

RULES OF MEANING AND PRACTICAL REASONING

ABSTRACT. Can there be rules of language which serve *both* to determine meaning and to guide speakers in ordinary linguistic usage, i.e., in the production of speech acts? We argue that the answer is no. We take the guiding function of rules to be the function of serving as reasons for actions, and the question of guidance is then considered within the framework of practical reasoning. It turns out that those rules that *can* serve as reasons for linguistic utterances cannot be considered as normative or meaning determining. Acceptance of such a rule is simply equivalent to a belief about meaning, and does not even presuppose that meaning *is* determined by rules. Rules that *can* determine meaning, on the other hand, i.e., rules that can be regarded as *constitutive* of meaning, are *not* capable of guiding speakers in the ordinary performance of speech acts.

0.

On a widely shared conception of language use as an essentially rule-governed activity, linguistic meaning is taken to be *determined* by rules. It is also part of that conception that speakers are *guided* by the rules that determine meaning, and, moreover, guided precisely in ordinary linguistic usage, i.e., when performing speech acts. In this paper, we are going to consider whether these two ideas can be coherently combined. Since an agent is guided by a rule only when that rule is a *reason* for acting one way rather than another, the question of guidance by rules of meaning will turn on their capacity to serve as reasons for linguistic acts. And in order to serve as such a reason, the rule must, in a relevant way, enter into the practical reasoning that provides the true reason explanation of the act.

We are going to argue that the answer is negative. To the extent that rules concerning meaning *can* serve as reasons for speech acts, they are not meaning determining. Conversely, the rules that can serve as meaning determining rules do not fit into the relevant practical reasoning. Rules that do fit into practical reasoning leading to speech acts require, since they are not themselves meaning determining, that facts of meaning are independently established. Moreover, this is all they require. They do not require that facts about meaning are established by rules at all. So this is our main thesis: to the extent that meaning rules guide speakers in ordinary



linguistic usage, by way of being reasons for speech acts, it is just irrelevant whether meaning is determined by rules.

The conflict between the two functions of meaning rules can be brought out by considering two conceptions of so-called constitutive rules. The intuitive idea of a constitutive rule is that of a rule which somehow constitutes, or makes possible, a new type of action, or a new practice or activity. We shall proceed as follows: In section 1, we shall sketch how the idea that rules guide action can be fitted in with, and understood in terms of, practical reasoning. In section 2, we shall present the standard model of constitutive rules, and show, first, that on that model, rules cannot guide speakers in the ordinary way. Second, it turns out that they nevertheless can enter into practical reasoning, though not in the way typical of rules. In section 3 a discussion of constitutivity leads to the suggestion of an alternative conception of constitutive rules. On this conception meaning determining rules are capable of guiding in the ordinary sense, but they do not fit into the kind of practical reasoning that underlies speech acts, and therefore cannot guide ordinary linguistic usage.

1. RULES IN PRACTICAL REASONING

For an action φ (or the agent) to be guided by a rule R it is clearly not sufficient that φ is in accordance with R. That might be a mere accident. It is not even enough that the agent always conforms to R, for the explanation of that may have nothing to do with R or any rule whatsoever (most people invariably proceed in corridors by alternatively putting each foot in front of the other). What more is required? At least it seems that R must enter into the *motivation* of the agent for acting one way rather than another. Simply put: the agent did φ (in C) because R told him to. In other words, R must be a *reason* for acting. Moreover, as a reason for a particular action φ , performed by the agent, R must have played a role in the practical reasoning that led to φ . The relevant practical reasoning is the one that provides, or corresponds to, the true reason explanation of φ .¹ Something is a reason, in this sense, irrespective of whether it is a good reason. Something I don't know of might give a better reason for doing what I do than the reasons I actually have, but it does not guide me if I am not aware of it. Therefore, we are not here interested in the question of when a rule gives a *good* reason for a particular action; rather, the question is the more basic one in what sense a rule can be a reason for acting at all. The reasons we are interested in are reasons which can be used to *explain* an action, but not necessarily to justify it, i.e., justify it objectively or from a third person perspective.

Now, an agent might have reasons, of which he is aware, to do various things, including actions that are mutually incompatible, i.e., cannot be jointly performed. He might also have reasons both for and against the same action. In the primary explanatory sense, not all of the reasons an agent has at a given time explain what he actually does. For instance, if he has a reason to go to the bank, but decides on something else, unrelated to his bank business, this reason does not explain his action. More precisely, the reason doesn't explain his action if he does not *believe* that he is going to the bank, even if he actually does go, believing, e.g., that he will end up in the bar. Moreover, as pointed out by Davidson (Davidson 1963, 11) an agent can *have* a particular reason M for performing an action, and perform it, but for a *different* reason, and in that case M does not explain the agent's course of action either. In the primary explanatory sense, it is only *the* reason for which the agent acted that explains his action. In this sense, a reason is an *operative* reason, i.e., an operative factor leading to the action. In a derivative sense, however, something is a reason if it is a *candidate* for being an operative reason, i.e., if it is something that can play a role in motivating an agent to perform a particular action, whether or not that action is eventually performed, and whether or not for that reason. Let us call such reasons *basic* reasons.

