John R. Searle

Social Ontology and the Philosophy of Society

Abstract: This lecture was originally given at the Einstein Forum in Berlin. It contains a summary of some of the themes in my book *The Construction of Social Reality* and continues the line of argument I presented there.

Tonight I want to make a plea for a branch of philosophy that in English speaking countries does not yet exist. I don't know if it exists in German speaking countries. It is what we might call the 'philosophy of society' and I see this subject as centering essentially around questions of social ontology. We already have something called 'social and political philosophy' and we have the 'philosophy of social sciences'. But if you look closely at those disciplines you find that social and political philosophy really tends to be political philosophy. And the philosophy of social sciences really tends to be about questions of methodology in the social sciences. What I'm going to urge tonight is that, in the sense in which there is a separate branch of philosophy that we now think of as the philosophy of language, but which didn't exist at the time of Kant and Leibniz; so we should think of the philosophy of society as a separate branch of philosophy. And we should think of such topics as the nature of the appropriate methods in the social sciences, which you get in the 'philosophy of the social sciences', or the moral implications of social organizations, which you get in 'social and political philosophy', as special topics within the philosophy of society. Tonight I am going to make a plea for the central branch of the philosophy of society which deals with the question of 'social ontology'.

The question is: "what is the ontology of social reality?" To begin with, I am going to make an assumption which I won't justify today, though perhaps in tomorrow's discussion I can. The assumption is that we should abandon the traditional Western distinction between mind and body. We have in the Western intellectual tradition been cursed with this dualism, with the idea that there is a fundamental distinction between the mind and the body, between the mental and the physical. And in some respects the situation is made even worse by those philosophers who say that in addition to two worlds,

Analyse & Kritik 20 (1998), S. 143–158

we also live in a third world. Philosophers such as Frege and Popper, and more recently Jürgen Habermas have said we should think of reality as dividing into three different worlds. My own view is we should never have started counting. Descartes started counting and got up to two, Frege got up to three. I am saying "don't start counting". We live, as my colleague Donald Davidson likes to put it, in one world at most. That is enough for us. And the basic structure of this world is pretty much as described by physics and chemistry. Ultimate reality consists of entities we find it convenient to call 'particles'. They are organised in systems. These systems are defined by their causal relations, some of those systems are organic systems, some of the organic systems have consciousness. With consciousness comes intentionality, and when we have consciousness and intentionality we have reached the evolutionary stage of animals, mammals and especially primates like ourselves, who form social groups.

1. Some Puzzling Features of Social Reality

Now with all that by way of introduction, here is my problem. If you look at our social life it is a remarkable fact that there is a class of entities that have a very important role in our lives, but they only are what they are, because we believe that that is what they are. I will take an obvious economic example: I carry around in my wallet these sordid bits of paper. And they are really not very important as physical objects, but they matter to us. They are examples of 'money'. Now here is my puzzle: It is only money, because we believe that it is money, and yet it is an objective fact that it is money. That is, when I go into a store and I give them one of these, they don't say: "Well, maybe you think it's money, but why should we care what you think?". They accept it as money. So, here is the initial formulation of my puzzle: How can there be an important and objective class of entities that only exist because we think they exist?

I believe that when you start a philosophical investigation you have to start naively and I am just going naively to tell you some of the puzzling features about social and institutional reality. After having gone through a stage of naiveté we must become immensely sophisticated in giving answers to our puzzles. I have never found the algorithm for deciding when you have to stop being dumb and naive and when you start being smart and sophisticated. We shall just play it by ear as we go along. Anyway, here goes with half a dozen puzzling features of social reality.

Problem number one is that there is a kind of self referentiality in social concepts. Something is only money if we believe it is money, and it is only property, marriage, government, a cocktail party, tenure, a summer vacation, if that's what we believe it is. But now, if money is partly defined as that

which is believed to be money, then philosophers are going to get worried. If it has to be believed to be money in order to be money, what is the content of the belief? It looks like you are going to have circularity or infinite regress, because if part of the ontology of money is believed to be money, then part of the definition of 'money' is believed to be money, and consequently the belief that something is money, has to be in part the belief, that it's believed to be money. And that means, you are in trouble, because then the content of that belief is that it is believed to be believed to be money, and so on. So that's the puzzle: how do we avoid circularity or infinite regress in the definition of 'money' if the concept has this self referential component? And what goes for money, also goes for property, marriage, government, and all sorts of other social and institutional phenomena.

