

## Rule-governed Practices in the Natural World

*Abstract: I address the question of whether naturalism can provide adequate means for the scientific study of rules and rule-following behavior. As the term “naturalism” is used in many different ways in the contemporary debate, I will first spell out which version of naturalism I am targeting. Then I will recall a classical argument against naturalism in a version presented by Husserl. In the main part of the paper I will sketch a conception of rule-following behavior that is influenced by Sellars and Haugeland. I will argue that rule-following is an essential part of human nature and insist in the social dimension of rules. Moreover, I will focus on the often overlooked fact that genuine rule-following behavior requires resilience and presupposes an inclination to calibrate one’s own behavior to that of the other members of the community. Rule-following, I will argue, is possible only for social creatures who follow shared rules, which in turn presupposes a shared (first-person plural) perspective. This implies, however, that our scientific understanding of human nature has to remain incomplete as long as it does not take this perspective, which prima facie seems alien to it, into account.*

### **1. Forms of Naturalism**

Barry Stroud has famously noted that in the 20<sup>th</sup> century the term *naturalism* has deteriorated into an empty slogan – very much like *World Peace*: “Almost everyone swears allegiance to it, and is willing to march under its banner. But disputes can still break out about what is appropriate or acceptable to do in the name of that slogan” (Stroud, 2004, p. 22). In fact, in recent decades we have witnessed the birth of multifarious forms of naturalism, ranging from scientific to liberalized naturalism or to that of *second nature*. Naturalism is typically characterized on the basis of two theses that are held, either singularly or in combination, by all naturalists: a metaphysical thesis, according to which there are no supernatural entities, and a methodological thesis that postulates a continuity of method between philosophy and the natural sciences – but, again, both theses are so vague and imprecise that there are many different and mutually incompatible ways to spell them out. Our conception of naturalism will depend a lot on how we understand the term *nature*: in a broad sense, it refers to the realm of living organisms that populate our environment; nature, in this sense, consists of forests, flowers, birds, alligators, and mosquitos, etc. In a more restrictive and less biologicistic conception, which was advocated prominently by members of the *Vienna Circle*, natural entities are taken to be spatio-temporally extended entities. Finally, a scientific understanding of the term, as it was outlined by Wilfrid Sellars in his “Philosophy and the Scientific

Image of Man" (1991a), restricts the realm of natural entities to the non-observable entities that are postulated by the most basic natural sciences to explain the behavior of the observable ones. According to this form of naturalism, nature consists of atoms, molecules, DNA, etc. Relative to these different conceptions of nature, the metaphysical thesis of naturalism, according to which there are no supernatural entities, will find very different interpretations.

In the present paper I will discuss an argument that is pertinent more to the methodological than to the metaphysical thesis of naturalism. It focusses on aspects of our nature that we can get into full view only if we take a methodological stance that is not in continuity with that of basic natural sciences like physics or chemistry, but introduces a new element that is alien to them. In a nutshell, my argument runs along the following lines: human beings do not live in isolation, but in social communities that are constituted by rule-governed practices that evolve over time; it is, thus, one of the essential elements of human nature that we are able to discern rules and to engage in forms of rule-following behavior. My argument is not, however, concerned with the indispensability of the normative dimension.<sup>1</sup> Rather, I will focus on an aspect that concerns method. My goal is to show that the very attempt to understand rule-following behavior, even in its most basic forms, and in further consequence to understand the normative dimension, requires us to take a *methodological stance* that goes beyond the natural sciences: systems of rules and rule-following practices are possible only in groups of individuals who commit to the rules. Whether or not an individual merely manifests regularities or is genuinely following rules that she shares with other members of her community, however, cannot be established from a third person point of view; rather, we need to take the first person perspective into consideration. Unlike Husserl, however, whose critique of naturalism resulted in a reassessment of the first person singular, my suggestion is that we should acknowledge the centrality of the first-person plural, following a Wittgensteinian suggestion that one cannot follow a rule *privately* (cf. 2009, pp. 87f, §202). Genuine rule-following behavior requires an individual to show resilience. A shared system of rules is possible only if members of the community continuously calibrate their own behavior to that of other members of the community. These two aspects, as we shall see, are not abstract imperatives. They are natural facts and, thus, open for a naturalist explanation – as long as we understand naturalism in a large sense.

---

<sup>1</sup> The recourse to rules and rule-following behaviour is often taken to show the indispensability of the normative dimension and, in consequence, the irreducibility of a layer of reality that cannot be accounted for by basic natural sciences, such as physics or chemistry. This is not my concern here. Moreover, I do not need to presuppose that we can draw a strict distinction between the factive and the normative, or that the natural sciences cannot address normative issues, but can remain neutral on these issues.

## 2. A Husserlian Argument

In his article “Philosophy as a Rigorous Science”, Edmund Husserl developed an interesting argument against naturalism. He suggests that insofar as naturalism consists in “the naturalizing of ideas and consequently of all absolute ideals and norms” (Husserl, 1965, p. 80) it undermines its own basis or, as we read in Lauer’s translation: “without realizing it, naturalism refutes itself” (Husserl, 1965, p. 80). Husserl states:

He is, after all, going on presuppositions, to the extent that he theorizes at all, to the extent that he objectively sets up values to which value judgments are to correspond, and likewise in setting up any practical rules according to which each one is to be guided in his willing and in his conduct. The naturalist teaches, preaches, moralizes, reforms. ... But he denies what every sermon, every demand, if it is to have a meaning, presupposes. (Husserl, 1965, p. 81)

The basic line of Husserl’s argument could be summarized in the following way: naturalism attributes a central role to the theories of the natural sciences and insists that only natural entities, i.e. the entities postulated by the natural sciences, exist. But the theories of the natural sciences are not natural entities. Thus, according to the naturalistic maxim, they should not exist. Moreover, the formulation of scientific theories presupposes what naturalism denies: in Husserl’s words, “absolute ideals and norms” (Husserl, 1965, p. 80). Husserl does not elaborate this latter point in this article. His reference to the *Prolegomena* of his *Logical Investigations* (see 2001), where he has presented his critique of psychologism, makes clear, however, that for Husserl a comprehensive understanding of our scientific activity has to take into account a realm of reality that cannot be reduced to the regularities we find in physics or other natural sciences. Naturalism has to presuppose, but – according to his own maxim – has at the same time to deny that scientific theories can exist only in a normative dimension: scientific theories are *evaluated*, they have a truth-*value* (they can be true or false) and contain representations of scientific facts that can be *correct or false*.<sup>2</sup> The attempt to reduce the truth of a theory to facts of nature or to the survival value of the creatures who have developed it – as W.V.O. Quine has famously suggested<sup>3</sup> – is,

