

Why social sciences are natural, and why they can't

Por qué las ciencias sociales son naturales, y por qué no pueden

JESÚS ZAMORA-BONILLA

UNED, Madrid
jpsz@fsf.uned.es (ESPAÑA)

Recibido: 01.03.2011

Aceptado: 28.12.2011

RESUMEN

En este artículo se argumenta que no existen diferencias ontológicas ni epistemológicas fundamentales entre las ciencias naturales y las ciencias sociales, aunque el ámbito de lo social es tan complejo e inestable que dificulta la emergencia (y por lo tanto, el descubrimiento) de regularidades empíricas robustas y relevantes. Se presenta también, como una sugerencia acerca de cómo podrían ser «naturalizadas» las ciencias sociales, un modelo abstracto de conducta-basada-en-criterios-normativos, modelo que es coherente con los descubrimientos empíricos en ciencias cognitivas, y también puede ser implementado en simulaciones informáticas

PALABRAS CLAVE

Naturalismo, positivismo, reduccionismo, inferencialismo, normatividad, racionalidad, hermenéutica.

ABSTRACT

It is argued that there are no fundamental ontological nor epistemological differences between the natural and the social sciences, though the social realm is so complex and unstable that it makes it difficult the emergence (and, hence, the discovery) of significant robust regularities. As a suggestion of how social sciences might be 'naturalised', an abstract model of normative based behaviour is also presented, such that it is coherent with empirical discoveries in cognitive sciences and capable of being implemented in computer simulations.

KEYWORDS

Naturalism, positivism, reductionism, inferentialism, normativity, rationality, hermeneutics.

1. INTRODUCTION

Once upon a time, in a world conceptually very far from ours, there were a group of researchers devoted to the study of bacteria, or prokaryotes, who called themselves 'biologists', and a different group of researchers, who studied eukaryote organisms, and who received the name of 'organismologists'. Organismologists were very proud of the obvious differences between their discipline and that of the biologists, and flagged those differences every time they needed to show that organismology was a science wholly independent of biology, completely 'irreducible' to it, and who demanded a totally different set of concepts, principles, and research methods, far more significant than the more or less trivial similarities between both disciplines and their objects. After all, this was a matter of common sense: for example, the concept of an 'organ' or 'digestive system', or the principles of sexual reproduction, do not apply to bacteria save in a very metaphoric sense; individual organisms of the same species (for example, siblings) have usually great and important differences between them to an extent that bacteria of the same type have not; and also most multicell organisms, contrarily to prokaryotes, can often be observed and manipulated without the need of very sophisticated instruments. Organismologists accepted, as a matter of course, that eukaryote organisms had evolved millions of years ago from bacteria, that both types of beings shared a common biochemistry and genetics (though much more complex in the case of eukaryotes), and that both types of beings engaged in lots of mutual relations, so strongly that, in general, bacteria are fundamental for the life of eukaryotes, either as prey, as parasites, or as symbionts. But all this was not relevant enough to turn into ashes the obvious truth that the systems in which multicellular organisms (and even unicellular eukaryotes) consisted were so complex and radically different from the simplicity of bacteria, as to deserve a completely independent field for their systematic study. 'We would not understand anything important about organisms', leading organismologists used to proclaim, 'if we were forced to describe them as mere aggregates of individual bacteria... what very probably they are, of course'.

I think that nearly everybody that knows something about biology in our own world would easily realised that organismologists were right in some sense (eukaryote, and particularly multicellular organisms, are terribly more complex than bacteria, and 'irreducible' to these), but wrong in other sense (their fields of study are not really 'heterogeneous'). I shall argue in this paper that the situation of social scientists resisting the 'naturalisation' of their disciplines is similar to the resistance of organismologists to be considered as just a 'branch' of a deeper

and more general discipline called ‘biology’: they are right in pointing to the significant differences between ‘social systems’ and ‘(simpler?) living systems’, but these differences do not ground the conclusion that social science is not a ‘natural’ science. I will also try to argue, however, that, due to some features of social systems, the unavoidable ‘naturalistic nature’ of social sciences gives little comfort to those optimists who might have thought that social sciences could turn ‘more scientific’ by recognising their own naturalness. So, my conclusion will be something like acknowledging that, it is not that social sciences ‘should’ be naturalised, but that *they are ‘natural sciences’ as they are*¹, and that they probably cannot become *significantly* more ‘successful’ in an empirical or theoretical sense as they are now.

