

1-1-2004

Collective Intentionality, Complex Economic Behavior, and Valuation (book chapter)

John Davis

Marquette University, john.davis@marquette.edu

Published version. "Collective Intentionality, Complex Economic Behavior, and Valuation, in *Elgar Companion To Economics and Philosophy*, edited by John B Davis; Jochen Runde; Alain Marciano; John Bryan Davis. Cheltenham : Edward Elgar Publishing 2004: 386-402. [Publisher link](#). © 2004 Edward Elgar Publishing. Used with permission.

19 Collective intentionality, complex economic behavior, and valuation

John B. Davis

'We think, therefore we are.' (Shaftesbury, 1900 [1963], vol. 2, p. 275)

In this chapter I depart from the standard view of the individual in economics as an atomistic being to consider the individual as a socially embedded being. There are of course many different ways of understanding individuals as socially embedded; the conception I employ, however, is based on collective intentionality analysis, particularly as formulated by Raimo Tuomela. There is an advantage to economic analysis in doing this. Whereas other views of social embeddedness are holistic, and reason mostly in terms of social entities, collective intentionality analysis is explicitly an account of individuals, albeit in a particular kind of social setting. This makes it possible to compare the understanding of economic behavior that emerges from a collective intentionality analysis of individuals with the understanding of economic behavior associated with the standard rationality view of individuals as atomistic beings. Further, as an account of individuals, collective intentionality analysis also offers a way of understanding the seemingly paradoxical idea that individuals can be socially embedded and yet remain distinct beings. The basic idea derives from our understanding of first person plural intentions, or we-intentions. Only individuals form such intentions, just as only individuals form first person singular intentions, or I-intentions, but we-intentions effectively embed *social relationships in individuals*. This contrasts with holist accounts of social embeddedness that rather run the risk of eliminating individuals when they embed *individuals in social relationships*. Collective intentionality analysis thus allows us to both talk about socially embedded individuals specifically as individuals, and compare their behavior to that of atomistic individuals. Finally, since individuals form both kinds of intentions, combining accounts of behavior understood in collective intentionality terms – what I characterize as deontologically rational behavior – with accounts of behavior understood in instrumentally rational terms, offers foundations for a complete account of individual economic behavior. I suggest that economic behavior in such accounts should be considered complex.

Determining the extent to which individuals are deontologically rational rather than instrumentally rational in economic life seems to be in part an empirical

question of the extent to which individuals are active in social settings in which they express themselves in we-intention terms. In the discussion here, I restrict my attention to smaller, relatively cohesive, institutionally well-structured social settings – social groups – on the assumption that shared intentions are more likely to have specific consequences for individual behavior in these sorts of circumstances than in larger, more diffuse, loosely organized social settings.¹ My argument for this assumption is *not* that smaller social groups more effectively monitor or discipline individual action – this would reduce deontologically rational to instrumentally rational behavior – but rather that smaller groups have stronger prospects of producing determinate outcomes, and this reinforces individuals' commitment to their shared intentions. Compare, for example, the need an individual feels in a place of employment to act upon relatively well-defined intentions shared with other employees ('we need do our respective jobs to meet our production deadline') versus the lesser need an individual feels in larger, more loosely organized social settings to act upon vague intentions that may only be weakly shared by others ('we need to do something about inflation'). Focusing on mid-sized social groups, then, essentially operationalizes collective intentionality analysis for economics by emphasizing the kinds of behavioral consequences that have been the subject of atomistic individual rationality analysis.

Section 1 briefly discusses social groups. Section 2 then reviews Tuomela's contribution to collective intentionality analysis. In section 3 I turn to how the socially embedded individual conception explained in collective intentionality terms involves a view of individual economic behavior distinct from that involved in standard rationality theory. A different view of normative reasoning associated with a collective intentionality analysis of the socially embedded individual conception is discussed in Section 4. Finally, section 5 offers concluding remarks.

Social groups

Social groups that are relatively cohesive and well-structured have been extensively investigated in sociology, anthropology, and social psychology. A social group may be characterized as (1) a plurality of individuals tied to one another by (2) some principle of membership that implies (3) a system of individual rights and obligations. A social group is not the same thing as a social category (such as income class, ethnicity, gender, etc.), which researchers employ to classify or group individuals according to a set of characteristics which the researcher selects. Groups, rather, are collections of individuals whose shared characteristics derive from their interaction with one another. Of the characteristics of social groups surveyed in the recent literature, I emphasize the following as particularly relevant to the analysis of mid-sized social groups in economic life: that individuals engage in repeated interaction, that they define themselves as members of a group, that they are defined by others as belonging

to the group, that they share and observe group rules and norms, and that they participate in a set of interlocking roles that are central to how the group functions (Cartwright and Zander 1968, p. 48).

