

FUNCTIONAL EXPLANATION AND EVOLUTIONARY SOCIAL SCIENCE

Harold Kincaid

1 INTRODUCTION

From their conception to the present, the social sciences have invoked a kind of explanation that looks suspect by the standards of the natural sciences. They explain why social practices exist by reference to the purpose or needs they serve. Yet the purposes invoked are generally not the explicit purposes or needs of any individual but of society or social groups. For example, Durkheim claimed that the division of labor in society exists in order to promote social solidarity and Marx thought that the state served to promote the interests of the ruling class. Social scientists have found these explanations as irresistible as their critics have found them mysterious. This chapter traces the controversies over these explanations — generally called functional explanations — and argues that they are widespread in some of our best current social science and that they can provide compelling information in some cases, despite the many doubts about them.

Section 2 surveys the general usage of functional explanations from classic sociologists to the present and common doubts about them. Section 3 surveys past accounts of how functional explanations work and their problems and develops two distinct models of functional explanation, selectionist explanations that show how a practice exists in order to bring about its effects and functional role analysis which explain systems in terms of component parts. Sections 4, 5 and 6 discuss specific instances of functional explanation as well as some general mechanisms thought to underlie them. Section 4 looks at norms, institutions and rational choice mechanisms, and 5 and 6 at inequality, organizational ecology, behavioral ecology, and varieties of selectionist mechanisms. Section 7 turns to functional explanation as functional role analysis. Section 8 draws morals and points to open questions.

While the overwhelming focus of this chapter is on functional explanation, I think functional explanation also raises a number of broader issues in the philosophy of the social sciences and philosophy of science more generally. These will surface briefly throughout, and I want to outline them here. I make no pretense to provide a thorough or conclusive discussion, only enough mention to show possible connections.

A long-standing dispute in and about the social sciences is the extent to which large scale social phenomena can be explained in individualist terms.¹ Variants of

¹For further development of this debate, see the essays by Little and Zahle, in this volume.

various individual traits. Other variants might argue against individualism by seeing groups as the units of selection or by seeing social structure as in some sense prior to individual traits. A related issue concerns the explanatory power of evolutionary psychology and human behavioral ecology: its defenders [Tooby and Cosmides, 1992] and Smith [1981] claim that all of culture, including social organization, can be explained in terms of psychological modules selected for in the Pleistocene.

Not far away from both individualism, evolutionary psychology, and behavioral ecology are questions about the unity of science. Individualism is often put forward as a thesis motivated by the goal of integrating the social and natural sciences; evolutionary psychology and behavioral ecology are given a similar justification. Similarly, those such as Gintis who advocate evolutionary game theory see it as providing a “common language” to unify the social sciences [Gintis, 2000]. An alternative picture depicts scientific unity as a piecemeal affair and doubts that functional explanation in the abstract is likely to suffice to integrate the social sciences.

Finally a broad set of questions about theories, models and explanation surface in debates over functionalism. A deep rooted tradition in philosophy of science and in much science itself sees science as essentially about producing theories and sees explanatory power as resulting from their ability to unify and generalize [Kitcher, 1993]. Models with idealizing assumptions are often an essential part of that process. Skeptics doubt that theories have such a central role, that unification explains, and that models with unrealistic assumptions explain.² They advocate an alternative picture that focuses on piecemeal causal explanation which downplays the role of abstract theories and places tight constraints on successful models. Many defenders of functional explanation in some or the other form defend it on the grounds of its ability to provide theories, to unify and to model. Hence these issues are joined in any discussion of functionalism.

2 HISTORY AND CONTROVERSIES

No definitive history of functionalism in the social sciences has been written. This deficit is perhaps due in part to two factors. First any such account would be nearly coextensive with much of the history of the social sciences in total. Moreover, from its inception, functionalism has been ambivalent between two different claims (detailed in the next section) — namely, that society is composed of component parts with specific functions as opposed to the claim that there is some kind of selective mechanism ensuring that those social parts exist in order to serve their functions. Different social science traditions emphasize one or the other of these claims for different reasons at different times in their history, making a cohesive

²See Kincaid [1996; 2004].

narrative difficult. In what follows I impose an order on this messy history.

Most of the key founding figures in sociology and anthropology — Comte, Spencer, Marx, Durkheim, Weber, Radcliffe-Brown, and Malinowski — espoused versions of functionalism. Comte [1974], who coined the term sociology, had a fundamental interest in the foundations of social order combined with the common 19th century belief that useful analogies can be drawn between biological organisms and societies. He also had to fight for the legitimacy of sociology as a discipline, and he may have done so by

appealing to similarities between the new social science and the much more established science of biology. Organic analogies and a concern with stability thus led naturally to the idea that society was composed of a set of interlocking component social parts — families, classes, and cities were his favorite elements — and that these parts contribute in different ways to the stability of the whole. That contribution was their function.

Missing from Comte is the functionalist emphasis on selective processes, though it was implicit in the idea of contributing to stability. This is not surprising given the pre-Darwinian times. However, another pre-Darwinian, Spencer, brought the selection through competition idea to social explanation (and indeed to biology through his influence on Darwin). Spencer [1975] adopted Comte's organicism. Societies were composed of interdependent parts analogous to those of organisms. He put this in terms that survive to this day: societies have structures (morphologies) and structures have functions; those functions are to meet the needs of the overarching society. He combined this distinction between structure and function with the claim that societal forms follow an evolutionary pattern from simple and undifferentiated to complex and specialized functions as they developed.

These ideas are in Comte, but Spencer develops them in more detail. Grafted on to them, however, was Spencer's defense of laissez-faire. Competition among free, self-seeking individuals produces efficient structures. Implicit here is the claim that those structures not only have a function but exist because they bring them about. Comte might have agreed that social structures exist in order to realize their functions, but the causal process was unspecified. Spencer is the first inkling of a mechanism, though he admittedly was vague on how his selection mechanism actually connected to his other functional claims.

Standard histories (cf. [Turner and Maryanski, 1979]) trace the next developments in functional thinking to Durkheim who then influences Radcliffe Brown and Malinowski in anthropology and then all three together set the stage for Talcott Parsons, the grand functionalist theorist of the 20th century. That narrative is too simple, I believe, and results in a diminished appreciation of the place of selection through competition in the functionalist story. To anticipate, a standard criticism of functionalism in its mature Parsonian form is illegitimate teleology, purposes without design. But since the 19th century there have been persistent attempts to develop a naturalized social selectionist account that would undergird functionalist claims in the social sciences. Some of the best contemporary work, I shall argue later, continues that tradition.

Missing from traditional histories of functionalism are Marx and, more controversially, Baldwin and Weber. Marx's functionalism pervades his core ideas. The "superstructural" aspects of society — the state, ideology, etc. — exist in order to promote the basic class relations characterizing the mode of production. Those class relations in turn exist in order to promote the productive forces. Marx is clear in each case that institutions have a characteristic effect or function that in some sense explains why the institutions exist.

Marx is not altogether clear on the mechanisms behind these explanations, but certainly he is groping toward some sort of competition and selection process. This explains Marx's reaction to Darwin: though Marx's core views were firmly established before the publication of *Origin*, he found in Darwin "a basis in natural science for the historical class struggle" and rational explanation for teleology [1922, 245]. Marx provided the most developed causal mechanism for the claim that class relations exist in order to promote the forces of production. He claimed to see a process whereby further

development of the material means of production become limited by the existing social relations, resulting in crises and an ensuing change in social relations. Of course the details, plausibility, and generality of that story is up for grabs, but Marx certainly does try to describe a roughly selective mechanism undergirding this functional analysis.

When it comes to the functional explanation of institutions like the state, things are fuzzier. We do not have the simple fettered growth — crisis — revolution model. Presumably the claim is that those regimes and state forms persist that are consistent with basic class relations and those that are not are eliminated. How this is all supposed to work for Marx is left undeveloped, though subsequent followers have provided more detail.

Another interesting figure if we want to explain functionalism as it is practiced today, especially in anthropology, is James Baldwin. Known now for the possible effect named after him — the Baldwin effect where mutations occur to make an acquired trait have a genetic basis — Baldwin made novel efforts to combine Darwinian natural selection, learning, and cultural change. Organisms learn by imitation, particularly humans. Imitation is not random, but shows a process of differential adaptation. This leads to a form of “social heredity” of culture that is different from genetic inheritance, though it too supports a selective process. This is one of the first clear, relatively rigorous discussions of gene-culture coevolution.

When cultural selection returns again later, it will be in the work of Leslie White, who implicitly invokes Marx (he wrote during the heyday of McCarthyism) and explicitly refers to Baldwin.

Max Weber [1978] was openly critical of organic analogies and the unilinear evolution found in the early functionalists. Yet, demonstrating once again the appeal of functionalist accounts and the Darwinian mechanism of variation and selection, Weber’s concrete empirical analyses of social phenomena extend the kind of selective mechanism found in Marx and Baldwin [Runciman, 2001]. For example, Weber’s account of religious development invokes various kinds of selection and retention processes. The exemplary prophet passes on beliefs and values that compete with others in the ensuing generations and are selected by their fit with the needs of individuals. Then competing organized practices compete for power, with success determined by features of the social environment. Similarly, Weber’s account of the protestant ethic and its role in the rise of capitalism has at least one reading that calls on similar processes: it was not the upward mobility of individuals from Protestant families per se that lead to change but the set of practices that Protestant family firms and workers brought with them that was crucial to the development of capitalism. They were, as Weber, puts it, the “carriers” of beliefs and values that were more or less adapted to their conditions of life.

Durkheim [1965] is the most influential of the early functionalists. He borrowed Spencer’s structure/function distinction, but put it to an empirical use that Spencer did not. In *The Division of Labor in Society* [1933] and *The Elementary Forms of Religious Life* [1961] Durkheim argued that particular social institutions serve to promote social solidarity. The division of labor does so by increasing interdependence among individuals. Religious beliefs symbolize the society in which they are embedded and thereby promote social integration. He backed up both with extensive empirical discussions.

Durkheim’s functionalism sometime extended beyond identifying functions to explaining institutions as existing because of their functions. He explicitly separates

causes of social phenomena from their functional effects. Yet he appeals to a Darwinian competition mechanism (explicitly invoking Darwin) in discussing the division of labor: increasing population causes increasing competition which results in increasing specialization. It looks like the division of labor exists because it allows societies to successfully function in the face of individual competition. Arguably we have here again the ongoing tension between functionalism as analyzing the role of parts in the whole and as providing a causal mechanism based on those roles.