2. NATURALISM, POSITIVISM AND REALISM

Naturalism is usually seen by many social scientists and philosophers of the social sciences as a frightening menace, and often as a completely wrong and misguided approach. It is usually amalgamated with other philosophical monsters, such as positivism or reductionism, and at times even with ‘realism’ or ‘rationalism’ (forming with them a very poisonous combination often referred to as ‘scientism’). I will discuss in this section the problem of the connection between naturalism, reductionism and positivism, and why I think that social scientists should not worry at all at any of these three supposed Behemoths. Let’s start with positivism. So many horrible things have been said about positivism in the last fifty years, that now it passes as a matter of fact that one should not be a positivist; particularly, the story is often told that those sciences, both natural, social or hermeneutic, that tried to rigidly follow the precepts of positivism, ended in something similar to a ‘brain death’ state, with nothing interesting to discover about their realm of study. For many of these critics, as well as for some visionary scientists that tried to apply those precepts, ‘positivism’ seemed to be reducible to the boring inductivist strategy of collecting as many single data one can, and then finding out by mere statistical recounting all the ‘regularities’ those data might happen to exhibit. Of course, no real science actually works in this naively Baconian way, nor most self proclaimed (and much less those merely referred to by others as) ‘positivist’ philosophers, have defended such a caricaturesque idea of the scientific method. But even when a more liberal and benevolent description of the positivist admonitions is adopted, most people in the social sciences still rejects that positivism is a sensible methodological approach. Just to take an example of this criticism, consider the recent paper by Isaac Reed (2010). Reed starts by making a distinction between the ‘context of investigation’ (i.e., what the social scientist *does*) and the ‘context of explanation’ (i.e., what the social scientist is trying to explain, the object she *studies*). According to

¹ For other arguments in favour of naturalism, see Rosenberg (2005) and Ross (2011).

Reed, the basic mistake of what he calls ‘positivism’ is the normative relation that is assumed to exist between both contexts:

Positivism begins with a reduction of the context of investigation to the rules of logic and to a precise language whose meanings are transparent to those in the scientific community who use it. It also proposes that the context of explanation be reduced to that which can be observed—and, to a great degree, that which can be measured. Then, through the process of inductive generalization, covering laws are *imported* from the context of explanation to the context of investigation. (Reed, 2010, 24; original’s italics).

Besides the fact that it seems to be inconsistent to assert that «positivism reduces the context of investigation to the rules of logic and to a precise language» just a few lines after having asserted that «the context of investigation would include both of Reichenbach’s contexts (of *discovery* and justification)» (my italics), for positivists’ ‘context of discovery’ was explicitly taken as not reducible to logical rules, Reed’s explanation is rather obscure about what the ‘importation’ from one context to the other would consist of. Perhaps this simply means that the regularities that are observed in the ‘context of explanation’ are ‘accepted’ or at least ‘proposed’ in scientific papers or books (but, if it is so, what’s wrong with it?). And Reed makes a caricature of the process of ‘importation’ by reducing it to ‘inductive generalization’, overlooking the essential role of hypotheses and theories, as well as of abductive and analogical methods (that only in the most extreme and unrealistic views of science of some positivist philosophers were considered as illegitimate, if at all). The consideration of hypotheses also allows to see that it is an oversimplification to assert that in the context of explanation only ‘what can be observed’ counts: if by ‘observation’ it is understood *exclusively* what is observable ‘by the naked eye’, then the assumption is false about how science, even social science, works (for all ‘observation’ is ‘theory laden’); if, instead, ‘observation’ is taken to mean ‘anything that is intersubjectively detectable by *some* means’, one does not see what’s wrong in taking into account only ‘observable’ things in what you want to scientifically describe and explain: is Reed suggesting that social scientists must assert things about which there *is absolutely no way* to empirically determine whether they are as we are asserting? I doubt it.

Reed contrast this caricature of positivist social science with his own preferred, hermeneutic proposal, one in which the two ‘contexts’ overlap to a certain extent, in a way that allows the scientist to ‘share’ the meanings that constitute the basic stuff the social world is made of:

the context of investigation is structured culturally—by a set of meanings we know as «social theory.» The various contexts that the researcher encounters are also structured culturally—by meanings that are different in different times and place, but that tend to be much less abstract and much more pragmatically organized than «theory.» The contexts are brought into contact with each other when the researcher uses concepts from the context of investigation

to describe and interpret the context of explanation she is studying (Reed, 2010, 34-35).

