

Grounding Assertion and Acceptance in Mental Imagery

Christopher Gauker

University of Salzburg

Abstract: How can thinking be effective in enabling us to meet our goals? If we answer this in terms of representation relations between thoughts and the world, then we are challenged to explain what representation is, which no one has been able to do. If we drop the appeal to representation, then it is hard to explain why certain inferences are good and others are not. This paper outlines a strategy for a nonrepresentationalist account of the way in which the structure of reality may drive cognition. It begins by restricting our attention to certain paradigm cases of linguistically-mediated cooperative exchange. Second, it introduces a kind of nonconceptual thinking -- thinking in mental images. Third, in terms of imagistic cognition, the nonnormative conditions under which sentences are asserted and accepted is explained. Fourth, in terms of these conditions, the possibility of successful linguistic exchange in the paradigm cases is to be explained. Fifth, the strategy calls for an explanation how a language can be internalized and become a medium of intrapersonal thought. Finally, the norms of discourse can be conceived as norms that supervise the processes of linguistic exchange so explained. In this paper only the first three steps of the overall strategy will be carried out in any detail. But a strategy for completing the last three steps will be sketched.

1. The dilemma of effective cognition and a sketch of its resolution.

Here is a simple-minded theory of human cognition. First, the mind perceives an object or some objects before it. For example, the object might be a ripe mango. Then the mind observes that the object has a property or that the objects stand in some relation. For instance, the mind might judge that the mango before it is ripe fruit. Then the mind applies some general knowledge to that observation. For instance, the mind might apply the knowledge that ripe fruit is edible. Finally, the mind draws a possibly useful conclusion: The object before it is edible.

If we flesh this out in terms of mental representation, then we seem to be able to explain how thinking is successful in meeting our needs and desires. Roughly, thoughts are in

touch with the properties of things, with the relations between things and with the relations between those properties and relations in the sense that it *represents* them. By virtue of being in touch with the world in this way, the mind can be guided by what is really the case in the world. Changes in the world are registered in the mind in such a way that, by means of its general knowledge, the mind can adapt to those changes. The structure of reality is partially duplicated in structures in the mind, which enables the body that contains these structures to adapt to the structure of the world.

Details aside, one major problem with this account is that no one has ever explained what it consists in for the mind to be in touch, in the requisite sense, with one property rather than another. Why should we say that, in having a certain thought, the mind is in touch with the property of being ripe fruit, rather than the property of being a mango or the property of being a yellow and red thing? In other words, what is the representation relation between a mental representation and the things and properties that they represent? Well, we have waited a long time for an answer to this question, without ever getting even a promising start.¹ Accordingly, I think that some of us should think about how we might explain successful cognition in a way that does not pose the question.

Having imagined that success in cognition is explicable given a correspondence between representations and representata, one might engage in the following train of thought: The chain of causes and effects mediating between sensory inputs and behavioral outputs is only the sequence of representations. The properties represented play no causal role. So, really, we can drop them from our explanation and appeal only to the representations themselves considered as meaningless structures. If we find that we cannot explicate the

¹ The best known attempt that goes beyond mere hand-waving is Jerry Fodor's asymmetric dependence theory (Fodor 1987). For my critique of this, see my 2003. A more recent attempt is Sider's concept of reference magnets (Sider 2011). A cursory glance through the nouns and adjectives across a few pages of a dictionary should quickly persuade one that very few words stand for properties that qualify as reference magnets in his sense. Here are some from the beginning of the "p"s: *pace*, *pachinko*, *pacified*, *packing*, *paddle*, *padlock*, *paeon*, *pagan*. Which of these stands for a reference magnet or something more like a reference magnet than what other similar, but merely possible words might have stood for?

representation relation, that does not matter. We can simply leave it out of the account. If we need to talk about the meanings of representations at all, then we can think of talk of meaning as a more or less indirect way of talking about the proper role of the representations that have those meanings.²

The trouble with this anti-representationalist strategy is that it has a hard time explaining why we should expect the whole system of representations to be successful in meeting our needs and desires. If we no longer think of the system as tracking states of the world, then we can no longer describe it either as adjusted to reality or as achieving a desirable result. We might still say, of the system as a whole, that it was designed by natural selection to enhance our survival. But that explanation does not really answer every question we might like to ask. Regarding some particular transformation of symbolic representations into other symbolic representations that the system performs, we would like to know, why is that a valid or reasonable or justified inference? Since a system can ensure its reproduction well enough, even if it does not do everything as well as can be imagined, we cannot expect to derive an answer to this question from the observation that the system was selected by nature to promote its own reproduction.

