

What is an institution?

JOHN R. SEARLE

University of California, Berkeley

1. Economics and institutions

When I was an undergraduate in Oxford, we were taught economics almost as though it were a natural science. The subject matter of economics might be different from physics, but only in the way that the subject matter of chemistry or biology is different from physics. The actual results were presented to us as if they were scientific theories. So, when we learned that savings equals investment, it was taught in the same tone of voice as one teaches that force equals mass times acceleration. And we learned that rational entrepreneurs sell where marginal cost equals marginal revenue in the way that we once learned that bodies attract in a way that is directly proportional to the product of their mass and inversely proportional to the square of the distance between them. At no point was it ever suggested that the reality described by economic theory was dependent on human beliefs and other attitudes in a way that was totally unlike the reality described by physics or chemistry.

Some years ago, when I published *The Construction of Social Reality*, I was aware that it had implications for the ontology of economics, but I was not aware that there had already been an important revival of the tradition of institutional economics. It would be an understatement to say that I welcome this interest in institutions; I enthusiastically support it. But I think that in the institutional literature there is still an unclarity about what exactly an institution is. What is the ontology, the mode of existence, of institutional reality? This article tries to add to this discussion.

Economics as a subject matter, unlike physics or chemistry, is largely concerned with institutional facts. Facts about money and interest rates, exchange and employment, corporations and the balance of payments, form the very heart of the subject of economics. When Lionel Robbins (1935), in a classic work, tells us that ‘Economics is a study of the disposal of scarce commodities’, he takes for granted a huge invisible institutional ontology. Two dogs fighting over a bone or two schoolboys fighting over a ball are also engaged in the ‘disposal of scarce commodities’, but they are not central to the subject matter of economics.

This article grew out of my participation in a conference on Institutional Economics at the University of Hertfordshire in 2004. I am grateful to the participants for helpful comments and I especially want to thank Geoffrey Hodgson and Tony Lawson. I also want to thank two anonymous reviewers for JOIE and most of all I thank my wife Dagmar Searle for her help.

For economics, the mode of existence of the ‘commodities’ and the mechanisms of ‘disposal’ are institutional. Given the centrality of institutional phenomena, it is somewhat surprising that institutional economics has not always been at the center of mainstream economics.

One might think that the question that forms the title of this article would long ago have been answered, not just by economists, but by the enormous number of social theorists who have been concerned with the ontology of society. I am thinking not only of such foundational figures as Max Weber, Emil Durkheim, Georg Simmel, and Alfred Schütz, but of the whole Western tradition of discussing political and social institutions that goes back to Aristotle’s *Politics*, if not earlier. You would think that by now there would be a very well-defined and worked-out theory of institutions. One reason for the inadequacy of the tradition is that the authors, stretching all the way back to Aristotle, tend to take language for granted. They assume language and then ask how human institutions are possible and what their nature and function is. But of course if you presuppose language, you have already presupposed institutions. It is, for example, a stunning fact about the Social Contract theorists that they take for granted that people speak a language and then ask how these people might form a social contract. But it is implicit in the theory of speech acts that, if you have a community of people talking to each other, performing speech acts, you already have a social contract. The classical theorists, in short, have the direction of analysis back to front. Instead of presupposing language and analyzing institutions, we have to analyze the role of language in the constitution of institutions. I am going to try to take some first steps toward this goal in this article. It is a continuation of a line of argument that I began in other works, especially *The Construction of Social Reality*, but I will draw also on my book *Rationality in Action*, as well as several articles.

In the twentieth century, philosophers learned to be very cautious about asking questions of the form, ‘What is . . .?’, as in, for example, ‘What is truth?’, ‘What is a number?’, ‘What is justice?’. The lessons of the twentieth century (though these lessons are rapidly being forgotten in the twenty-first century) suggest that the best way to approach such problems is to sneak up on them. Do not ask, ‘What is truth?’, but ask, ‘Under what conditions do we say of a proposition that it is true?’. Do not ask, ‘What is a number?’, but ask, ‘How do numerical expressions function in actual mathematical practice?’. I propose to adopt this method in addressing the question, ‘What is an institution?’. Instead of coming right out and saying at the beginning, ‘An institution is . . .’, I propose to start with statements reporting institutional facts. If we could analyze the nature of institutional facts and how they differ from other sorts of facts, then it seems to me we would be well on the way to answering our question, ‘What is an institution?’.

In some intuitively natural sense, the fact that I am an American citizen, the fact that the piece of paper in my hand is a 20 dollar bill, and the fact that I

own stock in AT&T, are all institutional facts. They are institutional facts in the sense that they can only exist given certain human institutions. Such facts differ from the fact, for example, that at sea level I weigh 160 pounds, or that the Earth is 93 million miles from the sun, or that hydrogen atoms have one electron. Of course, in order to *state* the fact that the earth is 93 million miles from the sun, we need the institution of language, including the convention of measuring distances in miles, but we need to distinguish the *statement* of this fact (which is institutional) from the *fact stated* (which is not institutional). Now, what is it about institutional facts that makes them *institutional*, and what sorts of things do they require in order to be the sorts of facts they are?

2. Observer independence, observer dependence and the objective/subjective distinction

I want to begin the investigation by making certain general distinctions. First, it is essential to distinguish between those features of the world that are totally independent of human feelings and attitudes, observer independent features, and those features of the world that exist only relative to human attitudes. Observer independent features of the world include force, mass, gravitational attraction, photosynthesis, the chemical bond, and tectonic plates. Observer relative features of the world include money, government, property, marriage, social clubs, and presidential elections. It is important to see that one and the same entity can have both observer independent features and observer dependent features, where the observer dependent features depend on the attitudes of the people involved. For example, a set of movements by a group of people constitutes a football game, not just in virtue of the physical trajectories of the bodies involved, but also in virtue of the attitudes, intentions, and so on of the participants and the set of rules within which they are operating. Football games are observer relative; the trajectories of human bodies are observer independent. I hope it is obvious that most of the phenomena we discuss in economics, such as money, financial institutions, corporations, business transactions, and public offerings of stock are all observer relative. One can say that, in general, the natural sciences are concerned with observer independent phenomena and the social sciences with observer relative phenomena.

