Salience and Questions Under Discussion

Eric McCready
Aoyama Gakuin University

September 2012

Abstract

The idea of questions under discussion is highly useful. Still, it remains unclear exactly how speakers can go about determining what question is at issue. This paper proposes a way to recover the question under discussion, conceptualized as a decision problem, by first determining a set of salient questions by using tools from contextualist epistemology, and then selecting one question from that set via a notion of maximal utility change.

1 Introduction

A *question under discussion* can be understood as the current goal of a linguistic interaction: what one of the conversational participants is trying to find out or ought to do in order to realize her goals. This notion has proved to be a highly useful tool in (formal) pragmatics since its introduction by Roberts (1996). For instance, among other applications, it has been used (more or less directly) to help understand relevance in the Gricean sense, which can be viewed as the extent to which a given utterance helps resolve uncertainty with respect to the question under discussion (van Rooij, 2003a); to give a characterization of alternative sets (Roberts, among many others); and in defining a notion of at-issueness as that content which directly contributes to answering the current QUD (Simons *et al.*, 2011). It is thus safe to say that it is a crucial piece of many insightful analyses.

However, some foundational issues surrounding the idea of a QUD remain ill-understood. The main problem is simply that we have no independent way to determine what the particular question is that speakers are addressing in a conversational situation. There is no theory on the market (to my knowledge) that allows

specification of the QUD in a given context. In the literature that makes use of the QUD notion, it is nearly always assumed that the question under discussion is given. Once a particular QUD is assumed, it can of course be used in various ways. But the question remains: how do we know what the QUD is?

One reason this issue remains unaddressed is, I think, that most linguists feel it to lie beyond their proper domain of study. The idea seems to be that how the QUD is determined should follow from general psychological processes, which are not in the purview of linguistic theory. However, this puts us as working linguists in a bit of a worrisome position. The foundations of the QUD remain a black box. Without some explication of the workings of this mechanism, it is hard to feel justified in using QUDs, or at least in assuming that we are genuinely able to take them as given. For what bars us from choosing QUDs which are convenient for our current theories, and thus diminishing the predictive power of such theories? Worries like this make it look imperative that we develop, or at least explore, a theory of selection for questions under discussion. Doing so is the goal of this paper.

We can then ask the following question as a preliminary: What is a QUD, in general? In the literature it is generally assumed that it is a question (perhaps this is obvious), and therefore, semantically, a set of propositions. Here I will take a somewhat different and more general approach. I will assume that a QUD amounts to a contextually salient decision problem. The idea is that a semantic question is in essence a decision problem defined over which element of the question to believe.[1] This assumption is not solely my own; it is also made, for example, by Davis (2009), for whom it allows a generalization of his account of the Japanese particle *yo* to both recommendations to action and recommendations to believe a particular answer to a salient question. The main question to be addressed in this article then becomes the nature of the metric which decides what counts as a salient DP.

But we are far from out of the woods. It is notoriously difficult to say exactly how it is determined what counts as salient here. The problem is very general. It arises outside of the (direct) context of QUDs in many places: for instance, in attempts to specify sets of relative alternatives in focus semantics or for Gricean implicature or in some contextualist approaches to knowledge attributions, in anal-

[1]There are possibly problematic aspects of this analysis relating to the intentionality of belief: can I choose to believe a particular answer, or does such belief happen automatically? This question is too large to address here, but I (choose to) believe that belief has at least a partial active component. For instance, one may choose to believe an instance of testimony, or not. Thanks to Ede Zimmermann for discussion on this point.

yses of so-called unarticulated constituents, in contextually based restrictions on quantifier domains; indeed, the issue seems to crop up in some form in almost all areas of formal pragmatics. Without a theory of salience, though, we are in some trouble, for the predictions are not clear, as stated above. In a slogan: without a theory of what counts as a salient alternative, we don't have a complete theory.

The plan of this first step toward such a theory is as follows. The goal is a theory which can specify exactly what decision problems are under consideration. I'll therefore start out in the next section by spelling out what decision problems are and, more interestingly for our purposes here, how one can go about altering the optimal solutions to them, which leads immediately to a characterization of 'useful talk.' This characterization opens up a line which I will develop into a means for speakers to – normatively – select from a set of salient decision problems. The problem then is to make that set available. My strategy will be to borrow some tools from contextualist epistemology, in particular versions of such epistemologies which make use of relevant alternatives. After exploring one such theory in detail, I modify aspects of it for the case of decision problems. Here, I will specify some baseline parameters on salience; unfortunately not enough is known about how salience works and about what constraints exist on it to provide a full theory at this point (cf. Hawthorne 2004). Still, what we will end up is enough to work through a few examples, with which I close the paper, after discussing some loose ends and future prospects.

