Playing the Game of Language: An Introduction to Contemporary Pragmatism

Ryan Simonelli

Originally *posted on my personal blog*, *Absolute Irony*, on October 28, 2014

Intuition Turned on Its Head

Perhaps the most peculiar trick that we've developed as a species is the ability to make noises with our mouths that somehow end up being *about* things in our world. We can talk to others about the things that we stumble upon in the world. We can express our views on what those things are like. We can agree, or disagree, or agree to disagree. When you think about it for a moment, this is a very befuddling ability of ours. How should we go about trying to explain how it works? Let's start with a suggestive thought experiment:

Imagine you're a toddler in sitting in a crib. You constantly see a bunch of adults moving around you, making all sorts of strange noises with their mouths. Eventually, seeing that they often make these noises while looking at or gesturing towards various objects, you realize that these noises *signify* different things around you. "Hat," for example, signifies to those things that the adults put on their heads. Gradually, you learn how to make these noises with your mouth as well, giving you the ability to talk about things just like the adults. Perhaps you want the adults to give you the hat, and so you say, "Hat!" in the hopes someone will give it to you. Slowly but surely, you learn how to string these noises together to *fully* express your thoughts and desires, and, like that, you come learn language.

This picture of language implies that words mean what they do in virtue of the things they represent in the world. We string them together to explicitly express our experiences and thoughts, our feelings and desires, and so on. Clearly, this is the most intuitive way to think about language, and, until the middle of the 20th century, most philosophers thought that it worked this way. That was until a philosopher named Ludwig Wittgenstein published a book called *Philosophical Investigations,* now widely considered the most important philosophical work of the 20th century.

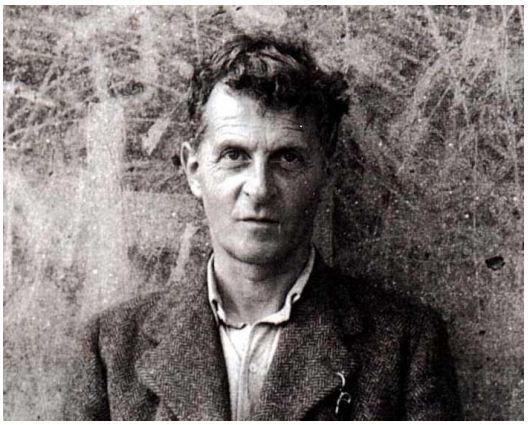

Wittgenstein famously started the *Investigations* by laying out this intuitive picture that most of us take for granted. He then goes on to argue that it's completely, utterly wrong. If we want a fundamental model for thinking about language, he argues, we shouldn't think of it as a means for explicitly *representing* our thoughts or experiences, our beliefs or desires, or *anything like that*. It is really just a *practice,* something we *do* that's woven seamlessly into our other activities.

The striking thing about Wittgenstein's suggestion is that, on his view, only once we're thoroughly immersed in this practice are we able to possess the sort of understanding that will allow us to know *what we're doing* by speaking a language. If we don't understand what we're doing when we begin to speak a language, the intuitive way I've just described a toddler's learning of language is impossible.

On Wittgenstein's way of thinking about things, one is *blindly conditioned* into the social practices of language, and it is only by way of competence in these practices that one can be said to have any concepts or understanding at all. Here's one passage from the *Investigations* that tries to illuminate this idea:

> *"A child has hurt himself and he cries; then the adults talk to him and teach him exclamations and, later, sentences; They teach the child new pain-behavior.*
>
> *'So are you saying that the word 'pain' really means crying?'—On the contrary: the verbal expression of pain replaces crying, it does not describe it."*

When reading this, it's easy to get tempted into thinking that crying is just a way of expressing the thought that one is in pain, and, when we learn language, we're just learning a more efficient way of expressing this inner emotion. Wittgenstein, however, is saying the exact opposite of this. Crying is just *what we do* when we're in pain, and *saying* we're in pain is, at root, the same sort of thing. If we think of language in this way, the notion of representation entirely drops out as a fundamental explanatory feature. In the first and most central instance, there is no thought or experience that these verbal expressions attempt to *capture*; one just *does them.*

The issue of whether or not we ought to think of language in this broadly Wittgensteinian fashion is one of the great dividing lines in philosophy. One way of drawing this division is to say that *representationalism* stands on one side, and *pragmatism* stands on the other. We can think of this distinction between representationalism and pragmatism in terms of two opposite *orders of explanation*. For the representationalist, we should explain our *use* of language in terms of how we employ various words and sentences to *convey meaning*. For the pragmatist, on the other hand, we shouldn't explain the use of

language in terms of meaning. Quite the contrary, the *meaning* of words and sentences is to be understood in terms of how these things are *used* in language.

