

Haugeland on Biological and Social Conceptions of Intentionality
Joe Rouse, Wesleyan University
ISPS Asilomar, 2009

John Haugeland's ground-breaking work on intentionality and understanding over the past three decades has pursued a twofold strategy. The critical side of his work has been a series of arguments to show that the nature and difficulty of the task of understanding intentionality have been seriously under-estimated by most philosophers. Most notably, the phenomenon to be understood, namely the kind of intelligent responsiveness exemplified *inter alia* by playing chess, doing scientific research, understanding a natural language, or perceiving objects as objects, has been mis-described. Haugeland concludes from his critical arguments that many prominent approaches to the intentionality question at best succeed in characterizing a different phenomenon. Thus, many familiar programs in the philosophy of mind could only account for the "ersatz intentionality" that he finds characteristic of higher animals and some computer programs. Dennett's intentional stance discovers a genuine and still-broader pattern in the world that can be characterized from a third-person perspective, but misses the essentially first-personal stance-taking by genuinely intentional systems. Brandom's account of the game of giving and asking for reasons successfully explicates something akin to Heidegger's conception of *Gerede* (idle talk or bullshit), if that privative phenomenon were utterly disconnected from its orientation toward the possibility of any "telling," truthful disclosedness of entities as entities. And so forth.

Haugeland's constructive work, which builds upon his interpretation of Heidegger on truth and owned being-towards-death, has then sought to locate and hit the correct target for philosophical explication. The project has been to show what would be necessary for a system's comportment to constitute objects as norms, so as to let them actually govern what it does. Here we find Haugeland's emphasis upon the need to articulate and sustain an excluded zone of conceivable but impossible discoveries, through a double-edged, first-person-involved "existential commitment." Such a committed stance must be implicit in taking up chess, science, or any other genuinely intentional human comportment, on pain of not actually engaging in those activities. The crucial commitment is not to tolerate any apparent violations of an object-constitutive "excluded zone," a commitment maintained at first by revision or repair of the relevant skills and standards, and then if necessary, by giving up the entire enterprise to acknowledge failure in revision and repair. Only those systems capable of taking such a vigilant,

committed stance toward the ever-present possibility of their own existential death can genuinely understand and comport themselves toward entities as entities.

My paper is focused upon the critical side of Haugeland's project. In the first part, I highlight the importance of his critical arguments in a new way, by locating them at the intersection of two important but insufficiently recognized distinctions among the various philosophical approaches to intentionality. The upshot of Haugeland's critical arguments is to rule out three of the four possible combinations that can be articulated in terms of these two distinctions. In the second part, I suggest that seeing Haugeland's arguments in this light also allows us to consider whether he has somewhat mis-described their import. Specifically, he claims that his critical arguments rule out any attempt to explicate intentionality in terms of the biological or social character of genuinely intentional systems. I think he is right that his arguments rule out the most prominent contemporary philosophical conceptions of biology and social theory as the appropriate locus for an explication of intentionality. The proper conclusion to draw, however, is that these conceptions make for bad biology and bad social theory. A more adequate conception of these domains points toward the possibility of understanding intentionality as a social-biological phenomenon.

I — Locating Philosophical Conceptions of Intentionality

It is only sporting to place Haugeland's work within a classification of philosophical approaches to understanding intentionality. His well-known paper, "The Intentionality All-Stars" offers a compelling, even stellar, classification of the most prominent philosophical options on this issue as positions on the baseball diamond. Yet even he soon realized (and announced within the very same collection of published papers that presents the all-star team) that his own position, and Heidegger's, were playing a different game, and thus did not fit within this classificatory field.

Ia— Two Distinctions, Arrayed

My own proposal to track the most important differences among accounts of intentionality begins with a distinction that I introduced here two years ago. To help clarify the differences between Bert Dreyfus's and John McDowell's understanding of what it means for an intentional comportment to be conceptually articulated, I distinguished between descriptive and

normative approaches to understanding the domain of the conceptual (and *mutatis mutandis*, the domain of the intentional). A descriptive approach to intentionality seeks to articulate those features of intentional compartments that are operative in producing their directedness toward their objects. Some salient examples of descriptive approaches include Fodor 1998, for whom intentionality results from representational structures that play a functional role in cognition, Husserl 1982, for whom the structured correlations between noetic act and noematic sense constitutes the meaningful directedness of consciousness, Searle 1982, for whom intentionality is a complex biological property of organisms, Millikan's (1984) teleo-semantics, Carnap 1967 or the Marburg neo-Kantians on the logical structure of intentional performances, Jackson 2000, for whom conceptually articulated intentionality is established by a partition of possibilities that is operative in mental life and manifest in intuitions available upon reflection, or Dretske 1981, for whom intentionality is constituted by the primary information-bearing features of cognitive states.

A normative approach to intentionality, by contrast, identifies the domain as those performances and capacities that can be held normatively accountable in the right way. Such normativity might be conceived in terms of accountability to rational-conceptual standards, to objects, or abilities-to-be. The defining feature of a normative approach, however, is that the intentional domain is constituted by its accountability to assessment. There must be some way in which intentional performances or states can be held accountable to relevant standards, and they are intentional in virtue of how they would stand up to such accounting. Thus, for example, intentional systems might be taken to be those that can be interpreted as mostly rational in context. Unlike descriptive theories, normative accounts can allow that the constitutive means of holding intentional performances accountable (e.g., upon reflection or interpretation) need not actually be operative in all intentional states or performances. Thus, to adapt an example from my earlier paper, on a normative account of intentionality, chess grandmasters playing blitz chess need not have a concept "in mind" when they respond to the board position with a rapid move, so long as they could hold what they are doing accountable to the regulative, constitutive and strategic norms of chess play. The fact that they understand these concepts and norms, and that their play largely accords with them, is sufficient for them to be intentionally directed toward

rooks and knight forks rather than to plastic figurines on dark and light squares. Some characteristic normative approaches to the intentional include Brandom 1994 on the game of giving and asking for reasons, Davidson 1984 on radical interpretability, Dennett 1987 on the intentional stance, McDowell 1994 on conceptual understanding, Heidegger 1962 on care and the existentiell possibility of *Eigentlichkeit*, or Haugeland's 1998 own account of existential commitment.

A second dividing line among approaches to intentionality can most easily be drawn in terms of Husserl's distinction between empty and fulfilling intentional relations. An empty intending can be directed toward its object in its absence, including the modes of absence marked by the non-existence of the object, or by failure to satisfy its intentional manifestation under some aspect. By contrast, a fulfilling intending presents the object itself as directly given under some intentional aspect. The sense of fulfillment carried by those intentional performances need not be infallible, of course. Some intentionally directed states, performances, or entities can, for example, encounter objects in ways whose internal features cannot be distinguished from a perceptual or other intuitive givenness of the object itself, and yet which turn out not to be fulfilling presentations.

Husserl's distinction then highlights the opposing directions taken by two strikingly different approaches to understanding intentionality. The more traditional philosophical approach has been to start by understanding empty intending, and then to ask what it is for an empty intention to be fulfilled. These approaches have typically identified intentional directedness with some form of representation. The problem of how to understand non-referring intentional states and erroneous presentations seemed to dictate beginning with what it is to represent an object (possibly absent, mis-represented or non-existent) via an intentional comportment, and only then to ask how some of our representations can actually present their objects "directly." Such approaches have notoriously confronted problems of skepticism.

An alternative direction has been to start with a system's actual relations to entities, and then ask what it would be for those relations to be intentional (and thus meaningful, and conceptually or otherwise aspectually articulated). The challenge for these approaches has been to show what it is for an engagement or interaction with one's surroundings to open a space of

articulated engagement accountable to norms. The most common motivation for such an approach has been “baldly” naturalistic, in McDowell’s phrase. Intentionality is taken to be a feature of some entities, states or performances that are causally or physically interactive with their surroundings. The next step is ask what it is for a causal interaction to involve intentional directedness, under an aspect, such that what the system does could be appropriately understood as in error, i.e., as intending something other than what it actually interacts with, or as intending it under an aspect that it might not actually possess or display. Dretske’s appeals to information-bearing states or Millikan’s teleosemantic functional norms are familiar examples of such an approach. Yet Heidegger illustrates a different approach to intentionality that also begins with intentional fulfillment (an understanding of being exhibited in an ability-to-be) without construing fulfillment in causal or other narrowly naturalistic ways. Both Dreyfus on practical/perceptual coping as the “ground floor” of intentionality, and McDowell’s “direct realist” account of perception as rational second nature, likewise take up a fulfilling intentional compartment as their starting point, without construing that engagement with the world in a “baldly naturalistic” way.