In the present context we are interested in what it is for a rule to be a basic reason. We are not here concerned with the decision theoretic aspects. It is clear that in order to decide what to do, an agent needs to compare and weigh the basic reasons he has for and against a particular course of action, and this process typically leads to so-called "all-things-considered"-judgments about what is desirable to do, themselves typically leading to decisions. In describing this one needs to take account of what it is for one reason to be stronger than another, and this must be done in a full explanation of why one basic reason rather than another emerged as the operative reason. To come out as the operative reason, however, something has to be a basic reason in the first place, and what it is for rules to be basic reasons, within the pattern of practical reasoning, is our present concern. We shall say that an agent is *guided* by a rule only if that rule is a basic reason for the agent. We shall also say that an agent was guided by a rule R *in* performing a particular action φ only if R was an *operative* reason for performing φ .² Practical reasoning, according to the model we shall use here, does not conclude in normative statements in a strict sense, i.e., in conclusions about what the agent *ought* to do, or what would be (morally) *right* to do, or what would be the *good* thing to do. Rather, practical reasoning in the sense required here ultimately concludes in *intentions*. There

may be intermediate conclusions which are not intentions, but the ultimate conclusion is always to be an intention.

Now, *prima facie*, rules seem to be obvious candidates for being basic reasons. Rule R tells you to perform an action of type Φ in context C. Granted that you know R and recognize condition C to obtain, it seems that R does tell you what to do, viz. to Φ . And that in turn seems to give you a reason to Φ . This, we think, is indeed roughly right. But matters are a bit more complicated, and the simple view needs to be qualified.

An example of practical reasoning corresponding to the simple *prima facie* view would be:

- (R) When in context C, Φ !
- (B) I am in C
- (I) So, I shall Φ

Here ‘B’ indicates that the statement is something believed by the agent, and the ‘I’ indicates that the statement expresses an intention of the agent. Let’s refer to this as *the simple inference*.

Now the first thing to notice here is that the concept of intention is applied to a type of action rather than to an individual action. From the explanatory perspective, this is misleading. For a person may honestly declare his intention, e.g., to kick George and then go on to kick George, although not as a result of this intention, the reason being that the agent falsely believed the man he was kicking wasn’t George, but Arnold. However, the agent did certainly intend to do what he did, for his behaviour wasn’t involuntary. So there was another intention he had, which more directly corresponded to, and also explained, what he in fact did.

This already highlights the need to recognize, as has been particularly stressed by Davidson, that individual actions can be recognized under different descriptions, i.e., as belonging to different types. One and the same action can be of the types *kicking George* and *kicking the man in the corner*, and one and the same action can be of the type *eating something sweet* and *eating something poisonous*. Clearly you can be favourably disposed to an action under one description and unfavourably disposed to the same action under another description. Your intention to perform a *particular* action is always connected with some description(s) of the action and unconnected with many others. To the extent that there always is a particular intention behind, and corresponding to, the particular action, it is better in this context to restrict the use of the term ‘intention’ to what corresponds to the particular action.

As regards mental states directed at *types* of action we prefer the term ‘pro-attitude’. Desire is the paradigmatic pro-attitude, and to the extent

that ‘desire’ is used in a technical sense, in theoretical contexts concerning action explanation or practical reasoning, the two terms can be regarded as synonymous. With respect to the example above we would then say that the agent had a pro-attitude both to *kicking George* and to *kicking Arnold*, and that, since he believed he was kicking Arnold, it was the latter pro-attitude that motivated his intention, not the former.

Hence, we must recognize a form of elementary practical reasoning corresponding precisely to the transition from a general pro-attitude to a particular intention. This will have the form:

- (PA) I want to Φ
- (B) (This action) φ is a case of Φ :ing
- (I) So, I shall do φ

Here we have a general pro-attitude to a type of action Φ , a belief that a particular action φ is of that type, and a resulting intention to perform the particular action. The belief and the pro-attitude are both reasons for the action, and such reasons, the immediate reasons behind a particular action, have been called “primary reasons” (cf. Davidson 1963, 5ff).

An example of another (non-primary) form of standard belief-desire reasoning is the following:

- (PA) I want to turn on the light
- (B) By flipping the switch I will turn on the light
- (PA) So, I want to flip the switch

Here we see an initial pro-attitude and a belief as reasons for a resulting pro-attitude, while in a primary reason, a pro-attitude and a belief are reasons for an intention. In general, reasons figuring in standard belief-desire reasoning can be divided into two categories: *motivational* and *doxastic*.³ Both components are necessary for explaining action. While the guiding or motivational force of the conclusion is derived from the pro-attitude, the belief does auxiliary work either by identifying a particular action as a token of some type, or by representing an action of one type as a means of performing an action of another type. The belief thus serves as a theoretical transition between two practical, motivational states of mind.

In the simple inference the rule occupies the slot for the motivational reason, and that is basically right. What rules should do, intuitively, is to introduce to the agent a motivation, which need not be there initially, to perform, or to abstain from performing, an action of a particular kind. In performing φ , you were guided by R if you did φ because (and in extreme cases just because) R tells you that you must. By contrast, if something

occupies the *belief* position in a practical inference, then it *doesn't* operate as a rule, because then it doesn't contribute to the motivational part of the reasons behind an action.⁴ However, even in this respect the simple inference is misleading. The reason is that simply *knowing* a rule isn't enough to be motivated by it. You also have to take the rule to be *in force*, more precisely, to be in force *for you*.⁵ For one thing, rules, like propositions, are, or are reasonably conceivable as, abstract objects. Taken as such, knowing a rule amounts to knowing two things: what kind of rule it is (thus, e.g., whether a violation has a moral or semantic or legal or game-like significance) and what kind of actions accord with it.⁶ And knowing how to comply with a rule is different from wanting to, just as knowing the truth-conditions of a proposition is different from believing that it is true. The rule can only have motivational force if it is taken as being in force.