Durkheim's ideas had a strong influence on two of anthropology's founding fathers, Radcliffe-Brown and Malinowski. Both attributed their basic orientation to him, yet made important elaborations in the functionalist picture. Radcliffe-Brown [1948] emphasized two crucial ideas: the distinction between structure and function and the notion of functional prerequisites. The first step in explaining society is to identify social structure — "the sum total of all social relationships of all individuals at a given time" [1948, 55]. Kinship systems were a prime example of social structure for Radcliffe-Brown (not surprisingly given anthropology's early interest in small scale societies) and they also illustrated his belief that social structure can be further divided into subsystems or roles.

There is a clear sense in Radcliffe-Brown that this project of identifying structure was explanatory in its own right. It required showing how the structures of a given society met its "functional prerequisites" or its "necessary conditions of existence," viz. consistency — a clear definition of roles and rights — and continuity.

The continuing question of whether (1) institutions have functions or (2) have them and exist because of them likewise arises in Radcliffe-Brown's functionalism. He denies that he is explaining historical origins but does claim to account for "sociological origins." Perhaps Radcliffe-Brown is sliding from their positive effects to the conclusion that things exist for those effects. But perhaps not. The key issue may be what "accounting for" comes to. Radcliffe-Brown sees science as essentially about producing laws. He may then think that if he can relate functions and needs in a lawlike way, he has explained them without making a causal claim linking effects and existence. We will see that similar assumptions about explanation surface in more contemporary accounts of functional explanation [Cohen, 1978].

Malinowski [1944] added further important wrinkles and some differences in emphasis. Identifying social structure meant institutional analysis. That required identifying, among other things, relevant norms, material culture, and societal needs met by the practice in question. Institutions served needs that were largely the needs of *individuals* rather than the collective needs of society as in Radcliffe-Brown.

Radcliffe-Brown and Malinowski were important influences on the next and most elaborate functionalism, that of Talcott Parsons. Parsons functionalism aims to provide a full fledged theory of societies in general. He identifies what he believes to be the essential components or subsystems of any society. The working of these components is then explained in terms of their contribution to social equilibrium. Social equilibrium is maintained when the universal functional needs or prerequisites are met.

This grand vision is fleshed out with a plethora of further nomenclature that becomes increasingly baroque over Parsons' career. But in short form, Parson's theory is this. There are four basic subsystems in any society: the biological, personality, cultural, and social. The biological system concerns inherent needs, drives, etc. of humans. The

personality system involves psychological dispositions, cognitive states, etc. Ideas and symbols form the cultural system. A social system is the pattern of roles and statuses that structure society. Each of these systems must meet four functional needs:

- 1 adaptation – obtaining resources from the environment
- 2 integration – maintain coherent relationships among their component elements
- 3 goal attainment – setting goals and allocating resources to achieve them
- 4 latency – reproducing organizational structure and managing tensions between units.

Parsons is an "analytic realist" about these categories: they "correspond not to concrete phenomena, but to elements in them which are analytically separable." [1937, 730]

Running parallel to Parsons' functionalism in the late 50s were developments in anthropology harking back to the selectionist strain found earlier in Marx, Baldwin

³This is precisely the same kind of analyses advocated by Hempel in roughly the same time period [Hempel, 1970].

and Weber. Functionalism in the guise of Radcliffe-Brown and Parsons had little space for history or change; societies were integrated, self-contained systems. Both tenets were challenged by the rise of cultural ecology or ecological anthropology. Steward [1955] and White [1949] were the most influential early proponents. For both, key issues were how societies and the subsystems composing them adapted to ecological surroundings. The selective processes described by ecologists in biology were relevant to the form and processes of change that described societies. These ideas were the direct inspiration of contemporary functionalism of the sort advocated by Rappaport [1967] and Marvin Harris [1979], both of who provided numerous explanations of social practices in terms of some kind of selective processes.

At the present, various versions of functionalism are alive and well, making it a continuing active research program⁴ in the social sciences. I will sketch some representative examples in later sections.

I won't try to sketch the corresponding history of criticism that functionalism generated, for most of the standard criticisms are present in some form from the beginnings in the 19th century. Rather let me list some of the most important:

Questions about mechanism and illegitimate teleology: Many ([Elster, 1983]; [Vayda, 1987]; [Little, 1989]) claim that functional explanations do not provide mechanism connecting beneficial practices with their effects and are thus unconfirmed and/or inadequate explanations. Because there is no mechanism, functional explanations end up with a mystical form of teleology — the need for social institutions makes them exist.

Questions about change, diversity and agency [Layton, 1997]: In its Parsonian form, for example, functionalism sees society as a self-maintaining system where social control is emphasized at the expense of individual agency and societal conflict and cultural diversity downplayed in favor of universal societal components.

Questions about illegitimate use of biological explanations ([Ellen, 1982]; [Hallpike, 1986]; [Fraccia and Lewontin, 1999]; [Bryant, 2004]). Societies do not have clear boundaries, do not have the equivalent of genes, do not reproduce, do not merely react to

their environments but change them, have guided rather than blind change, and are not populations of traits. All of these differences show that biological analogies to natural selection and organismic organization have no legitimate place.

Questions about tautology and vacuousness [Hallpike, 1986]: A persisting charge against functional explanation is its lack of content. This flaw is seen in various aspects of functional explanation. The essence of functionalism seems to be the claim that stable societies survive, where to be stable is to persist and thus survive. Or societies are said to adapt, but adapting is given no more content than simply

⁴As a rough guide, contemporary work breaks down into traditional sociological theorizing influenced by Parsons ([Alexander, 1985]; [Luhman, 1995]), rational choice game theory, evolutionary game theory [Bowles, 2004], cultural selection theory ([Boyd and Richerson, 1985]; [Durham, 1991]; [Cavalli-Sforza and Feldman, 1981]), mimetics [Aunger, 2000], behavioral ecology [Smith, 1992] and evolutionary archeology [Lyman and O'Brien, 1998].

“try to cope with their situation.” Neither claim seems to be saying anything nontrivial. Functionalists' predilection for putting societies or their practices into complex and highly abstract classificatory schemes likewise seems nonexplanatory.

3 CLARIFYING ISSUES

This section clarifies what is involved in functional explanations as the prelude to assessing them. There is an enormous literature over many years attempting to analyze talk of functions in general — across the social sciences, the biological sciences, and common sense usages [Allen *et. al.*, 1998]. Though various useful ideas come from this literature, its main project seems to me misguided from the start, for it presupposes an impossible and unhelpful notion of what is to give an account of something. Thus it is necessary to first discuss what we are doing in giving an account of functional explanation.

Standard philosophical approaches have tried to give the individually necessary and jointly sufficient conditions for “has the function of” or “exists in order to.” Such a definition is supposed to hold for all uses of these expressions. The evidence for these definitions are linguistic intuitions, usually on the part of the philosopher doing the analyzing. In short, an account is a formal definition of the concept.

This project is suspect for many reasons: (1) There is good empirical evidence from cognitive science that we don't represent our concepts in terms of necessary and sufficient conditions but use prototypes and relative similarity instead [Rosch, 1978]. (2) Even if we did represent our concepts in this way, it is implausible to think that functional language must have the same content regardless of whether it is used in biology, the social sciences, or in every day talk about artifacts. (3) Even if we found necessary and sufficient conditions that matched linguistic intuitions, it is not clear what that tells us about the science. Why should philosophers' intuitions about what they would say tell us anything about the science? It is not even clear that capturing scientists' linguistic intuitions about language use would tell us much, for we want to know how functional explanations work and when they are legitimate, not something directly illuminated by even scientists' linguistic intuitions.

Thus in providing an account of functional explanation, I am not providing a definition

tested against linguistic intuition. Instead we want some kind of framework that catches the main features of at least some classes of functional explanations and to do so in way that helps us evaluate them. Thus we want to know: How do they work? What kinds of processes in the world do they commit us to? What is their relation to causal explanations? and so on.

A first useful distinction — one that the definitional project seeking a universal definition ignores — is between functional explanations as the citing of roles in a complex system contrasted with functional explanations that claim something exists in order to do something. If we ask why there is a carburetor on my old Toyota or what it does, we can answer that it serves to provide a proper mix of fuel and air to the pistons. We describe its typical causal inputs and outputs, its causal role in the system that is my car. However, knowing that it has a role does not mean it automatically exists in order to do so. A typical causal effect of my brake pads is to cause the rotors to wear, yet we know they weren't designed in order to cause wear. In general showing that A has the systematic effect B is not the same as offering the explanation A exists because it does B.

We saw in our historical survey that functionalism vacillates, often unconsciously, between identifying functions in the sense of describing social structure vs. explaining social practices as existing because of their effects. So much in Radcliffe-Brown or Parsons is only about describing social structure. Alternatively, some ecological anthropology explains particular social practices as existing because of their effects without any parallel commitment to arguing that there is a cohesive structure to the overarching society and certainly without describing it in systematic detail.

In general, the claim that societal functions exist in order to do specific things is the more controversial of the two, for it seems to commit us to a rather mysterious causal process that is somehow nonetheless seductive. The bulk of this chapter thus focuses on functional explanation in this teleological sense. Nonetheless, functionalism as the description of integrated systems has been quite influential among social scientists and at the same time little discussed by philosophers of social science. In the last section of this chapter I discuss some of conceptual and empirical issues these explanations raise.

There is a third sense of functionalism prominent in the philosophy of mind and cognitive science literature that I should mention to avoid confusions. Proponents of functionalism in these areas argue roughly that the material realization of mental states is not essential to them — it is the input/output relations and the relations to other functional states that make them what they are. A root idea here is presupposed by functional explanations in the social sciences referring to causal roles and to the selectionist accounts given later for “a exists in order to b.” The common idea is that causal explanation can proceed at a given level of aggregation without specifying the microlevel physical detail that realizes it. In other words, we can describe causal relations that abstract from details about the components realizing the entities that stand in those causal relations. Both explanations in the social sciences of causal roles and of “existing in order to” will depend upon this assumption, but they are not exhausted by it.

Let's look at several contemporary accounts of the idea that social practices exist in order to bring about their effects. Faia's [1986] *Dynamic Functionalism* makes causal feedback relations central. On this picture a functional explanation shows how A's causing B in turn influences A itself. For example, contacts outside ones social group reduce prejudice which in turn causes greater social contacts. While explanations in

these terms are no doubt given with some success in the social sciences, they still do not get at the stronger sense that *As* exist in order to *B*. Feedback relations for *Faia* are just circular or mutual causation. By Newton's laws of motion, every body in the universe causes and is caused by every other.