We can bring these two distinctions together in a 2x2 array that classifies various approaches to understanding intentionality according to their location in that space:

Accounts of Intentionality	1: Primacy of empty intending/representation	2: Primacy of fulfillment : causality, perception, being-in-the-world, etc.
A: Descriptive account	Husserl: essential structures of consc. Carnap: logical structure of language Jackson: a priori partitions Searle: intentionality as biological Minsky et al.: GOFAI	Dretske: information-bearing states Millikan: teleo-semantics Fodor: cognitive representations Dreyfus: practical/perceptual coping
B: Normative account	Sellars: we-intentions Quine: radical translation Dennett: intentional stance Davidson: radical interpretation Rorty: conversation of mankind Brandom: game of giving/asking reasons	Heidegger: Dasein’s disclosedness McDowell: perception and action as rational 2 nd nature Lance: language as sport Haugeland: existential commitment Rouse: discursive practices as causal phenomena

Moreover, this classification has a further advantage of allowing us to recognize differences between two different kinds of philosophical disagreement. Some of the most focused (but also narrow) philosophical disagreements about intentionality involve differences over what plays a more-or-less-agreed upon philosophical role. Thus, we can recognize such “intramural” disagreements between Husserl and Carnap over whether logical syntax or essential structures of consciousness constitute meaning, between Dretske and Millikan over whether evolutionary history is directly relevant to intentional content, among Quine, Dennett, Davidson and Brandom over how best to characterize rational interpretability in context, between Dreyfus and Fodor over whether our engagement with the world should be understood in terms of skilled activity or causal-functional representations, or between Minsky and Searle on whether syntactic structure or conscious awareness constitutes intentional content. On the other hand, in some critical philosophical encounters, what is primarily at issue is the more substantial difference between descriptive and normative accounts, or between taking empty or fulfilling intentions as the point of philosophical departure for understanding intentionality. Here we find how naturalists like Dretske, Millikan or Fodor would reject Carnap’s formalism or Jackson’s two-dimensionalism; Quine’s or Davidson’s criticisms of traditional theories of meaning; or McDowell’s criticism of Davidson’s or Brandom’s coherentism. Most strikingly, we find Heidegger’s criticism of Husserl in the Marburg lectures taking up both issues at once: in arguing for the primacy of categorial over eidetic intuition, Heidegger takes Husserl to task for starting with empty intending, while his criticism of Husserl’s naivete about the being of consciousness points toward Heidegger’s own normative conception of Dasein’s relation to its own ability-to-be (as concern for its own possibility rather than as self-awareness).

Ib— Haugeland versus A2

Part of what I find striking about many of Haugeland’s critical arguments is that their primary targets are alternative locations on this array. His critical arguments are not intramural disagreements over how to work out the fulfillment of an agreed-upon philosophical role, but instead take up what roles need to be played and why. Their intended upshot is not merely that a particular account is broken-backed, but also why no other account of that sort will do.

Moreover, Haugeland takes his arguments against the other possible positions on my array to be interlinked. In particular, his arguments against the positions in A2 and B1 are intended to show that their failures are reciprocal, in a very strong sense. Broadly speaking, his point might seem to be that the A2 positions (most characteristically represented by Dretske, Millikan or Fodor) are able to account for how intentional compartments are engaged with the world, at the cost of being unable to show how that engagement is meaningfully articulated by any genuinely normative accountability. Similarly, it may seem as if the B1 positions (characteristically, Brandom and Davidson, or else those self-described Wittgensteinians who seek normative bedrock in the accepted compartments of a community) are able to account for a richly articulated space of meaning and normative authority, at the cost of losing any accountability to (and thus directedness toward) the world. Yet to put the point this way is to overlook the interdependence of these two concerns. Because the socially-constituted articulations appealed to in B1 cannot be held accountable to anything beyond their own responsive dispositions, they do not actually achieve an articulated space of meaning, either. Reciprocally, because the causal or biological systems from which Millikan or Dretske start cannot meaningfully articulate the causal chains or lineages of descent in which they are implicated, they also fail to identify the intentional relation as causal or etiologically functional. Thus, Haugeland's arguments show that truth and meaning, or objectivity and aspectuality, are constitutively intertwined in the form of "norms of objective correctness" (1998, 317).

The canonical articulation of Haugeland's arguments against the A2 and B1 strategies are in sections 3-5 of "Truth and Rule-Following," supplemented by a parallel argument against Dretske in section 2 of "Objective Perception," by further reflections on Davidson in section 17 of "Truth and Rule-Following," and on all of the B1 positions in "Mind Embodied and Embedded." These arguments can be reconstructed in a sequence from physical to social accounts of intentional relations, in which the difference between philosophical approaches (A2 and B1) falls out as a difference in the kinds of normativity established by the systemic interactions that are involved.

The sequence starts with Dretske's effort to identify the intentional relation between a perceptual system and the object perceived by that system, in physical, information-theoretic

terms. The problem he sets is how, from the long and involved causal chain that culminates in (for example) my hearing someone at the door, we can identify the person at the door as the intended object of my perceptual state.¹ Dretske's account turns on considerations of perceptual constancy. The object of a perceptual state is whatever it carries information about in a "primary" way rather than via some more proximal or of some more distal object. Perceptual intermediaries such as a vibrating eardrum or oscillations in the air are not perceived objects, because different intermediaries could have produced a qualitatively indistinguishable experience, and hence the experience does not inform us which intermediary was involved. More distal causes, such as the button being depressed or the person pushing the button, are also not primary information-bearers, because their involvement is only indicated via the more proximal stimulus of the bell: without hearing the bell ringing, we would obtain no information about visitors, at least not by auditory perception.

Haugeland in effect argues that Dretske is able to pick out the bell as the supposedly proper object of perception only because he is already committed to its unity and significance as an object on other grounds. A perceived object remains constant across variance in the visual or auditory stimuli that can be experienced as presentations of it. Yet although these various presentations differ from one another in many respects, Haugeland points out that Dretske's account of primary-information-bearing requires that there be no relevant similarity among these various sensory presentations of one and the same object, and that he cannot do:

¹ In the text, Haugeland follows Dretske in the latter's assumption that the proper object of this perceptual experience is the ringing of the bell. Dretske's account, if it succeeded in its own terms, would indeed identify the proper object of auditory perception in that way. Yet Haugeland's own constructive arguments, which proceed from the constitutive standards that govern the perceptual situation, ought to indicate (rightly, I think) that we hear someone at the door by the ringing of the bell, rather than hearing the bell (in a sense similar to our also hearing someone at the door by vibrations in the air). Even when we speak of hearing the doorbell, we don't literally mean the bell (which is usually located physically away from the door to be more effective in communicating the information about someone at the door). If we were to say to someone else under normal circumstances, "I hear the doorbell, could you go take care of it?" we would look askance if they used a towel to muffle the bell in the hallway rather than answering the door. Moreover, if our perception were in question ("was that the doorbell?") we would go check for a person in the doorway rather than seek evidence that the bell had recently vibrated.

If there were any single kind of stimulus that mediated all and only the constant perceivings (same kind of perception of the same kind of object), then the perception would carry information that the stimulus was of that kind [rather than about the object]. ... What's worse, it seems that there must be such kinds, if sensory perception is to be possible at all. For if one can reliably recognize the squareness of the table from varying perspectives, then there must be something— something higher order, global, relative to context, or whatever— normally common to all and only the stimuli from such objects, on pain of rendering perception magical. (1998, 245-46).

Although Haugeland does not specifically note this point, such higher order stimulus-kinds are not unique in blocking any direct information-conveyance from perceived object to perceptual state. A similar argument could be posed for any of the relevant causal intermediaries: a common object as cause of recognized auditory similarities must equally well present higher-order similarities among the vibratory patterns in the air that cause the higher-order stimulus pattern, or the electrical excitations of the nerves that convey it. The grounds for picking out, from the sequence of causal intermediaries and precedents, one stage as the common “object” of its perceived effect cannot be causal or information-theoretic. Causal or physical interactions as such cannot meaningfully articulate the world.

Whereas the physical conveyance of information does not establish any norms for articulating its constitutive causal sequences into objects and their experiential effects, biological evolution does introduce normative considerations. An organism is a functional complex, and functional and evolutionary biology combine to account for why certain components or operations are typically present in organisms of that taxonomic group:

The normative force [of biological functioning] is part of and integral to a larger account of how individual organisms of that kind work as a whole on the whole.... The understanding is holistic and statistical: the norms governing the component functions are intelligible together in terms of their interdependent roles in enabling the whole system to succeed— that system-level success being understood in turn in reproductive and evolutionary terms. (1998, 308)

Haugeland nevertheless argues that norms of biological functioning cannot articulate intentional

content. The reason he offers is that biological systems cannot distinguish between proper functioning and objective correctness or truth. The outcome of the argument, however, is that biological function cannot confer meaning upon its functional performances.