2 Decision Problems and Useful Talk

The first step, then, is to give a more detailed definition of decision problems.

The traditional formulation of decision problems is a 4-tuple as in (1); details can be found in e.g. Jeffrey (1983); van Rooij (2003a). Here, relative to an agent A, W is a set of worlds, $\mathcal{A}$ is a set of possible actions, $\mathcal{P}$ is a probability function satisfying the usual postulates in (2), and U is a utility function. The set of actions $\mathcal{A}$ is externally determined. The postulates in (2) state informally that the probabilities of all possibilities sum to 1, meaning roughly that they must cover the entire set of worlds, and that non-overlapping possibilities have independently determined probabilities. Finally, U is a function from outcomes to values, defined here as natural numbers. As is well-known, the relationship between the values output by U is important, but the precise values given by U are not; indeed, as is also well-known, utilities are invariant under positive affine transformations (cf. Myerson 1991 or any other text on game theory).

(1) $D = \langle W, \mathcal{A}, \mathcal{P}, U \rangle_a$

(2) $\mathcal{P}$ satisfies:

a. $\mathcal{P}(W) = 1$

b. If $\varphi \cap \psi = \varnothing$, then $\mathcal{P}\varphi + \mathcal{P}\psi = \mathcal{P}(\varphi \vee \psi)$

The output of U applied to an action $a \in \mathcal{A}$ depends on both my preferences and what the circumstances are. My choice of eating strawberry rather than chocolate ice cream will give greater or lesser benefit depending on whether (a) I prefer strawberry to chocolate and (b) whether the strawberry ice cream has gone bad. This means that the action the agent should pick depends on what she takes the probability of a given circumstance to be. This observation is captured by a notion of expected utility.

(3) $$EU(a) = \sum_w U(a,w) \times P(w)$$

So the utility of action a is the sum of the utility of the action in all the worlds deemed possible by the agent (on the natural and common assumption that we are dealing with subjective probability).

We can rank possible actions (within a DP, or more precisely within the set $\mathcal{A} = 2(DP)$) according to their expected utility.

(4) $a \succ_{dp} a'$ iff $EU(a) \succ EU(a')$.

This notion in turn allows definition of a total ranking of available actions. This ranking is nothing more than an ordered list of the preferences of the agent across possible outcomes.

(5) $Ord_{EU}(D)$ defined as $\langle \mathcal{A}_D, < \rangle$, the set of actions in D together with a (total) ordering on them.

This ordering, in turn, is dynamic, changing as the agent gains information. In fact, there are three ways to change expected utility of a via the learning of new information. One may learn something that affects our estimate of the benefit of a, which corresponds to a change in the utility function U. One may learn something about the world that affects our estimate of the likelihood of the action leading to a positive outcome, which corresponds in a change in the probability function $\mathcal{P}$, possibly via a change in the set of available worlds W. Or, finally, one can learn something about our possible choices that causes us to change our views

of the best action available, which of course corresponds to a change in the set of available actions $\mathcal{A}$.

This observation leads to an obvious conclusion. Given a particular decision problem, any utterance useful to resolving that problem should induce one of the three kinds of change catalogued above. This is a definition of 'useful talk' (given that the other Gricean maxims, such as in particular Quality, are followed; changes induced by false information are obviously not useful). I want to suggest that this kind of relevance should factor into salience judgements, in particular into determining what problem one should select from a set of possibilities. But of course this selection requires the availability of such a set. I turn now to providing one. After doing so, I will return to the issue of usefulness.

3 Salience

What is an appropriate set of possibilities from which to select a decision problem? Presumably we need a set of decision problems that are salient enough to qualify as candidates for the decision problem at issue, i.e. the question under discussion. Making this available requires a theory of what decision problems are salient. I will provide the beginnings of such a theory by making use of a notion of salience that has had some currency in philosophy.