One advantage of construing groups in these terms is that it allows us to say that groups need not operate on a face-to-face basis. Much sociological, anthropological, philosophical, and social psychological research focuses on small groups that do exhibit regular face-to-face contact. Indeed, in the limit a relationship between two people can be seen as a kind of group if those individuals see themselves as being in some type of repeated interaction with one another and observe rules and norms which determine roles for them in the relationship. Margaret Gilbert uses as one of her main examples the idea of two people 'taking a walk together' (Gilbert, 1989). Were 'taking a walk together' a regular interaction between two individuals, on the understanding here they would constitute a group. More long-lasting relationships of all kinds between two individuals, then, would also qualify as instances of groups. In economic life, however, groups are generally seen as being larger, somewhat more impersonal, and not infrequently involve limited or even no face-to-face contact (firms, unions, cartels, governments, etc.), and this is the sort of case I focus upon in order to concentrate on individuals' social embeddedness. The two-person limiting case kind of group, however, would still be worth investigating in economics to explain recurring market interactions between individuals, where trust relationships are better modeled along the lines of group behavior than in standard atomistic individual terms. But I do not address this sort of case here, in order to avoid needing to include sympathy or empathy as elements or factors in the analysis.

Tuomela's analysis

The philosophical literature on collective or shared intentionality distinguishes we-intentions corresponding to use of 'we' language from I-intentions corresponding to use of 'I' language. We-intentions are explained as a structure of mutually reinforcing, reciprocal attitudes shared by individuals in a social group. Important contributors have been Bratman (1993, 1999), Gilbert (1989, 1996), Searle (1990, 1995), and Tuomela (1991, 1995). Others, such as Etzioni (1988), have also distinguished 'I' and 'we' thinking, without employing the idea that shared intentionality can be described as a structure of mutually reinforcing, reciprocal attitudes. An advantage of Tuomela's work is its specifically individual focus. Though he sometimes informally refers to intentions being shared, he also emphasizes that this is not meant to literally imply that we-intentions exist in society apart from individual we-intentions. Rather, a we-intention is defined as an individual's attribution of an intention to the members of a group to which the individual belongs, based on that individual both having that we-intention and also believing that we-intention is held by other individuals in the same group. That is, I can only use 'we' language that pertains to you and

I, if I think that you would similarly apply it to you and I. Thus, expressing we-intentions is a matter of whether there exists a set *reciprocal* attitudes, not whether there is an actual *sharing* of attitudes. Indeed, in the limit an individual could have a we-intention that no other individuals have, if that individual were simply mistaken about others' we-intentions. Thus, a we-intention is not a supra-individual group intention separate from the attributions individuals make to groups, and when people use expressions such as, 'the intentions of the group', this is just a shorthand device for referring to a collection of individual we-intentions on the part of individuals in the group.

Tuomela's analysis of what he regards as the 'standard case' is as follows. An individual expressing a we-intention assumes that it is mutually believed that the we-intention is held by other group members. Consider the case in which an individual's we-intention is rooted in an attitude ('fear'), which the individual believes other group members also attribute to the group. For an individual A who is a member of a group G, 'A we-fears that X if A fears that X and believes that it is feared in G that X and that it is mutually believed in G that X is feared in G' (Tuomela 1995, p. 38). 'X is mutually believed' if not only do I believe others believe X, but they also believe that I believe X.² On this basis, A might suppose that 'group G has some intention' reflecting 'G's fear of X' (say, whether the group will avoid some danger). Of course A can only surmise that others in G have the same fear and also that the fear of X is mutually believed by members of the group. The strongest case using the idea that X is 'mutually believed' (a shared belief) would involve saying that the fear that X is iteratively believed by everyone. But Tuomela allows 'mutual' to have strong and weak interpretations, because groups themselves have strong and weak criteria for supposing their members share a belief, attitude, or intention. The main point is that we-attitudes are a group attitude not in the sense that a group over and above its members has an attitude towards something, but in the sense that individuals 'generally' in a group have some such attitude that they express in 'we' terms. Thus saying that they 'generally' have a we-attitude depends not just on the mutual belief condition, but on both conditions which when combined provide us with a reason to suppose that individual members of a group are justified in saying what they (that is, 'we') intend.³

Tuomela uses this framework to distinguish between rules and norms, and I refer the reader to his work for a fuller account. Rules are the product of an explicit or implicit agreement brought about by some authority, and used to determine a distribution of tasks and activities to individuals. Rules may be formal and written, such as laws, statutes, regulations, charters, bylaws, etc., or they may be informal agreements between individuals, sometimes orally established and sometimes silently agreed to. In contrast, in the case of norms a network of mutual beliefs substitutes for actual agreements between individuals in determining distributions of tasks and activities across individuals. As with we-intentions

generally, mutual beliefs are beliefs reciprocally established between individuals, such that each believes that others have the same belief, and each also believes that others think the same about the others, and so on in a structure of reinforcing, mutually held beliefs.

Rules and norms are both understood to have motivational force, meaning that they constitute reasons for action on the part of the individuals who accept them. Indeed, rules and norms are typically framed as 'ought' principles, and impose requirements on individuals as members of groups in the form of specific prescriptions for individual action. Formally, individual A feels obliged to do X, because A is a member of the group with a we-intention whose consequence is a rule or norm to the effect, 'we believe members of the group should do X.' But rules and norms are different in virtue of the different means by which they enforce a distribution of tasks and activities among individuals (Tuomela 1995, pp. 22–24). The prescriptive force of rules derives from there being sanctions that apply, whether formal/legal or informal, to those individuals who do not observe them. In contrast to rules, sanctioning with norms takes the form of approval or disapproval on the part of others. Because norms are internalized by individuals, in that they themselves accept them as reasons for acting, individuals apply others' potential disapproval to themselves, as when feeling shame or embarrassment.