---

<sup>2</sup> Incidentally, a very similar argument was recently presented by Robert Brandom, according to who it is “due to Kant and was revived by Sellars” (Brandom, 2001, p. 600). This shows that Husserl’s argument is exemplary for a widespread line of reasoning against naturalism. My choice to present Husserl’s version of the argument is guided by the intention to show that in its core, the argument is accepted by philosophers who pertain to very different philosophical traditions – both by members of the analytic and the so-called continental traditions.

<sup>3</sup> See, for example, Quine (1998; 1969).

according to this line of reasoning, a mistake of a piece with the fallacious attempt to reduce ideal logical laws to psychological regularities.

Husserl's argument is, in my view, convincing and I am happy to accept the consequence that there are more things between heaven and earth than are dreamt of in this form of naturalism. However, a stubborn scientific realist could be tempted to counter this argument on the basis of Wilfrid Sellars's distinction between the *manifest* and the *scientific image of man*.<sup>4</sup> She could suggest that propositional representations of the world as well as the truth of scientific theories are only part of a highly deficient, yet useful strategy to grasp what is going on within and around us. This strategy, she could add, is characteristic for the *manifest image*, i.e., "the framework in terms of which man came to be aware of himself as man-in-the-world" (Sellars, 1991a, p. 6), which has been developed and refined over the centuries by our antecedents until we started to replace it, bit by bit, with the *scientific image*. As the latter emerges from the manifest image that already contains theories – yet only approximate ones that eventually will show deficient – the theories formulated within the scientific image should be able to account for the faculty of human beings to formulate scientific theories. Thus, by showing the limits of the old theories, some day in the future advanced natural sciences will shed a new light on the very nature of scientific theories that can liberate "ourselves from the burden of misconceptions and tunnel visions" (Churchland, 1989, p. 99) that are at the bottom of the anti-naturalistic argument that I have presented above. This shows that on the basis of Sellars's distinction between manifest and scientific image of man one could come to the conclusion – a conclusion, by the way, which Sellars would not have shared – that the normative concepts we use in epistemology, notions like "truth," "justified belief," or "rationality," will be "reconstituted at a more revealing level of understanding, the level that a matured neuroscience will provide" (Churchland, 1981, p. 84). My purpose in this paper is not to defend this line of reasoning, which I do not take to be overly plausible. I present it only to show that it might be desirable to complement Husserl's argument against naturalism with a line of reasoning that blocks Churchland's strategy and can so make the traditional anti-naturalistic argument stronger. In order to do so I propose to draw our attention to rules and their place in nature.

The current debate on rules (and in particular the debate on constitutive rules) is still dominated by a problematic conception of their nature, though. This, I think, results from the

---

<sup>4</sup> I do not want to suggest that Sellars did formulate or would have embraced an argument along these lines; I rather want to sketch a line of thought that was presented by students of Sellars and, in particular, by Paul Churchland.

choice of examples that are discussed. In the next two sections I will try to counteract this bias by drawing attention to aspects of rules and examples of rule guided systems that are familiar, but often overlooked in the debate. In particular, I will emphasize that rules are not just abstract systems, but need to be enacted (section 3) and draw attention to familiar aspects of the most complex and subtle of our rule-guided systems, language, that bring to the fore that our rule guided systems are highly dynamic. I will suggest that this is possible only because we have a natural inclination to continuously fine-tune our behavior to that of the other individuals with whom we share a practice (section 4).

### **3. The Rules in Question**

The idea that rules occupy a central place in human nature was expressed, with strong rhetoric force, by Wilfrid Sellars:

To say that man is a rational animal, is to say that man is a creature not of habits, but of rules. When God created Adam, he whispered in his ear, 'In all contexts of action you will recognize rules, if only the rule to grope for rules to recognize. When you cease to recognize rules, you will walk on four feet.' (Sellars, 1950, p. 298)

With this quote, Sellars emphasizes that discerning rules and rule-following behavior are deeply engrained in human nature. What conception of rules does Sellars have in mind? When we look at the philosophical debate, we find that rules are often conceived as principles that are or can easily be explicitly formulated and, thus, presuppose language and the ability to reason – which would make the capacity to discern and follow rules quite demanding and open a wide gap between humans and other animals, the behavior of whom would be reduced to mere regularities.

In contrast, the perspective on rules that Sellars, Wittgenstein and others have developed allows us to draw our attention to the traits of human nature that make language and reasoning possible in the first place: both language and reasoning are rule-guided practices and both, in consequence, presuppose the capacity to engage in rule-following behavior. However, the very fact that there are rational speakers who have never studied grammar and who are not familiar with the principles of rationality shows that with regards to language and reasoning the enactment of the relevant rules is logically and temporally prior to their formulation.<sup>5</sup>

---

<sup>5</sup> I will come back to this point in more detail below, see section 4.

This perspective, thus, allows us counteract the widespread tendency to identify rules with explicit principles – a conception that Robert Brandom has labelled “regulism”<sup>6</sup>. This conception is rooted in a strong intellectualist bias, which, I suspect, results from a one-sided diet of examples. Typically, highly complex systems like laws, lawlike regulations, or rulebooks of games like chess are taken as prototypical examples for rules. They invite for the idea that one can gain sufficient familiarity with the rules and learn to play a game just by studying the laws as they are stipulated in the rulebooks. And while this might be true for some games like chess, it is important to note that one can learn to play many games just by watching, i.e., without gaining any familiarity with the explicit formulations of the rules.<sup>7</sup>

Moreover, the regulistic conception obviously overlooks that most games require not only a familiarity with the rules, but also physical abilities, (bodily) skills, and technique that need to be trained. It is, thus, essential to take not only “formal” games like chess into considerations, but to supplement our diet of examples with games like soccer or rugby, where the physical enactment of the moves is decisive.<sup>8</sup> This will make space for a more basic and less intellectualistic understanding of rules that pays due attention to the fact that rules typically are enacted before they are stipulated; there are, in other words, rules that guide our behavior even when they are never spelled out in an explicit manner.<sup>9</sup> In this way we will be able to counteract the intellectualist bias that dominates the debate (in particular the debate on constitutive rules) and the sterile understanding of rules it has brought about.

