* it is still possible that Danny is not maximally healthy. If this is the case, and given that John doesn't know what keeps Danny from being completely healthy, John may be worried that participating in the competition will not be good for him. However, given p_{best}, John will be in a better position to decide exactly what to do. No matter whether John decides to invite Danny or not, John will now be more certain about making the right call, and hence the expected utility (EU) of the action with the maximal EU given p_{best} will be higher than the maximal EU given *p*.

(26) Context: Mary is looking to hire employees for her new publishing house, where most books deal with history, sociology, etc. Mary thinks we may want to interview Danny.

Mary: What does Danny do? What's his profession?
John: He is be-gadol a book editor (I'm not sure exactly which type of books he edits)

In this case, we take *p, q* (the rejected strong enrichment of *p*), and p_{best} to be as in (27):

(27) a. p: Danny is a book editor
b. q: Danny can edit all types of books (i.e. he doesn't specialize in any specific field)
c. p_{best}: Danny edits history books

As before, after learning *p* Mary will probably tend to interview Danny. However, Mary will not be completely certain this is the right action to choose, since it is possible that Danny specializes in editing a type of books which is not relevant for the publishing house. Even so, it may be still helpful to interview Danny. Learning p_{best} Mary would be in a better position to decide exactly what to do, and to decide whether to interview Danny or not. Mary would then be more certain of the chosen action. Here too the EU of the action with the maximal EU given p_{best} is a higher given p_{best} than it is given *p*.

4. An apparent counterexample and a suggestion for a solution

As mentioned above in section 1.1, *be-gadol* also yields a use which seems to pose a problem for our theory. An example is (28).

(28) Context: Danny tried to get into medical school. To do that one has to pass an exam. The passing grade for the exam is strictly 80 and, generally, if you pass you get accepted. Mary wants to know whether to congratulate Danny or not. She asks Sarah about his current situation.

a. Mary: Well, what about Danny? Did he pass the exam?
b. Sarah: be-gadol hu avar, aval hayu yoter miday muamadim tovim
in.big he passed, but were too many candidates good
"Be-gadol he passed, but there were too many good candidates."

According to the analysis we proposed above, *be-gadol p* is felicitous iff *p* is not the best answer to the QUD, i.e. not most informative, and not most helpful. This is so because there is another answer, p_{best} which IS the most informative and helpful answer.

The problem in (28b) is *be-gadol* is felicitous although its prejacent seems to be already the best answer: Most informative (i.e. does not imply some enrichment *q* which is negated by p_{best}), and also most helpful, since Mary is left with one action (i.e. congratulate Danny), with the highest EU.

We suggest that this kind of example can be nonetheless covered by our proposal for *be-gadol.* To explain this we use an independently motivated mechanism proposed in Simons et al (2010), which has been proposed in the literature regarding the connection between projective meanings and questions under discussion.

Simons et al (2010) suggest that what is projected is 'not at issue', where such 'not at issue' material is often, but not always, a presupposition. Projection of presupposition can be blocked when the content is at-issue relative to the QUD, i.e. is relevant to determine which of the answers is true.

However, they bring some examples where it seems that at issue content is nonetheless projected. Consider (29) below, a slight variation of the original example in Simons et al., which is odd without its context.

(29) Context: Chloe is writing invitations for her birthday party to kids in her class. Her mother notices that all of the invitations are to girls.

a. Mother: Are there any boys in your class?
b. Chloe: I don't like the boys in my class.

The direct answer to the mother's question is that there are boys in Chloe's class, and hence this answer constitutes at issue content. Importantly, this content is projected in Chloe's answer. Simons et al. (2010) claim that this is not a counterexample for their generalization that content does not project when at-issue. The claim they make is that (29b) is not an answer to the mother's explicit question, but rather to a broader QUD. Put in other words, the immediate QUD is stated explicitly in (29a), but there is also an implicit question, namely, *Why aren't you inviting any of the boys in your class to your party?* Making the accommodation of this broader implicit question would explain what is going on in the sequence in (29). Giv-

en this broader question the sequence is felicitous since "there are boys in my class" is not at issue with respect to this implicit broader question, and hence is projected in (29b).

Considering the shift from the immediate QUD to a broader QUD, the example in (28) with *be-gadol,* can be reconsidered. Although the explicit question asked by Mary is *Did Danny pass?* this is just a subquestion of a broader, implicit question *Did Danny get into medical school?* Interpreting the answer in (28b) as an answer to the broader QUD, *p* is once again a not-best answer, as required by our theory. The best answer to the broad QUD would be an answer along the lines of *Danny passed the test but didn't get accepted,* as demonstrated in (30).

(30) Rina: Danny be-gadol passed, but there were too many good candidates.
  a. p: Danny passed, i.e., he most likely got accepted, but maybe not.
  b. q: Danny got accepted.
  c. p_{best}: Danny passed but didn't get accepted.

In the context of the broader implicit QUD, if Mary wants to decide whether she should congratulate Danny or not on his acceptance to medical school, (30a) would not resolve this decision problem in the best way, since given p_{best} Mary will be more certain in her decision (namely not to congratulate Danny). In addition, as required, the best resolving answer (30c) negates the strong implication *q* (30b) of *p* (30a). Despite not being the best answer, *p*, namely *he passed*, is still 'close' to being such a best answer, since its EU is just a bit lower than that of p_{best}. This is because passing the exam almost always leads to getting accepted, so learning *p* Mary is close to being certain that she chooses the best answer (in this case, that she can congratulate Danny).

Also worth mentioning is the fact that this case is the only one where it is necessary to continue the *be-gadol* sentence with ***but*** (e.g. *but there were too many candidates*). *But* creates here a strong counterexpectational effect (see e.g. the QUD-based analysis of *but* in Toosarvandani 2014).