A rough test for whether or not a phenomenon is observer independent or observer relative is: could the phenomenon have existed if there had never been any conscious human beings with any intentional states? On this test, tectonic plates, gravitational attraction, and the solar system are observer independent and money, property, and government are observer relative. The test is only rough-and-ready, because, of course, the consciousness and intentionality that serve to create observer relative phenomena are themselves observer independent phenomena. For example, the fact that a certain object is money is observer relative; money is created as such by the attitudes of observers and participants

in the institution of money. But those attitudes are not themselves observer relative; they are observer independent. I think this thing in front of me is a 20 dollar bill, and, if somebody else thinks that I do not think that, he or she is just mistaken. My attitude is observer independent, but the reality created by a large number of people like me having such attitudes, depends on those attitudes and is therefore observer dependent. In investigating institutional reality, we are investigating observer dependent phenomena.

A second distinction we need is between different kinds of objectivity and subjectivity. Part of our puzzle is to explain how we create, out of subjective attitudes such as beliefs and intentions, a reality of corporations, money, and economic transactions, about which we can make objectively true statements. But there is an ambiguity in the objective–subjective distinction. Because objectivity and subjectivity loom so large in our intellectual culture, it is important to get clear about this distinction at the beginning of the investigation. We need to distinguish the *epistemic* sense of the objective–subjective distinction from the *ontological* sense. Thus, for example, if I say ‘Van Gogh died in France’, that statement can be established as true or false as a matter of objective fact. It is not just a matter of anybody’s opinion. It is epistemically objective. But if I say, ‘Van Gogh was a better painter than Manet’, well that is, as they say, a matter of opinion or judgment. It is not a matter of epistemically objective fact, but is rather a matter of subjective opinion. Epistemically objective statements are those that can be established as true or false independently of the feelings and attitudes of the makers and interpreters of the statement. Those that are subjective depend on the feelings and attitudes of the participants in the discourse. Epistemic objectivity and subjectivity are features of *claims*. But in addition to this sense of the objective/subjective distinction, and in a way the foundation of that distinction, is an ontological difference. Some entities exist only insofar as they are experienced by human and animal subjects. Thus, for example, pains, tickles and itches, and human and animal mental events and processes generally, exist only insofar as they are experienced by human or animal subjects. Their mode of existence requires that they be experienced by a human or animal subject. Therefore, we may say they have a *subjective* ontology. But, of course, most of the things in the universe do not require being experienced in order to exist. Mountains, molecules, and tectonic plates, for example, exist and would exist if there had never been any humans or animals. We can say that they have an *objective* ontology, because they do not need to be experienced by a conscious subject in order to exist.

It is important to emphasize that the ontological subjectivity of a *domain* of investigation does not preclude epistemic objectivity in the *results* of the investigation. We can have an objective science of a domain that is ontologically subjective. Without this possibility there would be no social sciences. In light of these two distinctions, we might say that one way to pose our problem for this discussion is to explain how there can be an epistemically objective institutional reality of

money, government, property, and so on, given that this reality is in part constituted by subjective feelings and attitudes and, thus, has a subjective ontology.

With these two distinctions in mind, the distinction between observer relative and observer independent features of reality, and the distinction between the ontological sense of the objective/subjective distinction and the epistemic sense of that distinction, we can place our present discussion within the larger context of contemporary intellectual life. We now have a reasonably clear idea about how the universe works, and we even have some idea about how it works at the micro level. We have a pretty good account of basic atomic and subatomic physics, we think we have a good understanding of the chemical bond, we even have a pretty well-established science of cellular and molecular biology, and we are increasing our understanding of evolutionary processes. The picture that emerges from these domains of investigation is that the universe consists entirely of entities we find it convenient to call particles (even though, of course, the word ‘particle’ is not quite right). These exist in fields of force and are typically organized into systems, where the internal structure and the external boundaries of the system are set by causal relations. Examples of systems are water molecules, galaxies, and babies. Some of those systems are composed in large part of big carbon-based molecules and are the products of the evolution of our present plant and animal species. Now here is our general question, and here is its bearing on the social sciences. How can we accommodate a certain conception we have of ourselves as conscious, mindful, rational, speech act performing, social, political, economic, ethical, and free-will possessing animals in a universe constructed entirely of these mindless physical phenomena? It is not obvious that we can make all our self-conceptions consistent with what we know from physics, chemistry, and biology about how the world is anyhow. We might, for example, in the end, have to give up our belief in free will. But since our self-conception is pretty well established and is pretty well substantiated by thousands of years of human experience, we are reluctant to give up any central portions of it without some very powerful reasons for doing so. The investigation in this article is focused on one small part of that larger problem. How can there be a social and institutional reality, including economic reality, within a universe consisting entirely of physical particles in fields of force?

3. The special theory of the logical structure of institutional facts: X counts as Y in C

I will be very brief in this section, because for the most part it will be a straight summary of material that I have previously published in *The Construction of Social Reality*.

Though the structure of actual human societies is immensely complicated, the underlying principles, I believe, are rather simple. There are three primitive

notions necessary to explain social and institutional reality. (There is a fourth, what I call the Background, that I will not go into here.)

Collective intentionality

The first notion we need is that of collective intentionality. In order to explain this notion, I have to say a little bit about intentionality in general. ‘Intentionality’ is a word that philosophers use to describe that feature of the mind by which it is directed at, or about, or of, or concerns, objects and states of affairs in the world. Thus, beliefs, hopes, fears, desires, and the emotions generally can in this technical sense be said to be intentional. It is important to emphasize that intentionality does not imply any special connection with intending, in the ordinary sense in which I intend to go to the movies tonight. Rather, intentionality is a very general notion having to do with the directedness of the mind. Intending in the ordinary sense is simply a special case of intentionality in this technical sense, along with belief, desire, hope, fear, love, hate, pride, shame, perception, disgust, and many others.