Thus if feedback warrants a functional explanation, then every physical process in the universe can be functionally explained. Obviously this conception is too wide. Moreover, circular causation is symmetrical — each component can equally be the cause of the other. However, functional explanations in the social sciences are usually not neutral in this way. For example, the state exists in order to promote the interests of the ruling class according to Marx, but that does not commit us to believing that those interests exist in order to promote the state.

Cohen [1978] makes some progress in narrowing the field of feedback relations to those that look more like functional explanations. According to Cohen, functional explanations are a kind of consequence explanation — explaining the existence of something because of the consequences it has. To show that social practice *A* exists in order to *B*, we have to show that it is a law that when *A* has a disposition to produce *B*, *A* comes to exist. Cohen uses this account to clarify Marx's functionalism. The relations of production exist in order to promote the forces of production in that it is a law that when new class relations would promote technological progress, they come to exist. This sort of feedback relation presumably is much narrower in application than *Faia*'s mutual causation.

However, Cohen's account is inadequate as it stands. Cohen takes functional explanations to consist in deducing the event to be explained from a specific kind of law. But there are well known problems with the general picture of explanation assumed here. There can be deductions from laws that do not explain because the laws are irrelevant: it is a law that men who take birth control pills do not get pregnant. Similarly, showing that *A* exists when it causes *B* does not show that *A* exists in order to *B* — the connection might just be accidental. A second problem with Cohen's account is that it requires too much. Larger corporations may exist because they can take advantage of economies of scale. Yet they do not automatically get larger when doing so would have that effect — lack of foresight, resources, etc may prevent it.

Stripped of its deductivist account of explanation and substituting causal notions in its place, a related possible account of functional explanations would be that *A* exists in order to *B* if:

- 1 *A* causes *B*
- 2 *A* persists because it causes *B*
- 3 *A* is causally prior to *B*, i.e. *B* causes *A*'s persistence only when caused by

A.

The first claim is straightforwardly causal. The second can be construed so as well. At t_1 , *A* causes *B*. That fact then causes *A* to exist at t_2 . In short, *A*'s causing *B* causes *A*'s continued existence.⁵

⁵Wright's [1973] account is a partial inspiration here, but it has to be stripped of its conceptual analysis pretensions. And the requirement here is to explain persistence rather than existence.

The third requirement serves to distinguish functional explanations from explanations via mutual causality. If A and B interact in a mutually positive reinforcing feedback loop, then A causes B and continues to exist because it does so. Yet the same holds for B vis-à-vis A. Functional explanations do not generally have this symmetry. Thick animal coats exist in order to deal with cold temperatures, but when cold temperatures are present there is no guarantee that thick coats arise. And surely, even if they do, they do not cause the cold to persist.

On the view advocated here, functional explanations are a unique subset of general causal relations. We can therefore now answer two general complaints about functional explanations in the social sciences: that they make an illicit appeal to biological analogies and that they are tautologous or vacuous.

The most general description of a causal system describes a set of individuals whose values evolve through state space. At this level of description we are told very little: current entities stand in some relation to past ones. Natural selection is inevitably an instance of this system, given that it is a causal system. Functional explanations as causal are also an instance. Every causal system is analogous in being a dynamical system. The point here is that whether one set of causal relations is analogous or disanalogous to another depends on the level of description we are using.

So at the most abstract level it is a trivial truth that functional explanations are indeed analogous to Darwinian evolutionary systems in so far as they are causal systems. They are disanalogous in that social entities have no DNA that replicates. But then the HIV virus has no DNA either (it is an RNA virus). We find analogous processes in DNA and RNA organisms despite the differences because we abstract from the details to identify abstract causal patterns.

So do functional explanations commit us to some illegitimate analogy to natural selection? No, because natural selection explanations are just one realization of the above schema which is thus the more general pattern ([Kincaid, 1986]; [Harms, 2004]). A's causing B may result in A's persistence by means that don't involve genetic inheritance, literal copying of identifiable replicators distinct from their vehicles or interactors, etc. In fact not all biological processes of natural selection require this level of analogy—differential survival can be caused by other processes (see [Godfrey-Smith, 2000]).

In this regard, Pettit [1996] notes that explanations of this general form do not even require a *past* history of selective processes. He argues that to establish a functional explanation we need only prove "virtual selection." Virtual selection refers to processes that would exist if some social practice with beneficial effects were to change. Suppose golfing may not be present now because of the positive benefits it had in the past, but if golfing were now challenged, then there would be pressures to maintain it. This virtual selection is just one way to make it true that A persists because it causes B, where B is a beneficial effect.

We can give a similar answer to the tautology or vacuity objection. Functional explanations in the social sciences certainly can be relatively uninformative as in "societies exist because they are stable" or "societies that cannot solve their problems are outcompeted by those that do." This is only slightly more informative than "the entities in dynamic systems persist so long as they are stable" or "organisms that are most able to survive, survive." However, it is only at this abstract level of description that content is lacking. When we pick a specific trait in a specific environment, "those that survive,

survive” becomes “those finches with long beaks gain greater access to seeds in specific environments, causing greater production of similar offspring.”⁶ The uninformative tautology has become an informative and highly contingent claim.

The point is thus that functional explanations in the social sciences are not inherently empty anymore than they are inherently analogous or disanalogous to biological processes. It all depends on the details — on the extent to which they can provide convincing evidence that the three conditions for functional explanations hold for specific social practices.

Seeing this point will allow us to dispense another standard worry about functional explanations, viz. namely that they provide no mechanisms [Elster, 1983]. “Mechanisms” is a buzz word that can cover quite different things. In particular, mechanisms might be vertical or horizontal. If the Fed influences GDP in raising interest rates because doing the latter increases savings, the savings are a horizontal mechanism, one between the two social entities. However, there is another sense of mechanism that refers to the lower level entities-individual agents — that realize social processes. These are vertical mechanisms.

Arguments for one kind of mechanism need not be arguments for the other. Thus Elster’s claim that we must have individualist mechanisms to prevent spurious correlations makes no sense for vertical mechanisms, for they are not intervening variables. Describing everything individuals did in bringing about the social event of the Fed’s raising interest rates will not control for other aggregate variables that confound our evidence that interest rates cause increased savings.

Furthermore, identifying horizontal mechanisms is not necessary to confirm causal claims. The standard clinical trial of a new drug, for instance, usually involves randomization to a control and treatment group with the aim of showing that the drug has an effect; the mechanism of action may not be known at all, but if the trial is conducted correctly we can make well confirmed causal claims. (And if the thought was that the problem is not one of confirmation but explanation, we explain the outcome without citing the intervening causal process as well.) In general we can check for all possible confounding causes (or mechanisms) by using Mill’s methods without know the precise mechanism actually at work.

Let me turn next to three important, confusion-causing complexities in functional explanation that we should note: their relation to functionalism as grand social theory, their relation to nonfunctional causes, and the kinds of questions they may and may not answer. Keeping these in mind can help avoid needlessly unproductive criticisms and defenses of functionalism.

Functionalism can either be taken in the grand theory sense as a full theory of

⁶ Sober [1984] defends Darwinian natural selection against the common charge of being a tautology in this way and in more detail.

society or given more modest aims. Parsons, of course, wanted a total theory as did Marx — a complete functionalist explanation of all aspects of society. Functionalism as social theory claims to be a full explanation of social organization, social change, etc. Every aspect of society is related to its ability to promote societal needs. Two points about functionalist social theory of such great ambition are worth making. Specific functional explanations might be compelling without committing us to explaining everything in the social realm by its function or even committing us to the possibility of grand social

theory in general. Given that many deny that functional explanations are ever successful, it will be a useful result to show how particular functional explanations can work well, even if the prospects for a grand functionalist theory is not defended. As a result, my focus here is on specific, low level functional accounts, though I will have some things to say about functionalism as grand social theory later when I discuss functional explanation as componential analysis.

Another point to note, to avoid some common confusions, concerns the relations between functional and nonfunctional explanation. Frequently the two are treated as mutually exclusive. They need not be for two related reasons. Functional explanations are a kind of causal explanation. Causal explanation involves picking out salient aspects of a complex causal field, i.e. picking out an element from the totality of factors causally involved. So we might have a functional explanation for why A persists and yet there might be further, nonfunctional causes involved in explaining A's persistence. In graphic terms, not all cases must be of type 1 rather than type 2:

[INSERT FIGURE 1]

Moreover, it is frequently helpful to think of explanations as answering questions ([Garfinkel, 1981]; [van Fraassen, 1981]). Questions, however, have their content determined in interesting ways by context. In particular, context sets implicit contrast classes for an appropriate answer. If I ask "why Adam ate the apple?" then there is an implicit contrast class for Adam (as opposed to whom), ate (as opposed to what other action), and apple (as opposed to what other object). Thus an answer to one question with a specific contrast need not be an answer to all other versions of the same question — hence the basis for the Willy Sutton joke: "Question: Why do you keep robbing banks? Answer: That is where the money is." So it is possible that there are good functional answers to a question with one set of contrast classes but not for another. In brief, functional explanations need not answer all questions to answer some.

It is also helpful to distinguish various kinds of evidence that might be advanced in favor of functional explanations. We can place those kinds on a continuum running from direct demonstrations that the three conditions hold to indirect plausibility arguments that infer that they do. Direct evidence would show that a practice, institution, etc. had a certain causal effect and that having that effect caused it to persist. If we have data measuring the persistence of institutions, practices, etc. along with data about presence and absence of institutions and their alleged effects, then standard causal modeling exercises are possible. ⁷ As we will see below, such data is sometimes available.

However, such decisive data often are not on offer, and social scientists then must resort to more indirect and contentious kinds of evidence. Let me mention several and their corresponding difficulties: trait environment correlations, optimality arguments, and stability arguments.

Anthropologists studying small scale societies sometimes claim that different social practices are correlated with differences in ecological environment. They then infer from this that the practices provide benefits in these environments and persist because they do so. Biologists have long used such evidence in support of natural selection explanations of animal traits. But such evidence has plenty of room for confounding. Showing correlation is not necessarily showing causation, and moreover, is not showing that the trait persists because of its effect. Inferences thus have to be made with care.

Optimality arguments try to establish that some practice would be the best way to solve a specific problem, then find instances of that practice, and then infer that it exists in order to solve that problem. So economists argue that the best way to maximize profits is to equate marginal revenue with marginal cost, and then conclude that when firms do so behave, they do so in order to maximize profits — either by rationally seeing that this works best or by economic survival of the fittest. Similarly, ecological anthropologists will argue that some foraging practice maximizes caloric intake and infer that it exists in order to do so.

Optimality arguments are only as strong as the inputs that go into them. We have to ask: have we measured all the relevant costs and benefits of the practice? Do we have good evidence about which practices are actually feasible? Do we have good reason to think that there is some process picking out the optimal? Optimality arguments are convincing when we can confidently answer a ratively to all three questions, questionable to the extent that we cannot.