Moreover, even when a rule is *stated* as being in force, and thought to apply to an agent, it still needs to be accepted by the agent himself. Again, from mere knowing what accords with a rule issued it does not follow that the agent has a reason to act one way or the other. The headmaster of a school may have issued a rule to the effect that the pupils are to take off their caps before addressing a teacher. But that cannot be a reason why this boy took off his cap, unless he did it just *because* of recognizing that the rule applied to him. It might have been the case that the boy indeed did know of the rule, i.e., as something issued by the headmaster, but just didn't care, and so didn't take the rule as something that should have any bearing on his behaviour. So in order for a rule to have motivational force for an agent, the agent must have an attitude of acceptance of the rule. And this is so whether or not the rule is in force for the agent from a general or third person perspective. The agent must have the relevant pro-attitude to the rule itself.

This, however, might look like shooting ourselves in the foot. For rules were supposed to provide motivation, and now it turns out that they in fact require motivation. So they may appear redundant. They are not redundant, however, in cases where a pro-attitude to a type of action *derives* from a pro-attitude to a rule. Therefore, a fully explicit exposition of practical reasoning involving rules would, because of the requirement of acceptance, start by introducing a pro-attitude to whatever the rule requires, add a belief about what the rule requires and derive another pro-attitude:

(PA1) I want to do what R requires

(B) R requires that I Φ

(PA2) So, I want to Φ

There might, in turn, be a variety of reasons for taking an attitude of the form (PA1) towards a rule. An agent can take a rule to be in force, e.g., because it is issued as a law, because it is observed in his community, or because it will save time to go by a rule instead of weighing all the available reasons each and every time. For present purposes, however, this issue need not be pursued.⁷

What is important here, is that to complete the rule-involving practical inference, another inference is necessary. Here, (PA2) together with another auxiliary belief identifying a particular action φ as of type Φ forms a primary reason for φ . The rule itself, however, does *not* figure in primary reasons at all. To understand how rules can figure as motivational reasons we therefore have to look beyond primary reasons. Rules, if taken to be in force by the agent, do inter into such reasoning as direct reasons for pro-attitudes, never as direct reasons for intentions. A fully explicit exposition of such reasoning, to sum this up, would consequently look like this:

- (PA1) I want to do what R requires
- (B1) R requires that I Φ
- (PA2) So, I want to Φ
- (B2) (This action) φ is a case of Φ :ing
- (I) So, I shall do φ

We shall not normally, however, make that explicit in the rest of the paper. Rather, we shall enter rules in motivational positions in practical reasoning with the tacit assumption that the rule in question *is* accepted by the agent, and so does provide a motivational reason for performing an action of the kind required by the rule.

A rule can be a reason, then, provided the agent takes it to specify actions he ought to perform.⁸ From this, a second condition can be immediately derived: since ought implies can, we can reasonably accept only rules that recommend something that we *can* do. Ought likewise implies that we can do *otherwise*. This holds for all practical reasoning; it would simply lose its point to say that a reason explained an action if that did not mean that there was an alternative. This, *a fortiori*, holds for rules taken as reasons as well; there would be no point in explaining a pro-attitude towards Φ by acceptance of R if R could not be violated. If, that is, there were no possible actions not belonging to Φ . The rule must make a distinction between possible actions if it can motivate action at all.

We can now go back to our original question: Can meaning determining rules guide action?

2. CONSTITUTIVE RULES: THE TRADITIONAL READING

What does it mean to say that rules determine meaning? Donald Davidson starts his famous attack on the view that meaning is conventional by making a distinction between activities that are essentially conventional and activities that are only contingently so. Paradigm examples of essentially conventional activities are games. Eating, on the other hand, is conventional only by accident. This distinction can be applied to rules as well. Davidson himself does so when explaining it as follows: “In explaining what it is to play tarot we could not leave out of account the rules that define the game; in explaining what it is to eat no mention of rules or conventions needs to be made” (Davidson 1984, 265). To say that rules determine meaning is to say that speaking meaningfully is an essentially rule-governed activity. Without rules, that is, there would be no meaning, either.

What is it, then, that makes essential rules essential? Davidson only hints at an answer when he says that the rules of tarot “define the game”. Rules like these determine what it is to play tarot. Not only can tarot not be played without rules, it cannot be played without these specific rules, the rules of tarot. Analogously, the idea that rules determine meaning implies that without rules there would be no meaning and that what something means is determined by the specific rules applying to it. Rules essential to the activity they govern are therefore commonly called “constitutive rules”. To explore the characteristics of constitutive rules any further is not a question of great importance for Davidson; he goes on to argue that communication by language is not an essentially rule-governed activity. On his picture of meaning and communication, rules are explanatorily superfluous; what it is to speak meaningfully can be explained without reference to rules.

In this paper we are concerned with a different point. On a standard understanding of constitutive rules, their very constitutivity prevents them from *guiding* action. This holds for all constitutive rules, not just those of language. Constitutive rules, on the standard model, do not present practical reasons in the way typical for rules.

In this section, we will be concerned with an understanding of constitutive rules that finds its roots in Wittgenstein and Rawls, and has been elaborated especially by philosophers such as Midgley (1959), von Wright (1963), Shwayder (1965) and Searle (1969). However, as we will show in the next section, the idea of rules constitutive of the activity they govern can be cashed in in different ways. On the traditional understanding presented first, constitutive rules are contrasted with prescriptive rules and

what is different about a constitutive rule is precisely that its relation to action is different.