#### **4. Grammar: A Complex Landscape**

The intellectualist bias results in a widespread view that we can take the rulebook of games as self-standing, (ideally) well-structured, and completely codified (or codifiable) systems of rules – systems, that is, that can exist independently of the rest of our rule-guided social practices. To show that this view is short-sighted, I would like to draw attention to a most complex and sophisticated, and yet most basic social phenomenon: language.

---

<sup>6</sup> See Brandom (1998, p. 18ff).

<sup>7</sup> This is likely true even in the case of chess, however. Even here, learning to play the game *well* requires not only familiarizing oneself with the rules, but also acquiring skills and strategies. Moreover, it is conceivable that one might learn to play chess just by observing two silent chess-players for a while (think of the fictional character of the chess champion Mirko Czentovic in Stefan Zweig’s novel *The Royal Game*).

<sup>8</sup> I am following here a suggestion that was also made by Jaroslav Peregrin (2014, p. 34ff).

<sup>9</sup> With this I do not want to suggest that the rules in question cannot be explicitly formulated *in principle*, but rather, that the statement of the rule is neither *temporally* nor *logically* prior to forms of behavior that can be regarded as compliance with the rule.

The conception of the rules of language has changed drastically over the last century. For a long time, philosophers have conceived of language as a formal system of symbols, the main task of which it was to describe an extra-linguistic reality. This understanding, which clearly rests on an intellectualistic conception of rules, was criticized by Wittgenstein and other ordinary language philosophers, who have drawn attention to the fact that language is a form of social activity. When we look at the current debate, however, one can easily have the impression that this critique has not been heard. The formalistic conception is still very present – with all the shortcomings that come along with it. Let me draw your attention to two points that seem important to me: First, a formalistic conception of rules cannot do justice to the idea that language is a system that is constituted by rules. This has led some philosophers, who were unwilling to revise their conception of rules and to open up to a less formalistic understanding – most prominently John Searle – to argue that the rules of language are not constitutive for language (cf. Searle, 1995, p. 72ff). Searle’s argument, which has been very influential in the debate, is based on an understanding of constitutive rules that risks falling into a form of regulism and that can be undermined once we come to appreciate that rule-following is a natural fact that belongs to our biological constitution.<sup>10</sup> Second, it cannot do justice to the fact that these rules are enacted by members of a community in concrete (worldly) situations – and will have to be enacted in future situations, unforeseen by now, where it might not be obvious how exactly the rule is to be applied. In these cases, future users will know to adapt the standards or set new ones – and, by doing so, they will keep the system dynamic. This aspect should be familiar: all living natural languages change over time. Even though we are all aware of this, however, the continuous evolution of living languages is not sufficiently appreciated in contemporary philosophy of language<sup>11</sup>, which reflects once more the predominance of the formalistic understanding of rules that I wish to undermine.

The differences between the rules that govern language, i.e., the rules of grammar (in a very large sense) on the one hand and laws or law-like regulations or the rulebooks of games, on the other, are salient: laws and the rules of chess, but not the rules of language, form a clearly defined and relatively small set of rules that can be explicitly formulated with minor effort. And

---

<sup>10</sup> Searle’s intellectualist bias reflects in the fact that he insists in the logical form of constitutive rules (“X counts as Y in context C”). For an interesting alternative, cf. Haugeland (1998a), who argues that constitution is a complex phenomenon that includes not only rules, but also capacities and skills as well as a basic form of commitment.

<sup>11</sup> There are interesting exceptions though, cf., for example, Ludlow (2014). While Ludlow discusses questions concerning word meanings, though, my point is concerned with the question of how the whole system of rules can remain dynamic and how an individual can come to conform to a dynamic, yet shared system of rules.

while we can easily imagine that chess has been brought into existence in the moment when the rules were explicitly formulated (for the first time)<sup>12</sup>, we cannot conceive of an analogous scenario for language, since formulating the rules that bring language into existence would already presuppose the possession of a language.

This argument, which was presented by Wilfrid Sellars in the form of a regress-argument<sup>13</sup>, explains why we cannot bring (our first) language into existence by stipulating the constitutive rules that govern it. It does allow, however, for the possibility that we could make the rules of language explicit *ex post*, by explicitly formulating – in a *descriptive* rather than *stipulative* manner – the rules that already govern our linguistic practices. After all, this is what grammarians do for a living: their ambition is to give us a systematic and at the same time an (ideally) comprehensive compilation of explicit rules to which competent users of language conform.<sup>14</sup> Moreover, this is also what arguably has happened with the rules of soccer – and likely with the rules of chess: the game has been played before the first rulebook was written; the explicit rules, thus, served to describe a practice that had evolved beforehand.

Rules of grammar differ from other systems of rules in characteristic ways, though. In what follows, I will draw attention to four aspects which, although familiar, are often overlooked in the debate, and which invite to adjust the received image of rules. Taken together, the four points undermine the guiding idea of grammar, according to which the rules of language can be comprehensively codified. Grammar, I will suggest, is best conceived of as a complex and dynamic system of enacted rules, a part of which will always remain implicit but which, nonetheless, is constitutive for language. If my argument is successful, it will require us to amend the received conception of constitutive rules and make place for the fact that there are implicit constitutive rules.