Stability arguments do not show that a practice is the best, only that once a trait comes to dominate, no other alternative practice is likely to replace it. There are forces in place that cause the practice to persist because of its e ffects. Put another way, we show that the practice is a local optimum.

Optimality and stability arguments are ubiquitous in some parts of the functionalist literature. There are two kinds of broad mechanisms that are sometimes invoked to underpin optimality and stability arguments: individual maximizing and selectionist processes. The former assumes in some fashion that individuals do the best they can, given their constraints and that in the process of doing so, create or sustain certain institutions, practices, etc. Those practices then can be said to persist because of their contribution to individual maximizing. Selectionist accounts describe the differential replication of social practices. Traits that produce differential replication are thus explained as persisting because of their e ffects. These mechanisms are spelled out in varying degrees and varying ways throughout the social sciences o ffering functional explanations.

⁷For discussion of causal modeling in the social sciences, see Woodward, in this volume.

With the account of functional explanation developed above and these various distinctions in hand I turn to the question where, how and with what success these sorts of explanations are employed in sociology and anthropology. There is no prospect here of answering these in blanket fashion as we have seen others try to do; we have to work case by case, asking if the three requirements are met. However, we can gain some generality in two ways: by considering functional explanations of broad social phenomena and by considering basic general mechanisms such as rational choice and selective processes. Throughout I focus on these two mechanisms as they are used to explain organizations, norms, and inequality.

4 RATIONAL CHOICE, NORMS, AND INSTITUTIONS

Let me turn first to functional explanations of norms, in particular explanations given by advocates of the "new institutional sociology," rational choice and game theory. These approaches use directly the functionalist language of "in order to." So Nee and Ingram

[1998] in a volume on the new institutional sociology argue that individuals jointly produce and uphold norms to capture the gains from cooperation. Aoki takes the same general perspective, but considers specific cases: he argues that property rights and community norms serve to regulate external diseconomies in a commons domain and that trading club norms in the medieval Middle East existed to promote efficiencies in the face of informational deficiencies [2001]. These are all functional explanations precisely because they claim that specific norms make some benefit possible and persist because they do so.⁸

Let's look at a paradigmatic example in more detail. In *Analytic Narratives*, Greif seeks to explain "the process of state formation, the determinants of political and economic performance, and the politics of resource mobilization" in late medieval Genoa [1998]. During the 12th century Genoa was formed as a self-governing city republic headed by elected counsels. Genoa's ruling elite came from clans formed around feudal landholders. Those clans sought out profitable enterprises and sought to use the consulate to advance their own interests. Despite these potential conflicts, for much of the 12th century there was no military confrontation between the clans. However, economic growth was slow as the clans did not collaborate in acquiring external acquisitions. In the late 12th century the consulate system broke down. After some years of civil war between the clans, the podestra system (basically a city manager) replaced the consulate system and led to peace and economic expansion based on joint external acquisitions.

Greif explains the cooperation of the 12th century and the establishment of the commune in terms of a repetitive complete information game in which the strategies are cooperate or challenge. The commune was a mutual deterrence equilibrium realizing the strategy set (cooperate, cooperate). However, the extent

⁸ There are complex issues about exactly what constitutes norms and institutions that I am skirting here. For further discussion of norms, see Rouse, this volume.

of cooperation was restricted by the incentives for investment in military resources. Thus, Greif ties Genoa's slow economic growth and lack of joint international acquisitions to the commune. The greater cooperation in the podestra system was made possible by increased external military threats (viz. from Barbarossa), which increased the incentive for cooperation over military investment and thus also led to joint international acquisitions and thus economic expansion. The podestra system is shown to be an equilibrium outcome in a collusion game between the clans and the podestra.

In our functionalist terms, then, Greif gives us multiple functional explanations. The consulate system—the set of norms that made it possible — existed in order to capture gains from cooperation. The podestra system — again a complex set of norms defining the functioning of the office itself and the interaction of the parties — existed for the same reason in the presence of an external threat.

How do these explanations fair? Many issues are at stake. In the remainder of this section I focus on two questions: (1) Are these explanations inherently superior as economists often claim because they do not take norms as given but explain them based on constrained optimization? (2) Are the optimality arguments given for these explanations as rational choice explanations compelling?

Advocates of explanations like those listed above think they have a virtue that many sociological explanations do not, viz. they do not invoke unexplained norms. They

explain norms themselves and do so in terms of individuals doing the best they can, given their situation and desires. In short, norms are explained, and they are explained in a way that does not make unwelcome assumptions about the malleability of individuals and the power of culture. These explanations provide a compelling mechanism behind the existence of institutions and norms. As a result, they avoid the common charge that functionalist explanations have no place for agency.

I argued earlier that appeals for mechanisms had to be evaluated in the context of what else we know. That seems to be the case here. Is every account of norms or institutions in terms of socialization, i.e. learning, automatically suspect? That seems a strong claim in the face of vast bodies of research in social psychology and elsewhere. It is likewise not obvious that the thesis of optimizing behavior has any automatic precedence. We have many well confirmed experimental results suggesting that individuals do not always optimize, because they follow norms of fairness, because they make mistakes in probabilistic reasoning, or because they have inconsistent preferences or objectives. So, the value of explaining norms in nonnormative terms is an open empirical issue.

Furthermore, it is doubtful that game theory models eliminate norms in any case. Games require that there be rules of the game defining payoffs, possible moves and their timing, the preferences of actors, and what about all of the above is common knowledge among the players. These are typically taken as given, not explained. Yet these elements are kinds of norms or institutional facts. Moreover, as we will see below, many games have multiple possible equilibria. Which equilibrium is hit upon is usually explained in terms of "salience" or "focal points," again assuming some preexisting norms or conventions.

The examples described above certainly do not explain in an institutional and norm free way. Greif takes the structure of feudal society as given. His agents are clans defined by their social position as landholders. This is a rich social structure with an accompanying set of norms, roles, etc. that spell out the institutions. The possible strategies of the clan are also taken as given, resulting from the previous history of Genoa. These descriptions are the framework on which his game theory account works. Greif explicitly acknowledges all of this [1998, 47].

I would again argue that this dependence on social structure has important lessons for the general philosophy of science debates mentioned in the introduction. Game theory is a useful technology or tool. But by itself it does not unify the social sciences or support individualism. It does not have that much content; content has to be filled in to apply it. Greif's actors are clans not individuals. The game theory analysis works by building on sociological understanding, not replacing it. The connections between economic, political and psychological accounts of the social world are not made simply by the use of a common technology, but by the substantive attempts to interconnect them in concrete explanations.

So Greif's explanations cannot be judged superior on methodological grounds alone; they have to be supported more directly. The explanations given here are rational choice explanations. They are supported by arguing that the norms in question are strategies in a game that would be adopted when a Nash equilibrium is reached. A Nash equilibrium is the situation where each person's choice constitutes the best reply to the choices of the others. Rational individuals will see this and act accordingly. In short, we have an optimality argument for a rational choice explanation of the institutions or norms in

question. This is standard fare for rational choice game theory accounts of social processes.⁹

Let me discuss the general problems that face these kinds of functional explanations and then ask how well Greif handles them. As in all optimality arguments, we want to know that these explanations are based on an accurate specification of the game that is being played—the costs and benefits, strategies, preferences, etc. of the actors—and on good evidence that there is some mechanism bringing about the optimal. I think there can frequently be doubts about both.

Determining what game is actually being played, if any, can be difficult because many of the defining components of the game may be unobservable. The classic question here is inferring the utility functions of the agents, when only behavior is observed. This is complicated by the fact that real agents occupy many roles at once and thus may be playing multiple games, including metagames—games about what games to play (see [Ross, 2005]). Needless to say, this can make presenting a convincing case quite difficult.¹⁰

Equally serious difficulties arise for the claim that there is a mechanism picking out the optimal. Traditional game theory accounts of institutions are rational

⁹For further discussion of rational choice models, see Pizzorno, this volume. ¹⁰Economists working in industrial organization — market structure and concentration issues — confront this problem often as they explicitly acknowledge [Sutton, 2001].

choice explanations. The institution is a Nash equilibrium where each player chooses his or her strategy as best response to all other strategies. The agent is assumed to know the possible strategies, payoffs, etc. and then to deduce rationally the best strategy.

One problem is that many games that seem relevant to explaining institutions do not have a single Nash equilibrium. In fact, for many situations, nearly any strategy can be an equilibrium. So rational deduction from knowledge about payoffs cannot be the whole explanation of the given institution — how did the agents hit on the same equilibrium? Obviously something essential is missing.

Suppose the claim is that, for example, a specific norm exists in order to capture gains from cooperation. Then we are claiming that the norm's function is the complete cause of its persistence. The mechanism supposedly realizing that process is alleged to be rational choices of individuals. However, what rational individuals should choose in the situation of multiple equilibria is underdetermined, so we have good reason to think our claim is false — that the norm's benefits are not the full cause of its persistence.

One standard answer to this problem from game theorists is that some equilibria are more salient — they are focal points. However, then the norm is not explained by its contribution to gains from cooperation alone. It may still be partially so explained, but the rest of the explanation comes from whatever explains the salience point. That is likely to be preexisting norms and institutions. This illustrates my earlier point that something can be both functionally and nonfunctionally explained without the explanations necessarily competing, depending on what we are trying to explain.

This can be put in terms of contrast class and background information as we suggested earlier. We need to distinguish between two questions: Why is there a norm (rather than none at all)? versus why this norm rather than that? The first may get a functional explanation when the second does not. We may have good evidence that there is a norm rather than none because individuals rationally saw that having some norm

would allow gains from cooperation. But we may also have evidence that the reason this particular norm was picked rather than others was not rational deduction from known payoffs, etc., but has to do with something else in the causal field, viz. the common culture that makes some strategies more salient than others.

Another route to avoiding the multiple equilibria problem is to deny that it exists. One way to do that is take the heroic route of Aumann [1986] and claim that rational agents facing the same information, etc. will have to reach the same conclusions. That strategy seems committed to a full logic of inductive inference, something we have good inductive grounds to think unlikely. An alternative route to eliminate multiple equilibria is to require more than just Nash equilibria. This is the project of the so-called “refinements” literature in game theory. The emerging consensus is that this project is not productive or, in other terms, too productive, for there are numerous extant proposed refinements.