While Wittgenstein famously said that the meaning of a word can be explained in terms of its *use*, he did not actually provide a theory about *how that explanation should go*. Wittgenstein was not particularly keen on philosophical theorizing. In fact, he thought there should be no philosophical theories at all! Many contemporary pragmatists, however, don't share this sentiment and believe that we can offer a systematic theory of meaning in terms of use. In the rest of this article, I'll give an outline of what this sort of theory looks like and why we should favor it over a representationalist one.

Speech Acts, Social Practices, and Pragmatic Force

According to Wittgenstein, speaking a language is more like playing a game than painting a picture. To get a clearer conception of this, Wittgenstein has us consider a society of builders who have a primitive language that they use when they're building. One builder will utter "Slab!" and, upon hearing it, the other builder will hand him a slab. The builder might then utter, "Block!" and, upon hearing this utterance, the other builder will hand him a block. Looking at this language, we can see that the function of these utterances is not to conjure up some picture in the heads of the ones to whom it's spoken. While it's *possible* that the utterance "Slab!" might elicit the mental image of a slab, it's clearly unnecessary in order for the utterance to do what it's supposed to do. Insofar as the utterance *pushes the hearer to get a slab*, so to speak, it has accomplished its goal.

The ability an utterance might have to "push its hearers" in various ways is called the "pragmatic force" of an utterance. Insofar as we think of utterances primarily in terms of their pragmatic force, it is apt to think of them as *actions* aimed to perform a particular function. Following Wittgenstein and thinking about language in this way, a field in philosophy of language emerged called *speech act theory*. Perhaps the most famous speech act theorist, J.L. Austin, gave some examples of how we should think of the main function of various familiar utterances in terms of their pragmatic force. Suppose you're getting married and the judge says, "I now pronounce you man and wife." Here, she's not simply making a statement about what she's pronouncing; she's *doing something*—namely, legally binding you together in marriage. The function of this is utterance should not to be thought of in terms of what her words *mean*, but, rather, what they *do*.

This notion of "pragmatic force" may sound somewhat magical, as if it's straight from Star Wars, but it's not really mysterious at all. All you need is a bunch of people acting together and censuring each other's behaviors such that they conform to certain patterns of actions. In the words of John Haugeland, you just need a bunch of *conformists*. Conformists *imitate* the behaviors of each other, such that patterns of behavior emerge, and *censure* behavior that deviates from these patterns. In the context of this conforming behavior, *social practices* emerge, and, in the context of these social practices, certain actions

will result in certain sets of behaviors. Speech acts are one class of actions that can have that effect, and that effect is what I've been calling their "pragmatic force." So, the reason the judge's words have the pragmatic force that they do is because there are certain socially enforced patterns of behavior in which the act is contextualized.

Traditionally, the pragmatic force of an utterance has been seen as *supplementary* to its meaning. However, some pragmatists have boldly claimed that, rather than merely being supplementary, the pragmatic force of an utterance is *all there is* to the meaning of an utterance. Much like Wittgenstein's primitive builders, the only thing our utterances do is "push people around" in various ways; there is no meaning they have over and above that function. This suggestion will likely be met with incredulity from skeptical readers. While it may make sense to think of the primitive "language" of Wittgenstein's builders purely in terms of pragmatic force, it is easy to point out that their language is far from our own. In fact, it is probably a stretch to call it a language *at all*. While the utterance "Slab!" is going to causally initiate an action involving slabs, there is no sense in which it is a statement *about a slab* which may be *true or false*. *Our* sentences, on the other hand, *are* about various things in the world and they *can* be true or false . . . or, at least, so it seems.

One route that some pragmatists have taken is to deny that our sentences *really are* about things in the world, or *really can* be true or false. Richard Rorty famously took this route, arguing that the idea that our sentences or thoughts *mirror* the world, *representing* the way things really are, is simply mistaken. Rorty's suggestion, however, is pretty hard to swallow. It seems that, if I say "The cat is on the mat," it's about the cat and whether or not it's on the mat. If the cat really is on the mat, then my sentence has *said it like it is*, so to speak, and so it's true. If pragmatism must come at the price of denying this obvious intuition, then wouldn't we be all the wiser to simply deny pragmatism? It seems so. We'll have to do better than that if we want to explain language in pragmatist terms.