Haugeland introduces a telling example to highlight the point. He imagines a species of bird whose normal functioning correlates its behavioral response to its perceptual input, in this case by refraining from eating most yellow butterflies, while eating most others when possible. These perceptual and behavioral mechanisms evolved in response to an environment in which most yellow butterflies are poisonous, and most others are not. Haugeland's claim is that such biological functions can genuinely account for the failures of those birds whose mechanisms do not make the normal discriminations (they are abnormal rather than merely variant, because their variance is functionally relevant). These functional norms cannot, however, make the birds' normal behavior meaningful. In particular, they cannot under any circumstances enable us to understand normal function as mistaken in either direction. A (normal-functional) refusal to eat rare, nonpoisonous yellow butterflies does not involve mistaking them for poisonous ones. Consequently, in the more common case, the birds' evolved mechanisms are also not intentionally directed toward an avoidance of poisonous butterflies. In the other direction, when normal function does not discriminate certain odd combinations of other colors from yellow, or certain odd patterns of yellow from non-yellow, it likewise cannot be mistaken about those butterflies' color. We, who understand the concepts of color and poison, can recognize such normal-functional responses as design limitations, but not as errors, because "correcting" those limitations has no role in the behavior (or its selective history).²

Biological functions can thus articulate only whatever patterns in the world that they

² Strictly speaking, they would only indicate design limitations if they resulted from a lack of relevant variance in detection- or response-mechanisms on which selection could operate, or the intrusion of non-selective "forces" such as drift. In other cases, the actual patterns exhibited could be functional, if a more fine-grained selective mechanism were too costly to the organism in energy requirements, discrimination-time, or some other selectively relevant feature. The moral here is that what is a selectively relevant biological trait, and what is a selectively relevant feature of an organism's environment, are holistically determined by its overall pattern of behavior and selection history, and not simply by more local correlations between perceptual input and behavioral output.

actually respond to (when functioning normally), even if those patterns are gerrymandered from the perspective of our conceptual understanding. Moreover, even if those response patterns were de facto co-extensive with conceptually significant features of the world, such that the birds always and only avoided yellow butterflies, those patterns would not display an intentional directedness toward the butterflies' color, for that result would merely be a de facto contingency. Haugeland does not spell out the underlying principle here, but the point is clear enough: intentional directedness must introduce a possible gap between what is meant and what is actually encountered, such that there is a possibility of error. Thus, in the case in which the avoidance response was actually co-extensive with the butterflies' color, a counter-factual query would be telling: if there were to be a shade or pattern of yellow that the birds' normal functioning would not lead it to avoid, would it have any means of self-correction?³ If there is no way for "it" (individually, collectively, or via a prosthetic phenotypic extension) to hold its own performances to account in some respect, then its behavior cannot properly understood to be intentionally directed in that respect⁴:

³ Haugeland himself uses such a counter-factual thought experiment to a similar end at the conclusion of "Objective Perception," by imagining a dog who seems able to recognize and distinguish different members of the same family, but is incapable of responding appropriately to the (impossible) counter-factual situation in which the family members' individual physiognomic properties were redistributed among them. So long as an ersatz intentional system is exposed only to the actual conditions to which its development is already adapted, it can seem intentionally directed, but its inability to respond appropriately under extraordinary circumstances exposes the illusion. Haugeland's strategy here parallels Dreyfus's (1979) earlier objections to the alleged intentional directedness of AI programs like Roger Schank's restaurant scripts, which cannot handle counterfactual circumstances for which they were not already designed (e.g., Schank's restaurant scripts' inability to answer questions about whether the waiter is wearing clothes).

⁴ The notion of an extended phenotype (Dawkins 1982), which originated in Dawkins's genic-selectionist program, but has been adapted to a more biologically adequate conception of natural selection as operating upon "developmental systems," incorporates such characteristic features of a biological life-pattern as beaver dams, bird nests, spider webs, or cities as part of the organism's phenotype. I include the possibilities of collective and phenotypically-extended means of correction within the scope of Haugeland's argument, because just as we cannot assume in advance what is a biologically relevant environmental feature or organismic trait, so we cannot assume in advance what is the relevant intentional "system." Biological intentionality

The trouble with the insectivorous birds is that there is no definition of that to which they are supposed to respond except as that to which they do respond when everything is functionally in order. ... The colors of the butterflies have no normative status at all apart from their involvement in that normal functioning. (1998, 314).

Embedded in this argument for the limitations of any selective-functionalist approach to intentionality, we can already see the need for some of Haugeland's characteristic constructive claims: that intentionality requires reflexive self-directed compartments that would constitute the possibility of self-correction, and that it must be directed not just toward actual entities, but toward a "space" of possibilities and impossibilities. Moreover, the argument suggests a different way of expressing what is characteristic of the "ersatz" intentionality that Haugeland acknowledges can be exhibited by animals or computer programs. Ersatz intentionality can mimic its genuine counterpart, because a system can be "designed" (selectively or programmatically) to produce actual patterns more or less co-extensive with a subset of those patterns that a genuinely intentional system would accept as possible. Those patterns of responsiveness are nevertheless not accountable to anything beyond their own functionality, because they only establish a de facto correlation rather than constituting a space of possibilities.

Let me sum up Haugeland's criticisms of those positions that occupy A2 on my table. These positions seek first to account for how intentionality is directed toward the world, by starting from a system's actual engagement with its surroundings, whether through causal interaction, or through a history of natural selection. They then try to account for how that system's engagement with the world is intentional and hence meaningful, in terms of its characteristic patterns of perceptual constancy or evolved functionality. Yet the patterns that show up in the causal transmission of information or the natural selection of functional performance cannot articulate a meaningful directedness. The constancy that Dretske hoped would identify an object as a unique target of recognition in fact must exist at every stage of causal transmission. We readily recognize it located at the object, and not at other stages, because patterns of stimulation or vibration are gerrymandered. Yet the difference between

could be a property not of individual organisms, but of communities of organisms, or of organism/environment complexes.

gerrymandered and coherent patterns is not recognizable at the level of information flow, but only via standards imposed upon that process from elsewhere, namely from the conceptual standards that govern meaningful experience. Dretske's account thus presupposes what it seeks to explain, namely the difference between our perceptual relations to perceived objects, and to their causal descendants or ancestors. Similarly, biological functionality (or artificial design, in the case of sophisticated AI computer programs) can seem to mimic patterns of intentional directedness, by creating patterns of actual response that under typical conditions mostly conform to patterns that a genuinely intentional system could recognize as meaningful. Yet the difference between genuine and ersatz intentionality becomes recognizable at the margins, either under actual conditions in which its discriminatory capacities fail to match an intelligible difference (as in the birds's response not always tracking real differences in color or toxicity), or under counterfactual conditions that expose the limits of its selective adaptation. Ersatz intentionality cannot distinguish between correctness and error, because whatever pattern it actually picks out under normal conditions is what it is "aiming toward." There can be no gap between normal function and correct identification, since biological normality is the only normative consideration in play.

The power of Haugeland's line of argument shows up especially clearly in its applicability in an unexpected direction. Bert Dreyfus's account of skilled practical coping has been a crucial influence on Haugeland's work, and important elements of Dreyfus's approach are embedded within Haugeland's own constructive account, most notably in the central role played by mundane and constitutive skills, and the irreducibility of such skills to explicable rules. Yet Haugeland's line of argument against Dretske or Millikan also applies tellingly to Dreyfus's efforts to understand skilled coping as a distinct "ground floor" of intentionality to which the conceptually articulated upper floors must be subsequently connected. Dreyfus shares with Dretske or Millikan a commitment to beginning with intentional fulfillment rather than empty intending, in ways that highlight continuities between human intentionality and that of animals. Dreyfus of course characterizes that fulfilling comportment phenomenologically rather than naturalistically, as a responsiveness to tensions and solicitations in bodily engagement with one's circumstances, rather than as information-flow or naturally selected functional responsiveness. Yet once it has been separated from higher-level normative concerns that require conceptual

understanding and reflective assessment, skilled coping encounters the same problem that emerged for Haugeland's imaginary example of an insectivorous bird.

Consider a grandmaster playing blitz chess. The grandmaster's ability to recognize and respond almost instantaneously to complex patterns on the chess board is the outcome of an extended "selective" regime (artificial selection involving study of past games leading to recognition of significant configurations, rather than natural selection operating on a population).⁵ To the extent that we are simply talking about a felt responsiveness to complex perceptual configurations experienced as tensions and solicitations, then the grandmaster's play allows no space for the occasional strategic or even constitutive error. Grandmasters playing blitz chess do make errors, of course. Yet Dreyfus's account of skilled coping as ground-floor intentionality cannot recognize them as errors, unless the conceptually-articulated regulative and strategic norms of chess play already constitutively govern the pattern-recognition capacities involved. Just as "there is nothing that the [bird's] response can "mean" other than whatever actually elicits it in normal birds in normal conditions" (Haugeland 1998, 310), so if Dreyfus were right, there is nothing that a "normal" grandmaster's blitz chess play could mean other than what grandmasters actually do in various actual board configurations without conceptually articulable norms already at play in their performances. Any patterns that characteristically trouble blitzing grandmasters could only be recognized (from without, by systems that actually understand and deploy chess concepts and standards) as design limitations in their trained cognitive orientation, rather than errors in play (and of course, the counterfactual case of obscure positions that might not be arrived at in the ordinary run of play would be relevant here as well). As we shall see, the problem in Dreyfus's account, which in the end makes it unable to account for grandmasters' performances as chess, or their coping with the board as skilled rather than just functionally adapted to the actual history of play, is the inverse of the difficulty that Haugeland attributes to the B1 accounts of Brandom, Davidson, and others, to which I now turn.