Aside from the multiple equilibria problem, game theory rational choice mechanisms also make quite unrealistic rationality assumptions, even when multiple equilibria are not present. Individuals are required to calculate equilibria from information about the game that many smart students cannot do with training. Sometimes there is no known algorithm for calculating equilibria and they can only be found by a trial and error process. Frequently they require players to be rational Bayesian agents, something the empirical literature tells us we seldom are. For a great many rational choice game theory accounts in the social science, these problems raise serious doubts that the phenomenon in question has been shown to exist in order to maximize player’s joint utility functions.

How does Greif’s work stand up in the face of such problems? Fairly well. By taking as given quite a bit of the existing social organization and by making extensive use of historical context and evidence, he can make a fairly convincing case about the games that were being played. And while he invokes the subgame perfect Nash equilibrium concept (a refinement of Nash), a charitable interpretation of his account would not take the Genoese as actually calculating their best replies in that equilibrium. Implicit here is a learning mechanism where the relative payoffs — the nature of the game being played — are adjusted to reach the equilibrium that allows the podestri system to survive. For example, to make collusion with one of the clans by the podestri unlikely, various incentives and information pathways have to be controlled, which the Genoese found ways to do. However, this mechanism is not modeled or even verbally described in any detail, and it is a mechanism that falls outside rational choice game theory — in the evolutionary or selective mode of analysis. It is to those sorts of instantiations of functional explanations that we turn to next.

5 INEQUALITY

The work discussed in this section looks at functional explanations of inequality. This research falls somewhere in the middle of the rational choice — pure selection continuum. Individuals doing their best are either explicitly invoked or consistent with the accounts given. Yet formal rational choice explanations, game theoretic or otherwise, are not given nor is there explicit appeal to differential survival processes. Nonetheless, the claim is made that inequality exists because of its effects and various kinds of evidence are given. These characteristics are representative of much functionalist work

in the social sciences.

A long tradition of sociological research seeks to explain inequality in income, wealth, status, and authority in terms of the traits of individuals. A well-known and influential paradigm of this work is Blau and Duncan's study of the American occupational structure. For Blau and Duncan, "a pervasive concern with efficiency" leads to outcomes where "achieved status becomes more important than ascribed status." [1967, 430]. Individuals are sorted by their better ability to do the job. Using path analysis and a large data set, Blau and Duncan support these claims by examining the relative influence of one's own education, first job, father's education and father's occupation on occupational mobility.

Blau and Duncan's findings were expanded upon by further studies, with other individual variables and new data sets added. *The Bell Curve* [Herrnstein and Murray, 1994] is one of the later installments. The functional explanation common to all this literature is the claim that inequality exists in order to promote efficiency. Mechanisms are not made explicit: there is no appeal to rational choice maximizing models to provide the mechanism nor is there any explicit appeal to differential survival processes. So what is the evidence? It is only various correlation between traits potentially connected to efficiency. It is hard not to conclude that the evidence is weak.

A second question about Blau and Duncan's account of inequality concerns its completeness. We noted earlier that functional causes do not necessarily preclude other, nonfunctional causes from acting at the same time. These explanations of inequality are a case in point. While they put forth their account as if efficiency considerations explained all, there are good reasons to think the situation is considerably more complex. If inequality is functionally explained, it is so only partially. At most, inequality exists in order to promote efficiency, given preexisting social stratification.

There are two ways to see this. The first comes from the following useful analogy. Suppose we have a pen full of dogs of different sizes with differing abilities to capture bones thrown into the pit. Suppose also that dogs prefer larger bones to smaller ones and that they fight it out according to that preference when a bag of bones is thrown into the pen each day. When the fighting settles down, we will have canine inequality. Is that inequality explained by differences in ability to perform the job of bone snatching? It cannot be entirely. The distribution of bone sizes in the bag is a key component — bags with large variances will lead to large inequality, those with smaller variances, smaller inequality. For Blau and Duncan's case, the parallel is the preexisting set of occupations. Ability may determine which rung of the ladder you stand on, but it does not determine how far apart the rungs are or how many there are.

A second reason to think that these functional explanations are only part of the story — only explain given an already existing social structure — comes when we try to provide a relatively realistic picture of the process whereby individuals are assigned to jobs. When we say that the best individuals are assigned to the job, what is the reference group — best of which lot? In the models of perfect information and competition of neoclassical economics, the answer would be: best of all individuals. But any sociologically realistic account would realize that individuals are not matched with jobs that way. Much research [Granovetter, 1995] shows that individuals find jobs through social networks. Those networks are determined at least in part by preexisting social divisions of race, ethnicity, and the like. A realistic account of the mechanism behind the functional explanation appeals to other causal factors in addition to efficiency.

This is another instance of our general moral that functional explanation presupposes rather than replaces explanations in terms of social structure.

Explanations of inequality in the Blau-Duncan tradition are individualist in spirit — they aim to explain everything in terms of the traits of individuals. An alternative functionalist approach to inequality that avoids individualism is that of Tilly [1998]. Tilly argues that inequality cannot be explained in terms of traits of individuals, for there are “categorical” inequalities existing independently of individual differences in ability. Categorical inequality is the social grouping of individuals into types, backed up by systematic processes reinforcing those divisions. These divisions may be based on conscious discrimination, but need not be.

Tilly identifies several basic mechanisms behind categorical inequality: exploitation, opportunity hoarding, and emulation.¹¹ All involve using existing social categories such as race as a means to further individual interests. Exploitation uses social categories to deprive individuals of their full contribution to a joint activity. Opportunity hoarding comes about when a resource is confined to individuals of one group. Emulation happens when individuals form new organizations based on existing social categories.

Tilly has given a functional explanation: “categorical inequality persists [because it] facilitates exploitation and opportunity hoarding [and] solves organization problems.” Unlike Blau and Duncan, Tilly’s functional explanation does not take social structure as given but seeks to explain it. Again, the mechanisms are not explicit rational choice models nor are they explicitly put in terms of differential survival.

There are two main questions to ask about Tilly’s functional explanations: is categorical inequality always the best way to solve problems? And does categorical inequality exist solely because of its functions? We can have doubts of the first kind because discrimination has costs — slaves rebel — and yet Tilly’s evidence almost entirely focuses on the benefits to individuals. Unfortunately, we also have good reason to think that categorical inequality persists for reasons other than its value in exploitation, etc. We have lots of evidence, much of it experimental and robust, that: humans naturally tend to ascribe environmental causes to individuals (the “fundamental attribution error” — [Ross, 1977]); tend to blame the victim [Lerner and Miller, 1978]; reason via stereotypes, including social stereotypes [Bar-Tal *et. al.*, 1989]; display in group bias [Tajfel and Turner, 1979]; and blame themselves when they are discriminated against [Milner, 1983]. In short, there seem to be natural response patterns in humans that make categorical inequality likely. Tilly’s functional explanation is unlikely to be the whole story.

6 SELECTIONIST MECHANISMS

The functional explanations of inequality discussed in the last section are not full rational choice explanations, in that they do not explicitly state and solve a maximizing under constraints model. And though they invoked competition, they were also not explicit selectionist models of differential survival. In this section we

¹¹Tilly actually adds a fourth — what he calls “adaptation.” Adaptation is solving problems using existing social distinctions. However, that seems to me simply to be a more general description of what is going on in the case of exploitation, opportunity hoarding, and emulation.

turn to explicit selectionist underpinnings for functional explanations. There has been an

explosion of research on such social processes in the last decade, research that is arguably some of the most promising social science around.¹²

What is a selectionist mechanism for a functional explanation? Obviously, it must be a mechanism that realizes the three basic conditions. As suggested in Section 3, selectionist mechanisms in the abstract have very broad scope ([Kincaid, 1996]; [Harms, 2004]). The root idea is competition — differential growth (positive or negative) of entities of different kinds. When we identify the traits responsible for differential growth, then we know that those traits persist because of their effects. At this level of abstraction, nothing commits us to the full details of a standard biological natural selection process involving genes, phenotypes and reproduction.

Because selection mechanisms can be described very generally, there are numerous different ways they can be instantiated in concrete causal processes and likewise diverse ways of formulating and evaluating claims about those processes. It will be useful upfront to distinguish the various parameters that can be combined to produce specific selectionist mechanisms and evidence for them. In particular, selective processes may differ on: the extent to which human reproductive fitness plays a role, the kinds of things that compete, the kinds of evidence supplied, and the extent to which process are modeled formally.

We need to distinguish several different claims about selective social processes and their relation to genetic inheritance. There are multiple parameters that we might want to consider, among them being the extent to which it is cultural selection that explains rather than biological, the extent to which it is current reproductive value that explains as opposed to past, and the extent to which the explanation is a complete one or not. Species without learning and under strong selective pressure that have traits that are fully explained by relatively recent biological reproductive value represent one pole. Species with extensive learning capability acquired through a distant history of biological selection and minimal existing selective pressure but intensive current cultural selection are on the opposite end of the spectrum. Various combinations in between are possible. Exactly where humans fall on this spectrum is controversial.

We also need to ask what kinds of things we are trying to understand. Selectionist mechanisms are used to explain at least three different targets: the traits of individuals, the traits of groups, and the traits of social systems, i.e. of individuals or social groups that stand in some defining relationship to one another. The traits of individuals might be a propensity to commit a crime or to hold certain beliefs. The traits of groups might be a marketing strategy of an organization. The traits of a social system might be its type of production relations.

The distinction between selection of individuals and selection of groups hides some ambiguities that we should address here. An emerging consensus on biological and social evolution [Sober and Wilson, 1998] sees natural selection as a

¹²For further discussion of selection and evolutionary explanations, see Whimsatt and Haines, this volume. multi-level process that can act at various levels in the biological and social hierarchy, claims that the selective effects at the various levels can be identified and combined into a unified explanation, and is optimistic that cultural selection can integrate social phenomena with the biological. A corollary of this consensus is that evolutionary game theory may provide a universal language for the social and biological sciences [Gintis, 2000].

This consensus is too quick. So-called “multilevel” selection models are multilevel in only one, relatively weak sense; the consensus account of group selection includes groups only in an anemic sense. Common to the current consensus is the following picture: group selection of a trait occurs when the trait is differentially distributed in different groups in a population and those groups with a higher frequency of the trait are thereby more fit in that group size increases relative to other groups. In this situation the frequency of a trait can increase in the population as a whole, even though it may be less fit in each group. If the effects on group productivity are strong, the trait can evolve.

Note two things about this notion of group selection. The fitness of the group is defined by ability to increase in size — to increase the number of *individuals* in the group. Thus the unit of measurement is individual organisms specified by trait. It is this choice of unit that makes an integrated multilevel account possible: the effects of genic, individual, and group selection are compared in terms of differential survival of individual organisms of specified types.