One possible explanation for this is that this is the only case where p_{best} must leave us with the opposite action (e.g. not congratulate Danny) than the one we are left with after learning *p* (congratulate Danny). However, the precise nature of the relation between *but* and *be-gadol,* as well as other interesting interactions / parallels between the hedging operation of *be-gadol* and the semantics of *but* still await further inspection.

To conclude, we call this use the "change your question" use of *be-gadol.* For this use to be available *p* should be the complete answer to the immediate QUD (e.g. *Did he pass?*). To allow the hedging effect on the propositional level there must be a higher, super-QUD in the context (e.g. *Did he get accepted*?), or such a question must be available for accommodation. (see also Büring 2003, Simons et al 2010 on other constructions where such accommodation is necessary).

5. Summary and open questions

We looked at a range of hedging effects exhibited by *be-gadol*, which cannot be easily reduced to a single operation. To account for the full range of data in a unified way we suggested that *be-gadol* lexicalizes a hedging operation over answers to questions. It indicates that *p* is not the best answer to the QUD, and that there is a better answer, p_{best}, which is true.

Crucially, the hedging operation is done on two levels. First, the propositional level where p_{best} is more informative than *p,* since it entails *p* and rejects a strong enrichment of it. Second, the discourse level, where p_{best} is more 'helpful' than *p* given the goals and information state of the participants (though *p* is 'close' to being most helpful). This aspect of *be-gadol* can be captured by resorting to notions from the literature on resolving answers, and using formal tools from decision theory. To the extent the analysis is on the right track, it supports the linguistic reality of these notions of tools in a new way.

There are several open questions that still need to be accounted for and directions to check. For example, we are still checking whether the analysis will work for all examples with *be-gadol,* across contexts and decision problems, whether there are other/better ways to define what 'the best answer' to the QUD is. In addition, perhaps there are also other possibilities to make *p* 'close' to p_{best}, e.g. by ending up with more, but not many more actions after learning *p* than after learning p_{best}.

Furthermore, we aim to explore whether there are other members of this 'discourse hedgers' family. Are there more hedgers that are sensitive to how helpful their prejacent is relative to the goals of the participants? Potential candidates for such hedgers are *theoretically / in theory / in principle / by and large / basically*…etc. If these are indeed members of the same family, it would be highly significant to define their core semantics and characterize the parametric differences between them.

Maybe the most central issue to understand is the reasons for which a cooperative speaker would use *be-gadol* at all. If the speakers knows the best answer to the question, why use *be-gadol p,* and indicate that *p* is NOT the best answer? Perhaps what *be-gadol* actually does is signal that while its prejacent *p* is the best answer to the immediate decision problem, there is a 'higher', more elaborated decision problem which still needs to be resolved. If this is so, we need to develop a model with hierarchies of decision problems, similarly to hierarchies of QUDs (cf. Roberts 1996, Büring 2003 on Contrastive Topics). We leave this issue for further research.

References

Benz, A. (2006). Utility and relevance of answers. In A. Benz, G. Jäger, and R. van Rooij (Eds.), *Game Theory and Pragmatics*, pp. 195–221. Palgrave Macmillan.

Benz, A. and R. van Rooij (2007). Optimal assertions and what they implicate. *Topoi 26(1)*, 63–78. doi:10.1007/s11245-006-9007-3.

Büring, D. (2003). On D-trees, beans and B-accents. *Linguistics and Philosophy 26(5*), 511–545.

Ginzburg, J. (1995). Resolving questions, Parts I and II. *Linguistics and Philosophy 18(5)*, 459–527, *18(6)*, 567–609.

Kennedy, C. and L. McNally (2005). Scale structure, degree modification, and the semantics of gradable predicates. *Language 81(2)*, 345–381.

McNally, L. (2011). The relative role of property type and scale structure in explaining the behavior of gradable adjectives. In R. Nouwen, R. van Rooij, U. Sauerland, and H.-C. Schmitz (Eds.), *Vagueness in Communication*, pp. 151–168. Springer.

Nakanishi, K. and M. Romero (2004). Two constructions with *Most* and their semantic properties. In K. Moulton and M. Wolf (Eds.), *Proceedings of NELS 34*, pp. 453–467. Amherst, MA: GLSA.

van Rooij, R. (2003). Questioning to resolve decision problems. *Linguistics and Philosophy 26*: 727–763.

Roberts, C. (1996). Information structure in discourse: Towards an integrated formal theory of pragmatics. In J. H. Yoon and A. Kathol (Eds). *OSU Working Papers in Linguistics 49: Papers in Semantics*, pp. 91–136.

Rotstein, C. and Y. Winter (2004). Total adjectives vs. partial adjectives: Scale structure and higher-order modifiers. *Natural Language Semantics 12*, 259–288.

Sauerland, U. and P. Stateva (2007). Scalar vs. epistemic vagueness: Evidence from approximators. In M. Gibson and T. Friedman (Eds.), *Proceedings of the 17th Semantics and Linguistic Theory Conference (SALT 17)*, pp. 228–245. Ithaca, NY: CLC Publications.

Sassoon, G. and A. Toledo (2011). Absolute vs. relative adjectives: Variance within vs. between individuals. In *Proceedings of SALT 21*, pp. 135–154. Ithaca, NY: CLC Publications.

Simons, M. J. Tonhauser, D. Beaver, and C. Roberts (2010). What projects and why. *Semantics and Linguistic Theory (SALT) 21*, 309–327.

von Stechow, A. (1984). Comparing semantic theories of comparison. *Journal of Semantics 3*, 1–77.

Toosarvandani, M. (2014). Contrast and the Structure of Discourse. *Semantics and Pragmatics*.