First, unlike laws and rulebooks of canonized games, grammar is highly *dynamic*. Language undergoes continuous change and knows many different (geographical and sociological) variations. Violations of single rules do not normally constitute a problem as long as the overall

---

<sup>12</sup> It is worth noting that this is a mere theoretical possibility; in fact, chess has evolved over centuries from other games. It is very likely that the contemporary game-industry has many examples of games that have come into existence the moment the rules were stipulated, but it can be doubted that these games will have a successful and lasting destiny as chess.

<sup>13</sup> See Sellars (1991b), the argument is present (though in a less explicit manner) also in Wittgenstein (2009).

<sup>14</sup> The discussion below will show that the grammarians' ambition to come up with a comprehensive compilation of rules that form the corpus of our grammar might be impossible to achieve. Our language allows for so many exceptions and variations that cannot possibly be forced into the rigid structure of such a system. Nonetheless, native speakers do not seem to have difficulties to conform to the rules of language in their complexity. I take this to be a most remarkable fact.



consensus is not threatened. If a speaker commits a specific error regularly, it will be considered an idiosyncrasy. If it catches on, it can become the rule. Note, however, that it is possible to deviate from the standard only locally, i.e., with respect to a limited number of rules. Should a person consistently ignore all or most of the rules of grammar, she will not be regarded a competent speaker in the first place. Thus, even the most non-conformist of language users must conform to these basic rules *by and large*.

Second, language differs from laws and the aforementioned games also in *complexity*. Moreover, grammars usually do not only consist in a substantial number of rules, they also contain an even longer list of exceptions to these rules. In many cases it is possible (in principle) to introduce additional rules that account for most of these exceptions, but even those typically do not cover all of them – a margin of exceptions to the supplementary rules remain. Eventually one would have to list all exceptions one by one, but this would, of course, undermine the very idea of a grammar as a system of rules rather than a comprehensive list of all possible uses – very much like a map in the scale 1:1, as described in Jorge Luis Borges' short story "On Exactitude in Science", would undermine the very idea of cartography. This point illustrates well that the rules we expect grammarians to formulate are best thought of not as exact and comprehensive descriptions of linguistic practice, but as "signposts" that guide, but do not prescribe in a detailed manner, our linguistic behavior.

Third, unlike most games and systems of law, language does not require a privileged authority that is entitled to sanction violations. All competent speakers are at the same level (though some might feel more, others less confident); any competent speaker has to be able to detect (and point out to others) errors or violations. Unlike games or legal systems, however, language does not need a sanctioning authority or a "language police" to work properly.<sup>15</sup> Language is a game we play together – and we play it without referee.

Fourth and most important, there seems to be an interesting mismatch between linguistic competence and familiarity with the (explicit) rules that is particularly salient in the case of

---

<sup>15</sup> There are, of course, institutions that guide or watch the right use of language, like the *Académie française* for French, the *Accademia della Crusca* for Italian, or the editorial department of the *Duden* for German. These institutions play an important role in making language uniform over a larger territory and over time. Moreover, they assist competent speakers with their occasional doubts. They do not, however, enforce correct use, nor do they sanction deviations. To my knowledge, grammar is object to legal regulation only in German-speaking countries. In the context of the reform of orthography from 1996, which, unlike other reforms, served not the unification, but the simplification of German grammar and orthography, respective laws have been emended. It is interesting, though, that the members of the German *Bundestag* felt it necessary to pass the resolution "Die Sprache gehört dem Volk" ["Language is owned by the people"] that acknowledged that language is evolving continuously and in an organic manner.

language. An expert player of a game like chess or soccer is likely able to explain the rules to a novice – just like a good lawyer must be familiar with the exact formulations of the laws that are relevant in her field of expertise (by and large). Native speakers of a language, on the other hand, who (all in all) have no difficulties in conforming with the rich number grammatical rules and exceptions, are typically not able to explicitly formulate the rules, whereas non-native speakers often face difficulties in applying the rules of grammar correctly even when they are able to state them explicitly. We have, thus, the paradoxical situation that those who habitually conform their behavior to the rules are often not able to state them, while those who are familiar with the codified rules have difficulties conforming to them in everyday practice. This unveils an interesting point: learning to conform one's own behavior to a rule and learning an explicit formulation of a rule are two very distinct things. It also shows that human beings are able to fine-tune their behavior to a very complex landscape of rules.

These four points illustrate that acquiring one's first language does not consist in learning a system of abstract rules, but rather in learning to conform one's (linguistic) behavior to that of the others in a shared environment. This does not require an intellectual accomplishment, it is an achievement that is based on features that are deeply rooted in our biological nature. We are *conformist creatures* who continuously calibrate and fine-tune their own behavior to that of other members of the species in their environment. Our conformist nature allows the newborn child to be acculturated to a social community and to interiorize the rules by which the latter is governed. Moreover, it allows mature members of the community to stay in tune with the others by continuously re-adjusting their rule-following behavior to the constant changes in the overall system.

The insight that we continually calibrate our own behavior to that of the other members of the community allows to explain how it is possible that systems of rules are dynamic and change over time. We tend to overlook this point in our philosophical reflection when we focus on codified systems of rules which, by their very nature, serve to provide stability and perpetuate established systems. Laws (especially basic laws like the constitution), rulebooks of games, but also grammar books aim at establishing stable frameworks on which the individual participants of the practice can rely and within which they can make their moves. Nonetheless, we witness that our laws, the games we play, as well as language continually change and evolve. They do so in different ways: changes in the legal system are typically brought about by legislature or, in drastic cases, by revolutions; changes in games can be deliberated by the players, but can also “happen”

by silent changes in practice that occur over time; while changes in language are typically the result of unpredictable processes that are beyond conscious deliberation. This undermines our image of stable rules that, once deliberated, determine all possible future scenarios of a given kind. Rules are not like rails that are extended to infinity<sup>16</sup> and do not leave any leeway. Rather, they are like guidelines that recommend certain forms of behavior, but will have to be interpreted, adapted, or even modified in ever new circumstances of application that cannot possibly have been foreseen by the legislator or the authors of the rulebook.