Though dominant, this is not the only extant notion of group selection. To see that there is another conception, consider the basic requirements of the Darwinian mechanism, viz. individuals that reproduce individuals similar to themselves, variability in the traits of individuals, and differential survival of some kinds of individuals into the next generation. It is perfectly possible that these apply to groups of individual organisms *per se*. So group selection can occur when there are different kinds of groups that produce new groups that resemble them, when groups vary in their traits, and those traits have varying influences on the next generation. This is group selection where the units of measurement are groups, not individual organisms. If a trait leads to more groups of one kind, there can be group selection regardless of what happens to the number of individual organisms in them. Arguably this notion of group selection is what various biologists and social scientists have had in mind. It was explicitly contrasted with the current consensus notion in the mid 1980s ([Kincaid, 1986]; [Damuth and Heissler, 1987]).

The complications introduced by group selection in the second sense have not received sufficient attention. Group selection in the multilevel sense studies a different dependent variable than that based on survival of groups. Thus, the claims of multilevel selection to integrate both group and individual processes and the biological and social are suspect. There are also complex issues surrounding the very idea of selection “acting at a level” that I cannot address here. But at the very least it is important to keep the two different senses of group selection — differential survival of individuals because of group membership and differential survival of types of groups — distinct.

Two more important parameters concern how explanations are formulated and the kinds of evidence given for them. In formulation, the difference is mathematics: are the causal selective processes given some kind of explicit mathematical treatment? In the case of evidence, the difference is between direct types of evidence — measuring the causal influence of A on B and then the influence of that causal relation on A’s persistence — and more indirect evidence that infers the presence of a causal mechanism from a model with various idealized assumptions that somehow fit the facts.

I want to look at four examples of research that span the range of the differences listed above and to say something about their relative success: Hannan and Freeman’s work in organization ecology, Bowles’ evolutionary game theory account of property

rights, Richerson and Boyd's cultural selection theory, and Smith's work in evolutionary or behavioral human ecology.

I look first at organization ecology which provides a mathematical formulation of the selective process, takes the traits of groups to be the relevant target, provides direct evidence, and provides an account where biological selection plays no role. While I use Hannan and Freeman's work as an exemplar, theirs is but a part of a much larger literature that shows, for sociology, an usual degree of continuous research building on past results [Pfeffer, 1993].

Hannan and Freeman develop a model of the differential birth and death of organizations, the kinds of environments they face, and the strategies that are likely to do best in those environments. If organizations change slowly and if they compete for resources, then there should be differences in the founding and death of organizations as the number of organizations increases, assuming a fixed resource base. In other words, selection of organizations should be observable. If environments differ in significant ways, then we might expect that different kinds of organizations — organizations with different traits or strategies — would predominate in different types of environments. Those traits would then contribute to organizational survival and persist in the population because they do so. They would exist in order to take advantage of the specific environment in question. They would be functionally explained.

Borrowing from the ecologist Levins [1968], they distinguish three types of environments: stable, coarse grained variable, and fine grained variable environments. Intuitively, an environment is stable if the resource is always present, coarse grained variable if the resource comes in "lumps," and fine grained if it comes and goes at short, repeated intervals. Hannan and Freeman then take the further concepts of "specialist" and "generalist" from ecology, where specialists rely on a single resource and generalists on a range.

With these distinctions in hand, there are some natural predictions about which strategies should be present in which environments, assuming that selection is operative. Generalists will do best in variable coarse grained environments and specialists best in stable environments. If environments are sufficiently fine grained, then specialists may dominate there as well. These distinctions thus support some fairly precise functionalist predictions at the theoretical level.

Hannan and Freeman are able to tie these theoretical predictions to some impressive data: large original data bases on environment types and organizational strategies for restaurants and semiconductor firms. They control for a number of possible confounding factors — age effects, size effects, etc — and find the correlations between environment types and strategies holds up well.

Such work in organization ecology is a vivid counterexample to the various skeptical claims about functional explanation in the social sciences. Selectionist mechanisms can be precisely formulated that are instances of the general pattern described in Section 3 and that do not require biological equivalents of genes. Organizations are sufficiently stable and defined to undergo birth and death. The tautology problem — organizations that survive, survive — is solved in precisely the same way it is in evolutionary biology: by turning the abstract selection claim into concrete, defeasible claims about the causal interaction between types of traits and the environment. Nothing in these explanations entails that society as a whole is a throughgoing, stable entity or that change does not happen.

It is also useful to note what this work says about controversies over individualism, the role of the biological, and levels of selection. These are macrolevel explanations that abstract from the causal processes at the individual level. The selective process is group selection in a straightforward sense that makes no commitment to individual success. There is clear recognition that there is more to society than the culture of the evolutionary psychologists — which is only ideas in individual brains — and no prospect of explaining these phenomena solely in terms of psychological modules developed in the Pleistocene. The models developed can be represented in evolutionary game theory, but that in no way gives us a general theory relating individual behavior to the phenomena at hand.

A second, large body of work emphasizing selective mechanisms is found in evolutionary game theory. Bowles [2004], for example, provides an interesting evolutionary game theory account of the emergence of property rights among hunter-gather early humans. The main elements of the model are:

1. Individuals of four types: *Sharers*, where two sharers split a good evenly when they interact *Grabbers*, where grabbers get the good when they meet sharers and fight when they meet each other, gaining the good or cost of defeat with equal probability *Punishers*, who divide equally with each other and sharers, and who try to punish the grabber and either gain the good or bear the cost of defeat
- 1 Cultural transmission process where those strategies that do best are most likely to be adopted by the next generation
- 2 Group selection, measured in terms of differential survival of individuals as influenced by the average payoff per group
- 3 Second order punishment: individuals who do not punish are punished

Bowles shows that in this model there are two stable equilibria that cannot be invaded: the Hobbesian, where the grabbers outnumber the sharers and there is frequent fighting and no punishers, and the Rousseauian, with a mixture of sharers and punishers and no grabbers. However, when the possibility of property rights are added to the mix, then a mixed strategy is possible: act like a grabber when you possess something, otherwise act like a sharer. This is the Bourgeois strategy from the chicken and egg game. Bowles runs simulations that show that Bourgeois invades Rousseauian equilibriums to the extent that there is not uncertainty about possession. The transition from hunter gathering to farming made possession much easier to determine than under the presumably Rousseauian equilibrium during the hunter gather period. The upshot is that property rights persist once settled agriculture comes on the scene because it maximizes the success of groups with property rights.

Two obvious problems confront Bowles' account: First, our evidence about the details of human environments, social structure, biology, etc. in such distant time periods is sufficiently thin that any account is likely to be underdetermined by the evidence. Secondly, Bowles' models depend on numerous unrealistic assumptions, some of which can in no way be seen as rough approximations or idealizations. I have nothing original to say about the first problem which has been discussed at length elsewhere and moreover, it is not applicable to contemporary uses of evolutionary game theory. The

second problem is, and it is thus worth saying a few words about the problem.

There is a considerable temptation among scientists and philosophers to advance simple universal criteria that tell us when a simplified model succeeds. If it is “insightful,” “provides understanding,” and “similar” to the system it is modeling, then it is a good model. These defenses are inadequate. Insight and understanding are subjective states of investigators; warm and fuzzy feelings are not well known signs of reliability. “Similarity” is just too easy, since everything is similar to everything else in some respect. But what is the right respect? A well defended model with idealizing or false assumptions requires showing that we have good reasons to believe the causal processes identified in the model are actually at work in the world in the way the models says they are. Bowles provides no such argument, and much of evolutionary game theory applied to social phenomena is in the same boat.

One route around this requirement is to weaken the conclusions of the model. It is commonly said that idealized models at least show possibilities. But there are possibilities and possibilities. To describe a logically possible world is not to describe a socially possible one. One needs evidence that the simplifications and falsifications do not produce a socially impossible world. Possibility in a model and real world possibility are not the same; getting from one to the other requires argumentation.

I turn next to the selectionist models of Boyd and Richerson [1985; 2005a; 2005b]. They have developed one of the more sophisticated accounts of selectionist mechanisms with close attention to the interaction of cultural and genetic selection. For Richerson and Boyd, the target of selection is culture, and culture is information stored in the head of individual humans. When information in the brain of one individual produces behaviors that increase the odds that related information will appear in the head of another individual, we have a selective processes.

Boyd and Richerson distinguish three “forces” affecting cultural evolution: guided variation, biased transmission, and natural selection. *Guided variation* is learning. Using some specific standard that might be either innate or learned itself, judgments are made about the ability of various cultural variants — ideas, beliefs, etc. — to promote that standard. Guided variation does not depend on the extent of existing cultural variation, for new variants are generated in the learning process. *Biased transmission*, like guided variation, is differential adoption of culture based on fit with some standard, but biased transmission works only on existing cultural variants. Technological diffusion is a widely studied example. *Natural selection* occurs when one cultural variant spreads more than another without any evaluative standard determining whether it is adopted. Cultural practices that lead to greater fertility would be a case point.

In terms of the model developed in Section 3, all three mechanisms are forms of differential survival and thus possible grounding for functional explanations. However, Boyd and Richerson are particularly concerned with identifying the circumstances under which cultural selection would evolve by natural selection, and distinguishing these different forms is essential for that task.

The paradigm case for Boyd and Richerson is selection for individual traits — the information in the head. However, Boyd and Richerson do at various places allow for the possibility of cultural group selection [2005]. As is typical, their discussions are ambivalent on the exact processes involved. At one point they discuss tribes with competing norms and use differential survival of the *tribe* as their metric. At other

points [2005], they envision group membership causing differences in survival of types of *individuals*.

Boyd and Richerson rely primarily on previous studies by others to support their claims, and much of their efforts are analytical — separating possible causes — rather than explanatory. They cite extensive literature arguing that culture matters. They cite an equally extensive literature describing the processes of transmission of culture from individuals to individuals.

Let's look at two possible criticisms of Boyd and Richerson's work: (1) culture is not particulate in the way the selection models requires [Kuper, 2001] and (2) they illegitimately treat societies as populations of *individuals* [Fracchia and Lewontin, 1999]. Culture is a complex integrated set of meanings; it does not come in natural units. Societies are not just a population or collection of individuals; they are organized and structured.

Both criticisms have force. For Boyd and Richerson, information is always "in the head." Yet there is plentiful work from recent cognitive science that suggests this picture is much too simple. Cognition is embedded in larger physical and social environments in essential ways [Clark, 1998]. That social environment does indeed have organization that makes it very different from the standard population of population genetics, their paradigm. Hannan and Freeman's work on the differential survival of organizations shows that if anything does. The moral is thus one we have seen before: selectionist models are not universal theories in and of themselves and they often depend upon social science to provide them with the framework within which selection works.