I have nothing to oppose to this view, only that it is completely coherent with a non-caricaturesque view of ‘positivism’ (I mean, one which included the sophisticated reflections of Popper, Lakatos, or even the late Hempel and Carnap). In the same way as ‘educated’ concepts in the natural sciences (as ‘matter’, ‘heat’, ‘measurement’, ‘nutrition’...) have a *history* that connects them with ‘common sense’ and ‘local’ notions, the same happens with ‘scientific’ concepts in the social sciences; real scientists start at the point they can with the concepts and meanings they have, and try to *critically improve* them in order to offer *better explanations* of social facts than what could be attained just by using the ‘common sense’, ‘bar talk’ theories and meanings that ordinary people employ in their ordinary lives. The Reichenbachian ‘context of discovery’ is simply, amongst other things, the door through which ‘common meanings’ enter into the realm of social research. But what is important in science is what is *done* with these concepts, how they are (successfully or unsuccessfully) *rationaly* transformed by the process of investigation in a *new theoretical network* sagaciously articulated with empirical observations. Of course, one peculiarity of the objects that constitute Reed’s ‘context of explanation’ is that they (contrarily to galaxies, molecules or trees) can also absorb the symbolic products of that scientific process (at any stage of their elaboration, and with any possible degree of ‘fidelity’) and transmogrify them again into ‘ordinary meanings’. But a honest positivist will recognise as a *real empirical fact* that people do this, and will try to develop a social theory that allows to explain how this absorption is *possible* and what theoretical and empirical *consequences* it actually has. Hence, Reed’s proposal to replace the positivist schema of how social science works by a ‘hermeneutic’ schema can really be read as a vindication of a (knowledgeable) positivist method². Even when he ends by claiming that his position «also has the advantage (or, for some, the disadvantage) of offering a clear break between natural-scientific and social-scientific knowledge (because...) in social science, the role of meaning in structuring human action takes place both in the context of investigation and the context of explanation» (op. cit., 36), we must simply point to the fact that meanings (and the behaviours triggered by them) are *as natural* as any other component of the world’s furniture; of course they have properties that other things don’t have, but this is exactly true of *every* type of entity. Perhaps what Reed and other hermeneuticists want to say is that usually social scientists do not need to assume any *particular* theory about how nature permits the emergence of something as ‘meanings’, only some understanding of how meanings ‘work’ (and very often perhaps not even the latter). This is right, of course, but it is in the best case a half truth: in some (or even most) research contexts it does not

² To be honest, I must recognise that Reed’s schema is explicitly proposed as a ‘middle ground’ between the positivist and the anti-positivist approaches.

matter *how* the symbolic capacity of humans is grounded on their biological cognitive architecture, but surely there will be research areas *in the social sciences* where the psico-biological ‘fundamentals’, or some of their observable effects, are relevant (e.g., in studying the *social* effects of cognitive biases, learning capacity, motivation, and so on). And, what is more important, it would be appallingly hubristic the pretension that *no* empirical discovery about ‘how the mind works’ (and the body, to say it all) might *ever* have the tiniest chance of affecting the fundamental hypotheses of social theorising! Just as a symptom of intellectual modesty, the door of social sciences should be permanently open to such a possibility, particularly in these times when interdisciplinary research is becoming the norm rather than the exception.

But the most important criticism that can be done to this attempt to demarcate the social from the natural science by resource to considering the former as essentially engaged to a symbolic reality is not, from my point of view, that it ignores the possibility of ‘naturalising’ symbolic behaviour and meanings (a topic to which I shall devote the next section), but that it seems to ignore an important part of what social scientists actually can *do* with the ‘meanings’ they ‘share’ with their objects of study. For social science has not just the goal of *interpreting* what ordinary people say and do, in such a way as if a piece of social science research could *only* consist in something like a ‘edited translation’ of the lives of individual men and women. Far from this, many of the work in the social sciences consists in the attempt to discover those *causal* structures, processes or mechanisms that, on the basis of the symbolic capacities of individuals and *other* relevant factors, might *explain* what is ‘observed’ to happen in the social world. Concepts like ‘market equilibrium’, ‘class structure’, ‘strategic games’, or ‘social networks’ may of course have been ‘derived’ (usually through a rather tortuous path) from *lebensweltliche* notions, though, in order to construct theories that are reasonably capable of offering respectable explanations of whatever they were intended to explain, social researchers have been forced to transform them so thoroughly that they are now as different from their common sense origins as nightingales are from their pre-dinosaur ancestors. The social world is much more than just what ordinary individuals believe and say about it (if it were not, one might start doubting of the very need of a social science); it also consists in the *ways* in which those beliefs and actions are mutually interdependent and dependent of other causal factors or constraints, and we ordinary people may be as *ignorant* about the real working of those *ways of interdependence* as we are about the chemistry of protein digestion, no matter how much we happen to eat. This is the reason why not only positivism, but also *some kind of realism* is necessary in the understanding of the functions and methods of the social sciences (I shall come back to this question in section 4). Of course, social theories are ‘symbolic -and interest laden- social constructs’, but at least one of their goals is to make us *tentatively know* how the social world works, how some things in it depend on other things. And regarding the relevance of this goal, there is absolutely no difference between the social and the natural sciences.