⁵Haugeland endorses Dreyfus's main concern in the latter's argument, which is to insist that the patterns recognizable by grandmasters may have no higher-order articulation than that constituted by grandmasters' ability to recognize them, and hence their skillful recognition is irreplaceable by any rule-governed system.

Ic— Haugeland versus B1

So now let us consider the accounts in B1. Haugeland nicely captures a key feature of all of these approaches to intentionality by characterizing them as “interrelationist” (1998, 207-8). These approaches (which include Dennett’s, Davidson’s, Brandom’s, or Rorty’s, among others) all begin by specifying a holistic pattern of comportment that, if all works out, will collectively constitute meaningful, intentionally directed performances and states. In some cases (e.g., Dennett), the pattern is only recognizable as such from outside; in others (e.g., Brandom, and I would argue, Davidson), the recognition of the relevant pattern is itself construed as part of the pattern. The pattern in question, however, whether construed in terms of the intentional stance, radical interpretation, giving and asking reasons, or the like, is only vindicated as genuinely intentional if it can show that how we encounter the world causally or perceptually has a normative grip upon this pattern. Advocates of this approach differ on how to account for the normative force of the world’s grip upon putatively intentional comportments. Dennett looks to empirical assessment of its predictive success; Davidson to the token identity of rational patterns and causal interactions; Brandom to the objectivity of conceptual norms; and so forth. Haugeland’s arguments are intended to show that the world could never actually get a grip upon these patterns, because the only grip such accounts can acknowledge is internal to a constitutive holistic pattern of “meaning,” and thus provides no external constraint upon that pattern.

Haugeland argues against these views from two complementary directions. On the one hand, he argues that they allow only for “mere coherence” among the comportments they recognize as rational, rather than genuine accountability to objects. Thus, for example, he endorses McDowell’s argument that Davidson illegitimately “helps himself to the idea of a body of beliefs,” and he concludes that Brandom’s supposed objectivity proofs

show that there is no legal move, in Brandom’s system, from ‘Everybody believes *p*’ (or ‘I believe *p*’) to ‘*p*’. But they don’t show anything at all about what *could* legitimate ‘*p*’ instead; in particular, they don’t begin to show how ‘*p*’ could ‘answer to how things actually are’. (1998, 358)

Dennett’s intentional stance does not even purport to constrain its own application: nothing in Dennett’s account (apart from the institutional constraints imposed within normal scientific

practice that a charitable reading may extract from his naturalism) provides a check on the interpreter merely claiming predictive success for her interpretations, or in the other direction, dismissing a putatively intentional system as irrational after the failure of only the most desultory efforts to make sense of its performances.⁶ Rorty even celebrates the abandonment of any “hankering after objective truth,” with the result that Rorty’s conversation of mankind is governed solely by the *verfallende* norms of curiosity and idle talk.

Haugeland’s other argumentative strategy is to emphasize the inadequacy of theirs (and other) constitutive appeals to social conformity as a normative constraint upon individual deviance.⁷ Social conformity can produce complex patterns of social institution:

It isn’t only the norms as such that are socially instituted, but also the respective behaviors and circumstances that those norms “connect.” Thus, what it is to greet someone, and what it is to be a circumstance in which greeting is appropriate, are nothing other than what the community members accept and deem as such— ... they are themselves instituted along with the normative practices in which they occur, [which] allows for considerable intricacy and interdependence, [instituting] statuses and roles, accrued over time, [including] those contingent on whose turn it is, who owns what, which water is holy, or how the teams stand in the league. (1998, 311-312)

Yet what it cannot do, Haugeland argues, is to constitute objects as authoritative standards to which community agreement is itself accountable. So an entire community can establish subtle

⁶The real force of Haugeland’s argument against Dennett, therefore, is his argument in “Truth and Rule-Following” that the social-institutional constraints of scientific practice are insufficient to secure accountability to the world without existential commitment.

⁷ While Sellars, Brandom, and Rorty explicitly make appeal to “we-intentions,” socially recognized normative statuses, and community agreement, other inhabitants of B1 eschew explicitly social accounts. Yet Haugeland nevertheless clearly finds social conformity at work implicitly in their accounts. Davidsonian radical interpretation relies upon a postulated “massive agreement” in beliefs as criterial for radical interpretation of what a speaker means, and closes the circle of social conformity with the reciprocal application of this criterion “at home” (and also in interpreting others’ interpretation of oneself). Dennett does not apply his account reflexively in this way, but then his account of predictive assessment tacitly relies upon acceptance by a scientific community.

norms that determine when a folk dancer has or has not grocked, in circumstances that were or were not grockworthy; moreover, it can even distinguish between appropriate responses to the sacred jonquil and to the merely profane jonquils normally found in gardens and florist shops. What it cannot do is to mean by that response (e.g., by uttering ‘scarlet’ rather than ‘yellow’) that the difference between the sacred and the profane is chromatic.

Haugeland’s arguments against the various B1 and A2 approaches are not merely an adventitious grab-bag of arguments against philosophical positions with which he happens to disagree. They are systematic attempts to rule out whole families of theoretical projects.⁸ Moreover, the considerations that account for the failures of these alternative approaches do constructive work in Haugeland’s own account of how intentionality emerged as a late-breaking development within (but not of) the biological history of *homo sapiens*. Since I find myself in considerable sympathy with these arguments, and take them as compelling reasons to understand intentionality in terms appropriate to B2, I shall not address them directly. I shall instead consider their implications for how to characterize a B2 approach. Part of Haugeland’s interpretation of the failure of A2 and B1 approaches is that those failures show that understanding intentionality in terms of biological functions or social institutions is untenable. I think that conclusion is correct and important, as far as it goes. It remains to be seen, however, whether these considerations prevent us from understanding intentionality in terms of human biology and sociality. To that question, I now turn my attention.

II— Rehabilitating the Biological and Social Dimensions of Intentionality

Haugeland’s arguments against the A2 and B1 strategies make it initially difficult to see

⁸ I have not discussed Haugeland’s criticisms of the A1 positions. Not many of these approaches receive explicit critical attention in his work, apart from his early forays against GOF AI (good-old-fashioned artificial intelligence). Partly, I suspect that is because he recognizes A2 and B1 as the more serious philosophical alternatives. More fundamentally, I think it is because his reasons for rejecting the A1 positions will turn out to be a fairly straightforward combination of the difficulties confronting A2 and B1 (they neither institute an articulated normative space, nor account for how intentional systems are accountable to norms external to their own performances). If I am right about this last point, then it is a considerable virtue of Haugeland’s lines of arguments (and perhaps of my classificatory array) that his arguments combine in this way. I shall leave that exercise to another occasion, however.

how our biological functioning or our social practices could play a constitutive role in understanding intentionality. Clearly any viable account must be consistent with understanding ourselves as evolved, functioning organisms who participate in complex, iterated social practices. Yet Haugeland's arguments may seem to rule out understanding intentionality and its normative accountability to the world as constituted by features of our biological functioning or social relations.⁹ Three steps are needed to see how one could accept Haugeland's critical arguments, and yet understand intentional normativity in social and/or biological terms. The first two steps involve re-thinking our understanding of biology and social life, respectively. An important reason why philosophers have failed to understand intentionality in biological or social terms is the inadequacy of our familiar conceptions of the biological and social domains. The final step, already prepared in the first two, is to recognize the artificiality of the separation of the biological from the social dimensions of human life and understanding. Neither biological nor social conceptions of intentionality will do, if these are taken to be alternative approaches. We need to grasp human social life as an integral part of our biology if we are to see how it opens a space of intelligible possibilities. Needless to say, working out these three steps is a large project that cannot be accomplished in a single paper (let alone the second half of a paper). What I hope to do is to sketch an overview of the key steps with an eye to how they enable responses to Haugeland's critical concerns.