This line of thought about rules takes off from Wittgenstein. In his middle period, he became interested in the rules of games because he came to think of the rules of language as in many ways analogous to them. Applied to what he calls the “rules of grammar”, i.e., the rules for the use of words in general, in *Philosophical Grammar* he summarizes the main points of the analogy as follows: “without these rules the word has as yet no meaning; and if we change the rules, it now has another meaning (or none), and in that case we may just as well change the word, too” (Wittgenstein 1974, 133).

This Wittgensteinian idea is taken up by von Wright. He distinguishes between three basic types of norms for action, two of which are of interest for us: rules and prescriptions. Again, paradigm cases for rules are the rules of games. Interestingly, however, for von Wright their analogy with rules of language is less straightforward; what he likes to compare them to are rather the “laws” of logic. Paradigm cases of prescriptions are the laws of state. Von Wright explicates: “Prescriptions are commands or permissions, given by someone in a position of authority to someone in the position of subject” (von Wright 1963, 7). Prescriptions aim at making people behave in certain ways. However, this can hardly be said of something like the “laws” of logic. These, von Wright argues, are more like the rules of games than like the laws of state, they “neither describe nor prescribe, but determine something”, namely, what it is to play a certain game or what it is to think logically (von Wright 1963, 6). The difference von Wright tries to capture here is a difference in the relation of the respective norm towards action; while a prescription tells you what you ought to do or are allowed to do, a rule tells you what it *is* to Φ .

This very difference is elaborated by Midgley and by Searle, and their suggestions are virtually the same, even terminologically. Here we shall follow Searle’s presentation. Searle distinguishes between “constitutive” and “regulative” rules. Approximately, his regulative rules correspond to what von Wright calls “prescriptions” and his constitutive rules to von Wright’s “rules”. Searle says: “Regulative rules regulate a pre-existing activity, an activity whose existence is logically independent of the rules. Constitutive rules constitute (. . .) an activity the existence of which is logically dependent on the rules” (Searle 1969, 34). Searle illustrates this by the following example: It is possible, he says, that twenty-two men go through all the physical movements as gone through by two teams playing football, but if there were no rules of football, what they do would not be football.

The sense in which the activity logically depends on the rules thus becomes clearer when explained in terms of descriptions: “Where the rule is purely regulative, behavior which is in accordance with the rule could be given the same description (...) whether or not the rule existed, provided the description (...) makes no explicit reference to the rule. But where the rule (...) is constitutive, behavior which is in accordance with the rule can receive specifications or descriptions which it could not receive if the rule (...) did not exist” (Searle 1969, 35). While you can describe what is going on at a sophisticated dinner party as “eating” independently of the rules of dinner etiquette, what the twenty-two guys of the football-example are doing independently of the rules of football could be described as “running after some melon-shaped object and collapsing into heaps”, but not as “football”.

Let’s look at an example. As a constitutive rule of football, Searle cites the following: “A touchdown is scored when a player has possession of the ball in the opponent’s end zone while a play is in progress” (Searle 1969, 34). Rules like this one are doubly constitutive: they are constitutive of the game as a whole, i.e., form part of this game’s system of constitutive rules. At the same time, they are constitutive of a certain move of that game, here touchdown. Without the rule, actions could be described as “a player having possession of the ball in the opponent’s end zone while a play is in progress”, but not as “touchdowns”. What a touchdown is, is determined (or “defined”) by the rule. In this way, rules, to use another of Searle’s formulations, “create new forms of behavior” (Searle 1969, 33, 35–36); without the rule it would not be possible to intend to score a touchdown. As pointed out in the last section, an action is intended under a specific description. Without the rule for touchdown, the description of an action as a “touchdown” would not be available; there would not be such things as actions intended as touchdowns.

What exactly does it mean to say that such a rule “determines” something? From Searle’s example, the rule for touchdowns, we can learn that a rule determines a type of action, i.e., determines what it is to perform an action of a particular (new) type Φ . This, it does by connecting Φ with a different type of action θ . The θ -type can be described as such without the rule, the Φ -type cannot. This identification holds for a specific context C . These characteristics, Searle brings out by giving a standard formulation for constitutive rules:

(CR) Doing θ in C counts as doing Φ .

In the case of meaning rules an example might be:

(RG1) Uttering ‘green’ counts as expressing the concept *green*⁹

What Searle calls “regulative rules”, on the other hand, involve one type of action only. They can be outright prescriptions like:

(RR1) Do θ !

or conditional prescriptions like

(RR2) If C, do θ !

Form of statement is not, however, distinctive here. This, Searle explicitly notes; any regulative rule can be formulated in accordance with (CR). For example, the rule “Officers must wear ties at dinner” can be paraphrased as “Non-wearing of ties at dinner counts as wrong officer behavior” (Searle 1969, 36). What matters is that a kind of action is identified which, because of the very rule in question, can be performed by means of another kind.

Now, in what way can agents be said to be guided by rules that are constitutive according to the Midgley–Searle model? In the previous section we tried to understand this in terms of the idea that rules are reasons, according to this pattern:

(R) When in context C, Φ !

(B) I am in C

(PA) So, I want to Φ

This inferential form cannot, however, accommodate the Midgley–Searle constitutive rules. The reason is that the validity of this form of practical reasoning depends on the fact that acceptance of the rule induces a pro-attitude towards a type of action. This, in turn, is made explicit in:

(PA) I want to do what R requires

(B) R requires that I Φ

(PA) So, I want to Φ

The problem with constitutive rules in this context is that there simply isn’t anything that they require. According to the representatives of this view that is precisely what is characteristic of constitutive rules.¹⁰ The constitutive rule identifies one type of action as generated by another, i.e., says that you count as performing the one by performing the other. Suppose that the rule is

Doing θ in C counts as doing Φ .

