

Social Ontology as Convention

Mark H. Bickhard
mark@bickhard.name
<http://bickhard.ws/>

Abstract

I will argue that social ontology is constituted as elaborate hierarchical and interlocking conventions of multifarious kinds. Convention, in turn, is modeled in a manner derived from that of David K. Lewis. Convention is usually held to be inadequate for models of social ontologies, with one primary reason being that there seems to be no place for normativity. I argue that two related changes are required in the basic modeling framework in order to address this (and other) issue(s): 1) a shift to an intentional model, among other reasons, in order to account for normativity, and 2) moving away from the belief-desire, propositional attitude, framework for understanding the intentional realm toward an interactive, pragmatic model of intentionality. These shifts provide natural approaches to: 1) understanding the normativities of social realities; 2) the sense in which social ontology is often constituted in implicit relations among the participants rather than elaborated and iterated explicit beliefs and desires; 3) and language.

Lewis (1969) defined a convention, roughly, as a behavioral regularity that constitutes a solution to a coordination problem, where a coordination problem, taken from Schelling (1963), is a problem in which there are two or more possible outcomes that are preferable to all others among all of the participants to the situation. The problem, then, is to coordinate among the participants so that one of those preferred outcomes is reached. A convention constitutes a solution to such a problem in the sense that the regularity is the basis for such “coordinated” activity arriving at one of the preferred solutions — e.g., a regularity of “When approaching an oncoming vehicle, pass to the right”.

There are many details of Lewis' model and arguments for his model, most of which I will not be concerned with — though, in some cases, they have been the focus of extensive criticism (such as the role of purported rationality on the part of the participants in establishing a convention: e.g., Gilbert, 1989). Instead, I will focus on one aspect of his definition, and show how, with appropriate modification, a central criticism of the model can be avoided and how this modified model captures the basic nature of social ontology. The aspect that I will focus on is Lewis' reliance on an ontology of behavioral regularities, and the criticism is that behavioral regularities are inadequate to capture the normativities that are involved in convention and in social ontology more broadly.

Intentionality: Capturing the Normativity of Convention

The basic critical point is that regularities of behavior do not necessarily involve any normativities at all, and that defining convention in terms of such regularities, therefore, does not capture and cannot capture the normativities of conventions — the sense (perhaps multiple and various senses) in which conventions are normative, in which, for example, conventions involve obligations and responsibilities.

There is one sense in which Lewis' conventions do at least *seem* to involve normativity, and that is the sense in which they are solutions to coordination problems, and coordination problems are constituted in terms of relationships among participants' preferences. Violation of a convention, therefore, would seem to violate the preferences of those involved, perhaps even of the preferences of the violator. There is, therefore, at least a kind of instrumental normativity: honoring a convention is a means to the end of satisfying various participants' preferences, and is normative in that sense.

Regardless of whether this point is accepted or not, however, it does not address nor capture other senses of normativity, such as that participants have rational warrant for an assumption of *commitment* that others will do their part in honoring the convention. I may, for example, be counting on you to take your usual walk on the usual path at the usual time for the sake of some other purpose of mine, and be instrumentally frustrated if you do not. But, unless my expectation regarding your walk is more than just a predictive expectation based on past regularity, I have no reason to be disappointed in *you* per se. If, on the other hand, we have a convention of meeting at a particular point

along your walk, then I might have grounds for more than instrumental disappointment. For the moment, I limit myself to simply pointing out that there are such realms of normative expectation that are not addressed by behavioral regularities per se.

The standard response to such considerations is to shift the ontological grounds for modeling social ontology, including conventions, to that of intentionality — e.g., beliefs and desires. Normativities can then follow from the basic normativities of intentionality. If an arrangement, perhaps a convention, for example, involves a commitment (for example) on the part of the participants, then violations are intrinsically violations of such commitments, and that form of normativity, at least, follows directly.

This move does introduce normativity in some fairly natural ways, and I will endorse the move to intentionality for this and for one other important purpose. But characterizing intentionality in terms of beliefs and desires and related propositional attitudes gives rise to its own set of problems — intractable problems, I will argue. In its place, I will present an alternative approach to intentional phenomena, a dynamic model of representation and other intentionalities that derives from considerations of pragmatic interactions with the world, rather than proposition based beliefs and desires about the world.

The kinds of problems that standard belief-desire psychology frameworks produce can be illustrated fairly simply: Social realities involve various kinds of commonality of “belief” and “desire” (and perhaps others), and modeling these commonalities encounters unbounded regresses. So, in order for it to be the case that we have a common understanding that the current situation is a “taking a walk together” situation, we seem to have to each have a belief that that is the nature of the situation, that the other believes that that is the nature of the situation, that each of us believes that the other believes that this is a “taking a walk together” situation, and the third level, fourth, and so on (Lewis, 1969; Gilbert, 1989; Schiffer, 1972; Schmitt, 2003). If these iterated beliefs about each others beliefs fail to hold at any level, then a counter example can, in principle (and often in practice), be constructed. Similar unboundednesses can occur for desires and other “propositional attitudes”. These unbounded iterations of propositional attitudes are not realistic for anyone — no one actually holds, or could possibly hold,

such hierarchies of beliefs explicitly — and are especially unrealistic for toddlers, infants, and pets, all of whom we do seem to have genuine social relationships with.