Both criticisms also overreach. Some aspects of culture are relatively particulate. Boyd and Richerson refer often to work on the diffusion of innovations. New inventions certainly look like individual units that can be counted. If the claim is only that they are bundled with other information, Boyd and Richerson can (and do) concede that sometimes that it is complexes that are transmitted and selected. They do not emphasize such situations and there may be an argument that they are much more important than Boyd and Richerson grant. But that would not seem to show they always are so or that when they are bundled, no selectionist model is possible. Runciman [1989], for example, explicitly invokes a selectionist style model in the style of Boyd and Richerson for cultural practices, which are precisely interlocking sets of roles, beliefs, etc.

If the social world is not simply a population, it does not follow that nothing about it can be described in population terms. Hannan and Freeman's work does exactly that, albeit with a population of organizations.

I turn finally to Smith's [1981] instantiation of human behavioral ecology. Human behavioral ecology is a trend used largely in the study of small-scale human societies. Like Hannan and Freeman, this approach borrows models from evolutionary ecology. The standard usage is of optimality models, frequently optimal foraging models. Smith studied Inuit hunting practices, in particular group hunting size. Assuming that groups maximize caloric intake, Smith develops a model of the optimal group hunting size, derives predictions from it, and tries to show that it matches reality. The result is supposed to support a functional explanation: Inuit hunting group size exists in order to promote caloric intake.

What is the mechanism for this explanation? It is not some kind of cultural selection for hunting practices. Human behavioral ecology in general claims that social practices

exist in order to promote current reproductive fitness. Caloric intake, Smith argues, is a good proxy for reproductive fitness, hence it is measured.

Smith represents the opposite pole of the continuum from Hannan and Freeman. For them, selection acted on social entities and biological selection was nowhere in sight. For behavioral ecology, selection acts on genes and social selection seems nowhere in sight. While it is certainly imaginable that some social practices persist in part because they promote current reproductive fitness — religious beliefs concerning fertility are a possibility — it seems highly implausible that all or even the majority do. This seems right even if you think that any cultural selection account has to be explicitly integrated with biological processes. Boyd and Richerson, for example, describe gene-culture coevolutionary models of some plausibility that show that maladaptive cultural practices can evolve by natural selection. There would seem to be rather large problems for the entire human behavioral ecology program.

So the upshot is that there are a variety of possible selectionist processes that could undergird functional explanations. They provide a range of different substantive takes on the relation of the biological to the social and the nature of social explanation. They vary in their ability to provide convincing evidence. Yet it is hard not to think that some versions of the selectionist still represent one of the more promising trends in modern social science.

7 FUNCTIONAL EXPLANATIONS AS SYSTEMS ANALYSIS

My focus throughout this chapter has been on explanations that claim social practices exist in order to fulfill their functions. However, much of functionalism past (Radcliffe-Brown, Parsons) and present (Luhman) has been about identifying the components of social organization and their interactions. In other words, functionalism has been about the analysis of social systems.

Put in these terms, functionalism is, as Davis [1959] argued long ago, not essentially distinct from social science in general. Even at this most general level of functional theorizing, there are many interesting and surprisingly unexplored issues concerning the extent to which there is a unique and meaningful way to divide the social world into societies [Tilly, 1984]. Functional accounts, however, generally add further constraints to the very general idea of analyzing social components and these raise further questions. Frequently, it is claimed that the components have to fit together in special ways — ways that realize or are compatible with emergent properties of the social system. There are different ways to pursue these claims and difficulties in each case of identifying clear theses. And always in the background are the prospects for functional explanations in the stronger sense — of identify what things exist in order to do rather just what they happen to do.

Obviously the issues here are broad and numerous, too much so to say anything compelling in a short space. However, this aspect of functional explanation is undeniable and thus cannot be ignored. I close with a sketch of some of the issues.

We can distinguish at least the five following questions about functionalism as systems analysis. There are no doubt interesting connections between these issues, but they are relatively distinct.

- 1 What are the systems being analyzed?
- 2 Are there universal components of any social system that suffice to explain its basic characteristics?
- 3 What are the components of social systems and how are they identified?
- 4 Is there one right way to divide a social system into components, and if not, what does that say about the reality of those components?
- 5 How does identifying components explain?

In its purest form, traditional functional analysis — with Durkheim, Radcliffe-Brown, and Parsons being prime examples — gives the following “organicist” answers.

1 System being analyzed: The system being analyzed is “society,” where that is identified with a clearly bounded unit and where the boundaries are spatial, cultural, and political. Small scale societies like the Neurer or modern nation states are paradigms.

2 Components: There are basic components that every society has and it is these that are fundamental to any explanation and comparison. Kinship systems or role and statuses are paradigm examples.

3. (a) What are the components? Here there is of course not complete agreement, but there is a basic set that is common to much functional explanation and much standard social science (note the relevance of Davis’ claim in this regard): kinship systems, social roles, norms, which identify, reinforce, etc. social roles, values, which might be equated with individual beliefs, public symbols, etc. and social structure, which describes at a minimum the relationships between social roles.

(b) How are components identified? The most explicit answer comes from Parsons (not a good sign!) who apparently thinks that the basic components of society are the necessary preconditions for human action (in some writings) or the necessary conditions for social stability or survival (in others). Seeing this Kantian a priori trend in Parsons makes greater sense of where he is generally coming from (see [Munch, 1988]). Other functionalists in this camp focus on societal needs more directly and see that relation as more contingent and empirically determined. Components are identified by their characteristic effect on social system functioning, usually described as stability or survival.

3 Unique way to divide: Since the components identified in (3) are universal, there are not other competing, equally legitimate analyses into components nor are there competing and equally valid ways of dividing societies into the universal components themselves.

4 How does componential analysis explain? Functionalists in this tradition vacillate in this regard. Often they claim that when components can be linked in a lawful relationship — relating for example social structure and kinship — then explanation has been achieved. At other points the explanation comes from showing the contribution of components to social stability, survival, or needs.

Associated with the organicist strand are several deep assumptions in the philosophy of science: that theories are essential and central to science, that theories describe lawful relations between elements, and that explanation comes from picking out those relations. These are a central part of the broad philosophy of science issues raised by functionalism that we identified in the introduction.

There are reasons, on my view, to doubt the organicist picture at nearly every step.

Tilly [1984] has argued persuasively that focusing on society as a whole gets the scale wrong. From the anthropologists we get a related concern: that societies are not in fact cohesive cultural units in the sense that there is one culture shared by all. Rather, culture takes on diverse forms in different contexts and social locations, with the actual practice being a fluid, ongoing matter of negotiation.¹³ It might seem obvious that every society has roles, norms, values, kinship systems, etc. But behind this truism are lurking some hard questions: What exactly do roles, norms, etc. come to? Social scientists use them in different and perhaps incompatible ways. Norms, for example, are sometimes taken as regularities in behavior, sometimes as requiring specific beliefs on part of the actor, and other times as requiring sanctions. So a first important task is simply getting a clear conception of these components.

A second related concern arises from trying to spell out the relations among these allegedly universal social components. Parsons, for example, distinguishes between the cultural, personality, and social subsystems present in every society. Yet elements of each seemed to be intertwined with all the others [Parsons, 1951], so it is unclear what the relation between them is. Are they real independent entities that can stand in causal relationships? The problem here parallels a long and relatively unproductive literature trying to specify what “culture” is and how it might separate anthropology from sociology.

A third set of problems concerns the uniqueness and reality of social components as well as their ability to fully explain. Are there multiple ways to divide up societies into “subsystems” as Parsons calls them? When are such divisions competing explanations and when are they complimentary? Should we take these entities as “real” or merely as useful instruments? These issues get surprisingly little attention in the social sciences or, for that matter, in the philosophy of social science literature. A natural route to approach them would be with functionalism of the third sort I distinguished early on — by asking if they represent real causal patterns that we can identify independently of their particular realization in individual behavior.

A further and fundamental issue concerns how the organicist tradition intends to explain by means of the components it picks out. We saw representatives like Radcliffe-Brown taking explanation to come from the citing of lawful relations between components. However, that presupposes a well refuted picture of explanation, viz. namely, the nomological deductive model. Organicists can move away from that conception of explanation by focusing on the causal contribution that components make to system survival or stability.¹⁴

¹³For development of these criticisms, see Risjord and Rouse, this volume.

¹⁴An alternative approach here would be to use cybernetic or homeostatic causal relations, something found in the functionalist tradition but not discussed here.

This suggests that componential analysis is perhaps likely to fare best with a selectionist framework already in place. The organicist strand in functional explanation is pre-Darwinian in that it treats societies as types with an underlying essence. Moving to a selectionist picture means dropping that assumption. It is then an empirical question what entities and at what level of aggregation and abstraction are sufficiently stable to compete over scarce resources. Answering that question would then be part and parcel of identifying social systems and their components or aspects. There may not be universal components with sufficient content to do much work in explaining social organization.

8 CONCLUDING MORALS AND OPEN QUESTIONS

The overriding moral of this chapter is that functional explanations in the social sciences work — not always and not easily, but sometimes. They can make illicit biological analogies or reduce culture to biology, but they need not. They can be empty or tautologous but that is not inevitable either.

Not surprisingly, assessing functional explanations raises numerous questions in the philosophy of science. Some have taken functionalism in its selectionist version as a universal language of social science that would unify the social sciences and the social and biological sciences. Some have thought that unification would vindicate some form of individualism. We have found reason to be skeptical about these claims.

The most promising version of functionalism is apparently some form of a selectionist process. Yet one of its most powerful instantiations — evolutionary game theory — runs directly into the problems raised by highly simplified models — the problems of when do they explain and how do we know? Keeping clear on which results are interesting intellectual exercises only and which are giving us well confirmed explanation of real world phenomena is an important and hard task.

A final pressing complex of questions concerns the scope functional explanations — how much of the social world can they account for? The selectionist variant faces this question in a form that actually arises for natural selection in biology — how much of development, life history, and large scale ecological organization can be explained in terms of natural selection? And to what extent are explanations of the former autonomous and necessary presuppositions of natural selection explanations themselves? We have seen that the same kinds of questions confront functionalism in the social sciences when it faces the issues of integrating selective processes using different metrics and in determining what games people are playing and why.¹⁵ There is clearly much interesting work to be done — conceptual and empirical — around the functionalist tradition in the social sciences.

¹⁵ For an interesting stab at seeing what some of the issues are, see [Laland *et al.*, 2000].

