

Rationality in action: A symposium

John R. Searle, Barry Smith, Leo Zaibert & Josef Moural
Philosophical Explorations 4 (2):66 – 94 (2001)

Abstract: John Searle's book *Rationality in Action* presents a sophisticated and innovative account of the rationality of action. In the book Searle argues against what he calls the classical model of rationality. In the debate that follows Barry Smith challenges some implications of Searle's account. In particular, Smith suggests that Searle's distinction between observer-relative and observer-independent facts of the world is ill suited to accommodate moral concepts. Leo Zaibert takes on Searle's notion of the gap. The gap exists between the reasons that we have for acting and our actions. According to Searle, whenever there is no gap, our actions exhibit irrationality. Zaibert points out a certain obscurity in Searle's treatment of the gap, particularly in connection with Searle's notion of 'recognitional rationality'. Finally, Josef Moural examines the interactions between Searle's theory of institutions and his theory of rationality, with emphasis on the connections between intentionality and Searle's notion of the 'background'.

Appended below is a preprint version of Barry Smith's contribution to this symposium (pp. 70-75) together with Searle's response (pp. 75-77)

The Chinese Rune Argument

Barry Smith

Abstract: Searle's tool for understanding culture, law and society is the opposition between brute reality and institutional reality, or in other words between: observer-independent features of the world, such as force, mass and gravitational attraction, and observer-relative features of the world, such as money, property, marriage and government. The question posed here is: under which of these two headings do moral concepts fall? This is an important question because there are moral facts – for example pertaining to guilt and responsibility – which hover uncomfortably close to the boundary between the observer-relative and the observer-independent. By means of a thought experiment involving an imagined Chinese society in which guilt is determined by the random throwing of sticks, I seek to show that moral concepts threaten the foundations of Searle's philosophy of social reality.

The Chinese Rune Argument

Barry Smith

Hard cases do not make good law, but they can, if they are well chosen, make for good philosophy. What I want to offer here is a hard case for the philosophy defended by Searle in his *Rationality in Action* and in *The Construction of Social Reality*.¹ I agree with very much of what John Searle has to say in both of these works. Like Searle, I am of the opinion that it is long overdue that ontology be restored to its rightful place as the central discipline of philosophy and I see Searle's two recent books as representing important contributions to a whole new way of doing ontology. I also agree with Searle that, in the attempt to understand culture, society, law, the state, and similar social formations, we do not need to invent new realms of being. All such formations exist – together with you and me, and our physical and biological parts and moments – within one single world, a world which contains both culture and biophysics. Thus I will here accept Searle's naturalism. There is one world and social reality is one (very special sort of) complex thread running through it.

Searle's tool for understanding culture, law and society is the opposition between brute reality and institutional reality, or in other words between:

1. observer-independent features of the world, such as force, mass, and gravitational attraction, and
2. observer-relative features of the world, such as money, property, marriage and government.

The question I want to pose here is: *under which of these two headings do moral concepts fall?* This is an important question because, as we shall see, there are moral facts which hover uncomfortably close to the boundary between the observer-relative and the observer-independent. Moral concepts thus threaten the foundations of Searle's philosophy of social reality. If, on the other hand, we can provide an account of moral concepts within the terms of Searle's dichotomy, then we will have solved at least part of the problem of explaining how humanity has been able to raise itself up by its bootstraps from out of the biophysical realm of brute reality and out into the deontic world of institutions (the problem in other words, of explaining how deontic powers can arise out of the biophysical reality of acoustic blasts, inscriptions, intentions and like phenomena).

I will concentrate my remarks here on the special case of guilt, a concept that is closely related to the notion of the self which plays so central a role in Searle's new theory of rationality. Thus if we can give a workable account of guilt, then we shall be well on our way also to a better understanding of this notion, too, as also of a series of related notions such as responsibility, blame, punishment, praise, desert, reward, condemnation, and so on.