The point that no one holds such hierarchies explicitly provides a hint of what is needed: some way in which such characterizations of situations could be *implicit* rather than explicit. But noting the need for some model of implicitness does not constitute such a model, and certainly just using the term provides no such model. I will argue that it is on this point that standard belief-desire psychology, the propositional attitudes, fails, and that the alternative model of intentionality — the interactivist model — that I will outline provides a solution.

Intentionality: The Nature of “Aboutness”

There are two relevant aspects of the interactive model of intentionality: the positive model and the critique of standard approaches. I claim that standard approaches *create* the sorts of problems mentioned above and that the interactive model avoids them, so both aspects are important here. I begin with an introduction of the positive model.

Interactivist Representation

Consider a relatively complex animal in interaction with its environment. At any given moment, it is faced with a selection problem of what to do next. Selection of what to do next only makes sense within a range of “possible things to do next” that are possible in the current situation: it does no good to select “open the fridge for a beer” if you are in the middle of the woods. A frog, for example, might have several possibilities available: flick its tongue at a fly “there”, flick its tongue at a fly “over there”, at a worm “down there”, and jump into the water in response to a hawk shadow. The actual selection depends on various considerations of effort and outcome¹, but my focus at this point is that the frog (and you in the woods) must have some indications of the range of possibilities within which to make such a selection.

The crucial property of such indications (for current purposes, as well as for many others) is that they might be true or they might be false — they have truth value. So, the frog's indication of the possibility of flicking its tongue in this particular manner in order

¹ Such considerations open into the realm of motivation (Bickhard, 2000).

to eat may be false because what the visual system has detected is in fact a pebble tossed by some student of frog behavior rather than a fly.² The normativity of truth value is the most important characteristic of intentionality, of aboutness, and it emerges in anticipations of what is or might be possible.³

Anticipations of simple interactive possibilities, however, do not look much like canonical models of representation, such as of objects. So, even though this model of anticipated interactive possibilities might capture the emergence of truth value, it still seems lacking with respect to any kind of replacement for belief-desire propositional attitudes. Some elaboration of the model is required to show that it might have the resources for handling more complex sorts of representation.

There are two resources in the model that are most important for current purposes. The first has already been introduced: more than one possibility can exist at a given time. The frog in the above example had four possibilities to select among. So, interactive possibilities can branch or split into multiple possibilities.

The second is that such interactive possibilities might iterate. Consider the frog when there are no flies around. There will be no detections (absent devilish students of frog behavior and other unfortunate conditions) and therefore no indications of interactive possibility. But, nevertheless, if there were to be a fly, it would in all likelihood be detected and the appropriate indication set up. The point is that the conditional from detection to indication is present in the frog even when the conditional has not been satisfied. In sufficiently complex organisms, this yields the possibility that an initial interaction might yield or detect the condition or conditions for a second or further interaction possibilities. I may not be able to open the fridge for a beer in the woods directly, but I do know how to get home (itself possibly a rather complex endeavor) where I *can* open the fridge for a beer.

The realm of indicated interactive possibilities, then, can be potentially quite complex (in sufficiently complex organisms) with branching indications connecting to

² Note that what is represented here (in this model) is the potentiality for tongue flicking followed by eating, not (necessarily) the fly itself at all: detections or differentiations (e.g., of flies) do not in themselves constitute representations of that which has been detected or differentiated.

³ This normativity, in turn, derives from a normativity of biological function. I will not develop that part of the model here (see Bickhard, 1993, 2004, in preparation; Christensen & Bickhard, 2002).

the conditionals for still further possibilities, which might yield still more, and so on. Such a web of interactive possibilities can cycle onto itself and extend as far as the capacities and knowledge of the organism permit. I call this organization of indicated, conditional, interactive possibilities the organism's *situation knowledge*, its knowledge of what it could do in the current situation, much of it conditional on various other mediating interactions.

Object Representations?

Consider now a child's wooden toy block. The block offers multiple possible interactions, including various visual scans, manipulations, dropping, chewing, throwing, and so on. Furthermore, any one of these possibilities is reachable from any other, perhaps via appropriate mediating interactions. Any visual scan can be reached, for example, assuming the block is being held, by appropriate manipulations to bring that aspect of the block into view.

The interactive potentialities of the block concerning perceptual interactions and manipulations, then, are all mutually reachable. Furthermore, this entire organization of mutually reachable potential interactions is itself invariant under a major class of other interactions, including manipulations, dropping on the floor and going into another room, putting in the toy box, and so on. It is not, however, invariant under, for example, crushing or burning.

This is "just" Piaget's model of the representation of small objects in terms of interactions translated into the language of this model (Piaget, 1954). I can steal such models because both Piaget's model and interactivism are interaction based rather than passive input processing based as for most models of representation. I similarly borrow other aspects of Piaget's model, though in other cases I argue that either more modification is required or that he got it wrong (Bickhard & Campbell, 1989).