The agent accepts this rule, and now the question is what counts as complying with the rule and what as violating it. And the answer in both cases is:

nothing. Doing θ does not count as complying, for the rule doesn't require θ to be performed, and for the same reason not θ :ing isn't a violation. Rules of this type cannot be violated since they do not even make any distinction between correct and incorrect actions.¹¹

A possible answer would be that Φ :ing by θ :ing does count as complying, while θ :ing without Φ :ing would be held to be a violation. But this cannot be right, for when the rule is in force the agent cannot possibly *avoid* Φ :ing when θ :ing in C. So this suggestion would violate the principle that for a rule to be guiding, it must be possible to violate it.

Because of this problem about compliance and violation, no pro-attitude to any type of action is induced by acceptance of the rule. And because of this rules of the Midgley–Searle kind do not fit into practical reasoning in the manner suggested in section 1. Rules of this kind cannot, apparently, guide speakers to do one thing rather than another. But then, the question is: How can they be meaning determining? It seems clear that to the extent that meaning is determined by rules, it must be by rules that make a distinction between correct and incorrect actions. Where no such distinction is made, however, there is no meaning determined either. So there is reason to think that Midgley–Searle rules are *not* in fact meaning determining rules.

For all said so far, however, it is still possible that Midgley–Searle rules can function as reasons for actions. Here is an example:

(PA) I want to say that p

(CR) Making an utterance of s counts as saying that p

(PA) So, I want to make an utterance of s

Clearly, a Midgley–Searle rule identifies one type of action as a means for performing actions of another type. And this is typically made use of when performing speech acts. In fact, this is an example of the practical reasoning that *standardly* underlies speech acts. We want to say something, i.e., perform a particular illocutionary act, and we need to utter an expression as a means of performing that act. Hence, such rules do seem to serve as reasons for actions, and in virtue of that they might perhaps be guiding, after all.

That is not the case, however. For note that the Midgley–Searle rule does not occupy a *motivational* position in the practical argument. It occupies a *doxastic* position. That is, it functions just as an ordinary belief, in effecting a theoretical transition from one pro-attitude to another. The agent's motivation comes from his pro-attitude to the illocutionary type, and all that is required for moving on to a pro-attitude to the utterance type is that the agent *believes* that he can perform the illocutionary act by means of making the utterance. We could, alternatively, enter an instrumental rule

saying “if you want to say that p , utter s ”, but accepting such a rule is again equivalent to believing that the means-end relation holds.

Because of this, the classification of the doxastic premiss as a *rule* is quite inessential. All that is required is that there are *facts* about meaning, like the fact that s means that p , which the speaker can know or believe to obtain. Since the speaker does believe that s means that p , he believes that he can say that p by means of uttering s . It is completely irrelevant whether those meaning facts in turn are established by means of rules or in other ways, and *a fortiori* irrelevant whether they could be established by means of rules of the Midgley–Searle kind.

So, for the Midgley–Searle rules to serve as reasons for speech acts, it isn’t necessary that they are rules, or that they determine meaning, or that meaning is determined by rules at all. We can still, however, ask whether it is *possible* that Midgley–Searle rules serve to determine meaning by means of a very non-standard form of guiding, i.e., by functioning as doxastic kinds of reasons for utterances. But, as far as we understand, taking this to be possible in effect means to give up the idea of constitutivity. For what I would be “guided” in by Midgley–Searle rules are the ways of *expressing* my thoughts (I would not be guided as to what thoughts to express). That, however, means that it would in fact be possible for me to express them either in this way or in that, and to use this expression either for expressing this thought or for expressing that thought. The rule would guide me in selecting one possible way among others. But this goes against the grain of constitutivity, for that is the idea that it is only because of first accepting some rule that it is possible for me to express my thought at all, and given the acceptance of that rule, it would not be *possible* for me to use that expression for expressing any other thought. So the Midgley–Searle rules cannot in fact be meaning determining at all.

3. CONSTITUTIVE RULES: A NON-TRADITIONAL READING

Part of the problem encountered above, however, is a consequence of the specific Midgley–Searle conception of constitutive rules rather than of the intuitive idea. The intuitive idea is that a particular practice or activity, like playing chess, or a particular action type, like castling in chess, could not in principle be engaged in or performed were it not for the rules. That does not, however, imply that constitutive rules must have the overt form of identifying different types of action, such as in “ θ :ing in C counts as Φ :ing”. Intuitively, if we changed half the rules of chess concerning how pieces may be legitimately moved, as distinct from rules stating what counts as what, we would have a quite different game, and therefore quite

a different practice of playing the game. Rules of other forms, simply laying down what is allowed, obligatory or disallowed, are, intuitively, as constitutive as the others.

As noted above, Searle realizes this. According to him, rules do not have to appear in the counts-as form in the rule books of a game to be constitutive. Neither is it sufficient for a rule's being constitutive that it *can* be phrased this way. As already mentioned, rules of all kinds can be brought into that form. However, Searle contends, when twisting a non-constitutive rule into the counts-as form, the Φ -term is not a specification, but a term of *appraisal*. He concludes, that – at least usually – constitutive rules are those that “naturally” can be brought into the counts-as form and where the Φ -term is used as a term specifying a new type of action, a type of action being defined by the rule (Searle 1969, 36).

However, bringing rules into the counts-as form in clear cases makes use of the existence of specially introduced verbs like ‘castling’ or ‘check-mating’, and their existence is quite inessential to the game; the rules of chess, for instance, could be stated without them. And for a lot of rules there are no such ready-made terms at hand. Moreover, by Searle's own requirement the intuitive constitutivity of e.g. a rule that rooks cannot move diagonally, cannot be captured by the counts-as form criterion. For bringing it to that form, like in “moving a rook diagonally counts as making an illegal move” makes crucial use of the term of appraisal ‘illegal move’, and it will therefore not be classified as constitutive.