IIa— Better Biology

I will be relatively brief about the requisite changes in our understanding of human biology. I have discussed some of them here before, as has Mark Okrent. Moreover, while they are necessary for an adequate understanding of intentionality as a bio-social phenomenon, they

⁹ Strictly speaking, Haugeland's own account of intentionality in terms of constitutive skills, standards, and (existential) commitments does include a constitutive role for social practices and aspects of human biology. Haugeland postulates a biological disposition to social conformity, and assigns a crucial role to social practices in instituting a complex field of social roles and instituted norms. Yet these only provide a background to intentional normativity rather its constitutive basis. Only the (existentiell) possibility of taking a kind of individual, existential responsibility for the conformity and intelligibility of all constituents of our practices (including what others do and how things show themselves) allows for the kind of understanding that Haugeland takes to be constitutive of genuine intentionality.

are not sufficient. There are three points I want to emphasize:

1. Haugeland's arguments responded to the predominant conception of biological normativity in terms of functions, and that was appropriate at the time they were written. Yet Okrent (2007) has now convincingly displaced that conception. We should think of a living organism as a way of living in an environment, whose ongoing maintenance and reproduction is the goal of what the organism does. Moreover, neither the organism's way of living, nor its environment, can be defined independently: the "environment" of an organism is not the totality of entities and their interactions in its vicinity,¹⁰ but only those that intra-act in significant ways with its characteristic life patterns.¹¹

2. Classical genetics and its role in the modern synthesis have encouraged the reification of organismic functioning and evolution in terms of "genes." Molecular genomics, developmental genetics, and the emergence of evolutionary-developmental and ecological-developmental biology have exploded much of that reification, suggesting a shift toward developmental cycles and patterns of co-development, with evolution occurring in multiple dimensions (Jablonka and Lamb 2005). One consequence of this recognition is that the characteristic goals of an organism's activities are a shifting target, especially through niche construction, by which an organism's way of life modifies its own defining conditions by modifying or even constructing its environment, and also through other forms of phenotypic plasticity. Genes for the most part cannot be identified with biologically relevant traits, but are

¹⁰ Part of the problem is that one can only define its relevant "vicinity" via the characteristic life patterns of the organism. A salient example, discussed in an earlier talk here, is provided by desert-dwelling populations of Drosophila pseudoobscura, which are more rather than less hydrophilic than their con-specifics inhabiting moister regions. That seems paradoxical, until one realizes that the relevant environment for the desert Drosophila is the very moist micro-environment within a millimeter of a leaf.

¹¹ Many of those intra-actions are indirect. Features of an organism's physical surroundings to which its own development and reproduction are completely indifferent may nevertheless be indirectly part of its developmental and selective environment if they belong to the biological environment of other species that are integral to that organism's development and evolution. Of course, the "strong holism" of biological properties (Okrent 2007, ch. 3) means that these features still have different significance for those two species of organisms.

instead highly important developmental resources that can take up different roles in various ways of living.

3. A third change of special importance to human biology results directly from the first two. Language is a central feature of human biology that must be understood in developmental terms. Languages are integral to the developmental environment in which human organisms become what they are, and thus are a pre-eminent example of niche construction. Moreover, one cannot think of the evolution of language in terms of an evolved capacity selected by an “external” environment for its functional role in a relatively fixed human way of life. The human phenotype has substantially co-evolved with the emergence of languages, such that languages have driven important aspects of human evolution and development.¹²

It is worth noting why the first two points are not sufficient to challenge Haugeland’s arguments about the difference between ersatz and genuine intentionality (and the third does not help without incorporating the revisions of familiar conceptions of social practices; language and its crucial role in human understanding shows especially clearly why the separation between human life and human society is so misleading). If anything, these points indicate more clearly why biology “by itself” cannot account for intentional normativity. The goal of an organism’s activity is its “normal functioning,” i.e., the way of life typical of organisms of that kind. One can account for abnormal behavior, and one can also understand innovative behavior, as what was once abnormal becomes incorporated within the organism’s characteristic way of life. Moreover, one can recognize that the organism’s way of life is “transcendent” in a limited respect, since it implicitly articulates its environment as differentially relevant to what it does. Yet while one can explicate the normativity of the organism’s goal-directed activity (does it succeed, or contribute to the possibility of such success, in maintaining and reproducing itself as *energeia*?), one cannot understand its accountability to anything beyond how it actually goes about trying to maintain its characteristic way of life, that is, as obstacles or affordances. In that context, Haugeland’s argument still applies. The organism’s normal way of life as the goal of its constituent performances indicates whichever features of its normal environment that it

¹² For discussion of these aspects of the biology of human language, see especially Deacon 1997, Lloyd 2004, Dor and Jablonka 2000, and Jablonka and Lamb 2005.

characteristically responds to, however gerrymandered they are from our conceptually articulated point of view. Even when the organism's behavioral repertoire is highly flexible, such that it is capable of novel responses to novel situations, its instrumental rationality toward the determinate goal of its own biological maintenance and reproduction provides the only normative authority in the vicinity. We can, of course, explain in many cases why that behavioral or physiological repertoire is as successful as it is at self-maintenance and reproduction (and sometimes why it fails to maintain itself under changing selective pressures). We nevertheless have no grounds for understanding that repertoire as striving toward a more conceptually adequate articulation of the world, but falling short.¹³

Iib**— Better Social Theory**

I now turn to a reconsideration of the social dimensions of intentionality. A widespread philosophical strategy appeals to the common behavior, presuppositions, or commitments of a community of agents: the convergence of various performances into communal regularities supposedly institutes norms governing the performances of individual members of that community. Haugeland's arguments target precisely such a conception of the "social world." Haugeland himself has constructively supplemented such accounts by postulating a biological disposition to "conformism" (modeled on Heideggerian *Abständigkeit*) to explain how such commonalities in performance might arise, and how they have normative force. I nevertheless think that this philosophically influential model of social life is unrealistic and distorting, in ways that parallel the better understood inadequacies of neo-classical economic models of rationality (in "idealizing" away important determining factors in the processes they seek to understand, and in tacitly presupposing what they seek to account for). A more adequate model of social interaction, however, when situated within a more adequate account of human biology, promises

¹³ Mark Lange's (2000, 2007) important work on the role of laws in scientific practice, including in the life and other so-called "special" sciences, provides another way of seeing this point. Laws (including laws concerning what Okrent would call an organism's way of living) provide norms of inductive reasoning, concerning what it would be for some body of evidence to go on in the same way, without unmotivated bends or scope restrictions. To take the normal discriminations made by an organism's way of life to be governed by a conceptually cleaner norm would be to introduce an "unmotivated bend" in our inductive reasoning in cognitive ethology or other relevant fields of functional biology.

to fit within the commitments of a B2 approach to intentionality.

Part of what make neo-classical models of economic behavior so powerful is that they get some things importantly right about their underlying conception of exchange relations. After that, a combination of simple, intuitively plausible assumptions (which nevertheless define away important determinants of market behavior), theoretical tractability,¹⁴ and slippage between descriptive and normative considerations (along with a helping of intellectual inertia) help account for the continuing influence of neo-classicism. I think similar considerations account for the persistence of communal conceptions of social life despite their evident empirical implausibility.¹⁵

So what is right about communalist models of social institution? Social practices do depend upon the interdependence of socially-articulated roles, with a consequent need for communication and mutual co-adaptation of their constitutive performances. Moreover, social-institutionalist accounts rightly recognize the role of both censoriousness and self-monitoring in sustaining that co-adaptation. They rightly highlight the “groundlessness” of social practices, in two crucial senses: social practices do not exist apart from their ongoing reproduction in new performances of those practices (which also accounts for the localizability of social practices, which only exist “where” they continue to be performed), and nothing “outside” of the interactive play of those performances determines their form or content.¹⁶ These models also rightly

¹⁴ The specific form taken by the theoretical tractability of neo-classical economics, namely its straightforwardly calculable mathematical form (aided and abetted by the cultural prestige of mathematical theorizing, reinforced in turn by its opacity to the uninitiated), is nevertheless disproportionately important to the persistent centrality of neo-classicism in economics and other social sciences.

¹⁵ During the brief heyday of D-N explanation in the philosophy of the social sciences, my late Wesleyan colleague Vernon Dibble put forward the Laws of Sociology to satirize this conception of social scientific understanding. The pervasive social phenomena captured by the First Law of Sociology (“Some do, some don’t”) succinctly express the underlying empirical difficulty confronting communal social-institutionalist models of normativity.

¹⁶ Physics and biology articulate important boundary conditions for social practice, but these conditions accommodate extraordinary social and cultural differentiation, despite a long history of failed efforts to ground social life in our postulated biological “nature” (as essence).

emphasize the iteratively anaphoric character of social practice:

We can envisage a situation in which every social practice of [a] community has as its generating response a performance which must be in accord with another social practice”
(Brandom 1979, 189-90)

Such iteration accounts for the rich and complex articulation of the social world, and its ontic fecundity.

Yet it seems to me that the most notorious false-idealizations of neo-classicism (commensurability of desires/goods, perfect information, a psychology of utility maximization, and market clearance/equilibrium) have specific counterparts in social-institutionalist accounts of normativity. Consider these parallel false-idealizations in turn. Social-institutionalist accounts of normativity usually only talk about the normative relations between an individual agent engaging in one activity within a single authoritative community. Yet real human agents belong to many communities at different scales, and participate in many practices. Their performances in one context typically bear upon their performances in others, yet those interactive effects work differently for the various agents whose performances matter in the original context. Thus, even if one accepted a communalist account, the actual normative force bearing upon the performances of any single agent would not be the same as those bearing upon others. Hence there is reason to expect that even under genuine social pressure to conform to what others do, people’s performances need not converge on a single common standard. Here we have a social-institutionalist parallel to the problem of commensurability of goods in neo-classical theory.