I will let this brief outline, then, suffice to indicate that the interactive model is at least a candidate for capturing the nature of representation, thus intentionality, and turn now to some fundamental criticisms of standard approaches.

Encodingist Models of Representation

I generically call standard models of representation encodingist models because they all assume that representation is constituted as some form of encoding — as some special kind of correspondence between the brain or mind and (something in) the world that somehow captures for the brain or mind what that correspondence is with. Real encodings do exist, and are in fact quite important, especially in our age. Encodings can be illustrated with Morse code, in which, for example, “...” encodes “s”. Morse code is useful because dots and dashes can be sent over telegraph wires, while characters cannot. Similarly, binary codes are ubiquitous in contemporary communications and computers.

But such an encoding relationship only exists insofar as one or more people know about it — unless they represent the dots and dashes and the characters and the relationships between them. All of the crucial representations have to already be in place in order for the encoding to exist. Encodings are representational stand-ins, not emergent representations themselves. They are useful as stand-ins, but not for modeling basic representation. Nevertheless, standard assumptions are that *all* representation is encoding, and this is what I call *encodingism*.

Lest this be attributed to the conventionality of Morse code, consider a neutrino count from deep in a mine that encodes properties of fusion processes in the sun. Here the relationship is completely natural, but it is still the case that no encoding relationship exists unless physicists know about all of the relevant processes, in the mine and the sun, and the relationships among them — unless they already represent them.⁴

Encodings are stand-ins for representation, thus presuppose representation, and, therefore, attempts to model representation strictly within an encodingist framework is committed to (multifarious) forms of circularity: for such purposes, encodings presuppose what they are supposed to be modeling (Bickhard, 1980, 1993, 2004, in preparation; Bickhard & Richie, 1983; Bickhard & Terveen, 1995).

⁴ Note that an *informational* relationship, in the sense of a factual covariance, might exist even though no one knows about it, and, therefore, there is no encoding relationship (see, e.g., Fodor & Pylyshyn, 1981; Fodor, 1987).

Critiques of Encodingism

The special correspondences that are supposed to constitute encodings are almost always factual correspondences⁵. The models constitute attempts to capture the normativities of representation and intentionality in factual, causal, terms. This fails, and yields a realm of powerful criticisms of the approaches. There are many other criticisms, some equally central, but I will focus on two normativity criticisms for current purposes: 1) the possibility of representational error, and 2) a strengthened version of this, the possibility of *system detectable* representational error.

The argument concerning the possibility of representational error can be introduced straightforwardly: if the favored correspondence relationship exists, then the representation exists, and it is correct (if the object did not exist, then the relationship would not exist), while if the favored relationship does not exist, then the representation does not exist. But the relationship either exists or not — there is no third possibility. Nevertheless, there is a third condition that must be modeled: the relationship exists, but it is incorrect (Millikan, 1984, 1993)⁶. The kind of model involved does not have the resources to address all three conditions. This is hardly surprising since two of the conditions are normatively differentiated, and the model attempts to capture all three of them in factual terms.

This problem has certainly been noticed, and addressed in differing ways by multiple authors. I argue elsewhere that none of these attempts succeed (Bickhard, 1993, 2003, 2004, in preparation; Bickhard & Terveen, 1995), but will not pursue the details of those arguments here because the second criticism mentioned above, that of accounting for system detectable representational error, suffices to reject *all* of these models. The attempts to account for the possibility of representational error ubiquitously assume an external perspective on the organism and its environment, and attempt to characterize the

⁵ Causal, informational, and so on. For more extensive criticisms of , for example, Fodor (1987b, 1990, 1990b, 1998, 2003), Dretske (1988), Millikan (1984, 1993), and Cummins (1996), see Bickhard (Bickhard, 1993, 2003, 2004, in preparation; Bickhard & Terveen, 1995).

⁶ Millikan's model is not subject directly to this problem, but does encounter other problems. In different senses, both Dretske's and Cummins' models could be argued to escape this problem as well, but, again, they fall to other considerations (Bickhard, 2003, 2004, in preparation). The argument above applies to models in which the content of the representation is determined by that which is to be represented — this category includes, for example, information semantic approaches.

conditions under which it would be appropriate to assert that the organism is in error about its environment. None of them even *attempt* to account for the organism itself being able to detect its own errors.

This might seem an inconsequential problem, best set aside at least for now, until it is realized that, without organism or agent detectable error, there is no error guided behavior or learning. We know that error guided behavior occurs (even in very simple animals) and that learning occurs, therefore we know that any model that cannot account for this is already refuted.

Recognition of the importance of this problem is strengthened by recognition of the fact that it is a variation on the radical skeptical argument: we cannot determine the truth or falsity of our own representations because to do so would require that we step outside of ourselves (take up the external perspective mentioned above) and compare our representations with what is being represented to see if they match or not. This argument has survived centuries of attempts to solve or dissolve it, so it is perhaps not surprising that standard models do not even address the problem.