Salience and relevance have been extensively considered in contextualist epistemology. One of the major problems studied in epistemology is the nature of knowledge and the location of the dividing line between knowledge and 'mere' belief. The usual way to determine what counts as knowledge is, effectively, intuitions about knowledge attributions. The standard methodology asks whether a particular individual can be said to have knowledge in a particular situation.[2] Contextualism about knowledge is a way of spelling out the truth conditions of knowledge attributions. Many authors (a classic reference is DeRose 1992) have argued that the conditions under which a knowledge attribution is true are sensitive to some contextual feature(s). One implementation of this view makes knowledge dependent on the space of possibilities in awareness: roughly, the idea is that something may not be known when certain possibilities are taken into account, but if such things are not under consideration, it may be known after all. The truth-conditions of knowledge attributions are thus dependent on the available (counter)possibilities. Specifying these possibilities amounts to specifying what

[2]Whether this methodology is appropriate is currently controversial. See Weinberg *et al.* (2001); Alexander (2012); Williamson (2012) for some discussion.

possibilities are salient. The starting point of my story about QUD selection is here; the idea is to move from a theory of salient, or relevant, possibilities to one of salient decision problems.

The particular implementation of contextualism I will be working with here is that of Lewis (1996). Lewis is one of the few authors who has seriously engaged with the problem of selecting a set of salient possibilities for epistemological purposes; most authors in the area take their existence for granted. Lewis provides seven rules for determining what counts as a salient world for knowledge evaluation. These rules have been controversial in the philosophical literature, mostly for reasons that have little to do with our purposes here. A bit of terminology before beginning: in Lewis's rules, he talks about 'ignoring' or 'not ignoring' possibilities; I will say that not-ignored possibilities are 'salient'. The essential idea of Lewis's analysis is that a knows p iff p is true in all salient (not-ignored) a-accessible worlds. Let us now proceed to look at the rules; as we'll see, some of them prove to be useful for talking about salience for formal pragmatic applications as well. However, some are not; some of the rules have a purely epistemological character, talking, for example, about the method by which a particular bit of information has been acquired. Specifically epistemological rules like these have no obvious application to questions under discussion, and I will leave them out here.

Here is the subset of the Lewis rules that I will consider, four in all. Lewis himself states the rules in the form of longer heuristics embedded in philosophical discussion. I have therefore rephrased and shortened them substantially.

(6)
a. Rule of Actuality: The actual possibility is salient.
b. Rule of Belief: If one believes p, then p-worlds are salient; further, if one ought to believe p based on available evidence and arguments, then p is salient too.
c. Rule of Resemblance: If two possibilities 'saliently' resemble each other, then if one is salient then they both are.
d. Rule of Conservatism: If the people around us find something (not) salient, then we may also find it (not) salient.

My strategy now is to make appropriate substitutions in the rules above, and see what happens. The aim of the substitutions is to shift the discussion from not-ignored possibilities to salient issues or decision problems. For instance, I will replace 'possibility' with 'decision problem' in the above rules, and 'believes' with 'is currently addressing'.

Under this sort of substitution, the Rule of Actuality comes to state that the actual decision problem is salient. I take the actual decision problem to be the problem that the conversational agents are actually engaged in solving. This is just the QUD, as long as we consider primarily communication in discourse (for some decision problems that agents may be concerned with may not be under discussion at all). I'll call this decision problem *ActDP* hereafter, for 'actual decision problem'. The result of this substitution seems sensible enough: the actual QUD should indeed be salient. Although the conversational agents may not be certain exactly what the QUD is, a baseline criterion for a set of decision problems which can serve as proper candidates for the QUD is that it include the actual QUD itself. This consideration is precisely that which motivates Lewis's Actuality rule.

The result of making similar substitutions in the Rule of Belief yields the following: if one is either trying to solve a given problem, or if one is not trying to solve that problem but should be, then that problem is salient. The first case yields the QUD or the problem an individual is actively engaged on. There is some obvious overlap with the Rule of Actuality here. For the second case, we find two distinct kinds of subcases: situations where a DP is under consideration in general but is not at present consciously available, and situations where someone faces a DP but has not recognized that fact at all. The first subcase can be exemplified by a case from Davis (2009), where someone is waiting for a train, which has come without her noticing; she should be engaged in the problem of how to get herself onto the train, but remains unaware of the relevant change in circumstances. The second case can be exemplified by considering some unforeseen consequences of an action, which nevertheless are obvious to some outside observer. Summarizing, we get this principle:

(7) AWARENESS.
If an individual is consciously trying to solve a decision problem, or should be trying to solve a decision problem, then that decision problem is salient.