Now given that we all have intentional states in this sense – we all have hopes, beliefs, desires, fears, and so on – we need to discuss the role of intentionality in human social groups. It is a remarkable property that humans and many other animal species have that they can engage in cooperative behavior. Obvious examples are playing in an orchestra or playing team sports or simply engaging in a conversation. In such cases one does act individually, but one’s individual actions – playing the violin part, for example, or passing the ball to another player – are done as part of the collective behavior. Sometimes there is even cooperative behavior across species as, for example, to take a simple case, when my dog and I go for a walk together. When I am engaged in collective action, *I* am doing what I am doing as part of *our* doing what we are doing. In all of these cases, an agent is acting, and doing what he or she does, only as part of a collective action. It is an extremely complicated question how exactly the intentionality of the individual relates to the collective intentionality in such cases, but I have discussed it elsewhere, and I will not go into it here (Searle, 1990).

Collective intentionality covers not only collective intentions but also such other forms of intentionality as collective beliefs and collective desires. One can have a belief that one shares with other people and one can have desires that are shared by a collectivity. People cooperating in a political campaign typically desire together that their candidate will win, and in a church, the people reciting the Nicene Creed are expressing their collective faith.

Collective intentionality is the basis of all society, human or animal. Humans share with many species of animals the capacity for collective intentionality and thus the capacity to form societies. Indeed, I will define a social fact as any fact involving the collective intentionality of two or more agents. Our problem, then, is to specify what is special about human collective intentionality that enables us to create special forms of social reality that go beyond the general animal

forms. Both the Supreme Court making a decision and a pack of wolves hunting a sheep are engaged in collective intentionality and, thus, are manifesting social facts. Our question is, what is the difference between the general class of social facts and the special sub-class that constitute institutional facts?

The assignment of function

A second notion we need is that of the assignment of function. Again, human beings have a capacity that they share with some, though this time with not very many, other species of animals, the capacity to impose functions on objects where the object does not have the function, so to speak, intrinsically but only in virtue of the assignment of function. Tools are the obvious case. Humans are tool-using animals *par excellence*, but, of course, other animals have tools as well. Beaver dams and birds' nests are two obvious examples. And in some cases animals are even capable of discovering useful tools, when the use of the object as a tool is not already programmed into the animals as part of their genetic endowment. Think of Köhler's apes, for example. Assigned functions are observer relative.¹

If you combine these two, collective intentionality and the assignment of function, it is easy to see that there can be collective assignments of function. Just as an individual can use a stump as a stool, so a group can use a large log as a bench.

Status functions

The third item we need, to account for the move from social facts to institutional facts, is a special kind of assignment of function where the object or person to whom the function is assigned cannot perform the function just in virtue of its physical structure, but rather can perform the function only in virtue of the fact that there is a collective assignment of a certain *status*, and the object or person performs its function only in virtue of collective acceptance by the community that the object or person has the requisite status. These assignments typically take the form *X counts as Y*. For example, such and such a move in a football game counts as scoring a touchdown. Such and such a set of procedures counts as the election of a president of the United States. Such and such a position in chess counts as checkmate. These exhibit the general form of the assignment of status function, *X counts as Y*, or, more typically, *X counts as Y in context C*. In all of these cases, the *X* term identifies certain features of an object or person or state of affairs, and the *Y* term assigns a special status to that person, object, or state of affairs. Human beings have a capacity which, as far as I can tell, is not possessed by any other animal species, to assign functions to objects where the objects cannot perform the function in virtue of their physical structure alone, but only

¹ I think in fact that all functions are assigned and thus all functions are observer relative, but the general point is not essential to this article; so I just state the obvious fact that assigned functions are relative to the assignment and hence observer relative.

in virtue of the collective assignment or acceptance of the object or person as having a certain *status* and with that status a function. Obvious examples are money, private property, and positions of political leadership. In every case, the object or person acquires a function which can be performed only in virtue of the collective acceptance of the corresponding status.

I like to illustrate the distinction between status functions and other kinds of functions with a little parable. Imagine a tribe that builds a wall around its collection of huts, and imagine that the wall keeps members of the tribe in and intruders out, since it is difficult to get over the wall without the tolerance of the members of the tribe. But imagine that the wall decays to the point where it is nothing more than a line of stones, yet let us suppose that the people involved continue to – and watch this vocabulary closely – *recognize* the line of stones as a *boundary*. They recognize that they are not *supposed* to cross unless *authorized* to do so. Now, we are supposing that the wall, though it is no longer a large physical structure but simply a line of stones, continues to perform the same function that it did before, but this time not in virtue of its physical structure, but in virtue of the fact that the people involved continue to accept the line of stones as having a certain status. It has the status of a boundary, and people behave in a way that they regard as appropriate for something that they accept as a boundary. The line of stones has a function not in virtue of its physical structure, but in virtue of the collective assignment of a status, and with that status, a function which can only be performed in virtue of the collective acceptance of the object as having that status. I propose to call such functions *status functions*.

As this example is intended to make clear, the transition from physical function to status function can be gradual, and there may be no exact point at which we can say, the status function begins and the physical function ends. The vocabulary is revealing. ‘You can’t cross that’ can mean either ‘It is too high’ or ‘It is not allowed’ (or both).

The general logical form of the imposition of status functions is, as I said, X counts as Y in C , though I will point out some exceptions later.

It might seem that this is a very feeble apparatus with which to construct institutional structures; surely the whole thing could come tumbling down at any moment. How can it do as much work as it apparently does? The answer, or at least part of the answer, is that this structure has certain purely formal properties that give it enormous scope. The first is that it iterates upward indefinitely. So, for example, when I make certain sounds through my mouth, making those sounds counts as uttering sentences of English; but uttering those sentences of English counts as making a promise; and, in that context, making a promise counts as undertaking a contract. Making that kind of contract in that context counts as getting married, and so on upward. Notice the logical form of this: X_1 counts as Y_1 . But $Y_1 = X_2$ counts as Y_2 . And $Y_2 = X_3$ counts as Y_3 , and so on upward indefinitely.

Secondly, the whole system operates laterally as well as vertically. Thus, I do not just own property, but I own property as a citizen of the city of Berkeley in the county of Alameda in the State of California in the United States of America. Locked into this institutional structure I have all sorts of rights and obligations. For example I have to pay taxes to all four of those entities I just named, and all four are under obligations to provide me with all sorts of social services. I acquire various rights and duties as a property owner, and these interlock with other social institutions.