Note, however, that an organism that has indications of future interactive potentialities can check the truth of those indications by engaging in the relevant interaction. The interaction proceeds as indicated, or not, and this is determinable strictly *internal* to the organism, and, thus, does not require an external perspective. In this way, the organism can detect, however fallibly, its own errors, and guide its behavior and induce learning on the basis of those detections.

The interactive model, then, as a future oriented pragmatic model rather than a backwards facing “spectator” model (Tiles, 1990), solves problems that encodingist models cannot even address. It is a much more satisfactory model for the nature of representation and intentionality.

So, if there is good reason to move from a “behavioral regularity” based model of social ontology to an intentionality based model, there is also good reason to do so from within an interactivist framework rather than an encodingist framework. I will argue, in fact, that this resolves several central issues.

The Power of Implicitness

The crucial contrast between interactivist and encodingist approaches for current purposes (aside from the fact that encodingist models are at root circular, cannot handle error, and so on) is that encoding representation is necessarily explicit, while interactivist representation involves a rich and powerful realm of implicitness. Encodings are constituted by some element or event that carries a content, and, absent that content, there is no representation at all. The content has to be attached to whatever is to carry it, and this necessity for some sort of “attachment” is the core of the sense in which encodings are stand-ins: the content to be attached must be provided from elsewhere, so the encoding ends up being a stand-in for whatever provided the attached content. If something that is attached is *not* content, then again there is no representation. So, the content of an encoding must be explicit — without it, there is no encoding.

Interactive representation is much richer than this. First, the content, that which is predicated of the environment, is most proximately that this environment is one that is appropriate for the indicated interaction. Some environments will be appropriate and some will not. Crucially, the properties that constitute an appropriate environment are not explicit, but *implicit*, in the indication that that interaction is a possibility. Those properties or conditions are dynamically or (normatively) functionally *presupposed* in the indications of interactive potentiality.

Those contents also do not have to be attached. They are *internally* related to the interactive indications in the logical sense of internal and external relations. An internal relation is one that is essential to one or more of the relata. An arc of a circle, for example, cannot be that arc of that circle without having a specific relationship with the point that is the center of that circle — that relationship is essential, internal, to the nature of being that arc. Similarly, an indication of a particular interaction potentiality cannot exist without presupposing those conditions that would support that potentiality (Bickhard, 2003, in preparation). The content is internally related. In contrast, cause, information, and so on are all external relations. They are not internal to the nature of the “carrier” or “vehicle” of the representation, and, therefore, must be attached to that carrier or vehicle, thus creating a stand-in, not an emergent representation.

What is explicit for an interactive representation is the indicated interaction or organization of indicated interactions. What is implicitly presupposed is a realm of supports for those indications.

Frame Problems

Note that, in being implicit, the presuppositions of an interactive representation are *unbounded* relative to any attempt to capture them explicitly. Even if in fact there is a small finite number of such conditions, their being implicit means that no attempt to explicitly exhaust them can be assured of being complete no matter how many have been explicitly set out. Even if all have in fact been captured, the implicitness prevents that from being detectable or determinable in terms of the explicit list *per se*.⁷ Therefore, any attempt to capture a realm of implicit presuppositions in terms of explicit encodings will inevitably encounter an unbounded expansion of possibilities. Forms of this unboundedness have been designated as the frame problems (Bickhard & Terveen, 1995; Bickhard, 2001). These play an important but unfortunate role in the attempts to capture social ontologies.

Note in particular that attempts to model social ontologies in terms of propositional attitudes are inherently committed to encoding models of the nature of the propositions, and, therefore, to the explicitness (and external relations) of the contents of those propositions. This, plus one further property, is what generates the regresses of beliefs about beliefs, and so on, within these approaches.

Situation Conventions

I am now ready to introduce a model of social ontology based on this interactionist framework. I will show that it provides approaches to multiple forms of social reality, that it intrinsically captures multiple forms of normativity, and that it does so without introducing unbounded regresses of (meta-)propositional attitudes.

⁷ A much more developed discussion of these issues can be found in (Bickhard & Terveen, 1995; Bickhard, 2001).

The Inter-Agent Coordination Problem

When an animal encounters a rock, the interactive potentialities that follow are largely determinable from the perceptual characteristics of the rock. A special kind of problem emerges, however, when two or more complex agents must deal with each other.

One agent attempts to characterize the interactive potentialities of the situation, but the situation centrally includes the other agent. Unlike the rock, the interactive potentialities of the other agent are *not* largely determinable just from their perceptual characteristics. In particular, the interactive characteristics of the other agent depend strongly on that agent's interactive characterization of his or her situation. So, the first agent's interactive characterization of the situation depends on the other agent's characterization of the situation. But the other agent's characterization, in turn, depends on how they characterize the interactive potentialities of the first agent. Thus, each agent's interactive characterization of the situation depends on the other agent's characterization, generating an unbounded regress of potential characterizations.

Each agent has an interest in arriving at a reasonably accurate characterization of the situation, but doing so cannot be done independently. There is a coordination problem here in that each agent would like to arrive at something like a fixed point with the other agents in their mutual characterizations of their mutual situation. They would like to arrive at joint interactive characterizations that are interactively compatible with each other.