In sum, rule-following behavior is not something that is fixed once and forever, but a highly malleable process: a newborn child is trained to conform to rules when she is acculturated to a community, but the fact that systems of laws, rules of a game, and language are object to continuous change shows that also mature members of a (political, sporting, or linguistic) community continue to calibrate their own behavior to that of others. Wittgenstein raised an interesting point when he insisted that “it is not possible to follow a rule privately” (2009, pp. 87f, § 202)<sup>17</sup>, which underlines once more the essentially social aspect of rule-following behavior: it is something we do together with others who we regard members of our community. If we take into account that the systems of rules are – and have to be – dynamic, we see that this is possible only if the members of the community continually calibrate and fine-tune their behavior to that of other members of the group – this not only makes changes in the system possible, it also ensures that all of us can adjust to the changes and remain members of the community. This aspect of continuous calibration, which seems so crucial, has not been sufficiently appreciated in the current debate. It is important to note that calibration is not an effort that members of a community have to make in an intentional or conscious manner. Rather, the impulse to continually adjust our own behavior to that of others is deeply ingrained in our biological – and provides the glue that holds the social community together.

### ***5. Two Levels of Propriety***

In the preceding sections I have focused on rules that are logically and temporally prior to their formulation, arguing that partaking in social communities that are constituted by rules and engaging in rule-following behavior form a central part of what we might call human nature.

---

<sup>16</sup> Cf. (Wittgenstein 2009, p. 91, § 218).

<sup>17</sup> The quote is somewhat taken out of context, for Wittgenstein continues: “otherwise, thinking one was following a rule would be the same thing as following it.” Wittgenstein, thus, insists that rule-following depends on social criteria that determine the correctness of an application of a rule. Though I share this concern, here I mean to draw attention to the point that we tend to calibrate our behavior to that of other members of our community.

Naturalism looms: would this not show that we should expect to find a biological explanation for this aspect of human nature? In fact, there are studies that suggest that many of the skills that make our partaking in social communities possible, skills that enable communication, cognition and cooperation, are essentially part of our biological constitution.<sup>18</sup> Michael Tomasello, for example, distinguishes different kinds of inborn *intention-reading* and *pattern-recognition* skills that make language acquisition possible (see 2003, 2019). How could, then, a reference to rule-following behavior complement the classical argument against naturalism that I have discussed above (section 2)?

To see how it does, we need to take the distinction between *pattern-governed* and *rule-following behavior* that was introduced by Wilfrid Sellars (cf. 1991b) into a consideration. In our biological environment we find many organisms whose behavior displays a complex pattern that is object to biological explanations. Sellars illustrates this point with the example of the *waggle-dance* performed by bees, which can be explained on the basis of evolutionary biology. It is conceivable that at some moment in the evolutionary history of the bees, when the waggle dance had not yet existed, some bees started to manifest new forms of behavior in a completely accidental manner: for some reason or other they began to waggle and turn and though their behavior might have displayed the same pattern as that of our bees today, they did engage in it in a completely arbitrary manner. Their behavior, consequently, “is *not* appropriately described by saying that the successive acts by which the pattern is realized occur *because of the pattern*” (Sellars, 1991b, p. 326). Since the behavior did show evolutionary success, however, it was inherited to future generation and spread in the “wiring diagram” of the entire population of bees.

“It is by a mention of these items that we would justify saying of the contemporary population of bees that each step in their dance behaviour occurs because of its role in the dance as a whole.” (Sellars, 1991b, p. 327)

The movement of the individual bees is not considered arbitrary any longer, it has become part of a complex pattern that governs the bees’ behavior, the waggle dance. This does not mean that bees engage in rule-following behavior, though. Bees are (as far as I can see) not committed to partake in a community of waggle-dancers, they do not correct errors or violations and do not

---

<sup>18</sup> Incidentally: already Wittgenstein hints at this point in his *Philosophische Betrachtung*, where he suggests that one can train dogs, but not cats, to retrieve a stick. – And he goes on: “And if the child does not react to our encouragements, like a cat one wants to teach to retrieve, it will not come to an understanding of an explanation; or rather, the understanding begins here with reacting in a certain way.” (1984, p. 131, my translation)

allow for variations in their waggle dance. Moreover, bees merely *display* this complex pattern in their behavior, but they do not seem to have the means to reflect on the pattern, to adapt it to new situations, or to modify it.<sup>19</sup> Bees, in other words, are able to make moves that are appropriate “in the game” in determinate circumstances, but do not (seem to) have the means to reflect on whether or not a given move is appropriate, nor whether a certain pattern is suitable in a specific situation.

Contrast this with our own situation: early in our ontogenetic development, we were trained to react in determinate ways in given circumstances. When a newborn starts to manifest complex patterns of behavior, this is not the result of an act of deliberation performed by the newborn, since she does not (yet) have the means to ponder over the *reasons* for behaving in a certain way. She is merely trained to realize a complex pattern of behavior, very much like the bees. Newborns simply conform their behavior to that of their caregivers on the base of innate processes. These basic (yet complex) forms of pattern governed behavior allow for the development of a higher level of complexity which makes it possible to evolve from pattern-governed to rule-guided behavior. Thus, while it is possible to account for pattern governed behavior on the basis of the fine-tuning mechanisms that I have mentioned above, an explanation of the transition to rule-following behavior needs to be more complex. In my view, it requires us to take the holistic character (in a broad sense) of rules into consideration.<sup>20</sup>

Remember that the waggle dance is a form of behavior that enables bees to communicate a clearly defined kind of information to other bees of their hive: the direction and distance of a source of food. As a form of behavior of simple insects, the waggle dance fascinates because it is simple but, at the same time, surprisingly refined; yet in an important sense it is mono-dimensional, as it allows to communicate a very limited range of information and can be applied in a very limited number of situations only. In comparison, the social practices human beings engage in are not only more complex, they also build up to a multi-dimensional system. While bees have a “wiring diagram” (to use Sellars’s expression) that allows them to engage in one form of regular behavior, human infants are trained to adopt different forms of behavior that are suited to different situations and serve different purposes.

---

<sup>19</sup> These are, of course, empirical conjectures that might be falsified by future melittologists. For my purposes, the empirical adequacy of these claims is not relevant, though. Please note that the claim I am putting forward has the form of a conditional: if bees are not able to engage in the activities described, they are not in a position to engage in rule-following behavior.