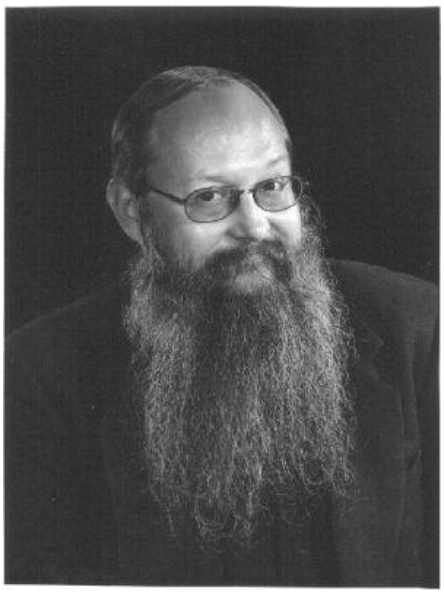

While Rorty wants to scrap the notion of representation altogether, Robert Brandom, one of Rorty's students, thinks we can have our cake and eat it too. Brandom thinks that we can account for the *representational* aspect of our language, the fact that we're able to speak *about* various things in the world, strictly in terms of the pragmatic properties of those speech acts. This is a bold and exciting prospect, and Brandom has spent considerable effort working out in detail how to do it. If successful, it's a genuinely revolutionary project for philosophy, dissolving longstanding problems about the nature of understanding. It takes Brandom 700 pages to lay out this theory in his book, *Making it*

Explicit, and even then there's unfinished business to attend to, so my short summary of his view here will be a rough outline at best, but here it goes:

Brandom's Game of Giving and Asking for Reasons

Let's return to the game metaphor for a bit, since it's one that Brandom finds particularly helpful. To start, let's take the example of chess. What do you need to do in order to play chess? Well, at the very least, you need to know the rules, play according them, and enforce play that accords with them (if you don't stop your opponent from moving a rook diagonally, you're no longer playing chess). We can talk about the state of game at a particular point in terms of what Brandom calls *deontic statuses*: what you *must* do in some circumstance and you're *allowed* to do in that circumstance. If your king's in check, for example, then you must make a move such that your king will no longer be in check, and there will be some set of moves that you're allowed to make that will accomplish this goal. The players of any game must be what Brandom calls "deontic scorekeepers," they *keep track of* and *enforce* these deontic statuses.

Our linguistic practices, Brandom thinks, can be articulated in much the same way. We are players in what Brandom calls the "game of giving and asking for reasons," the most fundamental game we can possibly play. This game can be centrally articulated in terms of two sorts of deontic statuses, *commitments* and *entitlements*. Commitments are what you're *obligated* to do according to the rules of the game, and entitlements are what you're *allowed* to do according to the rules of the game. In the same way that I alter the deontic statuses of the chess game when I move my rook to put your king in check, when I make a claim, I have altered the set of deontic statuses in a linguistic community.

Now, in chess, we understand all these deontic statuses *explicitly*. We're usually explicitly *told* the rules, and, when we correct behavior in others, we do so by telling them what the rules as well. However, when it comes to the linguistic game that Brandom describes, we must be playing before we're able to *say* that we're playing, since our ability to say things is the very thing this game is supposed to explain. So Brandom says that we start following the rules *implicitly*, without an explicit understanding of what they are. These implicit rules just are the norms of the social practices that emerged through conformist behavior that I described in the previously section (Well . . . Brandom is actually rather picky about how he wants to think about social norms, and he'd have a few issues with the way I've described them, but we can leave that aside for now). Within the context of these norms, speech acts are able to be socially significant.

Employing the terminology of Rebecca Kukla and Mark Lance (students of Haugeland and Brandom respectively), we can say that each speech act has a specific *output* and *input.* The output is what the speech act *does* to those to hear it (what I've been calling the pragmatic force), and the input is what *enables* us to perform it. Brandom claims that the output and the input of a speech act entirely accounts for the *meaning* of the speech act. What makes this claim interesting for Brandom is that he insists that these outputs and inputs be construed in *purely pragmatic* terms, in terms of the outputs and inputs of

other actions. This means, for example, the input for the assertion "There's a cat!" can't include actual *cats* as an irreducible component. Meaning is accounted for from entirely *within* the playing of the game, not from any things *external* to the game like cats.

Of course, Brandom does think that cats have *something* to do with the meaning of the sentence "There's a cat!" The cat is of course, *causally responsible* for my utterance of this sentence, and I'm disposed to make utterances of this sort only if cats are present. Furthermore, gameplayers are able to keep track of my reporting abilities such that they can infer from my *claiming* that something is a cat, *that it's a cat*. My own authority in making non-inferential claims like this is derived from others *granting* me such authority. On this account, even though what entitles me to the claim "There's a cat!" it has *everything to do with cats*, it's not the *cat itself* that entitles me to this claim. I am entitled to the assertion that there's a cat not because of the cat itself, but because I'm a member of a community that *recognizes me* as a reliable reporter of cats, and, accordingly grants me the authority to make such claims.