If any such solution to the coordination problem of interactive characterization should occur, it would constitute a convention about the nature of the situation in the general sense of constituting a solution to a coordination problem. It would be a convention about the nature of the (social) situation, a *situation convention*.

I am proposing such situation conventions as the basic ontology for social realities. In support of this claim, I will develop several properties of situation conventions, particularly with regard to: 1) how they can account for institutional social realities, 2) how they can account for various kinds of normativity, and 3) an outline of how they relate to language.

Institutionalized Conventions

My concern here is to show how this model of situation conventions can model more complex social ontologies such as institutions and roles. This will not be a complete model, but rather an indication of some of the resources available for addressing such issues — resources sufficient to make situation conventions a plausible candidate for modeling such social ontologies.⁸

Insofar as individuals have an interest in arriving at situation convention solutions to mutual agent coordination problems, they have an interest in ways to *invoke* such conventions. An invocation of a convention will not in general have an internal relation to the convention it invokes, so the power of some gesture or other interaction or condition to invoke a situation convention will itself constitute a convention — it will depend on the commonalities of interpretation of the gestural interaction or condition as inducing or invoking that convention.

One way for conventions to become established is via precedent and habituation. This was one of Lewis' fundamental points: he was concerned to show that language could be conventional even though there could not be a language in which the conventions of language were established. Precedent and habituation is an alternative. If two people accidentally meet, for example, at a particular restaurant on Tuesday for lunch, and each have an enjoyable time, they might meet again the next Tuesday, each at first wondering if the other might show up. The same might occur again on the third Tuesday, and perhaps the fourth. But soon each will develop the expectation that the other will be there, and a convention between the two to meet on Tuesdays at that restaurant for lunch will have been established, without any necessary explicit discussion. Thus could, in principle, the conventions of language have emerged.⁹

⁸ A more detailed discussion can be found in Bickhard (in preparation).

⁹ Gilbert (1989) develops detailed and strong arguments against Lewis' invocation of rationality in a game theoretic framework as grounds for such precedent and habituation origins of conventions, but her arguments turn on, among other things, an equating of rationality with strict deduction, a refusal to consider psychological assumptions as premises in such deductions, and other austere and desiccated notions. I will not examine the issues involved in these arguments as against Lewis' model, but simply point out that none of them even address more reasonable considerations of rational warrant that is not tied to strict deduction with maximum skepticism about premises. In general, invocation of a convention is necessarily presumptive, and there can be no guarantee that a presumptive gestural interaction will be understood or accepted. Establishing conventions is fallibilistic, and, even when apparently successful, there may be

Similarly, precedent and habituation is one way in which conventional means for *invoking* conventions might have emerged. A number of important properties follow from this point, but here I wish to focus on the implication that such conventions for invoking conventions will necessarily be repeatable across the relevant population in which the precedents and habituations have occurred and the situations in which the invocations might be appropriate. Such invocations might involve explicit interactions on the part of one or more people, or they might be conditional on properties of situations that arise in other ways. When encountering an oncoming driver, for example, it is conventional to pass on the right (or left). The condition of meeting an oncoming driver, thus, yields a conventionalized invocation of the convention of driving to the right. In appropriate situations, on the other hand, banging a gavel is a conventionalized way of invoking a meeting convention.

Conventionalization of means or conditions for invoking conventions requires some form of *typification* of the means or conditions. That is, a *kind* must be recognized as being the right kind for that invocation. There is nothing mysterious in this — all interactive representation is via differentiations into kinds¹⁰ — but it has interesting consequences.

First, it is this dependence on a multiply instantiable kind that makes conventional invocations of conventions *iterable* — they can be invoked (usually) some indeterminate number of times across (usually) multiple people. Furthermore, the conventions invoked, under such conditions, will themselves be of some general kind.

Conventionalized kinds or forms of interacting can come to be recognized as well as participated in — and recognized, represented, even from a third party perspective, when the recognizer is not a part to the conventional interaction. That is, they can become abstracted from particular occasions and participants: a cultural form of uncle-nephew

some failure of detail that may or may not be encountered in later interactions within the framework of that supposed convention. Maintenance of convention is an ongoing process, with invocation, maintenance, and repair always either involved or potential (Bickhard, in preparation). This is completely contrary to the kind of assumptions about rationality made by Gilbert's criticisms. So, even if Lewis arguably did make these assumptions that Gilbert criticizes, the cogency of those criticisms should have little weight against a more natural conception of rational warrant for and defeasible presumptiveness of conventional interacting.

¹⁰ This raises interesting questions about how reference is accomplished, questions that I will not pursue here (see, e.g., Bickhard, 1980, 1987, 2003, in preparation; Bickhard & Campbell, 1992; Bickhard & Terveen, 1995).

interaction might be recognizable just by watching an instance, or, similarly, a cultural form of clerk-customer interaction.