<sup>20</sup> For a discussion of a Wittgensteinian perspective on this topic, cf. Huemer (2006).

In most situations it is possible to react in the habitual way by reeling of forms of behavior that have been interiorized in the process of training. However, occasional tensions can arise between different forms of trained reaction that can be suited to the specific situations, but put contrasting demands on us. In these cases, it can become necessary to decide which of two conflicting inclinations one is to follow or which is most suited in the relevant situation. In these cases, it can be of advantage if an organism is able to distinguish two levels of propriety. In one sense, a given course of action is appropriate if it is in accordance with the pattern that constitutes the practice of which it is part; in this sense the course of action would qualify as a legitimate move in the game, as it were. It can happen, however, that a move that is legitimate in the game is not appropriate in a given situation. John Haugeland characterized this difference as “propriety of the performance as such, and the correctness of the result *vis-à-vis* their objects” (Haugeland, 1998a, p. 314).

Genuine rule-following behavior, it seems to me, requires an organism to be able to tell not only whether a move is appropriate within the practice, but also to evaluate whether it is appropriate in the concrete situation in which it is supposed to be enacted. This presupposes, in turn, that the organism is able to suppress a course of action that would be in accordance with the rules that govern the relevant practice, but inadequate in a specific situation. The organism must, in other words, be able to intentionally violate a rule.

Let me illustrate this point with an example that is borrowed from John Haugeland: imagine a species of birds who live on butterflies, but populate a habitat where most, but not all yellow butterflies are poisonous. For these birds, eating one of the poisonous yellow butterflies would have lethal consequences. The species has survived in this ecological niche only because in the course of evolution their antecedents had developed a strategy to spot yellow objects and avoid them. Whenever a bird is in front of a yellow butterfly and, rather than eating it, turns away from it, we could say that it has made a correct move in the game; it is following a “wiring diagram” that has become part of the biological constitution of the species. This is true even in the rare cases in which the yellow butterfly in question was not poisonous and the bird was threatened to starve from hunger. The bird’s move is correct according to the pattern of behavior it has inherited from its antecedents; from our perspective as observers, however, it seems obvious that eating this butterfly would have been the better option. After all, the butterfly in question was harmless. Moreover, it would have nourished the starving bird.

Now contrast this example with one that is inspired by John Dupré<sup>21</sup>: imagine a competent speaker of 17<sup>th</sup> century English who, when asked to give examples of different kinds of fish, included whales in her list. Should we consider her move correct or did she make an error? It depends. Her linguistic behavior was likely in conformity with the use of the words “fish” and “whale” at the time and none of her peers would have corrected her. If a contemporary speaker gave the very same answer today, however, she could hardly get away with it. Our use of the words has changed as biologists have convinced us that whales are mammals and that the classes of mammals and fish are distinct. The 17<sup>th</sup> century speaker has made a perfectly legitimate move in her language game, but in a sense she has made a mistake: she did not get things right. After all, in light of the taxonomy developed by contemporary biology, whales are mammals, not fish.

The very fact that there are conventions and patterns that can turn out to be inappropriate in relevant situations shows that they are not free-floating systems, they are rather essentially embedded in real-life situations. If an organism is able to appreciate the distinction between the two levels of propriety, it is able to perform a “reality check” of the habitual patterns that govern its behavior – which can be helpful when it comes to avoiding to learn about the inadequacies in one’s own pattern of behavior the hard way.

In sum, I want to suggest that one of the crucial differences between pattern governed and rule-following behavior is that the latter, but not the former, allows us to go beyond the blind realization of a given pattern and to intentionally decide whether in given circumstances it is more appropriate to play along and obey a given rule or to violate it.

### ***6. A Deeper Commitment and the First Person Plural***

The discussion so far has revealed that the difference between pattern-governed and rule-following behavior requires us to take a holistic perspective and an to distinguish two levels of propriety. We must not consider isolated practices one by one, but rather keep an eye on how they are interrelated (and possibly in conflict with one another) and how they are embedded in a physical, biological, and social environment. To fully appreciate this point, it might be helpful to consider a further example: imagine a group of anthropologists who discover (on a Pacific Island, in a hidden valley in the Austrian Alps, or on Twin Earth – choose your favorite scenario) a community that so far has had no contact with our civilization. The members of this community manifest forms of behavior that are very similar to ours. There is only one systematic difference,

---

<sup>21</sup> John Dupré discusses this example in his (1981, 1999).

though: as the anthropologists notice with surprise, the community does not know the practice of playing competitive games. They share a language and engage in social practices of all sorts (they collaborate in projects that require the joint effort of individuals, they engage in forms of economic exchange, they make music and dance together, etc.), but the idea that one could play games with others in order to win or meet an opponent in order to defeat her in a playful competition is completely alien to them. I don't know if such a community actually exists – nor whether it is possible, in principle, for such a community to exist – but invite you, for the sake of argument, to play along with imagining this fictional scenario which, at least *prima facie*, does not entail any obvious logical contradictions.

As time goes by the anthropologists learn to speak the language of this community and come to get to know them better and better. Of course, also the members of our community are curious and want to find out more about the anthropologists' background. One day, one of them finds a chess-set that the anthropologists have brought with them, along with the official rulebook of the *World Chess Federation*. Let's imagine that the anthropologists even manage to come up with a complete translation of the text. Could the members of the community learn to play chess by reading the rulebook? I have serious doubts. Even if the translation was excellent, the members of our foreign community would not know what to do with it. They likely would just stare and wonder why they should move these tiny bits of plastic on the chequered board. Getting acquainted with the abstract rules of chess does not suffice if you are alien to the practice of playing competitive games.

Of course, our anthropologists could teach them to play chess. In order to do so, they would have to give them more than the rules of the game, though: they would have to introduce them to the practice of playing competitive games. They would, in other words, have to find ways to explain to them why they should care to win and avoid to lose the game. It probably would not take much to do so, but note that this would substantially change the delicate fabric of social practices of this community.