Now that leads to a second question, and that is, what is the role of language in the constitution of social and institutional reality? It looks as if in the case of these institutional phenomena language doesn't just describe a pre-existing reality, but is partly constitutive of the reality that it describes. It looks like the vocabulary of money and government and property and marriage and football games and cocktail parties is partly constitutive of the phenomena. Otherwise, how do we account for the differences between animals that are incapable of language and consequently incapable of this sort of institutional ontology, and language-using animals like ourselves, where the words, in some sense we need to explain, seem to be partly constitutive of the social and institutional reality? Let me nail that down with an example. My dog Ludwig is very intelligent, but there are limits to his intelligence. Suppose I give him a pile of dollar bills and I train him to bring me a dollar bill whenever he wants to be fed. All the same, he is not buying anything, and it's not really money to him. His bringing me the money is not an economic transaction. Why not?

Let me nail down the problem about self-referentiality a little more closely as well. I said, it's only money, property, marriage, government etc. if we think that that is what it is. This actually has important consequences. Suppose, we decide we are going to give a cocktail party and we invite the whole population of Wilmersdorf, and we have a hell of a great cocktail party. But suppose things get out of hand and the casualty rate is worse than the battle of Gettysburg. All the same, it's not a war. It's not a war, unless people think it's a war. As long as they think it's a cocktail party, then it's a cocktail party, it's just a hell of a cocktail party. This feature of self referentiality is actually of some historical importance. I have always wondered, how could Cortez with 150 or so bewildered Spaniards beat the entire Aztec army? Not to mention Toltec, Mixtec, Aranhuac and other assembled tribes. Well, part of the answer is, they had a different definition of what they were doing. You see, the Aztecs were fighting a war according to their definition. That means

you get close enough to an enemy so that you can hold him without bruising him and later on you sacrifice him to the Great God Quetzacoatl by cutting out his living heart with an obsidian knife on the top of a pyramid. Well, that may be a great definition of warfare for Central American tribes; but it is very ineffective against Europeans on horses with metal weapons. So the sorts of phenomena I'm talking about actually have historical consequences. It isn't just that we are dealing with philosophical puzzles.

So far I have covered two sources of puzzlement. The constitutive role of language and self-referentiality. A third related source of puzzlement for me (and this has a special interest to me) is the special role of performatives in the creation of social and institutional reality. For a very large number of institutional facts you can create the fact just by saying you are creating it, provided you have the appropriate authority in the appropriate situation and the context is correct. So, you can adjourn the meeting by saying "I adjourn the meeting". You can declare war by saying "We declare war". You can pronounce somebody husband and wife by saying "I pronounce you husband and wife", and so on with a large number of cases. Now why is that? How can you create institutional reality just by saying you are creating it? You cannot do it with everything. You cannot score a goal in football by saying "I score a goal", or even "We hereby score a goal". So what is the difference? What is going on here?

Well, I will give you a couple of more of these puzzles and then we will start to try to solve them. Another puzzling feature of social reality is the complex interrelations among the elements. They seem to be systematic. So you don't just have money, but in order to have money you have to have a system of exchange, ownership, payment, debts and in general you have to have a system of rights and obligations. It might seem that games are an exception because games are self-enclosed in a way that money and property and marriage are not. But even in the game you understand the position of a batter and the position of a pitcher only in terms of understanding the notions of rights and obligations. And that already involves you in more general social and institutional notions. So I am struck by the pervasive interlocking character of the kinds of social and institutional phenomena that I'll be talking about.