We might at first suppose that it is obvious that guilt has to belong to the observer-relative features of the world. For how could something like guilt be part of brute physical

1. I thus continue the pursuit of hard cases for Searle's theory begun in "[The Construction of Social Reality: An Exchange](#)", *American Journal of Economics and Sociology*, 59 (2001).

reality? If guilt belongs to the realm of institutional reality, however, then it is subject to what we might call the Strong Appearance Injunction. As Searle himself expresses it:

In the case of institutional facts, the normal relationship between intentionality and ontology is reversed. In the normal case, what is the case is logically prior to what seems to be the case. So, we understand that the object seems to be heavy, because we understand what it is for an object to be heavy. But in the case of institutional reality, the ontology derives from the intentionality. In order for something to be money, people have to think that it is money. But if enough of them think it is money and have other appropriate attitudes, and act appropriately, then it is money. *If we all think that a certain sort of thing is money and we cooperate in using it, regarding it, treating it as money, then it is money.* (*Rationality in Action*, Chapter 6, italics added)

The Strong Appearance Injunction (or what I shall call ‘Strong AI’ for short), then, is the thesis to the effect that as concerns social reality what seems determines what is (and we note in passing that Searle is a vehement opponent of this thesis in the context of his own Chinese Room argument).

Strong AI can be a dangerous thing if it gets into the wrong hands.²

To many people it seems that God exists
Therefore (by Strong AI), God exists.

Consider:

People deal on a daily basis with computers.
To many people (perhaps to all people) it seems that computers bear blame.
Therefore (by Strong AI) computers do bear blame.

It will be rather easy for us to build a similar argument in relation to the phenomenon of guilt. Suppose, for the sake of argument, that everyone in China believes that the world is an organic whole of a very special type. Things behave as they do not because of their intrinsic natures but because of their position in this whole. As far as human beings are concerned, it

2. Recall the old Wittgensteinian saw:

People say ‘God exists’
(this language game is played)

is above all their position in the family (extending back, through many generations, to the original male) which is important in this connection. When an individual acts, then his action is, roughly, an expression of the spirit of his ancestors. Guilt and responsibility thus accrue not to the individual but to his entire family, according to patterns discernible via methods like plastromancy and scapulimancy (divination using tortoise shells and shoulder blades), and above all through manipulation of the runic hexagrams of the *Classic of Changes*.

And now comes the Chinese Rune Argument itself:


Everybody believes that guilt and responsibility are determined by the shaman via the manipulation of bamboo sticks

This is how it seems to everybody as far as guilt and responsibility are concerned

Therefore this is how it is as far as guilt and responsibility are concerned (by Strong AI)

Therefore guilt and responsibility depend upon the outcome of random manipulations of bamboo sticks

In order to establish guilt, therefore, the supplicant comes to the shaman, who sits in front of a mound of bamboo sticks. The shaman consults the spirits of the ancestors of the supplicant and then throws the sticks on the floor. The supplicant is then deemed to be either guilty or innocent according to the pattern which the sticks create, which might look something like this:



(We can imagine a series of variant forms of the argument. Thus in some parts of China, the sticks are used to decide what, from day to day, will count as crimes. In other parts of China, the sticks are used to determine the identity of the culprit, and so on.)

Responses to the Chinese Rune Argument

We can imagine a series of responses to the Chinese Rune Argument, some of which are more serious than others. At the head of the queue (I am beginning with the less serious responses) is what we might call the Systems Reply. Its proponents take no great issue with the notion that guilt might be determined by the random manipulation of bamboo sticks. They locate the error in the argument, rather, in the fact that it misidentifies the subject possessed of the ability to divine the presence of guilt. The shaman in the room does not

have this ability. Rather, he is merely part of a whole system, including bamboo sticks, ancestral spirits, and so on, and this *whole system* does have the capacity to divine guilt.

Next in the queue is what might be called the Multiculturalism Reply. This concedes that the argument is valid as it stands, but holds that Searle can save his position by writing different version of *Rationality in Action*, one for each different culture. Thus there will be one version for the Chinese, one for the Zoroastrians, one for the Californians, and so on. Western versions of the book would then talk about guilt, responsibility, the self, and so on, *as facts of our, Western natural history*. Perhaps this is what Searle is doing: offering us facts of our natural history. But then, given the grotesque relativism which is implied by the Multiculturalism Reply – which allows for facts about guilt, responsibility and so on to obtain simply in virtue of whatever views given social groups happen to hold, it is difficult to see how these different versions of *Rationality in Action* could then offer any profound contribution to our *philosophical* understanding.