Second, even social-institutionalists recognize an epistemic division of labor, such that not all participants within a community share the same level of understanding of what are supposedly the same practices. These differences are reconciled with the normative authority of a single community by positing norms of deference to experts, along with general convergence in the identification of relevant expertise.¹⁷ Yet such deference is often contested, especially in settings where multiple practices have a bearing: multiple forms of expertise may be relevant,

¹⁷ Consider Heidegger’s differentiation of genuine, average, and privative levels of understanding, as well as Putnam’s more widely recognized account of an epistemic division of labor.

different standards of expertise may be recognized, and deference itself may be contested. Perhaps most crucially here, however, the very norm of shared understanding has been increasingly contested. Peter Galison's (1997) substitution of "coordination of action and belief" within limited "trading zones" (modeled on the linguistic construction of pidgins and creoles), in place of translation or other models of convergent understanding, is just one example of the recognition that social life often proceeds via the localization, management and coordination of divergence rather than the construction of a genuinely common practice. The commonplace coordination of social life in the absence of shared understanding parallels the way real markets diverge from the neo-classical idealization of perfect information.

In a third false-idealization, this model of normative authority conceives the normative force of common performances as monolithic. While some deviance normally persists, it has no reciprocal effects upon what the community does.¹⁸ Even when shared practices change (whether by drift or persuasion), such change is not normatively accountable, but at most causally explicable.¹⁹ Resistance to community norms is not, indeed cannot be, treated as integral to the social-institutionalist conception of normative authority and force, but only as an external imposition that changes the content of what the community does without affecting its communal character. Here a Foucauldian account of power, which recognizes that power is always co-constituted by resistance, is both empirically and conceptually more adequate.²⁰ Understanding

¹⁸ Haugeland's account of "socially instituted normativity" is exemplary here in its parenthetical condition: "When an individual is deemed out of step, other community members intervene so as to restore compliance (in a way that, by assumption, is likely to be effective)" (1998, 311).

¹⁹ Haugeland's treatment of such cases is again instructive: "Imagine a charismatic florist who ... convinced everyone that a certain jonquil is sacred." Charisma and persuasive rhetoric are exemplary cases of non-rational causation that is not normatively explicable.

²⁰ It is conceptually more adequate, because it allows (indeed, makes central to the account) how power acts not merely via force applied to agents, but as power that re-shapes their field of possible actions. This re-shaping works in both directions, so as also to affect what it is intelligible for dominant agents (in this case, members of the "majority rule") to do. Instead of treating (causal) force as normatively inert, Foucault's concept of power explicates the normative significance of both force and resistance. For more detailed discussion, see Rouse 2005 and 1996 (pp. 180-194).

social relations in those terms, however, locates normative authority not in the dictates of a community to potential deviants, but in the ongoing, dynamic interactions between such imposition and resistance. Recognition of the normative significance of the interplay of power and resistance thereby challenges social-institutionalist conceptions of normative authority in ways comparable to the effect upon neo-classical models of individual behavior of denying that agents are perfectly rational utility maximizers .

These various challenges then also come together to undercut the most fundamental idealization involved in a communal-institutionalist conception of norms. Communal-institutionalist accounts only apply when there is a determinate fact of the matter about what some community normally does. In other situations, they supposedly explain why interacting agents tend to converge on a conformist “equilibrium,” but until that convergence occurs, the model cannot account for normative force in any particular direction. Yet just as markets may never actually reach neo-classical equilibrium (in which case, neo-classical economic theory does not characterize their behavior), so in many situations, “communities” may never settle on a determinate set of common practices. Indeed, that would be precisely the expected outcome once we recognize that agents participate in multiple overlapping practices, that their participation involves not a shared understanding with univocal deference to experts but only a localized coordination of divergent understandings, and that whatever normative authority is thereby constituted involves a dynamic interplay between power and resistance. Under those circumstances, difference irreducible to abnormality may be pervasive, with the consequence that a social-institutionalist conception of normativity would not apply.

The question that matters for my concerns, however, is whether a more adequate conception of social relations can circumvent Haugeland’s arguments against understanding intentional normativity as social-institutionally constituted. Much of my work for the past thirty years has been devoted to the articulation of a more adequate account of practices as the locus of intentional normativity.²¹ I conclude this section by highlighting a few key points that point toward a “non-equilibrium” account of the normative authority and force constituted through

²¹ For exemplary accounts, see Rouse 2007a, 2007b, 2002 ch. 5-9, and/or 1996 Part II.

such practices.²² The final section will then take up how I think these points could accommodate Haugeland's objections to understanding intentionality as social or biological.

Practices are not social regularities, and do not define communities of agents. They thus do not consist of various agents performing in similar ways, or sharing background beliefs or presuppositions. Practices instead are composed of performances that are mutually interactive in partially shared circumstances.²³ The intelligibility of various performances within a practice normally depends upon the anticipation and achievement of appropriate alignment with others' performances and their circumstances, toward some possible "end." Ends in this sense, however, are not something external to a practice for which it is merely instrumental. The end, the possibility toward which the practice is directed, is the practice itself, encompassing its present performances within a larger whole whose character is at issue and at stake in those performances.²⁴ In taking up various practices, then, agents are (in Heidegger's terms) "pressing into possibilities" that matter as partially constitutive of who and how they are.

Yet although practices are constituted as intelligible ends through the ongoing mutual alignment of various performances and circumstances, the performers and circumstances are usually only partly accommodating. In response to such mis-alignments, agents adjust what they do, both by changing their own performances, and by trying to affect what others do, or to re-

²² I talk about practices rather than social practices, because an adequate account of practices excludes any significant distinction between what is social and what is natural, or between social relations and their material embodiment and causal involvements. Practices are bodily performances in material circumstances, and the circumstances belong to the practices themselves. I nevertheless situate this account as part of a better social theory to highlight their historical engagement and continuity with social-institutionalist conceptions of normativity and eventually intentionality, as modified in light of my criticisms of these familiar false idealizations.

²³ Instead of practices being constituted by the doings of independently identifiable agents, agents come into being through being caught up in practices. Here, the vocative and cognitive dimensions of discursive practices that play a central role in Kukla and Lance 2009 are integral to any adequate account. See section IIc below.

²⁴ Practices can also be nested within more encompassing practices, such that (for example) what is at stake in a particular practice may be the appropriate resolution of something at issue in the larger pattern of practice to which it belongs.

arrange the circumstances. These patterns of mutual responsiveness and recalcitrance typically focus a practice on specific issues. Issues arise wherever some adjustment of performances or circumstances is called for to allow the practice to proceed coherently and intelligibly. Moreover, as I shall discuss further below, as discursively articulate beings, we often respond to those issues in part by saying what the issues are, and what are their inferential and practical consequences. Of course, these efforts to talk through what is at issue in a practice, including responding to divergent articulations of the issues, are themselves further performances belonging to the practice. Through these ongoing interactions, or intra-actions in Karen Barad's (2007) more perspicuous term, practices evolve and articulate themselves.

The normative "force" that binds us to practices, and makes us responsive to these issues, comes from their "ends," the very possibilities they provide for intelligible ways to understand and enact ourselves in the world. Brandom long ago described this binding potential in terms of "expressive freedom":

Expressive freedom consists in the generation of new possibilities of performance which did not and could not exist outside the framework of norms inherent in social practices. ... Expressive freedom, as the capacity to produce an indefinite number of novel appropriate performances in accord with a set of social practices one has mastered, is an ability which must be exercised to be maintained. (1979, 194)

What is at stake in a practice, and in the issues that divide its practitioners, are the very possibilities for who and what we might be through ongoing involvement in and submission to the practice.²⁵ Brandom's talk of "norms" is then misleading: norms are not already determinate

²⁵ Although I cannot develop this connection here in any detail, my discussion of the normativity of practices in terms of issues and stakes closely parallels Haugeland's (1998, ch. 12) distinction between inner and outer recognition of patterns, which is one way in which he expresses how there comes to be a gap between meaning and truth. "Outer recognition" involves noticing an intelligible pattern; "inner recognition" involves recognizing whether some element fits within that pattern. What is at stake is how some intra-active performances and responses matter (the relevant patterns are normative accountings); what is at issue is how one ought to continue those performances in light of what is at stake in them. Of course, in the case of practices, internal conflicts over what is at issue and at stake are critical components of the "patterns" through which we encounter our surroundings and make sense of ourselves as agents.