When the procedure or practice of counting *X* as *Y* becomes regularized it becomes a rule. And rules of the form *X counts as Y in C* are then constitutive of institutional structures. Such rules differ from regulative rules, which are typically of the form 'Do *X*', because regulative rules regulate activities which can exist independently of the rule. Constitutive rules not only regulate but rather constitute the very behavior they regulate, because acting in accordance with a sufficient number of the rules is constitutive of the behavior in question. An obvious contrast is between the regulative rules of driving, such as drive on the right-hand side of the road and the constitutive rules of chess. Driving can exist without the regulative rule requiring right or left; the rule regulates an antecedently existing activity. But chess cannot exist without the rules, because behaving in accordance with (at least a sufficient subset of) the rules is constitutive of playing chess.

Now I want to make a very strong claim. The institutional ontology of human civilization, the special ways in which human institutional reality differs from the social structures and behavior of other animals, is a matter of status functions imposed according to constitutive rules and procedures. Status functions are the glue that holds human societies together. Think not only of money, property, government, and marriage, but also of football games, national elections, cocktail parties, universities, corporations, friendships, tenure, summer vacations, legal actions, newspapers, and industrial strikes. Though these phenomena exhibit enormous variety, their underlying ontology reveals a common structure. The analogy with the natural world is obvious. Bonfires and rusting shovels look quite different, but the underlying mechanism that produces them is exactly the same: oxidization. Analogously, presidential elections, baseball games, and 20 dollar bills look different, but the underlying mechanism that produces them is the same: the assignment of status functions with their accompanying deontologies according to constitutive rules. (I will say more about deontology in a moment.)

We are now close to being able to give a provisional answer to the question which forms the title of this paper: 'What is an institution?' We have substituted for that question, the question: 'What is an institutional fact?' And I have claimed that these facts typically require structures in the form of constitutive rules *X counts as Y in C* and that institutional facts only exist in virtue of collective acceptance of something having a certain status, where that status carries functions that cannot be performed without the collective acceptance of

the status. This I am claiming is the glue that holds society together. There is a gradual transition from informal but accepted assignments of status functions to full-blown established institutions with codified constitutive rules, but in both cases the crucial element of deontology is present, as we will see. Furthermore, the notion of ‘collective acceptance’ is intended to be vague, because I need to mark a continuum that goes from grudgingly going along with some social practice to enthusiastic endorsement of it.

As a preliminary formulation, we can state our conclusions so far as follows: an institutional fact is any fact that has the logical structure *X counts as Y in C*, where the *Y* term assigns a status function and (with few exceptions) the status function carries a deontology.² An institution is any system of constitutive rules of the form *X counts as Y in C*. Once an institution becomes established, it then provides a structure within which one can create institutional facts.

Our original aim was to explain how the ontology of institutions fits into the more basic ontology of physics and chemistry and we have now done that: one and the same phenomenon (object, organism, event, etc.) can satisfy descriptions under which it is non-institutional (a piece of paper, a human being, a series of movements) and descriptions under which it is institutional (a 20 dollar bill, the president of the United States, a football game). An object or other phenomenon is part of an institutional fact, *under a certain description of that object or phenomenon*.

I am leaving out an enormous number of complexities for the sake of giving a simple statement of the bare bones of the ontology in question.

4. Status functions and deontic powers

How does it work, how does a set of status functions, deriving from systems of constitutive rules, function in the operation of society? The essential role of human institutions and the purpose of having institutions is not to constrain people as such, but, rather, to create new sorts of power relationships. Human institutions are, above all, *enabling*, because they create power, but it is a special kind of power. It is the power that is marked by such terms as: rights, duties, obligations, authorizations, permissions, empowerments, requirements, and certifications. I call all of these *deontic powers*. What distinguishes human societies from other animal societies, as far as I can tell, is that human beings are capable of a deontology which no other animal is capable of. Not all deontic power is institutional, but just about all institutional structures are matters of deontic power. Think of anything you would care to mention – private property, government, contractual relationships, as well as such informal relationships as friendship, family, and social clubs. All of these are matters of rights,

² One class of exceptions are honorific status functions, where the recipient has the honor or dishonor of the new status, but no real powers. Honorary degrees, knighthoods, presidential medals, and beauty contest victories are all examples.

duties, obligations, etc. They are structures of power relationships. Often the institutional facts evolve out of the natural facts. Thus, there is a biological family consisting of parents and their biological offspring. But humans have imposed on this underlying biology a rather elaborate formal and informal institutional structure, involving the respective statuses of the mother, the father, and the children. In so-called ‘extended families’ authority relationships and other status functions may include not only the parents and children but sundry other relatives. Furthermore, given the institutional structures, one may have families with parents and children where no one is biologically related to anyone else.

But that only forces the question back a bit: how exactly do these power relations function? The answer, which again is essential to understanding society, is that institutional structures create desire-independent reasons for action. To recognize something as a duty, an obligation, or a requirement is already to recognize that you have a reason for doing it which is independent of your inclinations at the moment.

It might seem paradoxical that I talk about institutional reasons for action as ‘desire-independent reasons for action’, because, of course, many of these are precisely the foci of very powerful human desires. What is more a field for human desire than money? Or political power? I think this question raises a deep issue: By creating institutional reality, we increase human power enormously. By creating private property, governments, marriages, stock markets, and universities, we increase the human capacity for action. But the possibility of having desires and satisfying them within these institutional structures – for example, the desire to get rich, to become president, to get a Ph.D., to get tenure – all presuppose that there is a recognition of the deontic relationships. Without the recognition, acknowledgment, and acceptance of the deontic relationships, your power is not worth a damn. It is only worthwhile to have money or a university degree or to be president of the United States if other people recognize you as having this status, and recognize that status as giving desire-independent reasons for behaving in certain ways. The general point is clear: the creation of the general field of desire-based reasons for action presupposes the acceptance of a system of desire-independent reasons for action. This is true both of the immediate beneficiaries of the power relationships (for example, the person with the money or the person who has won the election) and of the other participants in the institution.

5. Language as the fundamental social institution

I suggested earlier that one reason that traditional accounts of institutions, both in institutional economics and elsewhere, are incomplete is that they all take language for granted. It is essential to see in exactly what respect language is the fundamental social institution in order that you can see the logical structure of the other social institutions. It is intuitively obvious, even pre-theoretically,

that language is fundamental in a very precise sense: you can have language without money, property, government, or marriage, but you cannot have money, property, government, or marriage without language. What is harder to see is the constitutive role of language in each of these and, indeed, in all social institutions. Language does not just describe a preexisting institutional reality but is partly constitutive of that reality, in ways I need to explain.