There is one last puzzle I will mention. We could go on all night listing puzzles, but let's settle for five. The fifth puzzle that interests me is: though there exists a real institutional reality of elections, wars, property exchanges, stock markets and so on, nonetheless you can't have an institutional reality without an underlying brute physical reality. Here is an interesting fact. Money can take a very large number of forms. It can be in the form of gold or silver or paper or copper, it can be in the form of credit cards, and some primitive tribes use wampum or sea shells. By the way, most of your money underwent a dramatic physical change in the past 10 or 20 years that you

didn't even notice. It happened in the middle of the night. Most of your money is now represented by magnetic traces on computer discs in banks, and it doesn't make a bit of difference: You didn't lose any sleep at all over this, though there was a revolutionary change in the physical representation of your money. Now, here is the point. Almost anything can be money but at some point it has to have some physical reality. There has to be something whether it be gold or magnetic traces that counts or could count as money. Why is that? Why is the physical necessary, and why is there a primacy of the brute physical fact over the institutional fact?

2. Conceptual Tools Necessary to Account for Social Reality

Now we have a problem. Let's go to work to solve it. In order to solve it, I want to make another distinction that I have been presupposing and that I think is absolutely essential for understanding our position in the world. There are classes of objective facts in the world which have to be distinguished from certain other objective facts in the following regard. Many things that we think of as real nonetheless only exist relative to observers, in the form of reality that they have. We need to distinguish those features of the world that we might call 'observer-independent' from those features that are observer dependent. Observer independent features are those that, so to speak, don't give a damn about human observers, and here I am thinking of things like mountains and molecules and galaxies and processes like photosynthesis and mitosis and meiosis. All of those phenomena are observer-independent. But in addition to them, there are lots of other phenomena in the world whose existence depends on being treated or regarded in a certain way by human agents. Observer dependent phenomena would include such things as chairs and tables and glasses and money and property and marriage. So, we need a general distinction between those phenomena that are observer-independent and those whose existence is observer-relative.

Typically an observer-relative entity will have both sorts of features. So this object, which I carry around in my pocket, has a certain weight and that it has the weight that it has is observer-independent. It doesn't depend on me or anyone else, it depends on the gravitational relations between the object and the center of the earth. But this object is also a Swiss army knife and the feature of being a Swiss army knife is observer-relative. So we need a general distinction between those phenomena that are observer-independent and those that are observer-relative. Typically the natural sciences deal with phenomena that are observer-independent, phenomena like mountains and molecules and tectonic plates. Typically the social sciences such as economics, sociology,

and political science, deal with phenomena that are observer-relative. And here I am thinking of such things as political parties, elections, social classes and money. The question for this evening we can now state a little bit more precisely: we are discussing the ontology of a certain class of observer-relative social and institutional reality.

For the analysis of this social reality I need exactly three devices, three tools to try to analyse that ontology.

Here is the first one. We need to call attention to the class of entities to which we have assigned functions. Many of the most common concepts that we use in dealing with the world, for example concepts like "cars" and "bathtubs" and "tables" and "chairs" and "houses", involve the assignment of function. It is a remarkable capacity that humans and certain animals have, that they can assign functions to objects, where the object does not have that function independently of the assignment. And I want to make a strong claim about this assignment of function. I want to say: All functions are observer-relative. It is only relative to agents, only relative to observers that something can be said to have a certain function.

We are blinded to this fact by the practice in biology of talking about functions interchangeably with talking about causation. But there is a subtle difference. We do indeed discover such facts as the fact that the function of the heart is to pump blood. We do indeed discover, that the function of the vestibular ocular reflex is to stabilize the retinal image. But we discover those functions only against the background presupposition of certain norms. We have to assume that life and survival have a value, and it is against the presupposition of the norm, against the assumption that life and survival and reproduction are valuable, that we can say such things as that the function of the heart is to pump blood. If we thought that life and survival were worthless, that the only thing that really mattered was death and extinction, then hearts would be disfunctional, and cancer would have a useful function, it would hasten extinction. We don't think these things and it is crucial to our assignments of function that we don't. But it is only against the background of the presupposition of normativity, that we can discover such facts as the fact that the function of the heart is to pump blood.