Why does Brandom work things out in this fashion? One of the main reasons comes from Wilfrid Sellars, a philosopher who hugely influenced Brandom and Rorty before him. Sellars argued that we can't count something as serving as a *reason* for a belief if we consider it in complete *independence* of all our other beliefs. This is because, in order for something to count as a reason for a belief, we must already possess an understanding of the *significance* that thing has in informing other beliefs. In other words, we must *understand* that thing, having a *concept* of it, of the role it plays in the game of giving and asking for reasons. Only once something has *already been conceptualized* within this game, can it serve as the basis for belief; unconceptualized things *themselves* cannot. Now of course, once we *already* have the concept of something, and we understand how it factors inferentially with other beliefs, we can count the thing itself as entitling us to a belief. But note that this isn't *really* the thing itself; it's the thing *as we understand it*. If we want to explain the basis of all of this understanding, as Brandom does, we can't rely on unconceptualized things themselves to do the heavy lifting.

These considerations motivate Brandom's theory of meaning. Rather than thinking about the meaning of a speech act in terms of what it *refers* to, Brandom thinks about the meaning of speech acts in *inferential* terms, in terms of how it factors into the premises and conclusions of the various inferences that we might make. That is, the significance of a move is to be understood entirely in terms of how it affects other moves in the game, without relying on anything external to the game. *Commitment* to some claims will commit me to others, *entitlement* to some claims will entitle me to others, and, most importantly, *commitment* to some claims will *preclude* entitlement to others. So, for example, if I'm *committed* to the claim, "That's a cat!" I *can't* also be entitled to the claim "That's a dog!" This latter notion pragmatically explains the sense in which the claims "That's a cat!" and "That's a dog!" are *incompatible*. Saying what claims a specific claim is incompatible with, Brandom thinks, allows us to account for its *semantic content*.

So far, I've mainly been talking about claims as a whole. But what about *words*, the things that *make up* those claims? Brandom wants to accommodate words within his inferential framework, by thinking of

how the inferences which *license* a claim and which that claim *licenses* will change if parts of that claims are *substituted out*. According to Brandom, two words have the same meaning if you can substitute them in for each other in a claim and preserve the same inferences. So, "cat" and "feline" have the same meaning, just in case claims including the word "cat" will have the same inferential significance if we substitute the word "cat" with the word "feline." However, "cat" and "dog" have different meanings. While, "That's a Manx," licenses an inference to "That's a cat," it does not license an inference to "That's a dog." Quite the contrary, it *precludes* an inference to "That's a dog."

As I said at the beginning, Brandom promises to pull out a notion of representation out of this pragmatist framework. To see how he thinks he can do this, consider a case in which you and I are both out in the park. An animal is off in the distance and you say "There's a dog! I think it's a German Shepard." I, however, have better eyesight than you and I conclude that it's a cat. In this case, I might laugh and say, "You believe a *cat* is a German Shepard!" Here now, I've acknowledged that both of our beliefs are about the same thing, however, the thing your belief is about is not the thing you think it is. We can account for this effect of my speech act in pragmatic terms by showing how I've substitute *my own* commitments into your speech act in order to say what I take *you* to be committed to. This ability to substitute different commitments into claims across different perspectives is what allows us to see our claims as being *about* things in the world.

Though I have left out several details, this is the basic way in which Brandom attempts to explain the *meaning* and *representational* aspect of speech acts purely in terms of their pragmatic properties. If this theory is successful, Brandom can have his pragmatist cake and eat it too.

Why Pragmatism?

Brandom's inferentialist framework is a way of explaining language in pragmatist terms, but why would we embrace pragmatism in the first place? I've hinted at some of the reasons throughout this article, but I haven't explicitly stated the upside of adopting a pragmatist theory of meaning and representation. So, now I'll take a moment to do so. I believe there are two main upsides to this sort of pragmatism: an *explanatory* one and an *epistemological* one.