The counts-as form is simply inapt to capture the intuitive idea. The intuitive idea that some rules determine, or make possible, new types of action or new practices does indeed make it an essential feature of constitutivity that new action verbs *can* be defined by reference to the rules that constitute them. But again, that is not a property *distinctive* of what we intuitively conceive of as constituted by rules. For action verbs can be defined by means of quite irrelevant references to rules. For instance,

x queats $=_{\text{def}}$ x eats and the rule “Don't smoke” exists (somewhere).

This is clearly not constitutivity in the sense we have in mind, in particular when thinking about games and game rules. What is missing is an interesting relation between the rule and the action type. But which?

Searle does in passing offer another idea, however (misleadingly presented as a modification of the main one). About the system of rules constitutive of basketball, he says: “acting *in accordance* with all or a sufficiently large subset of the rules does count as playing basketball” (Searle 1969, 36, emphasis added). The idea, then, is that a rule is constitutive of, e.g., basketball if it is part of a system of rules such that an agent plays basketball

if, and only if, he is acting in accordance with all [sufficiently many of] the rules in the system.

This idea does provide an interesting connection between the activity and the rule. However, not the right one, for neither accordance with all the rules, nor accordance with sufficiently many, can be the right condition. It cannot be all, for that would make it impossible to violate a constitutive rule within the established activity. To be sure, there are systems of rules functioning this way. In chess, for instance, it happens to be the case that an (uncorrected) illegal move automatically terminates the game. This is not, however, a necessary feature of constitutive rules.¹² Spearing is a violation of the rules of ice hockey but does not terminate the game. Instead there are intra-game penalties. And as long as we are concerned only with the conceptual features of constitutivity, as opposed to the *raison d'être* of the constituted practice, not even penalties is generally needed. There are e.g. not in general any penalties for abusing language.

And acting in accordance with sufficiently many rules isn't right either, for that would mean that when, e.g., I am playing chess I am in fact playing indefinitely many games at the same time. For assume that they all, including chess, have $n + 1$ rules, that the first n rules of all the games are identical and that I act in accordance with *them*. Since there are indefinitely many candidates for the $n + 1$:st rule there are indefinitely many alternative games, or game variants, and by the "sufficiently many" criterion I play them all at the same time, which is just absurd.¹³

John Rawls was, we believe, more on the right track when he suggested that action types belonging to rule-constituted practices cannot be performed outside the "stage-setting" of the practice (Rawls 1955, 27). For instance, you cannot score a goal in soccer outside the context of a game. And this is clearly right, for no matter how you kick the ball, if there is no game no goal is scored.

However, what counts as a stage-setting cannot be simply physical, or physically specifiable, circumstances. For, again, 22 men may run around kicking the ball exactly as if there were a game, but if the game isn't on, still no goals are scored. And this does indeed indicate what is essential to the idea of stage-setting: that the relevant rules are *in force*. When we decide to start a game of soccer, what we decide is that the rules of soccer shall start to *apply*, i.e., be in force for us. We decide what to count as the field, the goal posts, the teams etc., and then, as we proceed to play, our actions are to be *evaluated* by the rules.

This, we believe, is the key to understanding constitutivity: a practice is constituted by a set of rules if it is possible to engage in that practice *only* insofar as the rules of that set are in force for the agent.¹⁴ I can make

a checkmate or castle only to the extent the rules of chess are in force for me as I make the move. If they are not in force, then I am doing something else. Once this is realized it is obvious that there is no *general* problem about complying with or violating constitutive rules.

Accordingly, it is essential, as regards constitutive rules as well as others, to distinguish between being in force and being complied with. This is not to deny that there is a substantial borderline area, where it is difficult to say whether a rule is in force at all because the violations are too frequent. Nonetheless, the distinction remains essential even in such cases.

Now, let us apply the present idea of constitutivity to rules of language. Taking up our old example of a rule determining the meaning of the expression ‘green’, let’s now suppose that instead of (RG1) we have the rule

(RG2) ‘green’ is correctly applied to an object x iff x is green.

As demanded above, (RG2) does make a distinction between correct and incorrect actions. More precisely, (RG2) determines which utterances of ‘green’ are *semantically* correct and which aren’t. If you know, that is, what it is for an utterance to be semantically correct (as opposed to being correct in any other respect), you can use (RG2) to semantically evaluate utterances of ‘green’.¹⁵

And here, too, we must distinguish between the rule’s being in force and its being complied with. It would be a mistake to think that a speaker must actually *follow* this rule, in the sense of always (intentionally) complying with it, in order to mean *green* by uttering ‘green’. In that case mistaken applications would not count as examples of expressing the concept *green*. It is not even required that the speaker often or even mostly complies with the rule. What is decisive is precisely that the rule is *in force* for the speaker. A speaker of English means *green* by ‘green’ precisely since an application of that term to an object is *correct* if and only if that object is green. And if the view that meaning is determined by rules is right, then this holds just in case the rule (RG2) is in force for the speaker. So, on the same assumption, we can say that (RG2) is *constitutive* of the act of meaning *green* by “green”, precisely because you cannot do that, i.e., mean *green* by “green”, unless (RG2) is in force for you.