When both the form of the conventional *interactions* and the differentiated forms of *participation* in those forms of interaction are abstracted and typified, there arises institutional forms constituted out of institutionalized *roles*. At this point we have institutionalized conventions, institutionalized forms, and institutionalized roles, arguably a sufficient range of resources for addressing complex societies as well as single relationships (Berger & Luckmann, 1966; Bickhard, 1980, in preparation).

I will not pursue here the elaboration of those models of various kinds of social ontologies, however, but turn to the issue of how such a framework can account for the normativities of social realities.

Normativities

There are many kinds and levels of normativity involved in situation conventions, including institutionalized conventions, which in most instances are implicit, presupposed, rather than explicit — unless they are made explicit, usually by some kind of failure or violation.

This model still captures the instrumental sense of normativity involved in social conventional interactions, but there is much more as well. First, engaging in conventional invocation, or accepting a conventional invocation, engages issues of the appropriateness of that person to be a participant in that convention. In some cases, such issues of appropriateness will themselves be issues of convention: do I in fact hold an institutional role that makes it appropriate for me to perform a marriage? Putting myself forward for such a role presupposes that I am appropriate, and, in that sense, makes the claim (however implicit) that I am appropriate.

More commonly, such issues of appropriateness have to do with issues about a potential participant that are closer to that person as person than to that person as holder of some institutionalized role. In engaging as a participant in some convention, I am putting myself forward as having other characteristics that constitute such appropriateness, such as the skills and reliability to be able to carry out my participation,

as well as the integrity to make the genuine attempt to fulfill the participation. I can and usually will be faulted for failures of any of these kinds. Defenses can be of several sorts: I might claim that I was not aware that I had indicated participation in some form of convention. Adequate support for such claims of ignorance *might* be accepted. As a variation, I might claim that you or I or both of us misunderstood various aspects of the supposed convention. I might attempt to shift characterization of my failure from an integrity failure to a reliability failure. I might attempt to reestablish my identity as a competent, reliable, and/or honestly intentioned social participant via some sort of penance or punishment.

In general, such failures and repairs, both of the conventional instance per se and of the social identities of those involved in that instance, is an ongoing process. Conventions, and trajectories of transformations of conventions, are sustained and repaired, not just invoked.

The central point here is that an agent participating in a convention is also a *potential* social agent more generally, and cannot be participating without actually being such a social agent. So, putting forward the claim (implicit or explicit) that I am a potential participant in some instance of a convention presupposes (usually implicitly) that I am a potential social agent more generally, including the normative issues of (at least) competence, reliability, and integrity.

To participate in social ontologies is to be a social person, and to be a social person is an intrinsically normative process. To put oneself forward as a legitimate, appropriate, competent, reliable, honest, etc. social participant is inherently normative in all of these, and additional, ways — but most of them are implicit in presuming to be a social participant at all.

Attempts to capture such implicit commitments and characterizations of social situations, therefore, will necessarily encounter their own version of the unboundedness of the frame problems: implicitness yields unboundedness relative to explicitness, and propositional attitudes commit to encoded propositions, and, thus, to explicit encodings. Finally, there is an inherent reflexiveness in mutual-agent coordination problems, and, therefore, in their solutions: some sort of resolution of my characterization of your

characterization of my characterization of.....of the situation is required. But, again, this hierarchy of potential reflections is implicit, not explicit (except for occasional special ascents up a few such levels, usually for purposes of deceit, such as for spies or con men). So, attempts to capture the structure of social ontologies in terms of explicit encoded propositional attitudes will not only encounter unboundedness, but unboundedness in this particular meta-level hierarchy form. It is no wonder that this literature is replete with such hierarchies.

Collectivities

There is another aspect of the implicit meta-level hierarchy of mutual agent coordination problems and their solutions, thus of situation conventions, that is worth mentioning. The problem for an individual agent in interactively characterizing his or her situation when one or more additional agents is present *is inherently a mutual, a social, problem*. It cannot be solved without reaching some sort of fixed point in the social relations among the individual characterizations.

So, the reflexiveness of mutual agent situations is intrinsic in the nature of the situation and does not require a prior *we* or *us* construal. Furthermore, the reflexiveness is in the nature of the problem, not just in the solution, but is necessarily reflected in the solution. Again, if this reflexiveness can only be captured explicitly, then we have unbounded “frame problem” type reflexivenesses attempting to capture the implicitness of the situation conventions. With respect to this issue, this becomes explicit propositional attitudes about some sort of inherently social *us** or *we** (e.g., Gilbert, 1989; Searle, 1983, 1995). This “socialness” is real, or so the interactive model would have it, but it is, again, initially and usually, implicit, with no explicit representations or indications or desires concerning an intrinsically social collectivity. Social situations are emergent, and so also are the representations appropriate to them, but they need not be explicit.

There is also a converse to these considerations: *any* social reality is emergent among the participants, and, therefore, *must* emerge out of the basic mutual agent coordination problem. Situation conventions, therefore, are not only sufficient, but also necessary to account for social ontologies.