3. MODELLING MEANINGFUL ACTIONS

There is another sense in which the fact that human action is based on meanings is taken to imply that social sciences cannot be naturalised. It has to do with the fact that natural sciences deal about ‘natural entities’, or ‘natural systems’, i.e., things that are *governed by some objective regularities* that we can try to know by empirical means (including, of course, devising hypotheses, testing them against empirical data, and modifying them according to the result of the test); human meaningful actions (as opposed to those facts derived from our mere physiology, so to say) are ‘free’, and hence not reducible to regularities. Of course, for a field to be classified amongst the ‘natural sciences’ it is *not necessary* that we discover a set of laws from which everything known *about* that field can be derived, only that *some* regularities can be known and that they form some more or less coherent and compact structure (constituting what is usually called a ‘scientific model’, or several of them). It is also *not necessary* that the regularities we happen to find in one field are logically derived from the laws of some ‘more fundamental’ field. We are far (or not so far?) from the times of the distinction between nomothetic and idiographic sciences, because the search and use of regularities and formal models in the social sciences is now commonplace, and because the relevance of hysteresis, path dependence, evolution and contingency is recognised in many fields of natural science, from biology to evolutionary cosmology; so, even if not all aspects of human action are apt to be captured by an scientific model, this does not mean that there are no important aspects that could be so apprehended. The fact is, of course, that different fields can be more or less indulgent than others in letting us discover a high number of significant regularities, and I shall just refer in the last section to the paucity that characterises the social world in this respect. Now, instead, I will try to offer a tentative, and very speculative conjecture about the way in which I think that many of the scattered and not-so-robust regularities that the social world has to offer us can be organised under what might be called (if the term did not sound so tremendously pretentious) ‘the fundamental law of the social sciences’, in the sense that it is a ‘law’ (or, more precisely, an abstract explanatory schema) that tries to express in a formal way what characterises a ‘hermeneutic’ explanation, and hence, the kind of *fundamental explanation* that is usual in most social sciences. This schema is not something intended to allow us to predict what will happen in every social circumstance (far from that!), for even this is not so common in the natural sciences, where basic laws do not really work as strict correlations between empirical variables, but as abstract frameworks for the construction of more specific models that need the addition of further, more restricted laws, as well as lots of tinkering with empirical details, simplifications and idealisations³. Indeed, what I want to present is a way of showing that ‘hermeneutic’

³ To use a very well known example, Newton ‘second’ law ($F = m \times a$) is useless if we want to derive from it any empirical prediction; to do that, we must build a more specific model helping us

theories are not qualitatively different from other scientific approaches in the sense that what they do can be explained as the construction of abstract formal models that can be empirically tested and also can be grounded on knowledge of natural facts, hence, nothing different of what is done in other sciences.

The most important reason why many think that human action being based on meanings entails that it cannot be 'naturalised' is because meanings are not only subjective, but *normative*. The presence of meanings in the explanation of human behaviour shows that our actions are based on 'reasons' (partially, at least), not on 'causes' (or not only), and the former cannot be reduced to mechanical laws in the sense that the latter can, or so it is argued⁴. I want to argue, instead, that contemporary science gives us enough instruments to *model* 'hermeneutic agents', so that this modellisation can be incorporated into still more complex theories that might, in principle, be subjected to empirical tests. This does not mean that the models of 'hermeneutic action' can be reduced to the content of other more basic natural sciences (biology, chemistry, physics), for, in general, every 'higher level' theory has concepts and principles that cannot be formally derived from the concepts and principles of 'lower level' ones. The only important thing is whether the model we present, or an alternative one, could be shown to be *coherent* with what we know about the biology and psychology of human beings, including their evolutionary history. The model, lastly, must not be taken as an unmovable 'law of nature', but *as a model*, and in this sense, as provisional and revisable.

Of course, I am not doing anything new in proposing a general hypotheses as a kind of 'basic regularity' in the social sciences. This is exactly what has been done from more than a century ago in at least a particular and powerful social science (economics), and more recently in others as well. I'm referring, obviously, to the assumption of expected utility maximisation, which is a mathematical principle that has served to construct literally millions of models or model applications. This principle asserts that agents have a coherent set of preferences (that can be mathematically represented by a utility function, u) and a coherent set of beliefs (that can be represented by a subjective probability function, p), and that the action performed by an agent will usually correspond to the one for which the following combined function gives the maximum value:

$$EU(a) = \sum_x p(x,a)u(x)$$

to identify the particular forces that apply to the circumstance, and apply that model to systems from which we have extracted some empirical data, and even all this can only be done after dramatic simplifications, idealisations, and other mathematical tricks are performed so that calculations become a practical possibility. So, the 'law' is more a *general recipe* for the construction of (more or less limited in scope) models than a transparent description of conspicuous correlations.

