Social Actions, Social Commitments

Herbert H. Clark

Social actions are the stuff of daily life. Walking on crowded sidewalks, working with colleagues, or eating with friends—these are activities we cannot carry out alone. It takes coordination to avoid collisions, negotiate business, and share food. Most of these activities are *joint activities*—activities in which two or more participants coordinate with each other to reach what they take to be a common set of goals. Without joint activities life would be impossible. We do more than merely work *around* each other. We work *with* each other, and on a range of common goals. Humans come equipped for joint action—with what Levinson (this volume) calls the "interaction engine"—and they engage in it from infancy on (see Boyd and Richerson, Gergely and Csibra, and Liszkowski in this volume).

Joint activities are managed through *joint commitments* (Clark 1996). I can commit myself privately to doing something and then act on that commitment. I may tell myself, "I'll have a beer when I get home," and when I get home, I have a beer. But for you and me to do something together—say, shake hands—it is not enough for me to commit privately to grasping your hand, or for *you* to commit privately to grasping *mine*. We must act on a *joint* commitment to shake hands. The argument here is that joint commitments are essential to all true joint activities. They are the guiding force inside Levinson's interaction engine.

When we take on joint commitments, we ordinarily do so for the benefits they afford—in avoiding collisions, negotiating fair contracts, and sharing food efficiently. But joint commitments also carry risks. Some risks come from ceding partial control over one's actions to others. Once you and I are committed to shaking hands, you might crush my hand or withdraw at the last minute. Other risks come from the

indeterminacy of joint commitments. When you and I agree "to talk," we may have only a vague idea of what about. Later, you may draw me into topics I did not anticipate or want to talk about. Joint commitments have moral and emotional repercussions. We may be happy and trusting when they benefit us, but angry and reproachful when they do not.

The goal here is to show how joint commitments are the driving force behind joint activities. I will illustrate with two joint activities, one quite ordinary and the other equally out of the ordinary. The ordinary one is of two people assembling a piece of furniture. It allows us to examine the normal, cooperative course of establishing joint commitments. The out-of-the-ordinary joint activity is a famous study of obedience by Stanley Milgram (1974). That activity, in contrast, illustrates the risks and the moral and emotional consequences of joint commitments.

Partitioning Joint Activities

People do not just happen to do things together. In shaking hands with you, I cannot grasp your hand without some sense, belief, or trust that you are going to do your part, and do it here and now. The idea is that people coordinate their parts in joint activities by means of joint commitments. To illustrate, I will begin with the cooperative assembly of a TV stand.

Two people I will call Ann and Burton were ushered into a small room, given the parts of a commercial kit for a wooden TV stand, and asked to assemble the stand from its parts.[1] They took about 15 minutes and were videotaped as they worked. Consider a 20-second segment in which they attached a crosspiece onto a sidepiece. They did this in a sequence of five paired actions, as represented here:

	Ann's action	*Burton's action*
1	A gets cross-piece	B holds side-piece
2	A holds cross-piece	B inserts peg
3	A affixes cross-piece	B holds side-piece
4	A inserts peg	B holds side-, cross-piece
5	A affixes side-piece	B holds side-, cross-piece

In line 3, for example, Burton holds a sidepiece steady while Ann affixes the crosspiece onto it. This is a joint action pure and simple. The two of them coordinate their individual actions—Ann affixing one board while Burton holds the other one steady—to reach a common goal, the attachment of the two boards. Ann does what she does contingent on

what Burton is doing, and he does what *he* does contingent on what *she* is doing. Note that these five joint actions together constitute a joint action at a higher level, "attaching the two side-pieces to the cross-piece," and that, in turn, is but one segment of a still higher level joint action, "assembling the TV stand." So, assembling the TV stand emerges as a *hierarchy* of joint actions, which is typical of joint activities (Bangerter and Clark 2003).

But joint activities take more than these joint actions. If we look again at the 20-second segment, we discover Ann and Burton talking about what they are doing:

(1) Ann   Should we put this in, this, this little like kinda cross bar, like the T? like the I bar?

Burton   Yeah ((we can do that))

Ann   So, you wanna stick the ((screws in)). Or wait is, is, are these these things, or?

Burton   That's these things I bet. Because there's no screws.

Ann   Yeah, you're right. Yeah, probably. If they'll stay in.

Burton   I don't know how they'll stay in ((but))

Ann   Right there.

Burton   Is this one big enough?

Ann   Oh ((xxx)) I guess cause like there's no other side for it to come out.