One way to put this point is to ask: what is the difference between saying that the heart *causes* the pumping of blood, on the one hand, and saying that *the function* of the heart is to pump blood, on the other. And it seems to me there is a crucial distinction, because once you introduce the notion of function you introduce normativity. Once you introduce the notion of function, you can talk about such things as heart disease, malfunctioning hearts, hearts that function better than other hearts. Notice, we don't talk about better and worse stones, unless we assign a function to the stone. If you think this stone will make a good projectile, then you can evaluate it. You can say

this one is better than that one. Or if you assign it an esthetic function you can say this stone is *an objet d'art trouvé*, and with such an assignment of function, you may think the stone has some artistic value. So that's the first point, we assign functions and all functions are observer relative.

The second notion I need is that of collective intentionality. All genuinely social behavior contains collective intentionality on the part of the participants. You can see the centrality of collective intentionality if you contrast genuine cooperative behavior from behavior which merely happens to be coordinated with other behavior. Suppose for example that we are playing in a symphony orchestra. Suppose I am playing the violin part and you are singing the soprano part, and together we are part of the performance of Beethoven's 9th symphony. We have to be able to make the distinction between me sawing away on the violin and you independently but by chance simultaneously singing "Freude, schöner Götterfunken", and us doing this intentionally together in concert. So a basic ontological fact about social and collective behavior seems to be collective or shared intentionality in the form of collective beliefs, desires and intentions.

But in my intellectual tradition the existence of collective intentionality creates a real problem. If all the intentionality I have, is in my head, and all the intentionality you have is in your head, how can there be such a thing as collective intentionality? There are a lot of ingenious efforts to try to solve this problem in philosophy. Basically they try to do it by reducing collective intentionality to individual or singular intentionality. They try to reduce we-intend, we-believe etc. to I-intend plus I-believe that you have such and such an intention. And then on your part it is I-intend plus I-believe that you have such and such an intention.

On the view that I am opposed to, the assumption is, We-intentionality must reduce to I-intentionality. Collective intentionality must reduce to individual intentionality. Otherwise you would have violated the "principle of methodological individualism". If you say that collective intentionality is primitive, then it seems you are in very bad company. It seems that you are postulating some kind of Hegelian Weltgeist that is floating around overhead, or something like that. Where I live you don't want to be caught doing that, otherwise you'll lose a lot of friends. Given that puzzle—how can there be collective intentionality, when all intentionality is individual?—it looks like we have to reduce collective intentionality to individual intentionality. An enormous amount of intellectual effort has been spent, in my view wasted, trying to do that. The analysis that comes out involves something called "mutual belief".

For example, consider a case where we are pushing a car together to try to get it started. Now that is a case of collective intentionality. So how is that supposed to be analyzed? The idea is this. When we are pushing the car

together, then I intend to push the car and you intend to push the car. And I believe that you believe that I intend, and I believe that you believe that I believe that you believe that I intend and so on up in an infinite hierarchy. And for you it is the same. It's "I believe that you believe," etc. on up. Now I think my poor brain will not carry that many beliefs and I want to suggest there is a very simple way out of this puzzle. The puzzle is, assuming that all intentionality is in the heads of individual human and animal agents, how can it be the case that it's all in our individual brains, if it is irreducibly collective? And the answer is, that we can have intentionality in your brain and my brain, which is in the form of the first person plural as much as we can have it in the form of the first person singular.

On my view there is a trivial notational solution to the puzzle. The irreducible form of the intentionality in my head, when we are doing something collectively is, 'we intend'. And I don't have to reduce that to an 'I intend' and a set of mutual beliefs. On the contrary, I have the 'I-intends' that I do have, precisely because I have an irreducible we-intend. To nail that down to cases, I am indeed playing the violin and you are singing the soprano part, but I am only doing what I am doing and you are doing what you are doing, because we together are collectively playing the chorale movement of Beethoven's 9th symphony. I hope everybody sees that point.

As I said, the problem I am discussing has a traditional name. It's called "the problem of methodological individualism". And the assumption has always been: either you reduce collective intentionality to the first person singular, to 'I intend', or else you have to postulate a collective world spirit and all sorts of other perfectly dreadful metaphysical excrescences. But I reject the assumption that in order to have all my intentionality in my head, it must be expressible in the first person singular form. I have a great deal of intentionality, which is in the first person plural.