It seems intuitively right to say that you can have language without money, but not money without language. But now we need to state exactly how and why language is essential. The general form of status functions is that we impose a status and with it a function on something that cannot perform that function in virtue of its physical structure alone. It can only function if it is assigned a status function, and in that respect it differs from other tools. Think of the difference between a knife and a 20 dollar bill. The knife will cut just in virtue of its physical structure. But the 20 dollar bill will not buy just in virtue of its physical structure. It can only function as money if it is recognized, accepted, and acknowledged as valid currency. The knife function can exist for anybody capable of exploiting the physics, but the status function can only exist if there is collective representation of the object as having the status that carries the function. *A status function must be represented as existing in order to exist at all, and language or symbolism of some kind provides the means of representation.* You can explore the physics of the X terms as much as you like, but you cannot read off the status function as you can read off physical functions, because there is nothing in the X term physically that by itself carries the status function. The piece of paper is only money, the man is only president, insofar as the piece of paper is represented as money and the man is represented as president. But now, if there are to be these representations, there must be some medium of representation, and that medium is language or symbolism in the broadest sense. We must have some means of representing the fact that this stuff is money or that that man is president in order that the stuff can acquire the status of money and the man can acquire the status of a president. No representation, no status function.

This is why pre-linguistic animals cannot have an institutional reality. My dog has very good vision, indeed much better than mine. But I can still see things he cannot see. We can both see, for example, a man crossing a line carrying a ball. But I can see the man score a touchdown and the dog cannot. We should reflect on this, because it is a very deep and important point. Why is it, exactly, that my dog cannot see a man score a touchdown? Is his vision not good enough? Well, we might train the dog to bark whenever a man crosses a white line in possession of a ball, but that is still not yet seeing a touchdown. To see a touchdown scored he would have to be able to represent what is happening as the scoring of a touchdown, and without language he cannot do that.

This also leads to very deep considerations about the ontology of institutional reality and its relation to cognition. In order to perceive the man score a touchdown, or to perceive that he is president, or to perceive that this is a

dollar bill, we have to think at two different levels at once. We have to be able to see the physical movements but see them as a touchdown, to see the piece of paper but to see it as a dollar bill, to see the man but to see him as a leader or as president of the United States. Now this looks like it is a standard form of *seeing as*, of the sort discussed by Wittgenstein, and of a kind which is common in Gestalt psychology; but in fact it differs sharply from them. It is not at all like the ambiguous duck/rabbit figure that can be seen either as a duck or as a rabbit. It is different because we have to think up a level. We have to think from the brute level up to the institutional level, and the capacity to think at different levels enters into the actual cognitive processes of our perception. I literally see a 20 dollar bill, I do not just see paper. I literally see a touchdown, I do not just see a man carrying a ball across a line. But the cognitive capacity to see these things requires a linguistic or symbolic capacity. To put it very crudely: no language, no status functions. No status functions, no institutional deontology.

Let us explore these ideas by going through some of the steps in which language is involved in the constitution of institutional reality.

We have the capacity to count things as having a certain status, and in virtue of the collective acceptance of that status, they can perform functions that they could not perform without that collective acceptance. The form of the collective acceptance has to be in the broadest sense linguistic or symbolic, because there is nothing else there to mark the level of status function. There is nothing to the line and the man and the ball that counts as a touchdown, except insofar as we are prepared to count the man with the ball crossing the line as the scoring of a touchdown. We might put these points in the most general form by saying that language performs at least the following four functions in the constitution of institutional facts.

First, the fact can only exist insofar as it is represented as existing and the form of those representations is in the broadest sense linguistic. I have to say ‘in the broadest sense’, because I do not mean to imply that full-blown natural languages with relative clauses, iterated modal operators, and quantificational scope ambiguities are essential to the constitution of institutional reality. I do not believe they are. Rather, I believe that unless an animal can symbolize something as having a status, which it does not have in virtue of its physical structure, then the animal cannot have institutional facts, and that those institutional facts require some form of symbolization – what I am calling language in the broad sense. The symbolization has to carry the deontic powers, because there is nothing in the sheer physical facts that carries the deontology by itself. No language, no deontology.

Secondly, and this is really a consequence of the first point, the forms of the status function in question are almost invariably matters of deontic powers. They are matters of rights, duties, obligations, responsibilities, etc. Now, pre-linguistic animals cannot recognize deontic powers because without having some linguistic means of representation they cannot represent them. Let me state this point with

as much precision as I can. Animal groups can have an alpha male and an alpha female, and other members of the group can make appropriate responses to the alpha male and the alpha female, but this hierarchy is not constituted by a system of rights, duties, obligations, etc. Indeed, the terms ‘alpha male’ and ‘alpha female’ are invented by ethologists from a third-person point of view to describe animal behavior, but the animal does not think, ‘I have to recognize his authority because he is the alpha male.’ What the animals lack is the deontology – the obligations, requirements, duties, etc. that go with the recognition of higher and lower status. For those obligations, requirements, and duties to exist, they have to be represented in some linguistic or symbolic form. Again, when a dog is trained to obey commands, he is just taught to respond automatically to certain specific words or other signals.

(By the way, I frequently make remarks about animal capacities. I do not think we know enough about animal capacities to be completely confident in the attributions we make, especially to the primates. But, and this is the point, if it should turn out that some of the primates are on our side of the divide rather than on the side of the other animals, in the sense that they have deontic powers and deontic relationships, then so much the better for them. In this article, I am not asserting the superiority of our species, rather I am trying to mark a conceptual distinction, and I assume, on the basis of what little I know, that where deontology is concerned we are on one side and other animals are on the other side of the dividing line.)

Third, the deontology has another peculiar feature. Namely, it can continue to exist after its initial creation and indeed even after all the participants involved have stopped thinking about the initial creation. I make a promise today to do something for you next week, and that obligation continues even when we are all sound asleep. Now, that can only be the case if that obligation is represented by some linguistic means. In general, one can say this: human societies require a deontology, and the only way they can have this is by having language. To repeat, no language, no deontology.