The example shows, in the first place, that the practice of playing chess is not isolated. It is part of a large system of practices which, over time, changes in a continuous and organic manner as new practices are included, others are modified, and others again disappear. It also shows, however, that a foreign person can become part of a community if she is able to conform her behavior to the practices that are constitutive for it – and starts to engage in them in a habitual manner. But why is it so easy to imagine that a foreign person is assimilated in a social



community? Why do we have no troubles imagining that the foreign community learns to play competitive games and can become chess champions? The reason, I want to suggest, is that in all these examples we think of persons who already do engage in rule-following behavior and who, in consequence, already possess all the necessary pre-requisites. In particular, they have a natural inclination to calibrate their own behavior to that of the other members of the community of which they are or want to be a part.

This might become more obvious when we contrast our foreign community with non-human systems: Can we teach computers to become chess-players? For sure, we can program them to make moves that are in accordance with the rules of chess. Moreover, *prima facie* it might seem that a good program does in fact actively pursue the goal to win and avoid to lose the game. But does it really *play chess*? John Haugeland took a skeptical stance with regards to this question, suggesting that computers miss a relevant element: “The trouble with artificial intelligence is that computers don’t give a damn” (Haugeland, 1998b, p. 47).

According to Haugeland, thus, computers do not play chess in the proper sense because they do not *care to win*. He admits that when we talk about computers and chess programs,

we construe chess-playing computers as ‘trying’ to make good moves and ‘wanting’ to win. At the same time, however, I think we *also* all feel that the machines don’t really care whether they win, or how they play—that somehow the game does not matter to them at all. ... It seems that what machines lack is a *reason* to win: some larger goal that winning would subserve. (Haugeland, 1998b, p. 55)

The question is not whether at some point in the future it will be possible to create computers that can be considered chess players in the full sense. The point, rather, is that in order to do so we would need to create computers that do play chess not because someone had double-clicked on an icon to start a program that is executed line by line by the machine. Rather, they would have to have a genuine desire to play the game and *care to win it*. We would, in other words, need to find ways to build machines that do not merely behave in accordance with a given set of rules, but who come to follow rules.

The distinction between two levels of propriety, that I have discussed in the preceding section, is of help here. While it is true that we follow most of the rules blindly, genuine rule-following requires that an organism is in a position not simply to realize a pattern, but also to have an understanding of why doing so is appropriate in a specific situation. The organism, in other words, needs to have an idea of why following this specific rule *matters*. This goes hand in hand

with an aspect which, according to John Haugeland, is crucial for genuine rule-following: an organism who follows a rule manifests in her behavior a sort of *resilience* or *determination to carry on* (Haugeland, 1998a, p. 341).<sup>22</sup> Not only does she need to be committed to the rules, she also has to insist if she does not succeed at the first attempt or correct her behavior when she becomes aware of having made a mistake. Moreover, she needs to care that her own behavior is “in line” with that of the other members of the community; she has, in other words, to be able and willing to calibrate and fine-tune her own behavior to that of the others in the sense discussed above (section 4).

These aspects might raise a methodological problem for the scientific study of rule-following, though. It is difficult to diagnose from a third-person perspective whether an organism cares about the rules, whether she understands herself as a member of a community and is willing to fine-tune her behavior to that of the others. Moreover, it is difficult to decide whether she merely realizes a pattern or is committed to a set of (implicit or explicit) rules. Engaging in genuine rule-following behavior requires that the agent has a sense of the situation in which she is enacting a certain pattern which, in turn, requires a proper perspective on one’s social and physical environment – and it is notoriously difficult to determine from the “outside”, as it were, whether or not an organism fulfills these criteria.

Wittgenstein at one point suggested that there were symptoms that indicated that a person was correcting an error – which might as well serve as symptoms for the kind of resilience in question:

There are characteristic signs of it [a mistake] in the players’ behaviour. Think of the behaviour characteristic of someone correcting a slip of the tongue. It would be possible to recognize that someone was doing so even without knowing his language.” (Wittgenstein, 2009, p. § 54).

Wittgenstein raises an interesting point, but his perspective seems anthropocentric: in fact, we usually have no difficulties to tell, from a third person point of view, whether someone (with who we share our biological constitution) corrects an error. This even holds when the person in question comes from a different cultural or linguistic background. It might be very difficult or even impossible to do so, however, when it comes to decide whether an organism (or system) that is very different from us corrects its behavior. But then they probably would not qualify to be

---

<sup>22</sup> For a more detailed discussion of Haugeland’s position, cf. Huemer (2003, 2005)

members of “our” community in the first place – for we had no way to calibrate our behavior to a shared standard.

From this discussion we can deduce two points that seem important to me: first, if we want to capture the difference between pattern governed regularities and genuine rule-following behavior, we must take the first-person perspective into account. This does not imply that it is impossible to study rule-following behavior scientifically, it rather means that we need to complement our scientific method in a way that allows it to address issues that are related to the subjective perspective of persons.<sup>23</sup> Second, the notion of rule-following that I have presented is essentially social. We have seen above that persons, who engage in rule-following behavior, calibrate their own behavior to that of other members of the community. Moreover, as Wittgenstein’s observation indicates, we normally have a very good understanding whether or not other members of our community follow a rule. It is, thus, not the first-person singular, but the first-person plural that we should take into consideration.

The shared perspective in question essentially relies on the mechanisms of calibration and fine-tuning. Only organisms who succeed in calibrating their behavior to a shared standard can be members of the social community. In turn, only membership in the community makes rule-following behavior possible. In consequence, the first-person plural perspective I have in mind does not require the members of the community to be able to share collective intentions<sup>24</sup> – they merely need to be able to conform to shared standards, which, in its basic form, seems to be a process that is deeply anchored in our biological constitution.

## **7. Conclusion**

In the present paper I tried to challenge intellectualistic conceptions of rules that treat them as abstract systems of principles. To contrast these views, I have insisted on the fact that rules are enacted by natural beings in their physical, biological, and social environment and who calibrate

---

<sup>23</sup> Moreover, taking the subjective perspective into account requires us also to take the subject’s temporality into account, which in turn can explain how rule-following practices can evolve over time. I am grateful to an anonymous referee for having drawn my attention to this point.