⁴ It is, however, an open question within philosophy of language whether meaning (as a semantic or cognitive concept) is 'normative', i.e., whether meanings entail some kind of obligations in the users of a language (see, e.g., Gibbard, 1994). I will take for granted, for the sake of the arguments, that they are.

where a is the action under consideration, and each possible x is a different and complete description of the consequences that might bring the realisation of a . Stated more clearly, the hypothesis just asserts that the option chosen will be the one with the higher *average utility*. In economics, and other areas of research, the principle of utility maximisation is even taken to be identical to ‘the’ principle of rationality, basically after the mathematician Jim Savage’s proof of the fact that the principle can be derived from a set of reasonable and *qualitative* assumptions about what is a ‘rational choice’⁵. There are, however, some strong reasons that recommend not to consider the principle of utility maximisation as a law capable of explaining all ‘meaningful’ action; I shall list a few of them:

1) The principle contains an unsolved tension between its prescriptive and its descriptive character. On some interpretations, it is assumed that rational people *should* behave according to the principle, not that real people actually obey it, but in this sense it is not an *empirical* model. On other (more common) interpretations, it is assumed that real people do follow the principle; but empirical violations are known to be very common, and so, in this sense, the principle has been *falsified* (at least if we interpret it as a universal generalisation).

2) The principle demands from the agents some cognitive capabilities that we know they don’t actually have. It is not only that real people does not know how to calculate probabilities and utilities (for that matter, they also don’t know how to calculate the trajectory of a freely falling object, but this does not prevent them to fall in a nicely Galilean way... so, our brains might in principle maximise expected utilities as ‘unconsciously’ as our bodies determine their falling trajectory). The problem is really that the principle demands that agents categorise the actions and states of nature in a way coherent with the concepts of probability theory, for example, that they think in terms of mutually exclusive final states of the world. It also entails that agents should be logically omniscient, i.e., that they know all the logical relations between the relevant propositions, for comparisons of probability must be such that, if A logically follows from B , then $p(B) \leq p(A)$, but surely real agents ignore many of these connections.

3) Perhaps more importantly, the maximisation principle is about the *result* of the agent’s decision, and not about the *reasoning process* he or she follows to arrive to that decision. In particular, the maximisation principle leaves out a lot of what is considered relevant to *understand* an entity as a rational agent, things like *language* and *deliberation* that are explicitly taken into account in hermeneutic approaches. Actually, if we take the principle as an explication of the concept of ‘rationality’, then we will leave out of this concept the very process of reasoning, i.e., the principle is not able of explicating what ‘reasoning’ (or ‘rational reasoning’) consists in. For suppose that, in order to decide what choice to make amongst the options within set S_ϕ , you should perform a reasoning process, but ought to *choose* first what reasoning process is best from the set of all those *possible* reasoning processes (or ‘arguments’) whose outcome is a choice between

⁵ Savage (1954).

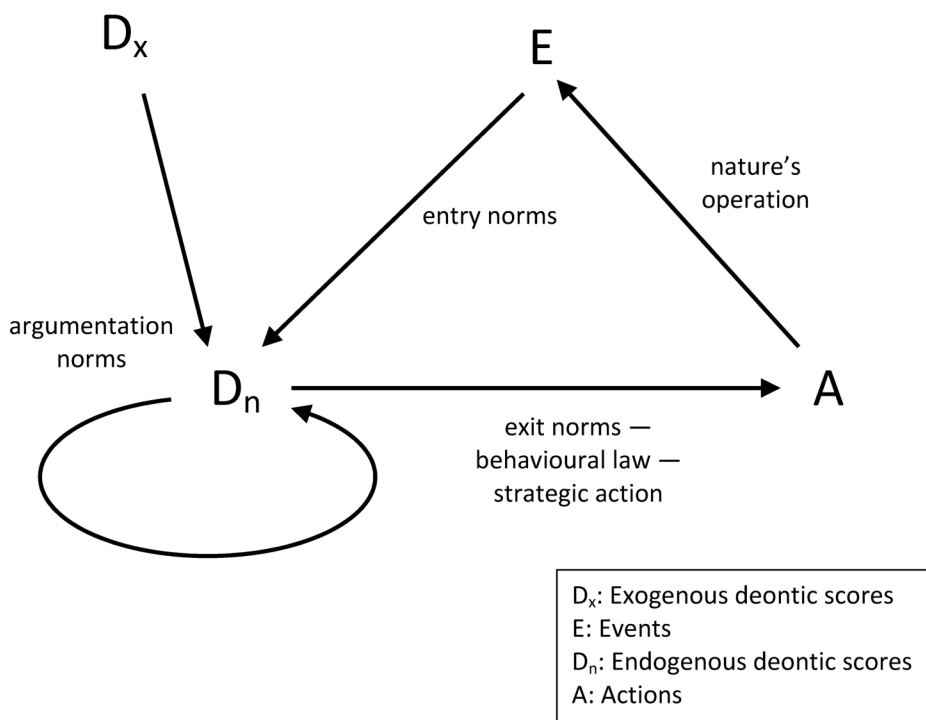
the elements of S_0 (let's call S_1 this set of arguments); now, in order to rationally choose one option amongst S_1 , you must make another reasoning process, choosing amongst the set of processes S_2 (the possible arguments whose outcome is a choice between the elements of S_1), and so on, and so forth. So, applying the principle of maximisation to the process of reasoning itself leads to an infinite regress. Hence, how is *rational reasoning* possible?