Fourth, a crucial function of language is in the recognition of the institution as such. It is not merely particular cases within the institution that this is my property, that that was a football game, but rather, in order that this should be a case of property or that a case of a football game, one has to recognize the institutions of property and football games. Where institutional reality is concerned, the particular instances typically exist as such because they are instances of a general institutional phenomenon. Thus, in order for me to own a particular item of property or to have a particular dollar bill, there has to be a general institution of private property and money. Exceptions to this are cases where an institution is being created *de novo*. But the general institutions, in which the particular instances find their mode of existence, can only exist insofar as they are recognized and that recognition has to be symbolic, linguistic in the most general sense.

6. Steps toward a general theory of social ontology. We accept (S has power (S does A))

I want now to discuss some of the further developments in the theory of institutional reality since the publication of *The Construction of Social Reality*. I want to mention two such developments. First, in the original statement of the theory, I pointed out that, in order for status functions to be recognized, there typically have to be some sorts of *status indicators*, because there is nothing in the person or the object itself that will indicate its status, since the status is only there by collective acceptance or recognition. Thus, we have policemen's uniforms, wedding rings, marriage certificates, drivers' licenses, and passports, all of which are status indicators. Many societies find that they cannot exist without status indicators, as, for example, the proliferation of identity cards and driver's licenses will attest. However, Hernando De Soto (2000) pointed out an interesting fact. Sometimes the status indicators, as issued by an official agency (where the agency is itself a higher-level set of status functions), acquire a kind of life of their own. How is this so? He points out that in several underdeveloped countries, many people own land, but because there are no property deeds, because the owners of the property do not have title deeds to the property, they are, in effect, what we would call squatters; they do not have status indicators. This has two consequences of enormous social importance. First, they cannot be taxed by the governing authorities because they are not legally the holders of the property, but, secondly and just as importantly, they cannot use the property as capital. Normally, in order for a society to develop, the owners of property have to be able to go to the bank and get loans against their property in order to use the money to make investments. But in countries such as, for example, Egypt, it is impossible for the vast amount of private property to be used as collateral for investments because so much of this property is held without the benefit of a property deed. The owners of the property are in effect squatters, in the sense that they do not legally own the property, though they live in a society where their status function is acknowledged and generally recognized and hence, on my account, continues to exist and generate deontic powers. But the deontic powers stop at the point where the larger society requires some official proof of the status functions. Thus, without official documentation, they lack full deontic powers. Collective recognition is not enough. There has to be official recognition by some agency, itself supported by collective recognition, and there have to be status indicators issued by the official agency.

A second and equally important development was pointed out to me by Barry Smith. He pointed out that there are some institutions that have what he calls 'free-standing Y terms', where you can have a status function, but without any physical object on which the status function is imposed. A fascinating case is corporations. The laws of incorporation in a state such as California enable a status function to be constructed, so to speak, out of thin air. Thus, by a kind of

performative declaration, the corporation comes into existence, but there need be no physical object which is the corporation. The corporation has to have a mailing address and a list of officers and stock holders and so on, but it does not have to be a physical object. This is a case where following the appropriate procedures counts as the creation of a corporation and where the corporation, once created, continues to exist, but there is no person or physical object which becomes the corporation. New status functions are created among people – as officers of the corporation, stockholders, and so on. There is indeed a corporation as *Y*, but there is no person or physical object *X* that counts as *Y*.

An equally striking example is money. The paradox of my account is that money was my favorite example of the ‘*X* counts as *Y*’ formula, but I was operating on the assumption that currency was somehow or other essential to the existence of money. Further reflection makes it clear to me that it is not. You can easily imagine a society that has money without having any currency at all. And, indeed, we seem to be evolving in something like this direction with the use of debit cards. All you need to have money is a system of recorded numerical values whereby each person (or corporation, organization, etc.) has assigned to him or her or it a numerical figure which shows at any given point the amount of money they have. They can then use this money to buy things by altering their numerical value in favor of the seller, whereby they lower their numerical value, and the seller acquires a higher numerical value. Money is typically redeemable in cash, in the form of currency, but currency is not essential to the existence or functioning of money.

How can such things function if there is no physical object on which the status function is imposed? The answer is that status functions are, in general, matters of deontic power, and, in these cases, the deontic power goes directly to the individuals in question. So my possession of a queen in the game of chess is not a matter of my having my hands on a physical object, it is rather a matter of my having certain powers of movement within a formal system (and the formal system is ‘the board’, though it need not be a physical board) relative to other pieces. Similarly, my having a thousand dollars is not a matter of my having a wad of bills in my hand but my having certain deontic powers. I now have the right, i.e. the *power*, to buy things, which I would not have if I did not have the money. In such cases, the real bearer of the deontology is the participant in the economic transactions and the player in the game. The physical objects, such as chess pieces and dollar bills, are just markers for the amount of deontic power that the players have.

In the early part of *The Construction of Social Reality* I said that the basic form of the institutional fact was *X counts as Y in C* and that this was a form of the constitutive rule that enables us to create institutional facts. But my later formulation in the book gives us a much more general account. I said that the basic power creation operator in society is *We accept (S has power (S does A))*; and that we could think of the various forms of power as essentially Boolean

operations on this basic structure, so, for example, to have an obligation is to have a negative power. What then, exactly, is the relationship between the two formulae *X counts as Y in C* and *We accept (S has power (S does A))*? The answer is that, of course, we do not just accept that somebody has power, but we accept that they have power in virtue of their institutional status. For example, satisfying certain conditions makes someone president of the United States. This is an example of the *X counts as Y in C* formula. But, once we accept that someone is president of the United States, then we accept that he has the power to do certain things. He has the positive power to command the armed forces, and he has the negative power, i.e. the obligation, to deliver a state of the union address. He has the *right* to command the armed forces, and he has the *duty* to deliver the address. In this case we accept that *S* has power (*S does A*) because $S = X$, and we have already accepted that *X* counts as *Y*, and the *Y* status function carries with it the acknowledged deontic powers.