In Tuomela's framework, then, it can be said that rules are the basis for institutions, and norms are the basis for social values. While it is true that many institutions also involve norms, as relatively settled social arrangements, institutions generally place greater reliance on rules. In contrast, since social values are rarely rooted in agreements, even informally, they usually place little weight on rules. Rather, social values reflect systems of mutual belief about individuals' interaction with one another. Thus when individuals create and/or change institutions, they adopt new rules, and produce new we-attitudes that define group action within an institutional framework that can be characterized in terms of agreements and corresponding sanctions. When individuals develop and/or influence social values, they adopt new norms, and produce new we-attitudes that define group action within a social value framework based on their mutual beliefs and systems of approval and disapproval. In both frameworks, rules/institutions and norms/social values, we-intentions are the foundation for understanding group action. Individuals thus influence institutions and social values as members of groups, and group action is the intermediate link between individual action and supra-individual institutions and social values missing from mainstream accounts of individuals' influence on institutions and social values.

Finally, it is important to emphasize that rules and norms can create different types of obligations – sometimes implicitly, sometimes explicitly – for individuals in terms of how different tasks, rights, and positions apply to different individuals

in groups. Tuomela characterizes an individual's position within a particular group in terms of that individual's tasks and rights within that group. An individual's tasks and rights are then further distinguished according to whether they flow from rules or norms operating within the group, that is, whether they are rule-based tasks and rights or norm-based tasks and rights. In contrast, across groups, individuals' social positions are understood in terms of the whole array of actions that individuals are required and permitted to do across various economic and social settings. These social positions assign individuals a variety of different tasks whose performance is in each instance protected by rights, where these tasks–rights combinations may themselves exist within established modes of implementation that are also understood in tasks–rights terms. The overall framework thus explains individual rights and duties within and across groups in terms of tasks–rights pairs that ultimately have *we*-attitudes in groups as their foundation.

A revised view of individual economic behavior

The behavior of atomistic individuals is understood in instrumentally rational terms, because individual objective functions are defined solely in terms of individuals' own preferences. With no basis for action other than their own preferences, and putting aside that they might act out of habit or behave irrationally, atomistic individuals can do nothing other than seek to satisfy their own preferences. In contrast, when we treat individuals as socially embedded, we no longer say that individuals act only on their own preferences, because they also act in accordance with those rules and norms which function as 'ought' principles – what I have termed a deontologically rational or perhaps a rationally principled type of behavior. But proponents of the atomistic individual conception understand rule-following and norm-observance in instrumentally rational terms. Are 'ought' principles operating in social groups then better explained in instrumentally rational terms? Is instrumental rationality a sufficient explanation of individual behavior? There seem to be three objections to saying that something other than instrumentally rational behavior is involved here. I respond to each objection.

First objection

One way in which to argue that observing rules and norms which have the force of 'ought' principles is still fully within the compass of instrumental rationality, is to maintain that the individual becomes subject to constraints additional to those usually assumed in standard constrained optimization analysis, namely, constraints associated with group membership. Though these additional 'social group' constraints further narrow individuals' choice sets, individuals would still maximize preferences, suggesting that socially embedded individuals are not significantly different from atomistic individuals. This argument, however,

ignores what is involved in saying that individuals observe rules and norms on account of their sharing intentions with others regarding those rules and norms. Shared intentions are those intentions which individuals ascribe to the groups of which they are members. But as intentions, they must stem from individual objective functions rather than constitute constraints on those objective functions. That is, shared intentions are like our ordinary intentions in expressing what individuals *choose* to do rather than what they are *limited* to doing. It is true that individuals in groups are more constrained in their behavior as compared to when they act outside of groups. But this type of constraint has an intended aspect, and is consequently different from the usual sort of constraint that is entirely external to the individual.

Second objection

A second argument for explaining shared intentions in instrumentally rational terms accepts that shared intentions stem from individual objective functions, but argues that they express individual we-preferences, just as ordinary individual intentions express an individual's own preferences (or I-preferences). We-preferences have been analysed by Sugden, and characterized as team preferences (Sugden 2000; also cf. Bacharach 1999). Sugden explicitly rejects collective intentionality analysis as developed by Tuomela, Gilbert, and others, on the grounds that it assumes individuals are bound by obligations or 'ought' principles, which he regards as inconsistent with an account of instrumentally rational behavior (Sugden 2000, pp. 189–90). To preserve the latter, he reasons, the former has to go. This implies that rules and norms are things that members of teams prefer to observe rather than believe they ought to observe. Moreover, if we-intentions are really the product of we-preferences, then it seems that it is no longer necessary to say that individuals in teams (or groups) need to be treated as socially embedded, since the obligations or 'ought' principles they observe are what they prefer. Sugden essentially draws this conclusion when he argues that the 'existence' question regarding whether teams and other groups exist (and therefore can act as agents) is independent of the theory of instrumental rationality enlarged to include we-preferences. Were groups thought to be agents over and above their members, there clearly would be a stronger case for saying that their obligations and 'ought' principles were not always preferred by their members.