Nothing comes without a price and we do pay a price for the solution that I am proposing to this puzzle. The price is this. It turns out that I can be mistaken, not only in what is happening in the world, but I can be mistaken about the very mental state that I have. That violates the Cartesian assumption that we cannot be mistaken about our intentions. But I think that is the right way to think of it. Suppose in the case where we are pushing the car I discover that you weren't in fact pushing? You were just going along for a ride, I was doing all the pushing. Well, then I was not only mistaken in one of my beliefs, but it turns out that in a way I also was mistaken about what I was doing. I thought I was pushing as part of our pushing and in fact that's not what was happening. I was doing all the pushing, you were just pretending. So that is a price that we have to pay. You can be mistaken about the nature of the activity you are engaged in, if you have an assumption about

the collective intentionality, which is not shared by your apparent cooperators. But that seems to be the situation we are in in real life.

The third tool is this. Years ago, when I first started working on speech acts I made a distinction between brute facts and institutional facts. Those facts that I said were 'institutional facts' presuppose a human institution for their existence, for example such facts as that somebody is checkmated in chess, or somebody is elected President of the United States. I wanted to distinguish those facts, which are called institutional facts, from 'brute facts' whose existence does not require a human institution, the fact, for example, that the earth is 93 million miles away from the sun. You need an institution in order to state or describe that brute fact; you need the institution of language and the institution of measurement in mileage, e.g. French and kilometers, to describe it that way and you could state the same brute fact using different institutions. But the point I'm making is, the fact of distance between the earth and the sun does not depend on a human institution, though of course you have to have institutions in order to describe or state the fact. Now here is the point. There is a class of facts that are institutional facts and another class of facts that are 'brute facts', because they do not require human institutions. And then, the question is, how are institutional facts possible?

I also made the claim that you need a distinction between two kinds of rules. One sort of rule regulates antecedently existing forms of behavior. Another sort of rule doesn't just regulate antecedently existing forms of behavior, but creates the possibility of new forms of behavior. I call the difference between these two sorts of rules—using a Kantian terminology here—'regulative' rules, that regulate antecedently existing forms of behavior, and 'constitutive' rules that constitute new forms of behavior. Examples are obvious: The rule "drive on the right hand side of the road" doesn't create the possibility of driving. Driving can exist without that rule. That is a rule to regulate the already existing activity of driving. But the rules of chess are not like that. It wasn't the case that there were a lot of people pushing bits of wood around on boards and somebody said: "Look fellows, we have to get some rules so we don't keep bashing into each other. You stay on the right with your knight and I go on the left with my bishop!" Rather the rules of chess are constitutive in the sense that they create the possibility of the activity in question. Playing chess is constituted by acting in accordance with at least a certain large subset of the rules of chess.

Now here is the bottom line of this discussion. Those rules have a typical form. The form is 'X counts as Y' or 'X counts as Y in context C'. That is, such and such a move counts as a 'legal knight-move'. Such and such a position counts as 'you being in check'. Such and such a position counts as 'checkmate', and checkmate counts as 'winning' or 'losing' the game. And what goes for chess goes for much more elaborate institutions: such and such

noises count as 'making a promise', such and such marks on the paper count as 'voting' in an election, such and such number votes counts as 'winning' an election, and so on with a large number of institutional structures.

3. Status Functions

We now have three tools to solve our problems. These are, first, the assignment of function, second, collective intentionality, and, third, constitutive rules, rules of the form 'X counts as Y'. With all this apparatus assembled, let's go to work. I will now try to put it all together. I want you to imagine a simple community of (let's call them) hominids, beasts more or less like ourselves. Now it's very easy to imagine that such organisms, such primates can assign functions to objects. It's easy to imagine that they use a stick to dig with, or they use a stump to sit on. They can assign a function of being a digging tool or a stool to sit on. But now it's not a big step to imagine that they do that collectively. That collectively they have a very big stick that they use as a lever, or they have a big log that they use as a bench to sit on collectively. So it's very easy to tie collective intentionality to the assignment of function.