However, even on the present conception of constitutive rules, there are problems when it comes to guiding linguistic activity. Midgley–Searle rules, to repeat, aren’t action guiding in the right sense, that of entering into motivational slots in practical reasoning. However, by stating that one

kind of action is a means for performing another, they do fit into doxastic slots. Especially, rules like

(R) Making an utterance of s counts as saying that p

enter into doxastic slots of exactly those practical inferences rationalizing speech acts. By contrast, on the present conception of constitutive rules, a semantic rule governing a sentence would have the form

(R*) uttering (or asserting) s is correct iff p .

(R*) is not concerned with any means-end relation, but with what it is correct to do on a certain condition. *Prima facie*, rules of this form therefore do enter into motivational slots if taken to be in force by the agent. Thus, (R*) would enter into inferences such as this:

(PA) If an utterance of s is correct, then I want to make an utterance of s

(R*) uttering (or asserting) s is correct iff p

(B) p

(B) So, an utterance of s is correct

(PA) So, I want to make an utterance of s

There are two problems here, both concerned with the idea of guiding. The first of these is a problem generally for rules that are constitutive according to the present understanding, not just for those constitutive of meaning. Here, it needs to be noted that their being in force, their being accepted by the agent as we have described it above, does not necessarily amount to always having a pro-attitude towards the type of action they determine as correct. In ice-hockey, it would, e.g., be a pro-attitude towards not spearing. But having that attitude is not necessary for playing ice-hockey. All you need to do is accept that spearing will be a violation of the rules, and *that* you can accept even with a pro-attitude towards spearing whenever not seen by the referee. In the case of meaning rules, it would be the pro-attitude towards saying something true, and having that pro-attitude is clearly not necessary for saying something meaningful. When I want to be ironic, for instance, I have a pro-attitude towards going *against* what the rule requires. It is still essential that I take, and that I am understood as taking, the rule to be *in force* even for that very utterance.

There is, however, a second, more serious guidance problem concerning rules constitutive of meaning. The kind of practical reasoning they can enter into is clearly *not* the kind of practical reasoning that rationalizes speech

acts. The initial pro-attitude of such an argument is that of performing a particular illocutionary act, such as asserting that p , or asking whether p , not that of making a correct utterance. Of course, we generally do want our utterances to be correct as well, but given this initial pro-attitude, I don't need the semantic *rule* for telling me what is correct. If I want to assert that p , then I don't need a rule for telling me that my assertion is correct if and only if p . All I need to know is what *expression* to use for making the assertion. But a constitutive rule of the present kind simply doesn't *tell* me that. It doesn't fit into the relevant piece of practical reasoning:

- (PA) I want to say that p
 (R*) uttering (or asserting) s is correct iff p
 (PA) So, I want to utter s

The result of putting it there, in place of a Midgley–Searle rule, is just incoherent. Therefore, it turns out that while the Midgley–Searle kind of rule *is* adequate for practical reasoning leading to linguistic utterances, constitutive rules on our model are not.

Of course, from the premiss that (R*) is in force, I can go on to infer that s does mean that p , and that uttering s does count as saying that p . But in this inference, (R*) does not fulfil any rule-like guiding function. It is just mentioned there as a meaning determining entity, in a premiss that states a meaning determining fact. It is only the meaning determining function, not the guiding function, that is relevant in this inference. For on some alternative theory of meaning we could have an analogous inference, with a corresponding premiss in which some other entity is mentioned, like a causal dependence relation, or whatever. If I know the meaning determining fact (whatever it is), and that it is meaning determining, then I can arrive at the corresponding conclusions about meaning and utterances, e.g. that s means that p . And knowing that s means that p is in fact all that matters to me as a speaker. I have no extra need for knowing the meaning *determining* facts as well, since that is not needed for speaking the language. Thus, even if meaning *is* determined by rules, there is no need for speakers to know about it.

It has turned out that, to the extent that the theory of practical reasoning is suited for understanding guidance by rules, the two ideas we mentioned at the outset cannot be coherently combined. Some rules, as exemplified in this section, can be regarded as determining meaning, and can guide speakers in odd circumstances, but they cannot be said to guide speakers in ordinary linguistic usage. Their relation to ordinary usage can only be of some indirect variety. Some other rules, like Midgley–Searle rules, or instrumental rules, can guide ordinary usage, but in that guiding capacity

they function just like ordinary beliefs about meaning. It is not essential to regard this guidance as guidance by rules at all. And to the extent it is so regarded, this is completely independent of the assumption that meaning is determined by rules.¹⁶ We would like to add, finally, that in so far as the basic view – that meaning *is* determined by rules – is *motivated* by the assumption that this can account for our practical competence as speakers, i.e., by explaining it as an example of guidance by rules, the motivation is simply misconceived. Some other justification is needed.

NOTES

¹ This is not, as will be clear later, a sufficient condition for being guided, but in the sense we are concerned with, it is necessary.

² Here we do not say “if and only if”, since the idea of being guided by a (normative) rule, we think, involves the assumption that the rule not only is a reason, but a motivational reason. See below.

³ In Raz’s terminology (cf. Raz 1990, 33ff) the belief is an “auxiliary” as opposed to an “operative” reason.

⁴ This doesn’t hold for rules like rules of thumb or rules of strategy. Such non-normative rules are called ‘instructions’, or ‘rules in the instruction sense’ in Black (1962, 110), ‘directives’ or ‘technical norms’ in von Wright (1963, 9–11) and ‘instrumental rules’ in Pagin (1987, 201). Such rules so to speak inform agents about good means of reaching certain ends, and acceptance of such a rule is normally just equivalent to a belief about means-end relations and doesn’t have any motivational role of its own. Sometimes, however, the belief serves as a reason for accepting a normative rule to the same purpose, viz. when you think that it is better in the long run to follow the rule than to trust your own judgement in each individual case, even if the rule sometimes gives bad advice, too. In such a case the rule can motivate me to do something, as a means to a certain end, even if I in fact don’t believe that it is, in those particular circumstances, a good means.