Sugden's argument, accordingly, depends on supposing that we-preferences do not really impose obligations or 'ought' principles on individuals. Why is it, then, that Tuomela and other proponents of collective intentionality analysis claim that this is a necessary dimension of we-intentions? The answer lies in their analysis of shared intentions as sets of reciprocal attitudes across individuals in groups. Though shared intentions are indeed individual intentions, unlike team preferences, which represent only what an individual independently prefers for

the team, an individual's shared intention is one element in a set of reciprocal attitudes. Thus, when individuals ascribe intentions to groups of which they are members, this represents not what they prefer to ascribe to the group, but rather what they believe to be the group's intention based on what they believe that they and other group members believe to be the group's intention. On this view, shared intentions imply 'ought' principles, because individuals share an intention over which they have very limited influence. Indeed, this combination of sharing an intention and having it stand over oneself helps explain the particular quality of 'ought' principles as binding precepts that individuals nonetheless embrace. Preferences, by comparison, have but one master, namely, the individual.

It is true that team preferences do have a shared aspect to them. But absent a set of feedback connections between individuals, such as Tuomela describes for shared intentions, the shared nature of team preferences is simply the result of an accidental alignment of individuals' *we*-preferences about teams of which they are members. Team members may happen to share preferences about the teams they are on. Yet if these preferences regarding the team begin to diverge, there is nothing in the interaction between individuals that brings about an adjustment in individuals' preferences regarding the team. Indeed, there are many examples of teams in the real world which operate on the basis of Sugden's team preferences, and as a result break down, simply because individuals are driven by what they prefer rather than by what they believe obligates them. The problem, basically, is that, with *we*-preferences, just as with ordinary individual preferences, *de gustibus non est disputandum*. That is, individuals retain their atomistic status, and the 'teams' of which they are members do not exist as teams in the customary sense of the term.

Third objection

These conclusions, however, suggests a third argument regarding how instrumental rationality might be sustained vis-à-vis collective intentionality analysis. Suppose that we treat rules and norms as conventions understood as coordination equilibria (Lewis 1969). Then using evolutionary game theory, individuals can be seen as instrumentally rational players who seek the best possible response to one another's individual strategies (a Nash equilibrium), and rules and norms can be explained as endogenously determined sets of reciprocal expectations. This would allow for a feedback/adjustment process, as operates in collective intentionality theory, but it would not explain this process in terms of 'ought' principles. Rather, following Hume's view of conventions, individuals find it in their *interest* to conform to rules and norms to which they expect others will conform. There are different ways of explaining why individuals would find this in their interest. Hume relied on sentiments of approval and disapproval, and indeed used this as the basis for his theory of justice. Since a system of justice implies 'ought' obligations, this game-theoretic/instrumental

rationality framework can also be argued to explain the 'ought' content of rules and norms, whether in moral or pragmatic terms. But in contrast to collective intentionality analysis, 'ought' principles in this instance derive from what individuals find to be in their interest.

In collective intentionality analysis, 'ought' principles derive from shared intentions, and shared intentions are explained in terms of individuals' use of first person plural 'we' language. Barring cases of deceit, first person plural 'we' language cannot be explained in terms of first person singular 'I' language, unless one denies elemental differences in human language, and engages in a reductionist sort of reasoning that has no support among linguists. In collective intentionality analysis, the reason that shared intention implies 'ought' principles is that they require a commitment on the part of the individual absent in the case of ordinary intentions that can be expressed in first person singular terms. Thus it seems clear that 'ought' principles that derive from collective intentionality analysis are not reducible to 'ought' principles that might emerge from a Humean framework. But this does not imply that the latter involves an unacceptable account of 'ought' principles, or that this account should be eliminated to produce one of 'ought' principles cast exclusively in shared intention and commitment terms. Rather it seems that both reductionist arguments should be rejected, because neither goes through, and because both are part of the view that thinking about individuals and society can be explained in terms of two inalterably opposed intellectual traditions: methodological individualism and methodological holism. Indeed, both accounts of 'ought' principles arguably have real world foundations. Just as there are teams that operate (often poorly) in terms of individual we-preferences, so there are 'ought' principles based on instrumentally rational behavior. Just as there are social groups that operate (usually more successfully) in terms of we-intentions, so there are 'ought' principles based on individual commitment.

My position is that individual behavior is complex in being rooted in both types of intentions. The challenge economists consequently face is in determining both the mix of types of behavior associated with different kinds of intentions, and in properly ascribing each kind of behavior to the correct real-world circumstances. Much mainstream economics, because of its adherence to the atomistic conception of the individual, imperialistically imposes instrumental rationality arguments on social settings where it does not apply. In using the wrong explanation in the wrong circumstances, mainstream economists impose 'thin' institutional explanations that overlook how the functioning of some social groups and institutional structures depends upon 'ought' principles stronger than can be explained in instrumentally rational terms. The holist economics tradition, in contrast, has at times been equally imperialistic, though in reverse direction, in using social whole-type explanations in circumstances for which they do not apply. These 'thick' institutional explanations overlook the extent to

which individuals are relatively free of shared intentional experience, as well as individuals' need for navigating across social structures.