standards to which performances are accountable, but are instead something ahead of us, toward which the various performances of a practice are mutually, but not always fully compatibly, directed.²⁶ Through participation in practices, we are always in the midst of “becoming who we are.” ‘Issues’ and ‘stakes’ are fundamentally anaphoric concepts. They permit reference to the scope and significance of a pattern, a concept, or a practice (what is at stake there), and what it would be for things to go on “in the same way” under other circumstances or more stringent demands (what is at issue), even though those issues and stakes might be contested or unknown.²⁷

With these considerations in mind, we can recognize that there probably are no practices that are merely socially instituted. Take Haugeland’s imagined example of “grocking” (a complex move within a folk dance that is only recognizable by skilled participants, and is called for only under similarly recognizable circumstances within a performance). The very description of it as a folk dance already partially indicates (and depends upon) what is supposedly at issue and at stake in correct performance (skilled display, collective identity, “entertainment” as counterpart to necessary activity, etc.). Moreover, responses to mis-performance in such an endeavor are not merely undifferentiated attempts to suppress divergence, but are themselves significant. Consider how one would respond differently to “mis-performance” due to clumsiness (as failure to “belong” to the group), clumsiness (as aesthetic failure), rebellion against community expectations, failed attempts at innovation, satirical critique, or momentary lapses in concentration. Consider in turn how the person being corrected, chastised, excluded,

²⁶ While I acknowledge some differences, this conception of comporting oneself toward issues and stakes is significantly indebted to Heidegger’s notion of understanding as pressing into possibilities (including the sense in which such possibilities can never be attained or “actualized”).

²⁷ Issues and stakes may be contested or unknown, but not non-existent. In this respect, my account of the normativity of practices meshes with a general feature of B2-conceptions of intentionality, which begin with an understanding of intentional fulfillment. Just as for Heidegger (for example), one can only intend something emptily because one has an understanding of what it means to be, so I will argue that one’s comportments can be directed toward non-existent or misconceived entities only because there is genuinely something at issue and at stake in the practices in which those comportments are caught up.

punished or ignored might respond in turn, and how others would take up the issue. How participants in a practice respond to inappropriate performances is itself part of a practice (which can itself be done correctly or incorrectly), including which mistakes are studiously ignored. It is only through the practice being caught up in an ongoing (and often changing) repertoire of nested possible responses that the practice actually has any normative authority, and any definite character as a practice. What the practice is goes beyond what is actually done, to constitute an articulated space of possibilities and how they matter. Moreover, it is possible that all of the participants are mistaken about what is actually at issue and at stake in the practices through which they enact their lives and make sense of themselves.²⁸

IIC— B2-Intentionality as Social-Biological

In this concluding section, I shall briefly point toward how I think the aforementioned revisions to philosophical conceptions of the biological and social domains enables a more adequate conception of intentionality. Only by developing such an account could we convincingly overcome Haugeland's reservations about understanding intentionality in those terms. I already indicated that a better understanding of biological normativity does not circumvent Haugeland's distinction between genuine intentionality and its "ersatz" simulacrum (even when the latter involves a flexible, rationally explicable responsiveness to more or less novel features of an organism's environment). An organism's normal ways of striving towards its characteristic goal cannot "mean" anything more or different than what organisms of that kind ordinarily do. Organisms (or species of organisms) can fail to maintain themselves in the face of environmental change, but they cannot be mistaken.

²⁸An account of normativity in terms of what is at issue and at stake in discursively articulated practices thus belongs to the post-Davidsonian tradition that emphasizes the "objectivity" of conceptual understanding. Yet the "objects" to which performances of a practice are held accountable are not "outside" our discursively articulated involvement in the world. Discursive practice cannot be understood as a "mental" or intra-linguistic structure or activity that then somehow reaches out to incorporate or accord to objects. The relevant "objects" are the ends at issue and at stake within the practice itself. "The practice itself," however, already incorporates the material circumstances in and through which it is enacted. Practices are forms of discursive and practical niche construction, in which organism and environment are formed together through mutually intra-active configuration. This point will be further developed in the next section.

A better understanding of “social” practices is also not sufficient. The problem with traditionally construed social practices is that they were not accountable to anything beyond what the community happens to do (or accept). To that extent, practices would be merely descriptively “normal” patterns of behavior rather than norms that constitute meaningful accountability to things aspectually intended.²⁹ In taking account of something supposedly at stake in a practice, which governs how one ought to respond to various (possible) performances of that practice by oneself or others (including nested sets of responses), however, I was implicitly appealing to something that social articulation alone cannot account for. Instead of instituting the relatively stable patterns of common activity postulated by social-institutionalists, it seems as if the complexity of social responsiveness would merely form shifting, drifting de facto patterns rather than relatively stable ones.

What does make a difference is to recognize a distinctive integration of socially and discursively articulated biological normativity (not just a combination of two distinct normative registers). The conception I am gesturing toward would join Haugeland’s, Heidegger’s, and other approaches to understanding intentionality that are located in B2 on my array. First, it is a normative approach, considering how our comportments are accountable to possibilities of contentfulness and correctness rather than a descriptive account of natural-selective, cognitive, behavioral, institutional, or other processes supposedly operative in producing intentional comportments as such. Second, it would start with a system’s actual involvement in its “surroundings,” and show how that involvement is meaningfully articulated, rather than beginning with empty (representational) intendings and then accounting for the difference between fulfilled and empty intentional directedness.

On this conception, our actual involvement in the world is conceived biologically, as organisms whose behavior aims toward the goal of maintaining and reproducing the dynamic pattern of its own simultaneous belonging to and differentiation from its environment. Yet the

²⁹ “Normal” supposedly has a different contrast class here than it does in characterizing normal biological functioning. Biological divergence from normal species patterns is abnormal; social divergence from accepted behavior is supposedly merely deviant. As we shall see, I think a social-institutionalist account cannot defend that characterization of non-conforming behavior, since deviance is a distinctive way in which behavior can matter. See the discussion below.

boundary between organism (as a reproducible, self-maintaining pattern) and environment is also ambiguous in several ways. As a “way of living,” we are individual organisms who participate in a larger pattern that constitutes that way of living as human.³⁰ Thus, other human organisms belong both on “our” side of the boundary between our (shared) self-maintaining way of living and our more or less shared environment, and also as integral parts of the environment in and against which we sustain our own individual existence. Languages are also both an integral and pervasive aspect of the environment in which we develop, and are appropriated within our own bodily patterns (literally so, since normal human neural and cognitive development only takes place within a discursively articulated environment, and the characteristic pattern of human life is a linguistically articulate one). Languages and their entanglement with other performances are thus the pre-eminent example of the evolutionary phenomenon of niche construction. Yet languages are also an exemplary case of a socially interactive practice, which only exist through human beings continuing to reproduce them in mutual responsiveness to one another in partially shared circumstances. Indeed, these are two different ways of describing the same point: it is language as a bio-social development that secured for us the possibility of intentional directedness toward entities as such.

That language is integral to human biological development and evolution is now a widely recognized (although still controversial) claim in evolutionary biology. We are linguistic/discursive beings, and not merely animals with an evolved capacity for language.³¹ I will not

³⁰ A similar point has played a prominent role in Haugeland’s (1982, 2008) own exposition of *Being and Time*, in the form of his insistence that Dasein is a singular entity, that is nevertheless also articulated into individual “cases.” Whether or not that point is correct as a matter of Heidegger interpretation (I think it is), I regard it as indispensable to an adequate understanding of intentionality.

³¹ The claim in question is not the more widely accepted view that a general capacity for language evolved at some point in the differentiation of *homo sapiens*, and was then a consequential factor in our species’s survival and demographic/geographic expansion. It is instead that languages, human neural and anatomical development patterns, and of course our distinctive patterns of neotenus development, co-evolved. The functional/anatomical patterns of human bodies are shaped by their development and evolution amidst discursive practice, while of course languages take the shape they do by selective reproduction suitable for human bodies and patterns of living. On such a conception, language and human cognitive functioning

discuss the neurological and anatomical evolution that has occurred in the human lineage under selection pressure for linguistic capacity, although it is a crucial part of the story. Rather I shall briefly highlight some other features of the linguistic/discursive way of life that such anatomical and neurological evolution have made possible. I think these features of our bio-social way of life together allow it to constitute the distinctively human phenomenon (so far as we know) that we philosophers have characterized as “intentionality.” There are five points I want to highlight.

1. The ability to recognize and produce articulated linguistic performances is itself a biologically evolved capacity, which also materially changes the environment in which humans develop and to which we respond. Spoken and written words, and the performances that deploy them, are among the most pervasive and influential features of the environment in which human beings develop biologically (and impairments of linguistic capacity, including the ability to respond appropriately to other speakers, are debilitating “abnormalities”).³² The specific capacities they enable, and the actual differential exercise of those capacities, become integral to the way of life that our ongoing activities aim to maintain and reproduce.

2. The emergence of language as a socially reproduced practice allows extraordinary capacities for the “internal” articulation/expression of patterns of intra-active performance, through its syntactically-marked semantic structure, its reflexive and recursive applicability, and the resultant ability to mark, track, assess and respond to an extraordinary range of differences and similarities in performance.