⁵ This is stressed in Glüer (1997), p. 197f.

⁶ Normally, what doesn’t accord with a rule is in conflict with it, but that depends on how you individuate rules. Nothing is in conflict with a pure permission, like an ordinary rule of inference, but a pure permission on its own has no regulating effect whatsoever. Pure permissions have an effect only against the background of a more general prohibition (no inference that is not directly or indirectly justified by the ordinary rules of inference is allowed at all), and so there is a question whether pure permissions are rules at all or more properly seen as parts of other rules.

⁷ Raz (1990) suggests that all rules are what he calls “exclusionary reasons”. An exclusionary reason is a second-order reason for *not* doing this: acting for some particular first-order reasons (p. 39). This solves the dilemma (p. 194) that rules are either redundant or unjustified. As first-order reasons, justified rules could not be distinguished from good reasons. This parallels our observation that for some rules, acceptance of the rule is equivalent to a belief. In general, however, our concerns are different from Raz’s, since we are concerned with the question whether a pro-attitude is *in fact* derived from a rule, not whether the attitude really needs, or can get, justification from the rule.

⁸ This does not mean that the reasoning concludes in an ought-statement. It is only the rule that tells you that you ought to Φ , and if you accept that rule, you take the corresponding pro-attitude.

⁹ This, of course, is an application of his idea of constitutivity to meaning rules that we do not find in Searle himself. And in one respect this does not capture the general idea, for a speaker of a different language than English can still express the concept green, although by way of another expression and another rule. Therefore we should rather say that the created type of action is that of expressing the concept *green* by means of the word 'green'. Although we think this does tell against the idea of rules as meaning determining, this is not the point we are making here.

¹⁰ Cf. Schnädelbach, 1990). Constitutive rules, according to Schnädelbach, "don't prescribe particular actions to us; they don't say that I should greet somebody or why, which move is to be made on the chess board or the soccer field or what one has to say on which occasions; one follows them so that what one does is a greeting or chess or soccer or communication" (p. 130, transl. by us).

¹¹ This is pointed out in Glüer (1997, 194ff. And Searle seems to agree: "it is not easy to see", he says, "how one could even violate the rule as to what constitutes checkmate in chess, or touchdown in football" (Searle 1969, 41).

¹² In fact, we have been told that it isn't true even of blitz chess, where an illegal move remains part of the ongoing game unless the opponent points out the incorrectness before making another move.

¹³ It is no good to reply here that the games really are the same (the reason being that just one rule cannot make a difference), since that would either make sameness of games non-transitive, or else have the consequence that wildly different games are really the same.

¹⁴ This account of constitutive rules is proposed in Pagin (1987), chapter 3. However, there Pagin also adds further conditions which are not relevant in the present context.

¹⁵ In Pagin (1987), chapter 2, it is argued that it isn't possible to explicate the concept of meaning by means of the concept of a rule. The reason is that if we only have an undifferentiated notion of correctness of actions, then we have no criteria for distinguishing rules that do concern the semantics of utterances from, e.g., rules concerning the social propriety of utterances. Not even structure is enough, for it is possible, even if implausible, that, e.g., rules of etiquette are recursively specified.

¹⁶ That meaning is determined by (normative) rules is, of course, one common reading of the more general thesis that meaning is normative. Both are investigated *in extenso* in Glüer (1997).

REFERENCES

- Black, M.: 1962, 'The Analysis of Rules', in M. Black (ed.), *Models and Metaphors*, Cornell University Press, Ithaca, N.Y.
- Davidson, D.: 1963, 'Actions, Reasons and Causes', *Journal of Philosophy* **60**, reprinted in D. Davidson, *Actions and Events*, Clarendon Press, Oxford 1980. Page references are to the 1980 edition, pp. 3–19.
- Davidson, D.: 1984, 'Communication and Convention', in D. Davidson, *Inquiries into Truth and Interpretation*, Clarendon Press, Oxford, pp. 265–280.
- Glüer, K.: 1997, *Ist Bedeutung Normativ? Versuch Über eine Sprachphilosophische These*, doctoral dissertation, Humboldt University of Berlin.

- Midgley, G. C. J.: 1959, 'Linguistic Rules', *Proceedings of the Aristotelian Society* **LIX**, 271–290.
- Pagin, P.: 1987, *Ideas for a Theory of Rules*, doctoral dissertation, preprint from Stockholm University.
- Rawls, J.: 1955, 'Two Concepts of Rules', *The Philosophical Review* **64**, 3–32.
- Raz, J.: 1990, *Practical Reason and Norms*, Princeton University Press, Princeton, N.J.
- Schnädelbach, H., 1990, 'Rationalität und Normativität', in: K. O. Apel and R. Pozzo (eds.), *Zur Rekonstruktion der praktischen Philosophie*, Stuttgart.
- Searle, J.: 1969, *Speech Acts*, Cambridge University Press, London.
- Schwayder, D.: 1965, *The Stratification of Behaviour*, Humanities Press, New York.
- von Wright, G. H.: 1963, *Norm and Action*, Routledge and Kegan Paul, London.
- Wittgenstein, L.: 1974, in Rhees (ed.), *Philosophical Grammar*, Blackwell, Oxford.

Department of Philosophy
Humboldt University of Berlin

Department of Philosophy
Stockholm University
106 91 Stockholm
Sweden