I do not attempt here to set forth a specific account of individual economic behavior as complex. There are a variety of different ways in which the two can be related.⁴ But consider an example. Suppose an employee in a business is assigned a set of rule and norm-based tasks associated with doing a particular job. If one rule is to invoice customers by the end of the month, and the norm for how this is to be done is to include in the invoice a complete description of all purchases made by those customers, the individual assigned these tasks is likely still free to perform them in a variety of ways (inquire as to customer satisfaction, pursue follow-up orders, institute new record-keeping practices, etc.). How well individuals do their assigned jobs, then, can be a matter of the extent to which they also act on their own preferences regarding the way a job is best done. They consequently act in an instrumentally rational way when already behaving in a rationally principled manner.

Normative reasoning and the conception of the individual as socially embedded

Deontologically rational behavior need not be normative in raising explicitly ethical issues, but it can be. Moreover, a specifically normative deontologically rational behavior might take on a variety of value forms according to the range of values operating in social groups. Thus in contrast to the more narrow normative framework standardly associated with instrumentally rationality analysis, valuation in collective intentionality analysis is complex and multi-dimensional. Following Amartya Sen, we might refer to this enlarged normative framework as a 'deontic-value inclusive consequentialist' framework (Sen 2001, p. 64). As he explains it,

It is neither that 'the good' comes first, and then 'rights and duties', nor that rights and duties congeal first, followed by the good, but that they are linked concepts that demand simultaneous consideration. While considerations of freedoms, rights and duties are not the only ones that matter (for example, well-being does too), they are nevertheless *part* of the contentions that we have reason to take into account in deciding on what would best or acceptable to do. The issue surely is *simultaneity*. (Sen 2001, p. 61, emphasis in original)

Here I address how normative values might arise and operate not just in social groups but in organizations and institutions generally, or, as it has recently been expressed, whether we may treat '*values as partly endogenous to the economic system, and economic systems and their performance as partly functions of people's values*' (Ben-Ner and Putterman 1998, p. xvii; emphasis in original). I begin by contrasting the standard view of how normative values operate in organizations made up of atomistic individuals. Essentially following Hume, the standard

account attempts to explain how conventions that lack a normative character in themselves can nonetheless come to acquire the status of moral norms.

'Moral' sentiments in organizations and institutions

Hume took a system of justice to be a set of conventions that arise when individuals come to expect one another to behave in regular and predictable ways (Hume 1739 [1888]). Individuals abide by a system of justice, because they find it in their interest to conform to its rules when they expect others to conform to them as well. The idea that such rules are 'conventional' comes from supposing that there are different possible rules of justice, and those that actually come about reflect a history of contingent interactions between people. Nothing a priori moral underlies actual systems of justice, making them for Hume not 'partly' but entirely 'endogenous to the economic system.' But why, then, should such rules be thought normative in nature? Why should they be thought to be anything more than simply persistent regularities in social behavior akin to other regularities that have no one believes have normative content? Hume's view, based on eighteenth-century Scottish-school psychology of sympathy, was that conformity with such regularities evokes sentiments of approval, and failing to conform with them evokes sentiments of disapproval. When these sentiments become widely shared and become attached to an idea of the general good, he believed they may then be characterized as a *moral* approval and disapproval. A system of justice, then, is ultimately nothing more than a relatively settled set of conventional expectations between individuals chiefly concerned with their own interest that is reinforced by sentiments of approval and disapproval.

This conception has been modified and redeveloped in recent years by Lewis and Sugden. For Lewis, conventions are coordination equilibria (Lewis 1969). Coordination equilibria can be explained in game-theoretic terms with players acting on individual strategies to achieve a common expectation regarding which individual strategies offer the best reply to one another (that is, they are Nash equilibria). Hume's psychology of sympathy is replaced by the characterization of individuals in terms of strategies, but any norms that emerge are still conventional and entirely endogenous to the economic system. Sugden similarly explains conventions in terms of individuals' expectations of one another conforming to regularities in behavior, but adds a concern individuals are said to have over incurring others' resentment as an emotion underlying conformity to conventions (Sugden 1986, 1989). When this emotion operates widely to reinforce individuals' adherence to conventions, Sugden suggests that *normative* expectations obtain among them (Sugden 1998). But against this it might be said that the emotion of resentment deserves the label 'normative' as much as Hume's approval and disapproval deserves the label 'moral.' Sugden argues in reply that this criticism misses the point behind providing a Humean *naturalistic* analysis of values. 'In

such an analysis, the definition of a moral sentiment has to be naturalistic; one cannot then object that some of the sentiments allowed by the definition are not "really" moral' (Sudgen, 1998, p. 84).