In short, we seem to face a dilemma — the dilemma of effective cognition. On one horn, we are obligated to do something that no one has been able to do, namely, to explain what it consists in for a mental representation to stand in the representation relation to an object, property, relation or state of affairs. On the other horn, we need to explain how the manipulation of meaningless symbols can reliably meet our needs.

There is a way out of this dilemma, I think. Suppose we can explain how the structure of the world drives our cognition without positing representation relations between mental representations and objects, properties, relations and states of affairs. When I speak of the

² This, I take it, is the core thought of the inferentialists such as Jaroslav Peregrin (2014). In saying that meaning characterizes “proper” role I mean to acknowledge that the functions that the inferentialist identifies with meanings are conceived as ought-to-be’s, not merely as a sum of the kinds of thing that actually tend to happen.

structure of the world driving our cognition, what I mean is that the events in cognition are sensitive to the structure of and the structural changes in the world. Then we can say that our thinking is successful because it is driven in this way by the structure of the world.

Here is an overall strategy for finding this way out of the dilemma:

First, we restrict our attention to certain cooperative paradigms. It is not the case that the cognition we wish to account for is only that which occurs in these paradigms, but these paradigms will be the cases that we can account for most readily and that will serve as our models when we attempt to extend our account to other cases.

Second, we posit and explicate a kind of nonconceptual imagistic cognition. By means of mental imagery we solve problems. This cognition is nonconceptual inasmuch as mental images do not represent the properties of objects or the relations between them.

Third, we explain in terms of imagistic cognition the conditions under which sentences are asserted and accepted.

Fourth, we explain how assertion and acceptance by means of spoken sentences could facilitate cooperation in the paradigmatic situations.

Fifth, we explain how the language can become internalized and constitute a medium of intrapersonal thought.

Sixth, we explain how norms of discourse, for instance, rules of logic, can be formulated in terms of structures constructed out of intrapersonal linguistic thoughts.

In this paper I will sketch some of the details of this strategy up through the first three steps. But I will also provide a kind of formula, the solution to which would take us through the fourth step. I will only briefly sketch the fifth and sixth steps.

2. First step: the paradigms of communication.

The cooperative paradigms that ought to be the objects of our initial explanations are situations in which linguistic communication facilitates the achievement of a collective goal. An example would be a hunting scenario. Imagine that there is a group of hunters, and the members of the group have various roles. One is a scout. He goes out ahead of the others to look for game. Another is a tracker. He works with the scout and is expert at recognizing disruptions in the soil and the foliation that indicate that game has passed by. Another is the arrow-maker. Another is the cook, who cooks the meals of the hunters. And so on. The goal of the hunters is to kill game. They eat the meat, for basic sustenance, and they make good use as well of the hide and the bones. All of this activity is mediated by language. The hunters agree on where and when to go, when to hunt and when to rest. They describe what they have seen, they state their needs, and so on.

Another paradigmatic scenario would be a case of boat-building in a pre-industrial age. Somebody has to design the boat. Somebody has to cut the logs from which the planks are fashioned, and someone has to fashion the planks. Somebody has to build the scaffolding from the which the boat will be built. Somebody has to make the pegs or nails by which beams of which the planks will be held to together. Somebody has to drill the holes that the pegs will be pounded into. If they are not accustomed to working together, somebody has to direct the activities of these several parties. All of this work is mediated by language. Somebody gives the orders. They tell one another when they will arrive. They describe the location where the planks are stored. And so on.

For a third example, consider barter trade. People from different parts meet in the market place. Some bring grain. Some bring animals. Some bring metal tools. Each declares to the others what he or she has and what he or she will take in return for a quantity of it. Perhaps there is a market manager who arbitrates disputes.

3. Second step: Imagistic cognition

All human beings and many nonhuman animals have a capacity to use mental imagery to solve problems.³ For example, if I need to replace a worn-out washer in a faucet, I can take the faucet apart, form mental images of the parts as I do, and record, so to speak, a mental movie of the parts as they come apart. After I have replaced the faulty washer, I can play that mental movie in reverse in order to put the parts back together again.