3. Yet language is not a general-purpose cognitive capacity. Its articulative possibilities are biased toward some domains of human expressive and responsive activity, and our linguistic abilities develop alongside other capacities that are not so readily articulable linguistically. Dor and Jablonka summarize this partiality well:

continue their co-evolutionary dance. For discussion, see Dor and Jablonka 2000, Jablonka and Lamb 2005 (especially ch. 6, 8), and Deacon 1997.

³² A striking example of the intertwining of the social and the biological dimensions of discursive practices are the widespread, ongoing efforts to transform discursively-articulated social life so as to incorporate as full participants those persons whose biological development impairs their ability to engage in linguistic interaction in the same ways others do.

The expressive envelopes of different languages are different in interesting and subtle ways, but they all share a common core. ... Interestingly, many of the messages which turn out to be very difficult to communicate through language seem to be very well suited for communication through other means of communication: we can mime and dance them, use facial expressions and body language to express them, paint and draw them, write and play music, prepare charts and tables, write mathematical formulae, screen movies and videos, and so on. (2000, 40)

Not only did we start out with non-linguistic cognitive and expressive capacities alongside the emergence of language, but those capacities have also proliferated and further developed.

4. Language is nevertheless not a self-contained practical/perceptual domain. It opens onto and “incorporates” other sensory/cognitive/performative capacities, via recognitive, demonstrative, anaphoric, and indexical locutions,³³ even while it is itself only intelligible as an integral part of our biological capacities for perceptual/practical interaction with our surroundings. I think the possibility of intentionality depends upon the development of such a semi-autonomous practice. The possibility of a gap between what our expressive performances “mean” or “intend,” and how those performances are involved with our actual surroundings, depends upon such semi-autonomy. We belong to a practice (“language”) that is both based in and accountable to our multiply intra-active involvements in the world, and yet which also

³³ The co-existence of language with other expressive capacities that are not readily expressible linguistically is part of what lends mistaken plausibility to the notion of “non-conceptual content.” Yet at least three mistakes must be made together to make this notion plausible. The first is to treat language as a self-contained practice, rather than one that depends upon both its incorporation within and its semantic inclusion of the whole of human bodily intra-action with our surroundings (including other discursively articulate human beings). The second mistake is then to conflate the conceptual domain with what is readily expressible in language. Language enables a distinctive capacity for conceptually-articulated normativity, but it is not co-extensive with it. We can articulate and express conceptual understanding through non-linguistic activities, although those activities are transformed by being caught up in discursively articulated normative accountability. The third mistake is to extract the result of their entanglement with discursive practice and identify it as an inherent feature of our various non-linguistic capacities in isolation. Here, the metaphorical connotations of “content” (suggesting distinctions between container/contained, and inside/outside) undoubtedly contribute to the ease of making this error.

enables the extensive articulation of patterns that are initially accountable primarily to other “intra-linguistic” performances.

5. Finally, but perhaps most crucially, the evolution of language draws upon, extends, and articulates our capacities for differential responsiveness to one another’s performances. The vocative (and recognitive) dimensions of linguistic activity,³⁴ and the consequent emergence of a more articulated mutual accountability, are what constitutes the possible divergence between our linguistic expression and our involvement in the world in other respects as a gap to which our various performances are accountable. Only because of the development of this vocative, second-personal dimension of our expressive and recognitive capacities (and their integration with the massive expressive/articulative capacities of language) can our responses to one another go very far beyond merely encouraging or discouraging various aspects of one another’s ongoing performances. Jablonka and Lamb’s (2005) account of the crucial, reinforcing entanglement of multiple dimensions of evolution (genetic, behavioral and discursive/symbolic in this case, since epigenetics may play a minimal role in vocative responsiveness) plays a crucial role here in my willingness to regard this aspect of the normative accountability that is constitutive of intentionality as biological .

Needless to say, I have barely gestured at the possibility of accounting for intentionality in terms of these aspects of our bio-social way of life as human beings. Nor have I considered at all where and how such an account would converge with and diverge from Haugeland’s own constructive approach to intentionality (as a modally-articulated existential commitment toward constitutive standards and skills). Yet I hope I have said enough to indicate why I think all five of these considerations would have to be worked out together, if one hoped to persevere with a broadly naturalistic conception of intentionality in the face of Haugeland’s critical arguments. Moreover, the specificity and complexity of the requirements that Haugeland’s arguments impose upon that task is one important measure of what those arguments have accomplished.

³⁴ Kukla and Lance’s (2009) ground-breaking book on the pragmatic topography of discursive practice is the pre-eminent discussion of the vocative and recognitive dimensions of language, and their role in constituting the normative accountability of discursive performances.

REFERENCES

- Aristotle. 1941. Metaphysics. In R. McKeon, ed., The Basic Works of Aristotle. New York: Random House.
- Barad, Karen. 2007. Meeting the Universe Halfway. Durham: Duke University Press.
- Brandom, Robert. 1979. Freedom and Constraint by Norms. American Philosophical Quarterly 16: 187-196.
- _____. 1994. Making It Explicit. Cambridge: Harvard University Press.
- Carnap, Rudolph. 1967. The Logical Structure of the World. Berkeley: University of California Press.
- Davidson, Donald. 1984. Essays on Truth and Interpretation. Oxford: Oxford University Press.
- Dawkins, Richard. 1982. The Extended Phenotype. Oxford: Oxford University Press.
- Deacon, Terrence. 1997. The Symbolic Species. New York: Norton.
- Dennett, Daniel. 1987. The Intentional Stance. Cambridge: MIT Press.
- Dor, Daniel and Eva Jablonka. 2000. From Cultural Selection to Genetic Selection: A Framework for the Evolution of Language. Selection 1: 33-55.
- Dretske, Frederick. 1981. Knowledge and the Flow of Information. Cambridge: MIT Press.
- Dreyfus, Hubert. 1979. What Computers Can't Do. Second Edition. New York: Harper & Row.
- _____. 2005. Overcoming the Myth of the Mental: How Philosophers Can Profit from the Phenomenology of Everyday Expertise. Proceedings and Addresses of the American Philosophical Association 79: 47-65.
- Fodor, Jerry. 1998. Concepts. Oxford: Oxford University Press.
- Foucault, Michel. 1977. Discipline and Punish. New York: Pantheon.
- _____. 1978. History of Sexuality, Volume 1. New York: Pantheon.
- _____. 2003. Society Must Be Defended. New York: Picador.
- Haugeland, John. 1998. Having Thought. Cambridge: Harvard University Press.
- Heidegger, Martin. 1962. Being and Time. New York: Harper & Row.
- Husserl, Edmund. 1982. Ideas, Volume I. Tr. F. Kersten. The Hague: Martinus Nijhoff.
- Jablonka, Eva, and Marion Lamb. 2005. Evolution in Four Dimensions. Cambridge: MIT Press.
- Jackson, Frank. 2000. From Metaphysics to Ethics. Oxford University Press.

- Kukla, Rebecca, and Mark Lance. 2009. Yo! And Lo!. Cambridge: Harvard University Press.
- Lance, Mark. 1998. Some Reflections on the Sport of Language. Nous-Supplement: Philosophical Perspectives 12: 219-240.
- Lange, Marc. 2000. Natural Laws in Scientific Practice. Oxford: Oxford University Press.
- _____. 2007. Laws and Theories. In S. Sarkar and A. Plutynski, ed., A Companion to the Philosophy of Biology. Oxford: Blackwell.
- Lloyd, Elisabeth. 2004. Kanzi, Evolution and Language. Biology and Philosophy 19: 577-588.
- McDowell, John. 1994. Mind and World. Cambridge: Harvard University Press.
- Millikan, Ruth. 1984. Language, Thought and Other Biological Categories. Cambridge: MIT Press.
- Okrent, Mark. 2007. Rational Animals. Athens: Ohio University Press.
- Quine, W.v.O. 1960. Word and Object. Cambridge: MIT Press.
- Rouse, Joseph. 1996. Engaging Science. Ithaca: Cornell University Press.
- _____. 2002. How Scientific Practices Matter. Chicago: University of Chicago Press.
- _____. 2005. Power/Knowledge. In G. Gutting, ed. Cambridge Companion to Foucault. Second Edition. Cambridge: Cambridge University Press.
- _____. 2007a. Practice Theory. In S. Turner and M. Risjord, ed. Handbook of the Philosophy of Science. Vol. 15: Philosophy of Anthropology and Sociology. Dordrecht: Elsevier. 630-681.
- _____. 2007b. Social Practices and Normativity. Philosophy of the Social Sciences 37: 46-56.
- Searle, John. 1982. Intentionality. Cambridge: Cambridge University Press.
- Sellars, Wilfrid. 1963. Science, Perception and Reality. London: Routledge & Kegan Paul.