I'll call the set of decision problems satisfying this principle AWAREDP $= \{dp | dp$ satisfies AWARENESS $\}$.

The third rule we are concerned with is the Rule of Resemblance. After appropriate substitutions, we get the principle that if two decision problems resemble each other in the right way, and one is salient, then they are both salient. The obvious question then is: what kind of resemblance is required? In epistemological contexts, Resemblance is meant to solve the problems of lotteries and of Gettier cases (cf. Hawthorne 2004; Gettier 1963). The salient commonality of both of

these cases is that whether an agent has knowledge or not depends on features of the situation that are inaccessible to the agent. We can then take the required notion of resemblance to be whether it is possible to distinguish the two possibilities for an agent internal to the situation.

In the context of decision problems, the goal is an optimal solution in the sense of utility maximization. I therefore define Resemblance with respect to optimal solutions. The basic idea is that two decision problems resemble each other iff, as far as the decision-maker can tell, a solution to one problem is always a solution to the other.

(8) MUTUAL SOLUTIONS.
If two decision problems have the same optimal solutions under all circumstances, then if one is salient, the other is as well.

We call such decision problems *similar* and write this relation $dp \sim dp'$. What kinds of problem pairs will satisfy this condition? At least pairs where $D = \langle W, \mathcal{A}, \mathcal{P}, U \rangle_a$, $D' = \langle W', \mathcal{A}', \mathcal{P}', U' \rangle_a$, where U agrees with U' on all $a \in \mathcal{A} \cap \mathcal{A}'$ and where either $\mathcal{A} \subseteq \mathcal{A}'$ or $\mathcal{A} \supseteq \mathcal{A}'$. That is just to say that, for the optimal solution to be the same in both problems, the optimal action must be present in both action sets and be payoff-optimal in both. Cases of this kind include the question of whether to have a beer or a coke versus a beer, a coke, or a martini where I prefer beer, or worries about whether to live in the city or the country, or live in the city, live in the country, or live on a sailboat in the open ocean, where the last possibility is presumably seriously dispreferred .

I will ultimately claim that the decision problem that should be selected with respect to an assertion of φ is the one that maximizes the change in expected utility induced by learning φ. Obviously, this claim is normative and defeasible; discourse considerations can easily override it, a point I will return to at the end of the paper. But if two DPs are similar, then we may have difficulties in picking between them by this criterion. I will define equivalence classes to get around this problem.

(9) $[dp] = \{dp' | dp \sim dp'\}$.

It is easy to see that this is really an equivalence relation: reflexive, symmetric, transitive. I interpret the Rule of Resemblance for decision problems as dictating the use of these equivalence classes of similar decision problems rather than individual decision problems themselves, and will follow this practice in what follows.

The final rule I will consider is Lewis's Rule of Conservatism. Roughly paraphrased, this means that what our neighbors consider salient is salient in fact. One can see that this rule has some overlap with the Rule of Belief, as it will have application to cases like the train case mentioned above, where the fact that the train has come is salient for the speaker but not (yet) for the addressee. In the context of communication, this rule seems obvious. One might compare the 'sense recoverability' requirement that various authors have proposed for rational communication (e.g. Roberts 2011; McCready 2012), to the effect that a speaker should seek to convey content that can be understood by the hearer; in this context, this can be taken to indicate that one aspect of (discourse-based) salience has to do with the availability of mutual reasoning about the beliefs of other agents about what is likely to be salient. As such, game-theoretic considerations are prominent. However, here, I will put these complex issues aside and propose only a much simpler rule based on something like group knowledge:

(10) DISTRIBUTIVITY. If a decision problem is salient for most members of a group, it is salient for all members of the group. The set of problems made salient in this way is DISTDP.

In the next section, I will use these rules and the (sets of) decision problems that satisfy them to put some constraints on what counts as a possible question under discussion.