Continuing with the example of the corporation, we can say that *so* and *so* counts as the president of the corporation and *such* and *such* people count as the stockholders. This is an example of the *X counts as Y in C* formulation, but, of course, the whole point of doing that is to give them powers, duties, rights, responsibilities, etc. They then instantiate the *we accept (S has power (S does A))* formula. But to repeat a point made earlier, the corporation itself is not identical with any physical object or any person or set of persons. The corporation is, so to speak, created out of nothing. The president is president *of* the corporation, but he is not identical with the corporation. The reasons for doing this are famous. By creating a so-called 'fictitious person' we can create an entity that is capable of entering into contractual relationships and capable of buying and selling, making a profit, and incurring debts, for which it is liable. But the officers and stockholders, are not personally liable for the debts of the corporation. This is an important breakthrough in human thought. So, what amounts to the corporation when we set it up? It is not that there is an *X* that counts as the corporation, but, rather, that there is a group of people involved in legal relationships, thus *so* and *so* counts as the president of the corporation, *so* and *so* counts as a stockholder in the corporation, etc., but there is nothing that need count as the corporation itself, because one of the points of setting up the corporation was to create a set of power relationships without having to have the accompanying liabilities that typically go with those power relationships when they are assigned to actual human individuals.

I regard the invention of the limited liability corporation, like the invention of double-entry bookkeeping, universities, museums, and money, as one of the truly great advances in human civilization. But the greatest advance of all is the invention of status functions, of which these are but instances. It is not at all necessary that there should exist status functions. Non-human animals do not appear to have them. But without them, human civilization, as we think of it, would be impossible.

7. Different kinds of 'institutions'

I have not been attempting to analyze the ordinary use of the word 'institution'. I do not much care if my account of institutional reality and institutional facts matches that ordinary usage. I am much more interested in getting at the underlying glue that holds human societies together. But let us consider some other sorts of things that might be thought of as institutions.

I have said that the fact that I am an American citizen is an institutional fact, but how about the fact that today is the 24 September 2004? Is that an institutional fact? What does the question ask? At least this much. Does identifying something as 24 September 2004 collectively assign a status function that carries with it a deontology? So construed the answer is no. In my culture there is no deontology carried by the fact that today is 24 September. In that respect, '24 September 2004' differs from 'Christmas Day', 'Thanksgiving', or, in France, '14 July'. Each of these carries a deontology. If it is Christmas Day, for example, I am *entitled* to a day off, and collective intentionality in my community supports me in this entitlement. Since every day is some Saint's Day, there is presumably a subgroup for which 24 September is an important Saint's Day that carries an institutional deontology, but I am not in that subgroup.

I think there is a sense of the word 'institution' in which the Christian calendar or the Mayan calendar are a kind of institution (both of them were, after all, *instituted*), but it is not the kind of institution that I am attempting to analyze. A calendar is rather a verbal system for naming units of time – days, months, and years – and indicating their relationships. Similarly with other verbal systems. Different societies have different color vocabularies, but that does not make the fact that the cloth in front me is magenta into an institutional fact. Similar remarks could be made about systems of weights and measures. The fact that I weigh 160 pounds is the same fact as the fact that I weigh 72 kilos, even though this same fact can be stated using different systems of measuring weights.

More interesting to me are those cases where the facts in question are on the margin of being institutional. I think that the fact that someone is my friend is an institutional fact because friendship carries *collectively recognized* obligations, rights, and responsibilities. But how about the fact that someone is a drunk, a nerd, an intellectual, or an underachiever? Are these institutional concepts and are the corresponding terms institutional facts? Not as I am using these expressions, because there is no collectively recognized deontology that goes with these. Of course, if the law or custom establishes criteria under which somebody is a recognized drunk and imposes penalties as well as entitlements for this status, then being a drunk becomes a status function. X counts as Y. And, again, I might personally feel that, as an intellectual, I have certain sorts of obligations, but this is not yet an institutional phenomenon unless there is some collective recognition of my status and of these obligations. When I pointed out in a lecture that being a nerd was not a status function, one of my students

told me that in his high school it definitely was, because as the class nerd he was expected to help other students with their homework. He was under certain sorts of collectively recognized obligations.

Another sort of ‘institution’ that I am not attempting to describe is massive forms of human practices around certain subject matters that do not *as such* carry a deontology, even though there are lots of deontologies within the practices. So, for example, there are series of practices that go with what we call ‘science’ or ‘religion’ or ‘education’. Does that make science, religion, and education into institutions? Well, we are using institution as a technical term anyway, and it is open to us if we want to call these institutions, but I think it is very important that we not confuse science, education, and religion with such things as money, property, government, and marriage. Within such gross human practices as science, religion, and education there are, indeed, institutions and plenty of institutional facts. Thus, for example, the National Science Foundation is an institution, as is the University of California or the Roman Catholic Church. And the fact that Jones is a scientist, Smith a professor, and Brown a priest again are all institutional facts. Why then are not science, religion, and education institutions? To ask of any word *W*, Does *W* name an institution? is to ask at least the following:

1. Is *W* defined by a set of constitutive rules?
2. Do those rules determine status functions, which are in fact collectively recognized and accepted?
3. Are those status functions only performable in virtue of the collective recognition and acceptance, and not in virtue of the observer-independent features of the situation alone?
4. Do the status functions carry recognized and accepted deontic powers?

So construed, ‘The National Science Foundation’ names an institution. ‘Science’ does not. The rules of scientific method, if there are such, are regulative and not constitutive. They are designed to maximize the probability of discovering the truth, not to create status functions with deontic powers. All of that is consistent with the fact that in my subculture to say that someone is a ‘scientist’ is to state an institutional fact, because it assigns a *Y* status, on the basis of meeting certain *X* criteria, that carries certain rights and responsibilities, a more or less specific deontology.

As I said before, I do not much care whether or not we want to use the word ‘institution’ for both those practices whose names specify an institutional deontology and those which do not, but it is crucial to emphasize the important underlying idea: we need to mark those facts that carry a deontology because they are the glue that holds society together.

8. Some possible misunderstandings

Each academic discipline has its own style, set of background practices, and habits. We inculcate these into our graduate students, and they are then passed

on, for the most part unconsciously, from generation to generation. There are certain special features of the cognitive style of economics as a discipline that I want to call attention to. I think these are probably, in general, very powerful intellectual resources, but they can also impede understanding when we are involved in the sort of interdisciplinary exercise in which I am currently engaged.