In my view, such a response is question-begging. The claim that 'moral sentiments' are just that, namely, somehow 'moral', needs a stronger defense than the suggestion that it should be possible to explain moral values naturalistically, and that therefore there must exist such things as 'moral' sentiments. Indeed, making this sort of argument seems to involve exactly what G.E. Moore famously labeled the 'naturalistic fallacy' (Moore 1903). Note also, that the program of producing a naturalistic account of normative values is closely associated with the aim of producing a positivist interpretation of moral life. Sudgen is explicit about this connection, asserting that economists 'trained in a positivist tradition' must seek to explain normative values without 'assuming the existence of moral facts' (Sudgen 1998, p. 76). A moral fact is a fact about something being right or wrong, good or bad; etc. For example, one might say it is a fact – specifically a moral fact – that it is wrong to needlessly harm another person. To deny that moral facts exist is to say there is nothing in society or nature that can be described in normative terms as a matter of fact. Normative values, rather, are subjective judgments, and must accordingly be explained in terms of some sort of 'moral' sensibility people exercise and impose on the world. Two obvious problems arise with this sort of approach. First, the idea that human society can or should be described naturalistically is highly questionable, and has not surprisingly, long been contested in the history of social science. Sudgen and others in this tradition generally do not explain why a natural science approach to social science might be plausible, and I am skeptical that any good arguments exist on this score. Second, this approach creates a very strong problem for making the transition from 'is' to 'ought' (thus Moore's naturalistic fallacy). Sudgen addresses this problem by simply insisting on a re-definition of 'moral' in naturalistic terms. Whether this is a reasonable re-definition of 'moral', however, depends on whether the account of moral behavior that emerges on these terms captures what we ordinarily think is bound up with the normative.

What is it, then, that is most characteristic of the normative domain? We can begin to answer this question by emphasizing the implied content of the moral 'ought.' When 'ought' appears in an expression in a normative way, it indicates the presence of a moral obligation. If individuals do something because they believe they morally ought to, they do so out of a sense of moral obligation operating upon them. There are many ways of understanding what this sense of 'ought' involves, but following Kant (1785 [1948]), I take the minimum essential idea to be a matter of doing something because it is required irrespective of one's inclinations or desires. But then the idea of acting out of a sense of obligation is not what is involved in acting on a 'moral' sentiment. If one is motivated to respect a norm or convention, because one fears others' resentment or disapproval for failing

to do so, it is not a sense of obligation but an inclination that operates as one's motivation. Rather fear of others' resentment or disapproval has *replaced* acting out of a sense of obligation. The closest Sugden comes to referring to a sense of obligation, then, is when he argues that one of the virtues of his analysis is that it 'allows us to consider cases in which normative expectations and self-interest pull in opposite directions ... cases in which individuals follow conventions even though this is contrary to self-interest' (Sugden 1998, p. 83). But this is not evidence of acting out of a sense of obligation, since one might well be inclined to observe a convention that was contrary to self-interest.

Thus if we take the idea of acting out of a sense of obligation to underlie what is involved in moral 'ought' thinking, a naturalistic, moral sentiments type of approach does not capture what we ordinarily think is bound up with the normative. This gives us good reason to conclude that the Lewis–Sugden type development of Hume regarding how values operate in organizations and institutions as conventions is not an account of distinctively *normative* expectations. This in turn suggests that a naturalistic approach probably cannot explain how normative values arise and operate in organizations and institutions. Thus, since the Humean tradition derives from its starting point in the notion that individuals are naturalistically described as isolated from one another, and generally acting in their own interest, let us rather begin at a different starting point by describing individuals as embedded in social groups in the collective intentionality sense, and then ask how normative values might arise and operate in organizations and institutions. Two questions can be addressed. (1) Does this alternative strategy enable us to talk about what is most characteristic of the normative domain, namely, a sense of obligation that individuals have when they use 'ought' language? (2) Does this strategy provide us an account of the variety and range of different types of relationships between normative values in social life along the lines of Sen's 'deontic-value inclusive consequentialist' framework?

Question 1: Moral obligation in organizations and institutions?

The first question can be answered by explaining what must be involved in saying that socially embedded individuals, understood in a collective intentionality sense, have we-intentions as well as I-intentions. The Humean framework, by taking individuals as fundamentally isolated from one another, operates exclusively with I-intentions. As previously argued, the use of 'we' language generally creates obligations for individuals – and not just in the moral sense. The collective intentionality framework, particularly as developed by Tuomela, shares with the Lewis–Sugden account of convention the idea that individual expectations are established within a system of reciprocal expectations between individuals. But Tuomela's account is different in that individual expectations have as their object reciprocal sets of we-intentions rather than have as their object the I-intentions implicitly involved in the Humean framework. We-

intentions, as previously argued, create obligations for individuals, because the successful use of 'we' by an individual needs to conform to how others use that same 'we.' Outside of a requirement of using language correctly, this obligation does not exist for the use of 'I.' Of course some obligations which individuals recognize are pragmatic, and consequently do not have moral content. But on the interpretation here collective intentionality analysis is not naturalistic, and certainly not motivated by positivistic aims. Thus it is as reasonable to suppose that moral facts exist as to suppose that they do not. From this it follows that some of the obligations individuals observe are indeed moral in nature. Though the dividing line between pragmatic and moral obligations may often be difficult to draw, and though it may change over time, it seems there are many clear cases of each, and thus fair to say that individuals who form we-intentions and use 'we' language often operate under a sense of moral obligation.