4 Salient decision problems

The process of specifying what is likely to be the question under discussion has two parts, as I construe it. First, we must specify a set of candidate decision problems; for me, those are the salient decision problems, as (initially) constrained by the rules given in the previous section. Second, we must select a single decision problem from this set. This issue is much more complex, but I will give one heuristic that allows a choice, after specifying in somewhat closer detail the set of salient decision problems.

We are now ready to define a set *DP* of discourse-salient decision problems. We can assume that such problems are part of information states in the sense of dynamic semantics (Muskens *et al.*, 1997). Standardly, information states are taken to be either sets of worlds or possibilities (when we are concerned with propositional information) or as sets of world-assignment pairs (when we are also concerned with anaphoric phenomena and quantification over individuals). It is

possible to extend the notion and take information states now triples of sets of worlds, assignment functions, and sets of decision problems. Thus a discourse-initial state will look like $\sigma_0 = \langle W, g, DP \rangle$, where all assignments are undefined. The idea is that *DP* will change over time as various decision problems become salient; this is the pool of potential questions under discussion.

The set *DP* must satisfy several constraints, imposed by the modified Lewis rules from the last section: the actual decision problem ACTDP must be included in *DP*,[3] and so must AWAREDP and DISTDP, those groups of decision problems satisfying the Rule of Belief and the Rule of Conservatism respectively. Note that these constraints are stated with respect to sets of equivalence classes of decision problems rather than the sets of decision problems introduced by our variants of the Lewis rules.

(11) a. $\{[dp] : dp \in \text{ACTDP}\} \subseteq DP$

b. $\{[dp] : dp \in \text{AWAREDP}\} \subseteq DP$

c. $\{[dp] : dp \in \text{DISTDP}\} \subseteq DP$

These constraints clearly do not exhaust the requirements on *DP*, but this is a first pass and it will do to illustrate how the theory is meant to work. In order to clarify what other constraints are active, much more empirical work is needed, not to mention some exploration of the psychological literature on salience. I will return to this point in the last section.

For now, with the notion of a set of potential questions under discussion embodied as *DP*, we may turn to the second question: how to decide what problem in *DP* is the one currently under consideration? The answer plainly depends on a range of factors, not least the content of the utterance itself. That content provides clues as to which decision problem is the right one: saying certain things is not relevant to some issues, but relevant to others. There can be no general ranking of objects in *DP*, statable in isolation from context and content, that will cover all available circumstances. Again, much more work is required here to itemize all the potential factors and let them play a role. Most likely, it seems to me, what is needed is something like a database of knowledge about how conversational interaction usually proceeds, a metalinguistic knowledge base, as it were.[4] Here

[3] Note that ACTDP is taken to be a set: this is the set $\{d\}$, where d is the decision problem actually at issue.

[4] It seems likely to me that at least some of this knowledge will take the form of the metalinguistic normality conditionals used in McCready (2012) to derive the interpretation of underspecified emotive expressions. This is not the place to explore the issue further, though.

I will simplify matters substantially and just propose a pragmatic principle that helps us choose.

This principle makes reference to a special type of relevance. An assertion can be identified as most relevant to a particular decision problem. If I tell you in the sushi place that a particular fish is fresh, it is most relevant to helping you decide what to order. Recall the notion of relevance proposed earlier in the paper, where relevance depended on the provision of information yielding a change in the ranking of expected utility $Ord_EU(dp)$. Now suppose that we relativize the ordering of possible actions in terms of expected utility to an information state σ, which now includes *DP* as one tuple element. We may then say that a sentence is relevant just in case its content is enough to change the utility ranking associated with some salient decision problem.

(12) φ is relevant wrt σ iff $Ord_{EU}(dp)(\sigma) \neq Ord_{EU}(dp)(\sigma')$, where $\sigma' = \sigma[\varphi]$.

On the basis of this notion, it is possible to define a principle of relevance maximization.

(13) MAXREL. When updating with φ, take the QUD/DP to be the most relevant decision problem in *DP with respect to* φ.

Unfortunately, the above is too underspecified: no notion of comparison of relevance between decision problems has been specified. How then to define relevance with respect to φ? A first obvious idea is to use the most salient problem. But we don't know how to specify this — doing so is the point of the present theory. But one heuristic seems obvious: we can say that the most relevant problem is that in which the most change is induced by processing of the new content. Then we pick up the problem to which learning the new information is the most useful. I thus spell MAXREL out as MAXCHANGE.