Language

Lewis' original motivation in addressing convention was strongly based on issues concerning language. Independently of Lewis' motivations, social reality is deeply intertwined with language. Can the notion of situation convention cast any light on language and the relationships between language and social ontology? I will briefly outline some considerations that suggest that the answer is yes, and that filling out that aspect of social reality yields some interesting consequences.

Consider first the critique of encodingism: If encodings must be stand-ins, then neither perception nor language can be fundamentally processes of encoding. In order for an encoding of something perceived, X perhaps, to be created, for example, some other representation of X would have to be available in order to provide the content that the encoding element (vehicle) can be attached to (Bickhard, 1993, 2004, in preparation; Bickhard & Richie, 1983). Similarly, in order for language to be constituted as encodings of mental contents, some other representation of those contents would have to already be available to an audience for that language in order that the encodings can be created via attachment of those contents (Bickhard, 1980; Bickhard & Campbell, 1992). Among other consequences, neither perception nor language can have the standardly assumed information flow and information processing form (Bickhard, 1980, 1993, 2004, in preparation; Bickhard & Richie, 1983; Bickhard & Terveen, 1995).¹¹

As an introduction to an alternative approach to the nature of language, consider first the conventional invocations of situation conventions that are discussed above. Some of these will be context sensitive in the sense that they will be appropriate only in certain contexts; some will be context sensitive in the sense that they will have differing evocational powers and consequences *depending on* the context. Of special relevance here is the possibility of being sensitive to the context of histories of prior evocations: equivalence classes of possible such histories constitute possible current states of the social situation. Those current states, in turn, constitute situation conventions concerning the coordination problem of how to understand or interpret the next evocations in the history. Engaging in appropriate histories of such evocational forms, then, can create the

¹¹ However much information in the strictly factual sense of covariation might exist, it does not and can not in itself constitute representation (Bickhard, 1993, 2004, in preparation).

context within which further evocational forms will have certain kinds of consequences in transforming the situation convention. That is, we can have the induction of various contexts within which further context sensitive inductions and transformations can occur. If this process of composing strings of situation convention transformational kinds is itself productive, in the sense of potentially generating an unbounded set of such strings, then we have a language, but a language in the sense of a productive conventionalized tool for manipulating and transforming situation conventions, not in the sense of transmitting encoded mental contents.¹²

Language, in this model, is a conventionalized means of productively constructing conventionalized transformations of situation conventions. Indirectly this involves transformations of people's characterizations — representations — of the situation, but, again, not via encodings and decodings.

Even more generally, social realities are constituted as commonalities of the participants' anticipations of what could and would be appropriate to ensue given the current situation. Much of what could thus ensue is constituted in further potential conversation or other form of language. Language, thus, is not only a tool box for transforming social realities, its possibilities are centrally (though not exclusively) constitutive of those realities.

In one more step, we find that language is not only a means for operating on conventions, it is itself a convention (for the productive generation of conventionalized transformations). Therefore, language can be used to discuss, analyze, and transform itself: meta-language.

I will end this indication of an interactive model of language with two further observations: 1) Human beings can keep up with the complex temporal trajectories of situation conventions in conversation, even though those situation conventions are not directly recoverable perceptually — they are hidden. This is an ability that has co-evolved with language, which creates/induces such trajectories, and is shared with no other animal species (there is some ability to track hidden trajectories, such as for hidden

¹² Note that not only is language conventional in this model, but the *object* of language as a form of interaction is itself convention.

movements of objects, in a few other species, but nothing close to the level of keeping track of a conversation). And 2) the trajectories of situation conventions that flow in a conversation are situation conventions that need not and likely will not ever recur. They are one instance, not repeated nor repeatable, social ontologies. At some point in a conversation, for example, everyone involved may know precisely how to resolve the reference for a pronoun, given the previous discussion, but that specific framework for understanding the rest of the sentence and sentences need not ever have occurred before nor occur again. Such situation conventions as the objects on which language “operates” could not possibly be modeled by Lewis’ behavioral regularities, because they do not involve any such regularities — but they do constitute a solution to a mutual agent coordination problem.¹³

Language, thus, participates in social ontologies in multiple deep senses. It is itself a social ontology; it functions to create, transform, and maintain social ontologies; and its potentialities constitute a primary constituent of social ontologies more generally. This is all in strong contrast to the mental encodings model, in which any involvement in social ontologies is at best peripheral.

Conclusions

I have defended Lewis’ basic notion of convention as solution to a coordination problem, though in importantly changed form. Among other changes, accounting for normativity requires moving to a realm of intentionality, as is common in the contemporary literature. But if intentionality is understood in terms of propositional attitudes, then there is a commitment to encodingism, and the necessary explicitness of encodings generates unrealistic hierarchies of those attitudes. If, on the other hand, intentionality is modeled in interactive terms, then there is an inherent domain of implicitness, of presuppositions in particular, that resolves these reflexivities without recourse to any explicit propositional attitude hierarchies.