So a collective intentionality framework, by operating with a conception of socially embedded individuals rather than atomistic ones, makes it possible to include a sense of moral obligation alongside individual inclination as a form of individual motivation. Turning to the second question above, then, what does the collective intentionality framework and the conception of individuals as socially embedded tell us about the range and variety of normative values in social life and the relationships between them?

Question 2: An expanded normative domain?

The emphasis on moral obligation thus far has rested on looking at moral obligation as something that particular individuals recognize. But a fuller characterization of the concept needs to see these obligations not just from the point of view of the individuals who have them, but also from the point of view of the individuals to whom they may apply. This suggests a concept of moral obligation which relies on an 'externalist' conception of the individual, where this is a matter of understanding individuals in terms of their relations to one another, in contrast to a concept of moral obligation which relies on an 'internalist' conception of the individual, where this is a matter of understanding individuals in terms of properties that apply to them independently of their relations to one another (Davis, 2003). An example of the latter is the Pareto efficiency standard, which employs an 'internalist' conception of the individual to explain normative recommendations that judge states of affairs according to whether one person is better off *ceteris paribus* all other individuals.

Externalist-individual normative concepts, it can be argued, just because they emphasize relationships between individuals, generally require that we give attention to a range of normative concepts that go beyond whatever particular normative concepts (say, regarding what is good) might constitute a particular individual's moral view. Thus to give any kind of detailed explanation of the moral obligations that one has to others, one typically also needs to have an understanding

of others' and one's own rights. But systems of rights are themselves generally embedded in broad social commitments to such ideals as freedom, equality, fairness, human dignity, community, justice, etc. Thus employing an externalist-individual type of normative concept typically commits one to examining an entire range of accompanying normative concerns. This also means that the connections between different normative concerns cannot generally be mapped out with any high degree of precision, making moral questions complex and often ultimately undecidable.

The idea of an expanded normative domain may be linked to a collective intentionality understanding of individuals' social embeddedness in terms of their involvements in social groups in the following way. Social groups generally have goals that help define them. Thus their members' we-intentions often concern a consequentialist type of moral reasoning, as when something is regarded as right on account of its helping bring about some outcome desired by the group. But this sort of consequentialist moral reasoning, when it is expressed in we-intention terms, also has independent concepts of moral obligation associated with it. Thus a particular individual using 'we' language in regard to what potential good consequences a group wants to bring about operates both with an idea that what is right is a matter of bringing about the relevant outcome, and also the idea that what is right is a matter of observing obligations upon oneself understood in we-intentions terms. This latter sense of right may draw in turn upon other ideals such as justice, equality, and dignity. This is one way of talking about a 'deontic-value inclusive consequential reasoning,' in which, '[i]t is neither that "the good" comes first, and then "rights and duties", nor that rights and duties congeal first followed by the good, but that they are linked concepts that demand simultaneous consideration' (Sen 2001, p. 437).

Concluding remarks

This chapter does not attempt to explain how instrumentally rational and deontologically rational economic behavior are coordinated. It does suggest, however, that this may in part depend upon the extent to which individuals are active in social group settings in which their behavior has an economic character, that is, where production, exchange, and consumption activities are engaged in. The standard view on the part of proponents of the atomistic individual conception is that behavior in groups can always in principle be decomposed into the behavior of instrumentally rational individuals. But this view has not stood up to scrutiny (cf. Kincaid, 1996), and in any case such arguments beg the central issue here, namely, that individuals act on we-intentions as well as on I-intentions. Thus the need to explain behavior as complex remains on the agenda. In closing I merely suggest a set of considerations that could figure in the way in which this issue might be addressed.

One possible view is that one of the two spheres somehow determines the boundaries of the other. Thus social groups might establish certain domains of activity in which individuals would act in an instrumentally rational fashion. As in the example at the end of the third section, the practices in a business firm, or in a department of one, might be to collectively delegate to particular individuals the responsibility of acting as they would find rational, subject to their observing the boundaries placed on that domain of activity by the group. Alternatively, instrumentally rational individuals might delegate domains of activity in which group considerations were regarded as primary. For example, were heads of households instrumentally rational in the market, they might nonetheless treat the household as a sphere in which customary relationships reflecting we-intentions would prevail. But this general model – one sphere determining the boundaries of the other – also suggests another model in which behavior in one sphere invades the boundaries and undermines the behavior of the other. For example, individuals may express we-intentions deceitfully, and act in ways that are contrary to them. Alternatively, groups may seek to impose rules and norms on individuals where mutual beliefs are absent.

One reason that instrumental rationality theory has been attractive in economics is that having a single model of analysis makes possible a high degree of logical and mathematical determinacy in economic explanation. But the consequences of achieving this precision are that certain types of behavior go unexplained, and possibly that the activity of individuals that is meant to be explained is misrepresented. Collective intentionality analysis constitutes one framework in which these risks might be avoided. The implication of this chapter is that a larger framework including that analysis which presupposes that economic behavior is complex is more likely to offer a more adequate account of economic behavior on both counts.