A response closer in spirit to the thinking of Searle might read as follows (we might call this the OJ Reply):

There are lots of words such as ‘guilty’ and ‘responsible’ which oscillate between an observer-independent and an observer-relative sense. Thus, consider the question ‘Was OJ guilty?’ In the observer-independent or brute fact sense, that is a question about whether or not he in fact intentionally stuck a knife into his wife. But there is another sense where it means ‘Was he found guilty by the courts?’ And the answer to the first question may be ‘yes’, while the answer to the second question is ‘no’. One set of questions has to do with the actual facts, what actually happened in the physical biological world. But then another set of questions has to do with what sort of status functions we assign to those functions.

The upshot of the OJ reply is this: that there is a sense of ‘guilt’ in which guilt is not an observer-relative phenomenon after all. We can pin down the relevant notion as follows:

(*) A subject S is guilty in the observer-independent sense of some deed F iff (defeasibly) x did F and F is a crime.

We say ‘defeasibly’ since S is not guilty if, for example, his life is threatened, or he is under military authority in time of war (or perhaps even if 99% of the people think it’s OK to kill your cheating ex-wife). But these matters need not concern us further here, since anyone who is interested in deontic powers faces a problem when it comes to giving a principled account of how a complete list of such conditions is to be formulated.

More important is to note that the sense of guilt which is at issue when we say, truly, that OJ is guilty, involves both an observer-independent and an observer-relative dimension. The former relates to the fact that, for the definition (*) to be applicable at all in a way that avoids circularity, we must be able to describe the deed F in observer-independent terms: we require an observer-independent specification of the brute facts upon which the pertinent status functions can then be seen as being imposed.

One problem for Searle (or for the proponent of (*)) is that there are many sorts of acts of which one can be guilty but for which it is very difficult and perhaps even impossible to give such an observer-independent description. Culpable negligence is one problematic example in this respect, since what turns an innocent mistake into a culpable one is a matter of cultural, observer-dependent conventions. But already we can learn from the phenomenon we call ‘bad manners’ that there can be examples of pieces of behaviour to which notions of guilt, blame and the like can be coherently assigned which yet cannot be described and relevantly classified in observer-independent terms.

It is however the final clause in our definition of guilt in the observer-independent sense – the clause relating to the fact that ‘F is a crime’ – which causes the most serious problems for Searle (if, indeed, he accepts the definition (*) at all). For of course it is the community which determines which deeds are classified as crimes. Whether a token act is a crime might be a brute fact, but deciding whether the act type is to be counted among crimes involves appealing to conventions and the like, and thus it has embedded institutionality. Some deeds – murder, theft – are classified as crimes in our society; others – blasphemy, petty theft between spouses, killing your enemy in a boxing match – are not so classified. Different communities have different repertoires or catalogues of crimes. Now, however, it seems that we are confronted once again with something like the Chinese Rune Argument.³ For if *being a crime* truly is an observer-relative property of a deed F, then we can imagine a society in which, for example, nodding your head is crime, and in such a society people would also be guilty whenever head-nodding took place.

Remember that our task here is to account for that feature of the world which goes under the heading of guilt as this term is used in our society. And surely this term is not used in our society in an observer-relative sense, which is to say: it is not used in such a way that it would be applicable, even in a society of the sort just described, to people merely nodding their heads. Rather we, in our society, would say that head-nodders in such a society merely *count as guilty*. (This suggests the charge that Searle runs together *guilt* and *what counts as guilt*, just as he can be accused as having run together in an earlier incarnation: *makes a promise* and *counts as making a promise*.⁴)

The issue here is a subtle one (and it has nothing to do with any suggestion that there is a specifically Western notion of guilt that is somehow superior to counterpart notions at work in other, non-Western societies). It comes down, rather, to the need to recognize what many important legal theorists have recognized through the ages, namely that there is a rationale

3. In chapter 8 of *Speech Acts*, Searle answers objections to his derivation of ought from is along the lines of there being two different senses of ‘ought’, one descriptive and one normative. Perhaps I am just offering here parallel arguments to the effect that there are two senses of ‘guilt’, one brute and one institutional. Then however we need to know what the relationship is between the two pairs (1) description and normativity, on the one hand, and (2) brute and institutional, in the other. And then if Searle is forced to claim that anything normative is institutional, then he is once again exposed to the Chinese Rune objection.