But now I want you to imagine the next step: Imagine—to take an example—our group of hominids live in a series of huts, and they build a wall around the huts. (This example wasn't designed with Berlin in mind, but anyway.) Imagine that they build a wall to keep intruders out and to keep their own members in. And now this is a case of the collective assignment of function, where the function is performed in virtue of the physics of the object on which the function is assigned. We just assume the wall is too big to climb over easily. But now imagine that the wall gradually decays to the point, where it is no longer able to keep the members of the community in, in virtue of its physical structure, nor to keep intruders out in virtue of its physical structure. But now let's suppose that, out of habit or whatever, the people involved continue to recognize the wall as a boundary i.e. they continue to acknowledge or accept that you are not supposed to cross the boundary. It is important to notice the vocabulary we use of 'acknowledge', 'accept', and 'recognize'. That is we imagine that the wall continues to serve its function, but no longer in virtue of its physical structure. It serves its function in virtue of the fact that it has a certain recognized status.

Now I wanted that to sound innocent, but I think that the move I just described is the basic move by which we create institutional reality of a specifically human sort. What happened was this. We imagine that an entity is used to perform a function, but it cannot perform the function in virtue of its physical structure. It can only perform the function in virtue of the collective recognition or acceptance of the entity in question as having that function.

And I want to say that is the underlying idea behind 'X counts as Y'. This line of stones, which is all that is left of the wall, now counts as a boundary. It now has a deontic status, it now has a form of power, which it exercises not in virtue of its physical structure, but in virtue of the assignment of function.

And I want to introduce a name for this sort of function—let's call these "status functions". A status function is a function that an entity performs not in virtue of its physical structure alone, but in virtue of the collective imposition or recognition of the entity in question as having a certain status, and with that status a function. And the structure of that—logically speaking—is the collective imposition of a function of the form 'this entity X counts as having this status and therefore this function as Y in this context C'. Now, I'm making a strong claim: this little device is the foundation stone of all institutional reality. So let's go to work and explain that claim.

I want to extend this account to the case of money. And just to nail it down to historical examples I want to talk briefly about the evolution of paper currency in Medieval Europe. (I love the Middle Ages, because it is, in a sense, the childhood of our civilization. In Medieval Europe, you see institutional forms that are growing and decaying, and the development of paper money is a very good example.) Initially people carried around gold and silver coins and the use of gold and silver was a form of barter. It was a form of barter, because the value of the coin was exactly equal to the value of the gold or silver contained in the coin, and the valuable coin was exchanged for other things. Now if you look in the text books they tell us, there are three kinds of money. There is 'commodity money', there is 'contract money' and there is 'fiat money'. But what they don't tell you is, what's the relation between them? The initial case we are talking about, where people actually had gold and silver, is a case of commodity money. Barter in gold and silver is both dangerous and inefficient, so people found they could leave the gold and silver with a group of people who worked on benches, and they were called "bankers", and the bankers would give them bits of paper on which it was said "we will pay the bearer of this note a piece of gold on demand".

With the introduction of the bits of paper we have now moved from commodity money to contract money, because the bit of paper is a contract to pay in gold or silver on demand. Later some genius discovered that you can actually increase the supply of money in circulation if you give out more bits of paper than you actually have gold in the bank. And as long as not everybody runs to the bank at once, it works. The bits of paper are still as good as gold. Much later on some genius discovered—and it took a long time to make this discovery: you can forget about the gold, and just have the paper. And that's the situation we are in now. We moved from commodity money, which is barter, to contract money, to fiat money. If you look at these bits of paper that I was waving around earlier, they seem to me good examples of the form

'X counts as Y', that is such and such bits of paper count as "currency". As it says on the piece of paper that I'm holding here, "This note is legal tender for all debts public and private". It counts as money in the United States. But that it counts as money, is a matter of collective acceptance of the status function in accordance with the structure: 'Such and such counts as so and so'. These bits of paper count as legal currency in the United States, just as these other bits of paper count as legal currency in the German Republic.