The role of mental imagery in cognition can be made salient by considering cases in which it fails us. I am wrapping Christmas presents. I cut a piece of wrapping paper from a roll, assuming that it is big enough to cover the box before me. But then I discover that it is not. This is a failure in imagination. I imagined the piece I was cutting entirely enveloping the box, and then discovered that reality did not conform to what I imagined. Or I might imagine using a wooden chopstick to prop open a window sash. But then I find that when I put the chopstick under the sash and let go of the sash, the sash, because it is too heavy, simply snaps the chopstick. Here again, I imagined – in the sense of mentally picturing – what would happen and then discovered that what I imagined did not happen. The fact that my imagination does not usually fail me in these ways shows that I have a reliable, but imperfect, ability to solve practical problems by imagining, i.e., mentally picturing, how things go.

For present purposes, we will need some distinctions between different kinds of mental imagery. One basic distinction is the distinction between what I call *committive (imagistic) representations* and what I call *(imagistic) representations of fictions*. Committive representations are those that are treated as if they represented something actually existing.

³ If there are people who claim not to, my guess is that they are simply unaware of it. That is possible, because it is no part of my thesis that mental imagery is necessarily conscious. For evidence that an operation on mental imagery (which I call *visual morphing*) can account for some results of experiments with monkeys and great apes, see my 2017.

That is, the effect of such a representation on behavior is that such that one would say that the bearer of the representation acts as if the thing represented actually exists. Representations of fiction are just the opposite; the bearer of the representation treats them as fictions.

Commissive representations may be perceptions or memories. But they may also be endogenously generated representations of that which the bearer expects to perceive a moment later (e.g., upon going around the corner). An imagistic representation of a fiction, by contrast, is a representation treated as if it represented a possibility. To say that a representation is fictitious in this sense is not to say that it is treated as if it represented an impossibility, just that it is not commissive.

If we want to compare mental images to pictures, then we have to compare them not to static pictures but to movies. They represent pictorially sequences of events. Among these mental movies, some represent realistically, others fantastically. A *realistic mental movie* is a movie that we treat as if it represents how things could actually go. For instance, I might represent a dancer dancing across the stage, making only moves such as a human being could actually make. That is not to say that I need ever have observed a sequence of events quite like that which I represent, just that it's a sequence of a sort that would not surprise me if I actually saw it. That's also not to say that what I represent in a realistic mental movie could actually happen, only that, perhaps through some error, I *treat* it as realistic. A *fantastic mental movie*, by contrast, is a mental movie that we treat as if what it represents could never happen. For example, if I imagine a glass of red wine falling and turning into a bird and flying away before it hits the ground, then that will be what I call a fantastic mental movie.

What I call an *imagistic hypothesis* is a realistic but invented mental movie that begins with a commissive imagistic representation. (It is not throughout its duration a perception or a memory.) For example, an imagistic hypothesis might begin with my current perception of my hand reaching toward a door handle. It might continue with images of my grasping the

door handle, my turning the door handle, and my pushing the door in. If one imagistic hypothesis extends another by adding representations of further scenes to it, then I will say that the former is an *imagistic elaboration* of the latter. In other words, if imagistic hypothesis X is an initial segment of imagistic hypothesis Y, then Y is an imagistic elaboration of X.

Further, I will assume that the mind is capable of a kind of imagistic decision-making. (See Nanay 2016 for a closely related hypothesis.) This consists of formulating imagistic hypotheses that may or may not culminate in an imagistic representation of a goal state and then acting so as to perceive the sequence of events represented in one of the hypotheses that culminates in an imagistic representation of a goal state. For example, I might entertain two imagistic hypotheses. In one I imagine myself grasping the door handle without first putting a key in the lock. In the other I imagine myself first turning a key in the lock and then grasping the door handle. If only the latter imagistic hypothesis culminates in an imagistic representation of the interior of my office, then I will turn the key in the lock before grasping the door handle.