Models and theories

Economists typically believe in models. In my experience in dealing with economists, they often talk about ‘your model’ as if one were not trying to give a factually accurate theory about the real world but to construct a model. And, indeed, of course, in classical economic theory one typically does construct models. One makes a set of assumptions about entrepreneurs trying to maximize profits and consumers trying to maximize utility, for example, and then one deduces certain conclusions. To the extent that the assumptions are true, the conclusions will be substantiated. To the extent that the assumptions are only partly true, or allow for all kinds of exceptions and interferences from outside the assumptions, then the applicability of the model to the real world will be to that extent limited. Economists in general are not worried by these limitations, because they think that as long as the model has important predictive powers, we need not worry about whether or not it is literally true in its details.

This methodological approach can be useful for lots of purposes, but it has impeded understanding of my own views. I am not trying to construct a model; I am trying to advance a theory that states an important set of facts about how society actually works. Just as when I say I have two thumbs, that statement is not a ‘model’ of my anatomy but a literal statement of fact, so when I say institutions generate status functions, that is not a model but, if I am right, it is a true statement of fact. It is not a case of constructing a model that ignores all sorts of complicating details.

Thought experiments

Economists, in my experience, typically confuse thought experiments with empirical hypotheses. Here is an example that has come up over and over. I point out that there are desire-independent reasons for action. A classic case of this is promising; when I make a promise to do something, I have a reason for doing it which is independent of my desires. When I point this out, economists often say, ‘Yes, but you have all sorts of prudential reasons why you would keep your promise; if you did not, people would not trust you, etc.’ These are familiar arguments in philosophy, but they miss the point. One way to see that they miss the point is to construct a thought experiment. Subtract the prudential reasons, and ask yourself whether I still have a reason for keeping the promise. The answer is not an empirical hypothesis about how I would behave in a particular situation, rather it is a thought experiment designed to

show the conceptual distinction between my prudential reasons for acting and the desire-independent obligation that I recognize when I recognize something as a promise that I have made. The point is that I am not making an empirical prediction about how I would actually behave under certain circumstances, rather I am giving a conceptual analysis where the concept of a prudential reason is a different concept from the concept of a desire-independent reason. The concept of promising, by its very definition, contains the concept of a desire-independent reason. To recognize something as a valid promise is to recognize it as creating an obligation, and such obligations are desire-independent reasons for acting.

Methodological individualism

It seems to me that there is a certain amount of confusion surrounding the notion of ‘methodological individualism’. Without going into too many details, I want to state the precise sense in which the views advocated in this article are consistent with methodological individualism. The sense in which my views are methodological individualist is that all observer-independent mental reality must exist in the minds of individual human beings. There is no such thing as a group mind or an Oversoul or a Hegelian Absolute of which our particular minds are but fragments. Another way to put this point, in light of the distinctions made in this article, is to say that all observer independent intentionality is in the minds of individual human beings. I want this sense of ‘methodological individualism’ to seem quite uncontroversial. It is perfectly consistent with the idea that there are predicates true of social collectives which are not in any obvious way true of individuals. So, for example, if I say that the United States government has a huge annual deficit, that statement has implications about the behavior of individuals, but it is not the individuals that have the ‘huge annual deficit’. A second issue that this definition of methodological individualism enables me to sidestep is that concerning ‘externalism’ in the philosophy of mind. I do in fact think that mental states are entirely in the head, but many contemporary philosophers think that the contents of mental states are not in the head but include, for example, causal relations to the real world and to the surrounding society. I do not think these views are true, but I do not need to refute them for the purpose of this investigation. I simply insist that all mental reality is in the minds of individuals. This is consistent with the theory that says mental contents and hence minds are not in heads, although I happen to think that theory is false.

9. Conclusion

I have now offered at least preliminary answers to the questions posed at the beginning of this article. At the risk of repetition I will state them:

What is an institution? An institution is any collectively accepted system of rules (procedures, practices) that enable us to create institutional facts. These

rules typically have the form of *X counts as Y in C*, where an object, person, or state of affairs *X* is assigned a special status, the *Y* status, such that the new status enables the person or object to perform functions that it could not perform solely in virtue of its physical structure, but requires as a necessary condition the assignment of the status. The creation of an institutional fact is, thus, the collective assignment of a status function. The typical point of the creation of institutional facts by assigning status functions is to create deontic powers. So typically when we assign a status function *Y* to some object or person *X* we have created a situation in which we accept that a person *S* who stands in the appropriate relation to *X* is such that (*S* has power (*S* does *A*)). The whole analysis then gives us a systematic set of relationships between collective intentionality, the assignment of function, the assignment of status functions, constitutive rules, institutional facts, and deontic powers.

The theory of institutions in this article is very much work in progress, as was the earlier work on which it is based. I see the theory of institutions as still in its childhood. (Maybe not in its infancy any more, but still its childhood.) Two methodological lessons for anyone wishing to pursue it further: First, because the institutional ontology is subjective, it must always be examined from the first person point of view. Institutional facts only exist from the point of view of the participants and for that reason no external functionalist or behaviorist analysis will be adequate to account for them. You have to be able to think yourself into the institution to understand it. Second, a consequence of this analysis is that society has a logical structure. Other parts of nature – the planetary system, mitosis, and the replication of DNA, for example – do not have logical structures. Theories about such parts of nature have logical structures but not the nature itself. But society consists in part of representations and those representations have logical structures. Any adequate theory about such phenomena must contain a logical analysis of their structures.

References

- De Soto, Hernando (2000), *The Mystery of Capital: Why Capitalism Triumphs in the West and Fails Everywhere Else*, New York: Basic Books.
- Robbins, Lionel (1935), *An Essay on the Nature and Significance of Economic Science*, 2nd edn, London: Macmillan.
- Searle, John R. (1995), *The Construction of Social Reality*, London: Allen Lane.
- Searle, John R. (2001), *Rationality in Action*, Cambridge, MA: MIT Press.
- Searle, John R. (1990), 'Collective Intentions and Actions', in P. Cohen, J. Morgan, and M. E. Pollack (eds.), *Intentions in Communication*, Cambridge, MA: MIT Press; reprinted in Searle, John R. (2002) *Consciousness and Language*, Cambridge: Cambridge University Press, pp. 80–105.