BIBLIOGRAPHY

- [Allen *et al.*, 1998] C. Allen, M. Beko and G. Lauder. *Nature's Purposes*. Cambridge: MIT Press, 1998.
- [Aoki, 2001] M. Aoki. *Toward A Comparative Institutional Analysis*. Cambridge: MIT Press, 2001.
- [Aunger, 2000] R. Aunger. *Darwinizing Culture: The Status of Memetics as a Science*. Oxford: Oxford University Press, 2000.
- [Baldwin, 1897] J. Baldwin. *Social and Ethical Interpretations of Mental Development*. New York: Arno, 1897.
- [Baldwin, 1895] J. Baldwin. *Mental Development in the Child and Race*. New York: McMillan, 1895.
- [Blau and Duncan, 1967] P.M. Blau and O. Duncan. *The American Occupational Structure*. New York: John Wiley and Sons, 1967.
- [Bowles, 2004] H. Bowles. *Microeconomics: Behavior, Institutions, and Evolution*. Princeton: Princeton University Press, 2004.
- [Boyd and Richerson, 1985] R. Boyd and P.J. Richerson. *Culture and the Evolutionary Process*. Chicago: University of Chicago Press, 1985.
- [Boyd and Richerson, 2005a] R. Boyd and P.J. Richerson. *Not By Genes Alone*. Chicago: University of Chicago Press, 2005a.
- [Boyd and Richerson, 2005b] R. Boyd and P.J. Richerson. *The Origin and Evolution of Cultures*. Oxford: Oxford University Press, 2005b.

- [Bryant, 2004] J. Bryant. An evolutionary social science? A skeptic's brief, theoretical and substantive. *Philosophy of the Social Sciences*, 34: 45-492, 2004.
- [Clark, 1998] A. Clark. *Being There*. Cambridge: MIT Press, 1998.
- [Cohen, 1978] G.A. Cohen. *Karl Marx's Theory of History: A Defence*. Princeton: Princeton University Press, 1978.
- [Comte, 1974] A. Comte. *The Positive Philosophy*. New York: AMS Press, 1974.
- [Davis, 1959] K. Davis. The myth of functional analysis as a special method in sociology and anthropology. *American Sociological Review*, 24:757-72, 1959. [Durham, 1991] W. Durham. *Coevolution: Genes, Culture and Human Diversity*. Stanford: Stanford University Press, 1991.
- [Durkheim, 1933] E. Durkheim. *The Division of Labor in Society*. New York: Macmillan, 1933.
- [Durkheim, 1961] E. Durkheim. *The Elementary Forms of Religious Life*. New York: Collier, 1961.
- [Durkheim, 1965] E. Durkheim. *Rules of the Sociological Method*. New York: Free Press, 1965.
- [Ellen, 1982] R. Ellen. *Environment, Subsistence and System: The Ecology of Small Scale Social Systems*. Cambridge: Cambridge University Press, 1982.
- [Elster, 1983] J. Elster. *Explaining Technical Change*. Cambridge: Cambridge University Press, 1983.
- [Faia, 1986] M. Faia. *Dynamic Functionalism: Strategy and Tactics*. Cambridge: Cambridge University Press, 1986.
- [Garfinkel, 1981] A. Garfinkel. *Forms of Explanation*. New Haven; Yale University Press, 1981.
- [Gintis, 2000] H. Gintis. *Game Theory Evolving*. Princeton: Princeton University Press, 2000.
- [Godfrey-Smith, 2000] P. Godfrey-Smith. The replicator in retrospect. *Biology and Philosophy*, 15(3): 403-423, 2000.
- [Granovetter, 1995] M. Granovetter. *Getting a Job*. Chicago: University of Chicago Press, 1995.
- [Greif, 1998] A. Greif. Self-enforcing political systems and economic growth: late medieval Genoa. In R. Bates, A. Greif, M. Levi, J-L. Rosenthal and B. Weingast, *Analytic Narratives*. Princeton: Princeton University Press, 1998.
- [Hallpike, 1986] C. Hallpike. *Principles of Social Evolution*. New York: Oxford University Press, 1986.
- [Hands, 2001] W. Hands. *Reflection without Rules: Economic Methodology and Contemporary Science Theory*. Cambridge: Cambridge University Press, 2001.
- [Hannan and Freeman, 1989] M. Hannan and J. Freeman. *Organizational Ecology*. Cambridge: Harvard University Press, 1989.
- [Harms, 2004] W. Harms. *Information and Meaning in Evolutionary Processes*. Cambridge: Cambridge University Press, 2004.
- [Harris, 1979] M. Harris. *Cannibals and Kings*. New York: Random House, 1979.
- [Heisler and Damuth, 1987] L. Heisler and J. Damuth. A method for analyzing selection in hierarchically structured populations. *The American Naturalist*, 130(4): 582-602, 1987.
- [Hempel, 1970] C. Hempel. *Aspect of Scientific Explanation*. Free Press, 1970.
- [Herstein and Murray, 1994] R. Herstein and C. Murray. *The Bell Curve*. New York: Free Press, 1994.
- [Kincaid, 1996] H. Kincaid. *Philosophical Foundations of the Social Sciences: Analyzing Controversies in Social Research*. Cambridge: Cambridge University Press, 1996.
- [Fracchia and Lewontin, 1999] J. Fracchia and R. Lewontin. Does culture evolve? *History and Theory*, 38: 52-78, 1999.
- [Laland et al., 2000] K. Laland, J. Olding-Smee and M. Feldman. Niche construction, biological evolution, and cultural change. *Behavioral and Brain Sciences*, 23: 131-175, 2000.
- [Layton, 1997] R. Layton. *An Introduction to Theory in Anthropology*. Cambridge: Cambridge University Press, 1997.
- [Lerner and Miller, 1978] M. Lerner and D. Miller. Just world research and the attribution process. *Psychological Bulletin*, 85: 1030-1051, 1978.
- [Little, 1989] D. Little. *Understanding Peasant China*. New Haven: Yale University Press, 1989.
- [Lyman and O'Brien, 1998] R. Lyman and M. O'Brien. The goals of evolutionary archeology. *Current Anthropology*, 39: 615-652, 1998. [Malinowski, 1944] B. Malinowski. *A Scientific Theory of Culture and Other Essays*. Chapel Hill: University of North Carolina Press, 1944.
- [Marx, 1861] K. Marx. A Letter to F. Lassal of 16 January, 1861. MECW 41.
- [Milner, 1983] D. Milner. *Children and Race*. London: Ward Lock Educational, 1983.
- [Munch, 1987] R. Munch. *Theory of Action: Towards a New Synthesis Going Beyond Parsons*. London: Routledge, 1987.
- [Nee and Ingram, 1998] V. Nee and P. Ingram. Embeddedness and beyond: institutions, exchange, and social structure. In M. Brinton and V. Nee (eds.), *The New Institutionalism in Sociology*. New York: Russell

- Sage Foundation, 1998.
- [Parsons, 1937] T. Parsons. *The Structure of Social Action*. New York: McGraw Hill, 1937.
- [Parsons, 1951] T. Parsons. *The Social System*. New York: The Free Press, 1951.
- [Pfeffer, 1993] J. Pfeffer. Barriers to the advance of organization science: paradigm development as a dependent variable. *Academy of Management Review*, 18:599-620, 1993.
- [Pettit, 1996] P. Pettit. Functional explanation and virtual selection. *British Journal for the Philosophy of Science*, 44: 291-302, 1996.
- [Radcliffe-Brown, 1948] A.R. Radcliffe-Brown. *A Natural Science of Society*. Glencoe, IL: The Free Press, 1948.
- [Rosch, 1978] E. Rosch. Principles of categorization. In E. Rosch and B. B. Loyd (eds.), *Cognition and Categorization*. Hillsdale, NJ: Erlbaum Seashore, R. H. (1947). ??????????????????
- [Ross, 2005] D. Ross. *Economic Theory and Cognitive Science: Microexplanation*. Cambridge: MIT Press, 2005.
- [Ross, 1977] L. Ross. The intuitive psychologist and his shortcomings: distortions in the attribution process. In: L. Berkowitz (ed.) *Advances in Experimental Psychology*, 10:173-220. New York: Academic Press, 1977.
- [Runciman, 1989] W.G. Runciman. *A Treatise on Social Theory*. Vol. II. Cambridge: Cambridge University Press, 1989.
- [Runciman, 2001] W.G. Runciman. Was Max Weber a selectionist in spite of himself? *Journal of Classical Sociology*, 1:13-32, 2001.
- [Smith and Winterhalder, 1992] E.A. Smith and B. Winterhalder. *Evolutionary Ecology and Human Behavior*. Hawthorne, N.Y.: Aldine, 1992.
- [Sober, 1984] E. Sober. *The Nature of Selection*. Cambridge: MIT Press, 1984.
- [Sober and Wilson, 1998] E. Sober and D.S. Wilson. *Unto Others: The Evolution and Psychology of Unselfish Behavior*. Cambridge: Harvard University Press, 1998.
- [Spencer, 1975] H. Spencer. *The Principles of Sociology*. Westport, Conn.: Greenwood Press, 1975.
- [Steward, 1955] J. Steward. *The Theory of Cultural Change*. Urbana: University of Illinois Press, 1955.
- [Sutton, 2001] J. Sutton. *Technology and Market Structure: Theory and History*. Cambridge: MIT Press, 2001.
- [Tajfel and Turner, 1979] H. Tajfel and J. Turner. An integrative theory of intergroup conflict. In W. Austin and S. Worchel, *The Social Psychology of Intergroup Relations*. Monterey: Brooks/Cole, 1979.
- [Tilly, 1984] C. Tilly. *Big Structures, Large Processes, and Huge Comparisons*. New York: Russell Sage, 1984.
- [Tilly, 1998] C. Tilly. *Durable Inequality*. Berkeley: University of California Press, 1998.
- [Tooby and Cosmides, 1992] J. Tooby and L. Cosmides. The Psychological Foundations of Culture. In J. Barkow, L. Cosmides and J. Tooby (eds.), *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*. Oxford: Oxford University Press, 1992.
- [Turner and Maryanski, 1979] J. Turner and A. Maryanski. *Functionalism*. Menlo Park, CA: Benjamin/Cummings, 1979.
- [Van Fraassen, 1981] B. Van Fraassen. *The Scientific Image*. Oxford: Oxford University Press, 1981.
- [Vayda, 1987] A. Vayda. Explaining what people eat: a review article. *Human Ecology*, 15: 493-510, 1987.
- [Weber, 1978] M. Weber. *Economy and Society*. Berkeley: University of California Press, 1978.
- [White, 1949] L. White. *A Science of Culture*. New York: Grove Press, 1949.
- [Wright, 1973] L. Wright. Functions. *Philosophical Review*, 82: 139-168, 1973.