Imagistic representations, I will assume, are nonconceptual. The term “concept” can mean at least two things. A concept can be a word-like component in a discursive representation (vehicle), or it can be a component of a propositional content. In the latter case, it is a component of a propositional content that stands to an individual word or phrase as a proposition stands to a whole sentence. In denying that imagistic representations are conceptual, I am denying the perceptual representations are conceptual in either of these ways. The vehicles of imagistic representations are not structured into concepts in the first sense, and the contents of imagistic representations (in whatever sense we may properly speak of contents here) are not composed of concepts in the second sense. I have defended this

claim in several other publications (Gauker 2011, 2012, 2017), and so I will not repeat that defense here.

Further, imagistic representations stand in relations of similarity to one another. Properly speaking, we might say that the things represented stand in relations of similarity to one another, but the relations of similarity between the things represented can be projected back onto the representations. The similarities I am talking are always three-place relations of the form *x is more like y than like z*. If we had to do with two-place similarity relations of the form *x is like y*, then we would need to specify a *respect* in which they were similar (since everything is similar to everything else in *some* respect). Thinking of similarity as a three-place relation however obviates the need to specify a *respect* of similarity (although we can still confine ourselves to certain dimensions of similarity). This idea that imagistic representations stand in relations of similarity to one another can be developed into a full-blown theory of imagistic representation, but this is another thesis that I have developed in detail elsewhere (Gauker 2011, 2012, 2017) and will not pursue any further here.

One last assumption that I have to make about mental images is that they may include auditory images of spoken sentences and that the images of singular terms may be *associated* with images of particular object-representing elements of the image. Alternatively, we can say that the represented occurrences of singular terms themselves may be associated the objects represented in the image. For example, if I form a mental image of ball in a box and a mental image of an utterance of the sentence, “This is inside of that”, then, “this” may be associated with the ball and “that” may be associated with the box. In view of the fact that an occurrence of “this” in a sentence of the form, “This is *F*”, may be associated with the object of a mental representation, we may also say that the object is *labeled* with the predicate *F*. In other ways as well objects may be labeled with predicative expressions.

4. Third step: Grounding assertion and acceptance in mental imagery

The next step is to explain in terms of imagistic cognition the conditions under which a speaker is disposed to utter a sentence and the conditions under which a hearer is disposed to accept a sentence. For this purpose, I will assume that we are dealing with a very simple language. It contains demonstrative singular terms, such as “this” and “that”. It contains some monadic and binary predicates. It contains a negation symbol and a disjunction symbol and scope indicators (such as parentheses). It does not contain quantifiers (such as “every”) or modal operators (such as “possibly”) or intensional verbs of any kind (such as “believes”).

While I will explain the conditions under which a speaker is *disposed* to utter a sentence, I do not want to suppose that such dispositions are always activated or are activated randomly. There is a general condition under which they are activated, as formulated in the following principle:

The Principle of Speaking Up. An agent will utter those sentences that he or she is disposed to utter in response to the things that he or she commitively represents in the course of imagining joint activity in pursuit of goals.

I will also explain what it means to *accept* a sentence. But first I need to lay down the assumption that people do typically accept what they are told.

The Principle of Acceptance. Except under conditions *C* (to be specified), if a speaker utters a sentence *P* in the course of a cooperative endeavor with a hearer *H*, then *H* accepts *P*.

There will be conditions under which hearers do not accept what is told to them (conditions *C*), but identifying those is not the most pressing challenge, and so I will not try to do so in this paper. Acceptance is the default. It is not typically the product of a complex reasoning process.

The basic schema for a production rule for the assertion of atomic sentences will be this:

The Basic Rule of Assertion. For each atomic sentence of L , under conditions D (definable in terms of imagistic representations), a speaker is disposed to utter that sentence.

In my 2011 book I attempted to flesh out this rule by specifying conditions D (see the Betweenness Rule, pp. 234–5). Since that attempt was rather complicated, and even there had only the status of an illustration of the sort of thing one might want to say, I will here confine myself to an example that pertains only to very specific circumstances.