On a deeper level, the recognition that any mutual agent situation is inherently a coordination problem situation, and, furthermore, one that is intrinsically social and

¹³ This requirement for non-repeated socially common understandings of a situation was one of the motivations for the original development of the situation convention model (Bickhard, 1980).

reflexively so, introduces issues of emergent socialities, with complex implicit structure, without having to introduce them either explicitly or ad-hoc. Socialness and social reflexivity are inherent in the nature of mutual agent situations. Social ontologies emerge as simple or complex solutions to those various and historically growing coordination problems.

References

- Berger, P. L., Luckmann, T. (1966). *The Social Construction of Reality*. New York: Doubleday.
- Bickhard, M. H. (1980). *Cognition, Convention, and Communication*. New York: Praeger Publishers.
- Bickhard, M. H. (1987). The Social Nature of the Functional Nature of Language. In Maya Hickmann (Ed.) (39-65). *Social and Functional Approaches to Language and Thought*. Academic.
- Bickhard, M. H. (1993). Representational Content in Humans and Machines. *Journal of Experimental and Theoretical Artificial Intelligence*, 5, 285-333.
- Bickhard, M. H. (2000). Motivation and Emotion: An Interactive Process Model. In R. D. Ellis, N. Newton (Eds.) *The Caldron of Consciousness*. (161-178). J. Benjamins.
- Bickhard, M. H. (2001). Why Children Don't Have to Solve the Frame Problems: Cognitive Representations are not Encodings. *Developmental Review*, 21, 224-262.
- Bickhard, M. H. (2003). Some notes on internal and external relations and representation. *Consciousness & Emotion*, 4(1), 101-110.
- Bickhard, M. H. (2004). Process and Emergence: Normative Function and Representation. *Axiomathes — An International Journal in Ontology and Cognitive Systems*, 14, 135-169. Reprinted from: Bickhard, M. H. (2003). Process and Emergence: Normative Function and Representation. In: J. Seibt

- (Ed.) *Process Theories: Crossdisciplinary Studies in Dynamic Categories*. (121-155). Dordrecht: Kluwer Academic.
- Bickhard, M. H. (in preparation). *The Whole Person: Toward a Naturalism of Persons — Contributions to an Ontological Psychology*.
- Bickhard, M. H., Campbell, R. L. (1989). Interactivism and Genetic Epistemology. *Archives de Psychologie*, 57(221), 99-121.
- Bickhard, M. H., Campbell, R. L. (1992). Some Foundational Questions Concerning Language Studies: With a Focus on Categorical Grammars and Model Theoretic Possible Worlds Semantics. *Journal of Pragmatics*, 17(5/6), 401-433.
- Bickhard, M. H., Richie, D. M. (1983). *On the Nature of Representation: A Case Study of James Gibson's Theory of Perception*. New York: Praeger Publishers.
- Bickhard, M. H., Terveen, L. (1995). *Foundational Issues in Artificial Intelligence and Cognitive Science: Impasse and Solution*. Elsevier Scientific.
- Christensen, W. D., Bickhard, M. H. (2002). The Process Dynamics of Normative Function. *Monist*, 85(1), 3-28.
- Cummins, R. (1996). *Representations, Targets, and Attitudes*. Cambridge, MA: MIT Press.
- Dretske, F. I. (1988). *Explaining Behavior*. Cambridge, MA: MIT Press.
- Fodor, J. A. (1987). A Situated Grandmother? *Mind and Language*, 2, 64-81.
- Fodor, J. A. (1987b). *Psychosemantics*. Cambridge, MA: MIT Press.
- Fodor, J. A. (1990). *Theory of Mental Content and Other Essays*. Cambridge, MA: MIT Press.
- Fodor, J. A. (1990b). Information and Representation. In P. P. Hanson (Ed.) *Information, Language, and Cognition*. (175-190). Vancouver: University of British Columbia Press.
- Fodor, J. A. (1998). *Concepts: Where Cognitive Science went wrong*. Oxford.
- Fodor, J. A. (2003). *Hume Variations*. Oxford: Oxford University Press.

- Fodor, J. A., & Pylyshyn, Z. (1981). How direct is visual perception?: Some reflections on Gibson's ecological approach. *Cognition*, 9, 139-196.
- Gilbert, M. (1989). *On Social Facts*. Princeton, NJ: Princeton University Press.
- Lewis, D. K. (1969). *Convention*. Cambridge, MA: Harvard University Press.
- Millikan, R. G. (1984). *Language, Thought, and Other Biological Categories*. Cambridge, MA: MIT Press.
- Millikan, R. G. (1993). *White Queen Psychology and Other Essays for Alice*. Cambridge, MA: MIT Press.
- Piaget, J. (1954). *The Construction of Reality in the Child*. New York: Basic.
- Schelling, T. C. (1963). *The strategy of conflict*. New York: Oxford University Press.
- Schiffer, S. R. (1972). *Meaning*. London: Oxford University Press.
- Schmitt, F. F. (2003). *Socializing Metaphysics*. Rowman, Littlefield.
- Searle, J. R. (1983). *Intentionality*. Cambridge: Cambridge University Press.
- Searle, J. R. (1995). *The Construction of Social Reality*. Free Press.
- Tiles, J. E. (1990). *Dewey*. Routledge.