Let's start with the explanatory upside of pragmatism. To explain this upside, we have to consider what the point of all this philosophical theorizing is in the first place. Here, I'll draw again from Wilfrid Sellars. The point of philosophy, Sellars says, is to "to understand how things in the broadest possible sense of the term hang together in the broadest possible sense of the term." He elaborates this claim by articulating to ways in which we might think about the world. The first is the *manifest image*. This is, roughly, the world as it pre-theoretically appears to us. The manifest image includes things like tables and chairs, stop signs (and the fact that we should stop at them), claims and arguments, and so on. On the other hand, there is the *scientific image*, the image that we come to know through rigorous scientific

investigation. It includes things like quarks and electrons, organic molecules, cells and circulatory systems, and so on. One of the fundamental goals of philosophy, then, in the quest of explaining how things "hang together in the broadest possible sense of the term," is to show how the manifest image and the scientific *connect*. Among the main challenges in connecting these two images is explaining how we have things like claims that are *about things in the world,* in a world that, as far as we know, is made up of things like atoms and molecules.

So, here's the explanatory challenge: How do we get things like claims that are *about things*, if such claims are ultimately just part of a world made up of atoms and molecules which *aren't about anything at all?* Pragmatism gives us a route to answering this question. The two main raw materials we need to account for representational language, on the pragmatist story I've offered, are *social practices* and *reliable dispositions to respond differentially to things in one's environment*. Presumably, we can give a *socio-biological* explanation of social conformity that is able to account for the former, and we can give a *neuro-biological* explanation for the latter. Both of these explanations will ultimately be rooted in *evolutionary* terms. The great thing about explanations that appeal to evolution, from a naturalistic standpoint, is that they always bottom out in *natural selection*, a non-rational, non-intentional, *purely causal* mechanism.

So, when it comes to explaining meaning and representation on pragmatism, we have a way of connecting the dots from the *scientific image* to the *manifest image*. If we take *representation* as primitive, however, no such explanation seems possible. Social practices and dispositions to respond to environmental stimuli are both unproblematic from the standpoint of the scientific image, but representation *itself* is not. Now, of course, representationalists have attempted to explain representation in naturalistic terms. Most of these attempts have been regarded as problematic in various ways, and I don't have the space to go over all of them here. The point is this: if we *can* get a notion of representation out of social practices and reliable dispositions, as the pragmatist thinks we can, then there's an explanatory path in sight, and we ought to follow it.

That's the explanatory upside of pragmatism. What about the epistemological one? In Rorty's book, *Philosophy and the Mirror of Nature*, he criticized representationalism on largely epistemological grounds. If one thinks of thought, in the first and foremost instance, as an attempt to *represent the world*, the following question arises: How will we ever know that we're representing things *as they actually are?* If you're a representationalist, this question is of crucial importance, and the classical representationalists struggled mightily with it. Descartes, quite famously, attempted to derive knowledge of the external world from the inner certainty of his own thought. But he didn't get much beyond this without invoking God. Locke thought that we could inductively infer from the coherence of the impressions that we have to the external world.

These traditional responses to the skeptical worry are *foundationalist* ones. They start with some set of indubitable inner impressions, and they infer outward to the things that those impressions are supposed to indicate. However, if the argument from Sellars that I mentioned in the previous section is roughly

correct, no such strategy can work. A mere internal sensation can't license an inference an external thing unless we *already have the understanding* that such a sensation is a reliable indicator of that thing. But, if inner sensations are *the only thing* we have from which to draw our concepts, then there is no way that we could understand these sensations in such a way that they license beliefs about the external world. In the previous section, I employed this argument with reference to external objects, but it works just the same for internal sensations as well. In fact, this was Sellars' initial target when giving this argument.

While the representationalist must struggle with this difficult epistemological problem, for pragmatist, no such problem arises. If our concepts of objects arise within a social practice in which we are conditioned to respond appropriately to such objects, then the fact that there is a reliable connection between the concept of an object and the object itself is built right into the nature of the concept. Our entitlement to employ the concept is not, in the first instance, derived from the thing *itself*, but from the authority we are granted by fellow gameplayers. With this way of thinking about things, the worry of epistemological skepticism is avoided.

While these reasons certainly do not *conclusively* justify this pragmatist project, hopefully they are able to provoke enough interest to justify further exploration into how the details of the project might be worked out. Further, I hope I've elaborated the theory enough to show that this it is at least a *plausible* way of thinking about language, one that might be able to accommodate the complexities and peculiarities of human language. If this project is on the right track, we may be well on our way to finally understanding the nature of understanding.

Further Reading

Ludwig Wittgenstein: Philosophical Investigations

Wilfrid Sellars: Empiricism and the Philosophy of Mind, and Philosophy and the Scientific Image of Man

Richard Rorty: Philosophy and the Mirror of Nature

Robert Brandom: Making it Explicit and Articulating Reasons