4. See my “[Reinach, Searle, Hume: An Essay on Material Necessity](#)”, P. Hanson and B. Hunter, eds., *Return of the A Priori* (*Canadian Journal of Philosophy*, Supplementary Volume 18), (1993), 301–322.

to the fact the some deeds are classified as crimes in given societies and others not – a rationale which makes such classification not a merely observer-relative (or better: culture-relative) matter. This rationale can be formulated as follows: crimes are necessarily such that they violate certain *basic norms of society*.

Now one might suppose that consideration of matters such as this is entirely alien to Searle's project (though given the ambitious nature of his philosophical project even if this were so it would not leave Searle completely off the hook). In fact however the matters here at issue are treated of by Searle himself in an important passage of *Rationality in Action*. For Searle, too, attempts to derive just such a rationale for himself, and he does this precisely in order to put himself in a position where he will be able to infer that (for example) murder and slavery are wrong, while other types of behaviour counted as crimes in other societies (Searle's preferred example is blasphemy) are acceptable. Searle thereby does indeed attempt (albeit unwittingly) to insulate himself against the Chinese Rune Argument. All of this occurs in chapter 5 of *Rationality in Action*, where Searle deals with the topic of what he calls 'strong altruism'.

The crucial passage, which bears more than a hint of Habermasian logic, is as follows:

Once I have consciousness and the self and I'm able to use language, I have already committed to the existence of other consciousnesses and selves on the par with my own.

Strong altruism can, Searle argues, be derived from the very existence of language and rationality.

I do not want to spend time criticizing Searle's argument here.⁵ Suffice it to note that, even if he can indeed prove that, when we use language, we need to treat other language users on a par with ourselves, this does not by any means suffice to prove that we have to treat *all* other language users or users of *all* other languages in just this same way. (We might, for example, suppose that the people over the hill speak quite the wrong sort of language to be treated as equals.) Once this is allowed, however, then the language-based argument for strong altruism opens up the floodgates once again to the Chinese Rune Argument.

Searle has set himself the task of describing in naturalistic fashion the way societies actually work and the way institutional reality is actually built up in this world of ours, while at the same time describing how deontic powers gain their force. This is a hard task indeed, and Searle has perhaps come closer to fulfilling it than any other contemporary philosopher of social reality. But while the arch-realist Searle has so heroically demonstrated the necessity of realism when dealing with the physico-biological world (and with the claims of the proponents of strong artificial intelligence), he has weakened his

5. Perhaps Searle does not want his remarks on strong altruism to be interpreted in a moral sense at all, so that Zaibert is right when he asserts that "in spite of the fact that Searle now cares about marriages, money, and property rights, he has as yet little to say about morality." ("Intentions, Promises, and Obligations", in B. Smith (ed.), *The Philosophy of John Searle*, Cambridge: Cambridge University Press, forthcoming.) Then, however, we must raise the question as to how these remarks are to be interpreted.

commitment to realism when dealing with the world of responsibility and guilt. 'Social reality and the space of reasons do not exist unless people have specific sorts of beliefs. This is the core of Searle's construction. What he still misses is that there are facts about the space of reasons which hold independently of and even in spite of how people believe.