Now notice, that once you have got that structure 'X counts as Y', then automatically certain forms of abuse become possible. If I go in my basement and produce a lot of things that look like these bits of paper, I will be producing counterfeit money. Thus money isn't just anything that looks like this but it has to be issued by the Bureau of Engraving and Printing under the authority of the Treasury. So one form of abuse is counterfeiting. The structure automatically makes it possible to have abuses, because you can present something as satisfying the X term even if it doesn't in fact, and that's counterfeit. Another form of abuse is if you get too many of the entities in question. Then you have inflation and in hyper-inflation the entities are no longer able to function as money. And what goes for money, goes for other forms of social institutions. You can have counterfeit lawyers and counterfeit doctors, that is, people who don't actually satisfy the conditions, but who masquerade as lawyers and doctors. I don't know how it is in Germany, but in the state of California we now have so many lawyers that there is a kind of inflation.

Now here is a puzzling question. If I am correct in describing the logical structure of status functions, if it is just a matter of imposing a status and with it a function, then how can the system be so powerful? How can these structures have such an enormous effect on our lives, when, as I have described it, it all seems so fragile? There are two parts to the answer to that. One is this: the structure can be iterated indefinitely. Let me give an example. I make noises through my mouth, I just emit this acoustic blast. But these count as sentences of English. And in a certain context making noises of that sort, uttering those sentences of English, counts as making a promise. Making that kind of promise in that kind of context counts as making a contract. Notice how we are going up in the hierarchy. The X-term at one level will have been the Y-term at an earlier level and you keep going with it. Making that sort of contract, counts as getting married. And in the State of California once you get married, all kinds of things happen. You are entitled spousal benefits, income tax deductions, all sorts of rights concerning property, taxes, and so on. So you get an indefinite iteration.

The second point is that you get interlocking structures. I don't just have money, but I have money in my bank account at the Bank of America, which is put there by my employer, the State of California and which I use to pay

my debts to the Pacific Gas and Electric Company as well as my federal, state and local taxes. Now, just about every word I uttered in that litany was an institutional notion. We are talking about interlocking institutional facts. The whole point of the institutional is often to structure the brute. For example, recently I went and stood in front of a woman at a counter. I made noises and she made noises, I gave her a plastic card, she gave me sheets of paper, and the next thing is I was on an airplane on my way to Berlin. The movement of my body was a brute fact. My body moved from California to Berlin. (My body is still complaining about the jetlag). But the institutional facts made the brute fact possible. We are talking about a structure whose point is not just to empower other institutional structures, but to control brute reality.

However, the structure is also fairly fragile and the amazing thing is how rapidly it can collapse. I will never forget the moment, when I saw the people coming over the Wall on television. It was an amazing moment, because I was of a generation that thought the two-power division of the world would go on indefinitely. But there came a point, when the system of institutional reality was simply no longer acceptable and it just collapsed quite suddenly. So you can have a collapse of the institutional structure, if it's no longer accepted, and you can have a decay of the institutional structure of the sort that I have been describing.

4. Solutions to the Puzzles

We have a couple of minutes left, so I am going to show you how we solve these puzzles I began with. First, how can there be self-referentiality without circularity or infinite regress? Well, the answer is, you don't have to use the word "money" in order to define money. The word "money" functions as a summary term or as a place holder for being a medium of exchange, a store of value, a payment for services rendered, a measure of value of other currencies and so on. And if something performs all of those functions, then it's money. So we do not have a vicious circularity or infinite regress. If I say in order for something to be money, people have to believe that it's money, there is no circularity because they can have that belief without having the word "money". The word "money" here just is a place holder for a large number of other functional expressions.

Now, what about our second and third points, the role of language and especially performatives, how can performatives create institutional reality? And the answer to that is, that where the X-term is itself a speech-act, then typically you can create the reality by performing that speech-act. So you can make somebody husband and wife by saying "I pronounce you husband and wife", or you can find somebody guilty in a court if you are a judge by saying

"I find you guilty as charged". And the creation of the institutional fact need not even take the performative form. It says on this 20 dollar bill "This note is legal tender for all debts public and private". Now, I am an epistemologist and my natural worry is "how do you know?", and I want to write to the Treasury and say "How do you guys really know that it's legal tender? Have you done a survey, have you done an empirical study?" And the answer is, it isn't an empirical claim. They make it legal tender by declaring that it's legal tender.