(14) MAXCHANGE. When updating with φ, examine the changes in $Ord_{EU}(D)(\sigma)$ for all $D \in DP$ and pick the one in which the change is maximal as the QUD/DP.

The intuition is that maximal change in the ordering corresponds to the maximal relevance, much in the way that van Rooij (2003b) takes the best answers to a question to be those that maximize the loss of entropy it induces in the question denotation.

What is the prediction then for the QUD? The answer depends how we define 'maximal change.' There will always be two *kinds* of decision problems in *DP* with respect to a sentence *S*. One could be characterized as the question of what action to take, given the information in *S*. I will call this the 'external decision problem.' In general, it is nonlinguistic in character, as it asks about the best 'external' action to take. The second asks whether the content of *S* should be believed. I will call this the 'internal decision problem;' it is linguistic in the sense that it asks what to do with the information in the sentence, rather than what to do on the basis of it. Note that both types are certain to be present: we need the second type if we wish to use decision problems to define questions under discussion.

Now a possible problem arises: if maximal change is understood as maximal resolution, the internal decision problem always wins, because once the content of *S* is believed (= added to the information state), the problem is solved. We don't want this to hold across the board, though we certainly would like it to 'win' sometimes; QUDs can be much more complex. We can consider alternative definitions. However, if maximal change is understood as maximal change in the 4-tuple defining the decision problem, then things get messy fast: there are too many dependencies between the elements in the tuple. Fortunately, we can simplify matters by taking maximal change to be maximal shift in expected utility. The result of this move is that we pick the decision problem from *DP* for which learning the content yields the most benefit for the processor. This sounds like an eminently reasonable pragmatic principle, though certainly a defeasible one.

Before moving on in the next section to show how the theory works for a few examples, I want to briefly mention an interesting issue. The action of the processor depends on his beliefs. Decision problems about action and belief are intimately intertwined. Does that mean that the the choice problem for these cases is a spurious issue? I think it does not. In some sense, we are interested in the intentions of the speaker of the sentence and on how the receiver can recover those intentions. For the speaker, the intent is *in general* to guide the hearer toward an action; belief change is a part of the process but is not the main goal. Of course, this is also case-dependent; it is easy to find scenarios where belief change is in fact the main purpose of the utterance, for example those involving the interpretation of any sentence appearing in this paper.

5 Examples

To see the working of the theory, I will focus on two kinds of cases. The first type is made up of cases in which a conversational participant raises an explicit question. This is obviously the easiest case. The second kind involves utterances in situations where intentions are clear.

I will need one key assumption for what follows, but I do not think it is a particularly problematic one. The assumption is that it is more useful to learn about a profitable external action than to learn about a profitable internal action. As support for this position, observe that external actions directly involving the external world, i.e., what to do in a particular situation. Conversely, internal actions involve the external world only indirectly via an agent's mental state, for they only concern what should be believed in a particular situation. Pretty clearly, in most cases only external actions carry direct payoffs—correct beliefs only help choose between the various options for external action. I take this to mean that external actions carry higher payoffs.

We can now turn to the cases. The first kind of case I will consider are those in which an explicit question is raised. Such cases are often cited in the literature on questions under discussion (cf. Buring 2003). Here is an instance. Suppose that we are at a conference and that there are three options for a restaurant for lunch: Japanese, Mexican, Thai. A asks B which is best, and B responds.

(15) A. Where shall we go eat?

B. The Thai place is good.

We can now ask: what is the QUD for B's utterance? Presumably the answer is obvious, but we would like our theory to derive it as a first step. The story will go like this. First, partitioning via A's question induces a decision problem $G = \{$go for J, go for M, go for T$\}$. Further, A's utterance introduces G to DP via Awareness. So the decision problem G is at least present in DP, as desired. Now, plainly, B's utterance essentially resolves G and guarantees some degree of utility for some choice of action, namely T. Assuming that it does not also resolve some higher-stakes DP, G is selected by MAXCHANGE. This is all straightforward enough.

For an example of the second kind of case, let us return to the train case mentioned above. Suppose that in this situation someone standing on the platform utters the following to you, who are sitting on a bench reading a newspaper.

(A) The train's here.