Response to Barry Smith

John Searle

I think Barry Smith supposes that I am trying to answer a question that I am not in fact trying to answer. I am not trying to answer such questions as, "What sort of moral society should we have," "What sorts of moral criteria should we have". I think those are legitimate questions, they just do not happen to be the ones I was trying to answer. The question that I was trying to answer initially was, "How can there be objective facts which are only facts because we think they are facts? How can there be facts where, so to speak, thinking that it is so makes it so?" It seems to me there are facts such that thinking them so does make them so. For example, we are all at something we call "The Wittgenstein Conference". Now ask yourself what fact about this makes it the Wittgenstein Conference? One could tell a whole story about how the institution was created and how it is sustained and so on, but it is essential to the operation of all this apparatus that people have to have certain attitudes, and those attitudes are in large part constitutive of the fact that this is the Wittgenstein Conference. So it could not turn out, for example, that unknown to us, scientific research will one day prove that though no one ever knew it at the time or before, but afterwards it was discovered that this was really the Ernest Hemingway Conference. Appearance is indeed one thing, reality something else. But there are cases where the appearances add up to the reality, and that is not an epistemic point. The point is, rather, that the fact that we treat something in a certain way creates a reality. And this I believe is true of institutional reality in general. The reality of money, property, marriage, government, universities, etc.

At this point Barry Smith's question becomes relevant: where does a fact such as "the Chinese guy is guilty" fall? Is this an observer-relative fact or an observer-independent fact? When I was writing my book on rationality I had a completely different set of questions. This question about the observer-relativity of guilt and innocence is not a question about rationality. In my book on rationality I find myself forced in chapter three to postulate the existence of a self in order to account for the gap. I will not spell out the whole argument here. You would have to read the book to get the whole argument. But the basic idea is that you cannot make sense out of rational explanations of human behavior without postulating a self because you have no other way to make intelligible the fact that we have adequate causal explanations that do not give causally sufficient conditions. The fact that we have such explanation bothers a lot of people - Galen Strawson, and Thomas Nagel, for example - and I am trying to answer their objections I want to claim that we can have adequate causal explanations which do not cite causally sufficient conditions because they have a quite specific logical structure which is different from the structure of explanations in the natural sciences.

What is meant by “the self” as I use it? It is really a formal notion. It does not mean “personality” or “character”. Having a self is not a matter of being a nice person, for example. The notion of the self as I am using it is just an entity that satisfies a certain set of formal conditions, namely it is the entity that is capable of agency, capable of rational thought including memory, perception, belief, and so on, and capable of making rational decisions. Once you have a notion of the self, you can make a lot of other notions intelligible, such as responsibility and guilt. But it was not my intention to offer an analysis of guilt.

Barry Smith points out that insofar as there is a set of institutional facts concerning guilt, the criteria for what counts as guilt varies from one culture to another. But that is not an objection to my theory; rather that is exactly what the theory should predict. Different sorts of people would count different sorts of things as guilt. But then what are we to say about Barry’s imaginary Chinese example? The way he has described the case, I think we would have to say that they have irrational criteria for deciding questions of guilt.

I believe that Barry’s account of my views is based on a misunderstanding. On my view there is, or can be, a simple brute fact about what someone did and what someone did intentionally. Let us suppose that it is a brute fact that a certain person committed a certain act and did so intentionally. Those brute facts will count as a “crime” in our society only if there is in addition to the brute facts a law which specifies that certain acts are criminal, others not. In such a case, then, you identify one institutional fact – the institutional fact of a crime having been committed - in terms of other institutional facts, the fact that certain laws exist. But the identification of one institutional fact in terms of other institutional facts is perfectly normal. Someone is a candidate for the presidency of the United States only if he is already a citizen of the United States. So “He is a candidate for the presidency” (one kind of institutional fact) presupposes “He is a citizen” (another institutional fact). So it is not an objection to me, to say that in order to say that such and such can be counted as guilty you have to make reference to another institutional notion, the notion of a crime. Sometimes Smith talks as if he thought crime and criminal guilt could be brute facts, but such facts as the fact that such-and-such is a crime and so-and-so is guilty of a crime are never brute facts. Such facts are only relative to some institutional structure.

So I do not really see that there is any serious difficulty here. There are non-institutional brute facts, and on top of these we create institutional facts by imposing status functions. The point is that in the U.S. since there is a law that says that smoking marijuana is a crime, to claim that Joe is smoking marijuana implies that (defeasibly) Joe is committing a crime. But the fact that smoking marijuana (the act-type) is a crime in the US is not a brute fact; rather, this fact is only the result of a convention, specifically, a law. So whenever someone says, “So and so committed a crime”, there is a brute aspect and an institutional aspect. The brute aspect would be the act in question, and the institutional aspect would be whether or not the society in question counts that act as a crime.