Our next point—and this is the most important—is the constitutive role of language. Why is language constitutive of institutional reality, in a way that it is not constitutive of other forms of reality? Why is it that money and property and marriage and government require a vocabulary in a way that tectonic plates and gravitational attraction and galaxies do not require a vocabulary for their existence? That is in fact a very hard question to answer and I spent a whole chapter on it in the book on which this lecture is based, but now I will just summarize the answer in one sentence: for institutional facts there has to be some form of symbolism because there isn't anything else to mark the transition from X to Y. We just count the X term as having a Y status. But if we so count it, there must be some way to represent that counting feature. My dog can see somebody cross the line while carrying a ball, but can't see him score a 'touch down'. Why not? Because in order to see him score touch down you have to have some way to represent the extra status function and that requires language.

Now you might ask "well why do you need words?" And the answer is in some cases you do not. Suppose we kept score in a soccer match by piling up stones. I score a goal so I get a white stone to put on my side and you score a goal so you put a white stone on your side, and these are points. I got a point and you got a point. But now here is the 'point': These stones now play a linguistic role. They are now symbolic. They now play the role of symbolising scoring in the game. So the language or some other symbolism has to be constitutive because there isn't an independent ontology. The move from X to Y is itself a symbolising linguistic move and there has to be some way for us to represent it, otherwise it doesn't function.

Well, our last questions had to do with systematic relations of institutional reality and also with the priority of the brute over the institutional facts. The answer to the first of these questions is this: The reason we have all this institutional ontology is to organise and regulate our lives. So there has to be a set of interlocking institutions. What I haven't had time to tell you is, all of this at bottom is about power. We are talking about how society organises power relations. It normally does it through the institution of status functions. Somebody is the boss and somebody else an employee, somebody is an elected president, somebody is defeated and so on. And all of this is

designed precisely to intersect with other elements of the society. So, in order to have money you have to have a system of rights and obligations. You have to have the ability to buy and to sell, to store value in the form of money, to receive money as payment for services rendered. So, that is the reason for the interlocking complexity. That's what we have the system for. It is designed and has developed to enable people to cope in complex social groups, in power relations.

The final question was, why is there this priority of brute facts over institutional facts? And the answer to that is, because the iterated structure of 'X counts as Y' has to bottom out somewhere. For instance, my making a contract can be derived from my signing my name, and my signing my name can be a matter of certain words being written on a page. But then you reach the point where there isn't any more 'X counts as Y'. You just have the brute fact, e.g. the marks on the page, as X-term. So institutional reality of ownership and obligations is built on top of physical reality, it has to bottom out in physical reality.

Now to conclude, I said I would like us to think of the possibility of creating a philosophy of society, where our first task would be to get an understanding of social ontology. If we got that, then, I think, it would give different cast to our political and social theories. I think that political philosophy in the West contains a large fantasy element about how we make social contracts with each other and about when people can violate or not violate the social contract. But in real life it isn't like that, in real life it's a matter of accepting or rejecting, or furthering or fighting against institutional reality. And one way to create institutional reality often is to act as if it already existed. This is how the United States was created. There was no way that a group of people could get together in Philadelphia, all of them subjects of the British Crown in a British Crown Colony and declare themselves to be an independent nation. There was no institutional structure to enable them to do that. Well they just did it. They did it and they got away with it. It helped that they had an army and had the support of the French and so on. But you can do this if you can get away with it. You can create an institutional reality just by acting as if it already existed.

One last thought I want to leave you with and that's this. In order to articulate this I have made it look much more conscious than it really is. Most of these things develop quite unconsciously, and indeed people typically are not even aware of the structure of institutional reality. It often works best when they have false beliefs about it. So there are a lot of people in the United States, who still believe that a dollar is only really money because it is backed by all that gold in Fort Knox. It's the gold in Fort Knox that makes the dollar money. This is a total fantasy, of course. The gold has nothing to do with it. And people hold other false beliefs. They believe someone is king

only because he is divinely inspired, or they believe that marriages have to be made by God in heaven, and so on. I am not trying to discourage them in these beliefs because often the institution functions best when people hold false beliefs about it. But I think as philosophers we must, as a first step in understanding social reality, and as our first step in creating a philosophy of society, understand the basic ontology of social reality.