These and other problems can lead to prefer a different explication of what rational ('meaningful') action consists in. Within mathematical social sciences, a famous strategy has been to substitute the idea of 'substantial rationality' (i.e., maximisation) for some computationally less demanding task, like the many versions of 'procedural rationality' (see, e.g., Laville, 2000). Within this line, I want to propose what might be formally understood as a kind of procedural rationality, but which is based on some intuitions from recent philosophy of language, and on an explicit consideration of the *normative* nature of meaning and reasoning (for further details, see Zamora Bonilla, 2011). The starting point of the model is to reflect on what makes a series of noises, lines of ink, or neural processes of verbal images an example of *linguistic behaviour*. Philosophers of language within the so called 'inferentialist' school⁶ assert that what makes of a behaviour a verbal (or 'rational') behaviour is that it is governed by inferential norms, i.e., that the production of an item of language behaviour *amounts* to making a move in an *inferential* game, so that the rules of the game determine the *consequences* that you and the other speakers are *entitled or committed to draw* from what you have said. So, the elements of a game like this are «normative positions», i.e., a description of what you *must* or are *allowed* to do in the next movements. We can easily extend the reach of this normative positions to behaviour which is not (purely) linguistic, i.e., the commitments or entailments defining your position in the game can be referred to the performance of extralinguistic actions, like paying something to someone, or taking children to the school. We can also extend the reach of the inferential norms to allow that an agent's normative position is determined not only but what he or she, or other agents, say, but also by what they have done in the past, or by other publicly observable events that may have happen.

So, a social situation is a set of positions which are normatively determined according to a network of rules (which need not be explicitly formulated), positions that are defined from what each agent has the right or the obligation to do if he or she were in that position of the 'game'. Of course, an essential aspect of the model is that the 'actions' that one can or must do, and what partially determine one's normative score in the next positions, are verbal actions, i.e., 'speech acts', so that agents can employ verbal communication and deliberation to *publicly define* the commitments or entitlements of each individual, group or institution, what I shall call their 'deontic scores'. Obviously, having a description of what each agent is allowed or entitled to do within a situation does not offer a

⁶ See, e.g., Brandom (1994).

full explanation of the agent’s behaviour; for this we also need the help of a couple of psychological theses. The first one is that we need to assume that deontic scores can serve as a *motivation* for the agent’s action, a motivation that needs not be the only one the agent has (in particular, the agent may have other desires besides to comply with her sense of obligation); i.e., deontic scores need to be acknowledged and ‘felt’, so to say, in order to have a causal relevance. The second thesis is that we have to assume some statistical correlation, or ‘*behavioural law*’, between an agent’s map of motives and her behaviour. In order to simplify, we need to assume at least that there is a probability distribution that expresses how likely it is that an agent performs action X under the assumption that she has a deontic score stating the her degree of obligation to perform X is such and such (and, obviously, this probability function can also depend on what the other desires and commitments of the agent are)⁷. Figure 1 shows how the structure and dynamics of social situations would look like from the point of view of this model.



In a social situation, the deontic scores are subjected to a dynamic process: we start with some antecedent (or ‘exogenous’) normative state, which is modi-

⁷ Actually, reasonable inferential norms will allow agents to take into account their own desires and those of the others in defining the deontic scores.

fied according to the *argumentation norms* that are accepted by the people in that situation, and on which the occurrence of other external events can also influence according to the situation's *entry norms*. The *exit norms* are those that determine, on their turn, what actions are commanded by each agent's deontic state; these norms, together with our behavioural assumptions, determine a certain probability of the *actions* that each agent will take, and these actions will have physical effects in form of *events*, that may influence again the agent's deontic states according to the relevant entry norms. In more sophisticated versions of this model, we can introduce a perspectival description of the deontic scores, in the sense that we don't need to assume that each agent's idea of what the commitments or rights of each other agent are is necessarily the same: disagreements about what one must or must not do, and even about what are the relevant applicable rules or their interpretations, are obviously possible. I want to insist in that this model is naturalistic in three important senses. In the first place, though it includes normative elements as an essential part of it, these are taken to be normal events in the natural causal structure of the world, for what makes 'normative' a state is the agents' psychological attitude of taking it as normative. In this respect, I have to mention the close connection of this approach with the notion of *homo suadens* that have been proposed in Castro Nogueira *et al.* (2008), esp. ch. 6-8, where a detailed evolutionary explanation is given of the fundamental role that the transmission of normative valuations from an individual to others has in the constitution and evolution of human cultures.