Suppose a child has observed some cats, and some of them have, in the child's experience, been labeled "cat". That is, in conjunction with the perceptual representation of a cat, the child has sometimes had a perceptual representation of a sentence of the form "That's a cat" or "There's the cat", accompanied perhaps by some gesture toward the cat, or has heard the question "Where's the cat?" followed by the perception of someone catching the cat. In accordance with my assumption that images of singular terms may be *associated* with images of particular object-representing elements of the image, let us suppose that the child associates such utterances with the cat. Similarly, the child has observed some dogs, and some of them have, in the child's experience, been labeled "dog". Suppose the child now encounters a novel animal and for some reason needs to label it either "cat" or "dog". Perhaps the child will react to the animal in one way if it decides to label it "cat" and in another way if it decides to label it "dog", and so the labeling is the child's method of deciding what to do. In this case, the child can adopt the following heuristic: If its imagistic representation of the novel animal is more similar to the imagistic representations associated with the label "cat" than to those labeled "dog", the child will be disposed to label the new thing "cat"; if its imagistic representation of the novel animal is more similar to the imagistic representations

associated with the label “dog” than to those labeled “cat”, the child be disposed to label the new thing “dog”.

My contention is that in this and similar ways, dispositions to label objects with words can be created by the similarity relations that a perceptual representation of the object stands in to other perceptual representations. This account of course will be limited to sentences that we might call *observation sentences*, by which I mean sentences of a kind such that a disposition to utter one may be grounded more or less directly in perception. It does not say that speakers will be disposed to utter observation sentences *only* in the manner specified. Even observation sentences may sometimes be theoretically grounded, in the sense that the disposition to utter them may sometimes be grounded in the utterance of other sentences, that, as we say, provide reasons to believe.

Acceptance, I have claimed, is a default. But I have yet to say what acceptance consists in. The answer will be different for different kinds of sentences, starting with atomic sentences. Here is where the concept of an imagistic hypothesis comes into play.

The Acceptance of Atomic Sentences: An agent *accepts* an atomic sentence if and only if, as a consequence of hearing or imaginatively entertaining that sentence, the agent confines his or her imagistic hypotheses to those such as would, if committive, elicit a disposition to utter that same sentence in accordance with the production rules for atomic sentences.

For example, if *A* says to *B*, “The door is locked”, then *B* will no longer entertain the imagistic hypothesis that depicts *B*’s opening the door just by turning the door handle. The reason is that that imagistic hypothesis is not one of those such that if it were committive and not merely hypothesized, then *B* would be disposed to utter the sentence, “The door is locked”.

Next I want to explain the conditions under which speakers of this language will be disposed to utter sentences containing logical vocabulary and the conditions under which they will accept such sentences. In order to do that, it will be useful to define another term:

Compatibility: A sentence s is *compatible* with an imagistic representation X for an agent A if and only if *some* of A 's imagistic elaborations of X are among the imagistic hypotheses to which the agent would confine him- or herself in accepting s .

It will aid understanding to spell out the dual notion of *incompatibility* as well, thus:

Incompatibility: A sentence s is *incompatible* with an imagistic representation X for A if and only if *none* of A 's imagistic elaborations of X is among the imagistic hypotheses to which the agent would confine him- or herself in accepting s .

For example, if A imagistically represents a sequence of events in which she unlocks a door, opens it and then closes it again without locking it, then the sentence "That is locked" is incompatible with her imagistic representation, because none of her imagistic elaborations of this representation is among those to which she would confine herself in accepting, "That is locked".

In terms of compatibility and incompatibility we can define a production rule for negations thus:

A Production Rule for Negations: An agent will be disposed to utter the negation of a sentence s in response to committive imagistic representation Z if and only if the agent has occasion to consider whether s is compatible with Z and finds that it is not.

For instance, if A committively imagistically represents herself approaching a door, turning the door handle and opening the door, without ever imagining herself using a key, then if A has occasion to wonder whether "That is locked" is compatible with the committive representation (with "that" associated with the door), she will find that it is not and consequently will be

disposed to utter, “That is not locked”. Correlatively, an acceptance rule for negations can be formulated thus:

An Acceptance Rule for Negations: An agent *accepts* the negation, *not-s*, of an atomic sentence *s* if and only if, as a consequence of hearing or imaginatively entertaining *not-s*, he or she confines his or her imagistic hypotheses to imagistic representations that are incompatible with *s*.

So if *A* says to *B*, “That is not locked” and *B* confines his imagistic hypotheses to such as are incompatible with “That is locked”, such as an imagistic hypothesis in which he imagistically represents himself approaching the door, turning the handle and opening it, then we can say that *B* accepts the negation, “That is not locked”.