Acknowledgements

Thanks for comments on a previous version of this chapter go to Jelle de Boer, Alain Marciano, Jochen Runde, John Searle and Raimo Tuomela. Errors are my own alone. This chapter also appears in *Protosociology*, vol. 18 (2003).

Notes

1. A similar argument regarding the economic consequences of social proximity was made by Adam Smith, though he relied on sympathy rather than shared intentions as the underlying motive force.
2. Gilbert explains 'we' language in a similar way: 'A person X's full-blooded use of "we" in "Shall we do A?" with respect to Y, Z, and himself, is appropriate if and only if it expresses his recognition of the fact that he and the others are jointly ready to share in doing A in relevant circumstances' (1989, p. 199). Gilbert holds that individuals use of 'we' language constitutes a 'plural subject' (1989, pp. 199ff).
3. Tuomela draws on an account of mutual belief that has become fairly standard among philosophers which relies on the idea of a hierarchical set of beliefs iterated across individuals (Tuomela 1995, pp. 41ff). See Shwayder (1965, p. 257) and Lewis (1969, pp. 52ff) for early formulations.

4. For one example of how such an explanation might be produced, see Minkler (1999), where a 'commitment function' is added to a standard utility function representation of individual behavior. The individual is said to engage in a two-step iterative procedure with the first step corresponding to a response to group requirements and the second step corresponding to an instrumentally rational maximization of utility.

References

- Bacharach, M. (1999), 'Interactive team reasoning: a contribution to the theory of cooperation', *Research in Economics*, **53**, 117–47.
- Ben-Ner, A., and L. Putterman (1998), 'Values and institutions in economic analysis', in A. Ben-Ner and L. Putterman (eds), *Economics, Values, and Organization*, Cambridge: Cambridge University Press, pp. 3–69.
- Bratman, M. (1993), 'Shared intention', *Ethics*, **104**, 97–113.
- Bratman, M. (1999), *Faces of Intention*, Cambridge: Cambridge University Press.
- Cartwright, D., and A. Zander (1968), *Group Dynamics Research and Theory*, New York: Harper and Row.
- Davis, J. (2003), *The Theory of the Individual in Economics: Identity and Value*, London: Routledge.
- Etzioni, Armitai (1988), *The Moral Dimension: Toward A New Economics*, New York: Free Press.
- Gilbert, Margaret (1989), *On Social Facts*, London: Routledge.
- Gilbert, Margaret (1996), *Living Together: Rationality, Sociality, and Obligation*, Lanham, MD: Rowman and Littlefield.
- Hume, David (1739 [1888]), *A Treatise of Human Nature*, edited by L. Selby-Bigge, Oxford: Clarendon Press.
- Kant, Immanuel (1785 [1948]), *Groundwork of the Metaphysics of Morals*, translated by H. Paton, New York: Harper and Row.
- Kincaid, H. (1996), *Philosophical Foundation of the Social Services: Analysing Controversies in Social Research*, Cambridge: Cambridge University Press.
- Lewis, D. (1969), *Convention: A Philosophical Study*, Cambridge, MA: Harvard University Press.
- Minkler, L. (1999), 'The problem with utility: toward a non-consequentialist/utility theory synthesis', *Review of Social Economy*, **52** (1), 4–24.
- Moore, G. (1903), *Principia Ethica*, Cambridge: Cambridge University Press.
- Searle, John (1990), 'Collective intentions and actions', in P. Cohen, J. Morgan and M.E. Pollack (eds), *Intentions in Communication*, Cambridge, MA: MIT Press.
- Searle, John (1995), *The Construction of Social Reality*, New York: Free Press.
- Sen, Amartya (2001), 'Symposium on Amartya Sen's Philosophy: 4 Reply', *Economics and Philosophy*, **17** (1), 51–66.
- Shaftesbury, A. (1900 [1963]), *Characteristics of Men, Manners, Opinions, Times*, edited by J. Robertson, reprint edn, Gloucester, MA: Peter Smith.
- Shwayder, D. (1965), *The Stratification of Behaviour*, London: Routledge.
- Sugden, Robert (1986), *The Economics of Rights, Co-operation and Welfare*, Oxford: Blackwell.
- Sugden, Robert (1989), 'Spontaneous order', *Journal of Economic Perspectives*, **3**, 85–97.
- Sugden, Robert (1998), 'Normative expectations', in A. Ben-Ner and L. Putterman (eds), *Economics, Values, and Organization*, Cambridge: Cambridge University Press, pp. 73–100.
- Sugden, Robert (2000), 'Team preferences', *Economics and Philosophy*, **16** (2), 175–204.
- Tuomela, R. (1991), 'We will do it: an analysis of group intentions', *Philosophy and Phenomenological Research*, **51**, 249–77.
- Tuomela, R. (1995), *The Importance of Us: A Philosophical Study of Basic Social Notions*, Stanford, CA: Stanford University Press.