In the second place, the model gives an extraordinarily simple way of reducing 'collective meanings' (intentions, judgments, decisions, functions, values, etc.) to natural events grounded on features of individual animals like human beings: the 'collective' arises just through the *systemic interdependence* of the individuals' thoughts, attitudes and behaviours, as well as through the fact that individuals can (according to the perspectival definition of deontic scores we have just mentioned) *attribute* deontic scores to *any entity* their inferential rules allow them to do it (like firms, states, families, churches, the dead, or the gods).

Last, but not least, the model is 'naturalistic' in the sense that it can lead to the construction of detailed computer models of artificial agents that are based on the dynamics of normative deliberation, and that can be empirically tested and improved (see, e.g., Castelfranchi *et al.* 1999; Boella *et al.*, 2006; Rahwan and Simari, 2008); these artificial models are not devised to offer something like a 'general theory' of society, but are limited formal constructions applicable to small world cases, a research strategy which is common in the application of computer simulations in the natural sciences.

4. WHY, IN SPITE OF ALL THIS, THERE WILL PROBABLY NOT BE A NATURAL SOCIAL SCIENCE

In spite of the fact that, as I have argued so far, social science fits within the

natural sciences both because of its field being a portion of the natural world obeying natural laws, and because of its general methods (the ways of knowing what's out there) being deeply identical to those in the natural sciences, the truth is that we cannot be very optimistic about the possibility of developing in the social sciences a *corpus of 'scientific knowledge'* as strong, general and robust as in the case of many natural disciplines, like physics, chemistry, or biology. The reasons for my pessimism come not from a suspicion about the capabilities of social scientists (that will not be, on average, smaller than those of any other human collective), but from some aspects of the *nature of the entities* social science deals about. These entities, which we might call 'social aggregates', are *complex systems constituted by the interactions of intelligent individuals* (something probably similar to what in Castro Nogueira et al. 2008 are called '*plektopoi*'), i.e., things like families, corporations, governments, markets, cities, tribes, etc. Both the fact that these systems are complex, and the fact that their constituents are intelligent actors, contribute to the (so to say) low epistemic efficiency with which scientific models can be applied to them. I will comment separately on these two groups of factors, but before that, I want to insist in that I am referring to a *low* cognitive efficiency, as compared to other scientific fields, so I am not asserting anything like that the social sciences can give no real knowledge: only that it is comparatively limited. Basically, what I mean is that it is unlikely that in these areas *robust empirical regularities* can emerge, and that this produces a very high level of *uncertainty* in the application of *any* model to most social situations, for we cannot be sure that the regularities on which our model is grounded are true with enough approximation to the case we are studying⁸.

Regarding those properties of social aggregates that can be responsible of the difficulty of finding out robust regularities, we may list the following ones:

1) Social aggregates, and particularly the bigger ones, are *scarce*, and hence they don't provide a sufficiently large basis for statistical inference most of the cases.

2) Their limits are *blurred*; it's difficult to determine where one market, or a family, or even a conurbation, ends and the next one begins. This means that *types* of social aggregates don't constitute, in general, what we might call a 'natural class'.

3) They are also very *irregular and unstable*: they do not behave always in the same, or even in similar ways, both comparing one aggregate with others of the same type, and comparing the temporal evolution of one single aggregate.

4) Last but not least, social aggregates are very *difficult to manipulate experimentally* (and not only for moral reasons).

As it is obvious, points 1-4 are more or less common in some other fields of science, basically those that are basically constrained to field research, i.e., big

⁸ For a more detailed discussion on the difficulties of empirical predictions in the social sciences, see Betz (2011).

portions of biology, geology, meteorology, and astronomy, but also in almost any science when it has to be applied 'in the wild'. The degree in which the listed difficulties appear in each case is variable, but it also has the consequence of 'infecting' with a higher or lower level of uncertainty the conclusions of research. Of course, in the case of social sciences these difficulties are also a matter of degree: we may have cases in which a nice experiment can be performed, and cases in which the statistical basis is big enough to justify sound inductive inferences or generalisations. But the norm is rather the opposite.

Regarding the reasons for scepticism derived from the nature of the social agents, here are the most important ones:

5) Individuals, and particularly their own cognitive apparatuses, are very *complex* systems in themselves (actually, much more complex than the social systems they constitute, if the complexity of brains is deducted). This entails that the behaviour of individual agents is very difficult to predict, and that they have also a very high level of creativity, i.e., the capacity of devising new forms of behaviour that the social scientist's models had simply not taken into account.