Finally, a production rule and an acceptance rule for negations can be formulated along the same lines:

A Production Rule for Disjunctions: An agent will be disposed to utter a disjunction *p or q* in response to committive imagistic representation *Z* if the agent’s imagistic hypotheses containing *Z* include at least one representation that would dispose the agent to utter *p* and include at least one representation that would dispose the agent to utter *q*, and for every one of the agent’s imagistic hypotheses containing *Z*, either it would dispose the agent to utter *p* or it would dispose the agent to utter *q*.

An Acceptance Rule for Disjunctions: An agent *accepts* a disjunction *p or q* if and only if, as a consequence of hearing or imaginatively entertaining that sentence or one of the disjuncts, he or she confines his or her imagistic hypotheses to the union of those to which he or she would confine him- or herself if he or she accepted *p* and those to which he or she would confine him- or herself if he or she accepted *q*.

On this account, an agent will not be disposed to utter a disjunction just because he or she is disposed to utter one of the disjuncts. However, the acceptance rule does allow for the acceptance of a disjunction simply as a consequence of accepting one of the disjuncts (since in confining oneself to imagistic hypothesis that would dispose one to utter p one is confining oneself to those that would either dispose one to utter p or dispose one to utter q).

In a fuller exposition, I would go on to formulate production and acceptance rules for conditionals, quantified sentences, modal sentences, indirect discourse sentences, and so on. An important further development is that interlocutors must learn to build, on the basis of what they hear, linguistic structures that I call an interlocutor's *take on the context*. Only relative to these structures can we define the production and acceptance rules for, for instance, conditional sentences. But again, the basic psychological mechanisms by which these *takes on the context* are constructed should be definable in terms of imagistic representations.

5. *Cooperation revisited*

I have now taken the first three steps of the strategy I announced at the beginning, which are the only steps I promised to work out in some detail. What I want to do now is explain how we might take the next three steps, although I will not actually take them.

The next step would be to show that a linguistic practice governed by principles of production and acceptance such as I have described could facilitate cooperative action. Here I will take as my example the hunting scenario that I described at the beginning of section 2.

Here are some assertions that might contribute to the conduct of a hunt:

There are tracks in the sand near the lake.

There are no tracks in the valley.

Balam has no spear.

Spears are in the cave.

The enemy is between here and the lake.

Namu stays home.

Of course, I am cheating already in that I am assuming that the reader knows what these sentences mean, but I will assume also that the reader's understanding is somehow justified by the same production and acceptance rules grounded in imagistic cognition that would pertain to the language of the hunters.

We may suppose that each of the hunters in the hunting party entertains a rich spread of imagistic hypotheses. Each such imagistic hypothesis is rooted in the hunter's current perceptual representations and develops from there. For instance, Balam may envision forming an advance scouting mission with Namu or staying back with the rest of the hunters; he may imagine going to the hut to get a spear or going to the cave to get a spear; he may imagine heading directly to the lake or may imagine an indirect route to the lake along the ridge over the lake, and so on. We can think of these imagistic hypotheses as forming a tree structure. When Balam imagines two different ways of elaborating on his mental movies, a branch of an imagistic hypothesis splits, to form two imagistic hypotheses that have a stem in common.

When a sentence is uttered by someone and Balam accepts it, whole tranches of his tree structure of imagistic hypotheses are erased. For instance, when someone says, "Namu stays home", all of those imagistic hypotheses depicting an advance scouting mission with Namu are erased. When someone says "Spears are in the cave", Balam retains only those imagistic hypotheses that depict him obtaining a spear from the cave. And so on.

What I would like to be able to prove, but cannot prove, is that under certain conditions, this process of paring away imagistic hypotheses through the assertion and acceptance of sentences will have the result that each party to the exchange will have an imagistically represented plan of action, such that if each of them engages in the kind of imagistic decision-

making of the kind I spoke of earlier, the result will reliably be successful coordination in the pursuit of some goal. In the case of the hunters, each will move his body in such a way that each one ends up with the tools and supplies that he needs and together they move in unison toward the game that they need to find.

I cannot show this, but I can point to some of the key variables that we need to tune so as to make the system work. First, we need to better understand the generation of realistic imagistic hypotheses. We cannot suppose that members of the linguistic community are confined to imagining courses of events that are exactly like those they have experienced. They have to be able to learn from their experience in a way that allows them to imagine what might happen in a way that we would consider realistic, rather than fantastic, as I explained that distinction above. Second, we need to flesh out the schematic basic rule of assertion that I set forth above with something more specific. We need to say, for each atomic sentence, exactly what the conditions are, definable in terms of imagistic representations, such that a speaker is disposed to utter a sentence under those conditions.