<sup>24</sup> It is, thus, not necessary to decide whether the group comes first or the individual, nor whether collective intentions are intentions that are shared by a number of individuals (cf. Bratman, 2014) or intentions that should be attributed to a plural subject (cf. Schmid, 2018). My goal is rather to draw attention to the fact that collective intentions are possible only on the basis of a deeper commitment – a commitment that is deeply rooted in our biological constitution – to partake in a social community, which, in turn, becomes manifest in one’s ability and willingness to adjust one’s behavior to that of the others.

their own behavior to that of other members of the community. This suggests that a scientific perspective cannot be comprehensive if it ignores rule-following behavior. Moreover, I have suggested that it is not possible to decide, on the basis of intersubjectively observable symptoms, whether an organism is merely realizing a pattern or engaging in genuine rule-following behavior. For that reason, I have concluded, a scientific study of rules and rule-following behavior has to take the first-person perspective into account. As rules and rule-following have an essential social dimension, though, it is the first-person plural that we need to take into consideration.

As long as naturalism remains attached to a received conception of scientific method, it will not be able to give a comprehensive view of the world. It needs to do justice to the fact that rule-guided social practices are as much part of the world as human beings who possess their own, subjective perspective that they share with others, and who calibrate their own behavior to that of other members of their community. Only forms of naturalism that can open themselves to this dimension can hope to achieve their goals.

### ***References of works cited***

Brandom, R. B. (1998). *Making it explicit: Reasoning, representing, and discursive commitment*.

Cambridge, Mass.: Harvard University Press.

Brandom, R. B. (2001). Modality, Normativity, and Intentionality. *Philosophy and*

*Phenomenological Research*, 63(3), 587–609. <https://doi.org/10.1111/j.1933->

1592.2001.tb00127.x

Bratman, M. (2014). *Shared Agency: A Planning Theory of Acting Together*. Oxford: Oxford

University Press.

Churchland, P. M. (1981). Eliminative Materialism and the Propositional Attitudes. *The Journal of*

*Philosophy*, 78(2), 67–90. <https://doi.org/10.2307/2025900>

Churchland, P. M. (1989). On the Nature of Theories: A Neurocomputational Perspective.

*Minnesota Studies in the Philosophy of Science*, 14, 59–101.

- Dupré, J. (1981). Natural Kinds and Biological Taxa. *The Philosophical Review*, 90(1), 66–90.  
<https://doi.org/10.2307/2184373>
- Dupré, J. (1999). Are Whales Fish? In D. L. Medin & S. Atran (Eds.), *Folkbiology* (pp. 461–476).  
Cambridge, MA: MIT Press.
- Haugeland, J. (1998a). Truth and Rule-Following. In *Having thought: Essays in the metaphysics of mind* (pp. 305–361). Cambridge, Mass: Harvard University Press.
- Haugeland, J. (1998b). Understanding Natural Language. In *Having thought: Essays in the metaphysics of mind* (pp. 47–61). Cambridge, Mass: Harvard University Press.
- Huemer, W. (2003). Husserl and Haugeland on Constitution. *Synthese*, 137(3), 345–368.  
<https://doi.org/10.1023/B:SYNT.0000004902.07390.18>
- Huemer, W. (2005). *The Constitution of Consciousness: A Study in Analytic Phenomenology*.  
<https://doi.org/10.4324/9780203641781>
- Huemer, W. (2006). The Transition From Causes to Norms: Wittgenstein on Training. *Grazer Philosophische Studien*, 71(1), 205–225.
- Husserl, E. (1965). Philosophy as Rigorous Science. In Q. Lauer (Trans.), *Phenomenology and the crisis of philosophy* (pp. 71–147). New York: Harper & Row.
- Husserl, E. (2001). *Logical investigations* (D. Moran, Ed.; J. N. Findlay, Trans.). London & New York: Routledge.
- Ludlow, P. (2014). *Living words: Meaning underdetermination and the dynamic lexicon*. Oxford: Oxford University Press.
- Peregrin, J. (2014). *Inferentialism: Why rules matter*. New York: Palgrave Macmillan.
- Quine, W. V. (1998). *From stimulus to science*. Cambridge, Mass.: Harvard University Press.
- Quine, W. V. O. (1969). *Ontological Relativity and Other Essays*. New York: Columbia University Press.

- Schmid, H. B. (2018). The subject of "We intend." *Phenomenology and the Cognitive Sciences*, 17(2), 231–243. <https://doi.org/10.1007/s11097-017-9501-7>
- Searle, J. R. (1995). *The construction of social reality*. New York: Free Press.
- Sellars, W. (1950). Language, Rules and Behavior. In S. Hook (Ed.), *John Dewey: A Philosopher of Science and Freedom* (pp. 289–315). New York: Barnes & Noble.
- Sellars, W. (1991a). Philosophy and the Scientific Image of Man. In *Science, Perception and Reality* (pp. 1–40). Atascadero, Calif.: Ridgeview Pub.
- Sellars, W. (1991b). Some Reflections on Language Games. In *Science, Perception and Reality* (pp. 321–358). Atascadero, Calif.: Ridgeview Pub.
- Stroud, B. (2004). The Charm of Naturalism. In M. De Caro & D. Macarthur (Eds.), *Naturalism in question* (pp. 21–35). Cambridge, Mass.: Harvard University Press.
- Tomasello, M. (2003). *Constructing a Language: A Usage-based Theory of Language Acquisition*. Cambridge, Mass.: Harvard University Press.
- Tomasello, M. (2019). *Becoming Human: A Theory of Ontogeny*. Cambridge, Massachusetts: The Belknap Press of Harvard University Press.
- Wittgenstein, L. (1984). *Eine philosophische Betrachtung (Das braune Buch)*. In: *Werkausgabe Bd. 5*. Frankfurt am Main: Suhrkamp.
- Wittgenstein, L. (2009). *Philosophical Investigations* (Rev. 4th ed.; P. M. S. Hacker & J. Schulte, Eds.; G. E. M. Anscombe, P. M. S. Hacker, & J. Schulte, Trans.). Malden, MA: Blackwell.