What is the QUD for this utterance? Presumably, one of the results of updating σ with the content of A's utterance will be to modify the action sets associated with at least some of the decision problems in $3(\sigma)$, namely those for which boarding the train would be a relevant possible action. Given that the addressee wishes to take the train when it comes, all such problems will have a new optimal solution in $3(\sigma')$, the information state resulting from the update with A's utterance, but not in $3(\sigma)$, the information state before that update. As a result, *MaxChange* tells us to select (the equivalence class associated with) a problem whose action set includes {board the train, not board the train}, for decision problems of this form will give a maximal payoff in this circumstance. (On the assumption that there are no other, more important, things at issue, of course.) Of course, I have not said exactly how the new utterance goes about altering the action sets of these decision problems; this is in itself a thorny issue. Still, this looks like progress.

One might wonder why I have claimed that a decision problem like {board the train, not board the train} is selected, as opposed to (say) $Bel = \{$believe p, not believe p$\}$? This problem is also newly introduced, and will be relevant in more or less just the same way as we say with the lunch case immediately above. This is where the assumption about the value of internal and external decision problems comes in. By assumption, since B is external and Bel internal, B is associated with a higher utility change, and so is more relevant by the proposed definition; it therefore is selected.

6 Conclusion and prospects

This paper has proposed a way of selecting questions under discussion, via a utility-based selection of decision problems from a set of possible such problems determined by variants of rules involving salience from contextualist epistemology. Plainly, this theory is only a first step. Two obvious deficiencies are the available constraints on the set of candidate decision problems and the simplicity of the rule selecting from that set. Obviously there are more constraints on the candidate set, as previously stated; deciding exactly what they are, and how they interact, is a delicate matter, and one which I must leave for future work. Here, reference to the psychological (and indeed economic) literature on awareness and salience appears a good starting point; I have begun this investigation, but am not yet in a position to propose a a fuller set of constraints. As for the selection rule itself, clearly MAXCHANGE isn't the only thing at issue: for instance, *persistence* is not factored in. One aspect of the QUD model is that one tends to keep with the

same question until it is resolved, but this is not really factored into the present model. Also, it seems problematic to link selection to utility maximization as directly as I have. MAXCHANGE is stated in a way that privileges decision problems where more is at stake. Perhaps a better notion would normalize payoffs and consider only changes across these relative values, rather than the current absolute change; I am not sure of the empirical consequences at this point. In sum, there are many avenues for further investigation.

References

Alexander, J. (2012). *Experimental Philosophy*. Polity.

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

Davis, C. (2009). Decisions, dynamics and the Japanese particle *yo*. *Journal of Semantics*, **26**, 329–366.

DeRose, K. (1992). Contextualism and knowledge attributions. *Philosophy and Phenomenological Research*, pages 913–929.

Gettier, E. (1963). Is justified true belief knowledge? *Analysis*, **23**, 121–123.

Hawthorne, J. (2004). *Knowledge and Lotteries*. Oxford.

Jeffrey, R. (1983). *The Logic of Decision*. University of Chicago Press, Chicago.

Lewis, D. (1996). Elusive knowledge. *Australasian Journal of Philosophy*, **74**, 549–67.

McCready, E. (2012). Emotive equilibria. *Linguistics and Philosophy*, **35**, 243–283.

Muskens, R., van Benthem, J., and Visser, A. (1997). Dynamics. In J. van Benthem and A. ter Meulen, editors, *Handbook of Logic and Language*, pages 587–648. Elsevier, Amsterdam.

Myerson, R. (1991). *Game Theory: Analysis of Conflict*. Harvard.

Roberts, C. (1996). Information structure: Towards an integrated formal theory of pragmatics. In *OSUWPL Volume 49: Papers in Semantics*. The Ohio State University Department of Linguistics.

Roberts, C. (2011). Solving for interpretation. Manuscript, Ohio State University.

Simons, M., Tonhauser, J., Beaver, D., and Roberts, C. (2011). What projects and why. In *Proceedings of SALT 20*, pages 309–327. CLC Publications.

van Rooij, R. (2003a). Quality and quantity of information exchange. *Journal of Logic, Language and Information*, **12**, 423–451.

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

Weinberg, J., Nichols, S., and Stich, S. (2001). Normativity and epistemic intuitions. *Philosophical Topics*, **29**, 429–460.

Williamson, T. (2012). Review of 'Experimental Philosophy' by Joshua Aalexander. To appear in *Philosophy*.