6) Agents act according *their own perception and understanding* of the social situation (which does not entail that it is necessarily an absolutely *right* understanding). The social scientist can often just *guess* what those perceptions are, even when she engage in 'participatory' research, and particularly when the goal of the social model or theory is, as it very often happens, to understand the complex *consequences* of the interaction of thousands of agents.

7) Human beings are able of *acting strategically* (perhaps other animals also are, to some degree), i.e., anticipating the actions of others. This entails that their own understanding of the situations (see point 6) can still be extraordinarily more complex because of strategic concerns, so that it can become nearly impossible for the social researcher to guess what is really the 'game' that agents are playing, or are thinking they are playing.

8) The fact that most social interactions have the form of strategic games also entails that many social situations, even if the 'real game' is correctly represented, can be intrinsically unpredictable, because that game may have *multiple equilibria*. Actually, normative deliberation, as depicted in the past section, often has the function of reducing this indeterminacy for the social agents themselves, helping them to find out 'the right' (in a normative sense) solution of the game, amongst all the ones that are possible. This is way many of the regularities that can be found in the social realm have the nature of *social norms*.

9) Lastly, this strategic capacity can also entail that the discovery of an empirical regularity in the behaviour of a social aggregate can lead to an intrinsic *instability* due to the *reflexivity* of social sciences, for when the agents learn the scientist's discovery, the former may use that knowledge to strategically change their own pattern of behaviour.

To this list of obstacles that the emergence of robust regularities encounters in the social world, we may add another one that in part derives from this intrinsic fogginess of social facts. It is that, in the absence of strong intersubjecti-

vely testable empirical laws serving as a harsh constraint to speculation, perhaps a too big space remains for the subjectivity of the social scientists, a subjectivity that, as in all cases, may contain both cognitive biases and non epistemic interests. When contumacious empirical data don't serve as a restraint of what one may assert about the world, then other considerations will determine what each social school prefers to say and imagine about the Rorschach test that social facts and data usually consist in.

If my diagnostic is nearly right, then it is unlikely that social sciences can become someday a branch of the 'natural sciences' in the sense of having discovered a set of laws that connect their objects with what we know about the biological constitution of individuals and that allow to make robust predictions about their individual or aggregate behaviour. But is this a bad thing? I don't think so. The perspective of sociologists, anthropologists or economists having the capacity to know with so much detail what will happen under such and such circumstances in the communities in which I live would be a rather terrifying possibility in many ways. The development of toy computer models of artificial communities, based on the ideas presented in section 3 or on other alternative intuitions, will surely lead to a higher and deeper amount of empirical knowledge about social facts, as the introduction of some profound theoretical insights and some statistical and other mathematical technics during the last century has allowed to increase our knowledge about the social world. I think that any addition in this corpus of knowledge (or in any other branch of science and technology, of course) should lead to a negotiation about who will have the power of applying it. But this will better be the topic of another paper.

ACKNOWLEDGEMENTS

Financial support from Spanish Government's research projects FFI2008-03607/FISO, FFI2008-01580/FISO and FFI2011-23267 is acknowledged.

REFERENCES

- BOELLA, G., *et al.*, (eds.), 2006, special issue on Normative Multiagent Systems, *Computational and Mathematical Organization Theory*, 12(2-3).
- BRANDOM, R., 1994, *Making it Explicit*, Cambridge (Mass.), Harvard U.P.
- CASTELFRANCHI, C., *et al.*, 1999, «Deliberative Normative Agents: Principles and Architecture». *Proceedings of ATAL'1999*, 364-378.
- CASTRO NOGUEIRA, L, *et al.*, 2008, *¿Quién teme a la naturaleza humana?*, Madrid, Tecnos.
- GIBBARD, A., 1994, «Meaning and normativity», *Philosophical Studies*, 5, 95-115.
- JARVIE, I., and J. ZAMORA BONILLA (eds.), 2011, *The SAGE Handbook of Philosophy of Social Sciences*, London, Sage.

- LAVILLE, F., 2000, «Foundations of procedural rationality: cognitive limits and decision processes», *Economics and Philosophy*, 16, 117-138.
- RAHWAN, I., and G. R. SIMARI (eds.), 2009, *Argumentation in Artificial Intelligence*, Dordrecht, Springer.
- REED, I., 2010, «Epistemology Contextualized: Social Scientific Knowledge in a Post-Positivist Era.» *Sociological Theory*. 28, 61-80.
- ROSENBERG, A., 2005, «Lessons from biology for philosophy of the human sciences», *Philosophy of the Social Sciences*, 35, 50-74.
- ROSS, D., 2011, «Naturalism: the place of society in nature», en Jarvie and Zamora Bonilla (2011), 121-136.
- SAVAGE, J., 1954, *The Foundations of Statistics*, New York, Dover.
- ZAMORA BONILLA, J., 2011, «Rationality in the social sciences: bridging the gap», in Jarvie and Zamora Bonilla (2011), 721-738.