The effect of a successful linguistic exchange will seldom be that the interlocutors all form the *same* imagistic hypotheses. The several interlocutors will typically have different roles and responsibilities in their exchange. Since their imagistic hypotheses will be what guide their actions through imagistic decision-making, we should expect that different interlocutors will form correspondingly different imagistic hypotheses. What I am describing is by no means a version of the traditional idea that in communication interlocutors share their mental states.

6. Glimpses of the last two steps

Many people have thought it hard to conceive of the possibility that we literally think in the very languages we speak. Language, they have assumed, may be conceived as a tool by

means of which we reveal the contents of our minds to others. But there does not seem to be any point in my revealing to myself the content of my own mind. What could I tell myself that I do not already know? At most the speech we hear could be an impetus to acquiring new concepts and learning new facts, and hearing the sounds of words in our auditory imagination might have a mnemonic function.

If we conceive of linguistic communication along the lines I have here described, then thinking in language makes literal sense. It is plausible that when we are engaged in imagistic decision-making our imaginations draw us in several different directions at once. There is what I expect based on what I observed two days ago, and there is what I expect based on what I saw yesterday. There is what I envision in my capacity as friend, and there is what I envision based on my capacity as citizen. We have various sources of information, and we play various roles in our community. These different streams of our selves can require some coordination. Spoken language may serve as a means by which coordination between these several streams is achieved. For further discussion of this conception of thinking in language, see my 2011 book and my forthcoming paper on inner speech (Gauker, forthcoming).

In consequence of both what other people say to others and what people say to themselves, interlocutors can build what I call a *take on the context*. I call it a *take* on the context, because we should think of it as an interlocutor's attempt to find a linguistic structure that really is *the context* that pertains to the conversation. So for each conversation there will be a *context per se*. The *context per se* for a conversation, as opposed to any given interlocutor's *take* on the context, is that ideal structure that the interlocutors would indeed all take to be the context if they were communicating most effectively.

The norms of discourse are rules that, at various levels, implicitly, unconsciously, or quite consciously, we strive to conform to in order to improve our cooperation by means of language. The rules of logic are norms of discourse in this sense. The rules of logic can be

defined in terms of the concept of context, as I have explained it here. We can formulate a recursive definition of the conditions under which the sentences of a language are true (or, as I prefer to say, *assertible*) relative to a context, and then we can define an argument, considered as a set of premises and a conclusion, to be *logically valid* if and only if for every context in which the premises are true (assertible), the conclusion is true (assertible) as well. My fullest elaboration of this conception of logical validity is in chapter 2 of my 2005 book on conditionals.

References

- Fodor, Jerry A., 1987: *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*, MIT Press.
- Gauker, Christopher, 2003: *Words without Meaning*, MIT Press.
- Gauker, Christopher, 2005: *Conditionals in Context*, MIT Press.
- Gauker, Christopher, 2011: *Words and Images: An Essay on the Origin of Ideas*, Oxford University Press.
- Gauker, Christopher, 2012: “Perception without Propositions,” in John Hawthorne and Jason Turner, eds., *Philosophical Perspectives 26: Philosophy of Mind*, Wiley, pp. 19–50
- Gauker, Christopher, 2017: “Visual Imagery in the Thought of Monkeys and Apes”, in Jacob Beck and Kristin Andrews, eds., *Routledge Handbook of Philosophy and Animal Minds*, Routledge.
- Gauker, Christopher, 2017: “Three Kinds of Nonconceptual Seeing-As”, *Review of Philosophy and Psychology*, DOI: 10.1007/s13164-017-0339-2.
- Gauker, Christopher, forthcoming: “Inner Speech as the Internalization of Speech”, in Peter Langland-Hassan and Augustin Vicente, eds., *Inner Speech: Nature, Functions, Agency and Pathology*, Oxford University Press.

Bence Nanay, 2016: "The Role of Imagination in Decision-Making," *Mind and Language* 31: 127–143.

Peregrin, Jaroslav, 2014: *Inferentialism: Why Rules Matter*. Palgrave Macmillan.

Sider, Theodore, 2011: *Writing the Book of the World*, Oxford University Press.