Burton   M-hm.
[8.15 sec]

Burton   ((Now let's do this one))

Ann   Okay

Ann and Burton's talk is not idle. It is what allows them to arrange, agree on, or coordinate who is to do what when and where. Here, too, Ann and Burton carry out paired actions, but the pairs are turn sequences like this:

(2) Ann   Should we put this in, this, this little like kinda cross bar, like the T? like the I bar?

Burton   Yeah ((we can do that))

In the first turn, Ann proposes that they attach the crosspiece, and in the second, Burton takes up her proposal and agrees to it. The two of them proceed this way throughout the TV stand assembly. They make agreement after agreement about which pieces to connect when, how to orient each piece, who is to hold, and who is to attach.

As this example illustrates, joint activities ordinarily can be partitioned into two activities: (1) a *basic* joint activity; and (2) *coordinating* joint actions. Consider the two parts of Ann and Burton's assembly of the TV stand:

The *basic joint activity*, or joint activity proper, is what Ann and Burton are *basically* doing—assembling a TV stand. It consists of the actions and positions they consider essential to their basic goal—the assembly of the TV stand.

The *coordinating joint actions* are what Ann and Burton do to *coordinate* their basic activity. They consist of *communicative acts* about the basic activity.

It takes both sets of actions to assemble the TV stand. The first set effects the assembly proper, and the second coordinates the joint actions needed to effect the assembly proper. Ann and Burton surely see these two activities as different. What they were asked to do was "assemble a TV stand." If asked, "But weren't you talking?" they might have replied, "Oh yes. That was to figure out who was to do what."

To complicate the picture, communicative acts are themselves joint actions (Clark 1996). For each utterance, speakers and addressees must coordinate the speaker's vocalizations with the addressee's attention to those vocalizations, the speaker's wording with the addressee's identification of that wording, and what the speaker means with what the addressee understands the speaker to mean. In earlier work, I have called the process of coordinating on these points *collateral* communication. So just as basic joint actions are coordinated by communicative acts, communicative acts are coordinated by collateral acts (Clark 1996, 2004). I will say no more here about collateral acts.

It takes coordination, therefore, to carry out joint activities, and communicative acts to achieve that coordination. But to agree on a joint course of action is really to establish a joint commitment to that course of action. How is that done?

Establishing Joint Commitments

The very concept of joint commitment is a puzzle. In common parlance, we can speak of an ensemble of people making a joint commitment, as in "The football team is determined to play better next year," or "The orchestra will now play Brahms." Yet it is only individuals who can make commitments (or have intentions). The members of an ensemble can each make up his or her own mind, but they can hardly make up each other's minds. How, then, do joint commitments get created from individual commitments?[2]

Varieties of commitment

Individual commitments come in many types. Consider four types of commitments to "go for coffee at Joe's Café at noon":[3]

1. *Private self-commitment*. Privately, without letting anyone know, I can commit myself to *myself* to go for coffee at noon.
2. *Public self-commitment*. I can make the same commitment *in front of* you—say, by telling you of my plan. It is not that I commit myself to you that I will go for coffee at noon. It is just that I make my self-commitment *public* between us.
3. *Simple other-commitment*. I can commit myself *to you* that I will go for coffee at noon—say, by promising you I will go. Not only do I make my commitment public between us, but I grant you the right to hold me responsible for fulfilling it. People make other-commitments for a range of social obligations.
4. *Participatory commitment*. Suppose you and I agree to meet for coffee at noon. I commit myself *to you* to taking part in this meeting *just so long as* you commit yourself *to me* to taking part in the same meeting, and vice versa. Our individual commitments are conditional on *both* of us being committed to the joint action.

A joint commitment is simply the sum of the participatory commitments of its participants.

These four types of commitment differ in how binding they are:

1. *Private self-commitment*. If I commit myself privately to going for coffee at noon, I can change my mind, or fail to get there, with no consequences for anyone else.
2. *Public self-commitment*. If I tell you about the commitment and then change my mind or fail to get there, these changes are public, perhaps to my embarrassment.
3. *Other-commitment*. If, instead, I commit myself *to you* and then change my mind or fail, I expect you to hold me responsible for the consequences. Perhaps you told a friend I would be there, and you are angry that I disappointed him or her.
4. *Participatory commitment*. If, finally, you and I are *jointly* committed and I unilaterally change my mind or fail, I expect you to hold me responsible not only for my individual failure but for subverting our joint action—what we would have accomplished jointly. If you turn up at noon for coffee, I will have wasted your time and abused your trust.

Plainly, private self-commitments are easier to alter midcourse than public self-commitments and other-commitments. But the most binding of all are joint commitments.

Meeting for coffee at noon is what David Lewis (1969) called a *coordination problem*, and our agreement to meet is a solution to that problem. With the agreement, you and I establish the mutual belief that we each expect both of us to go for coffee at noon.

Joint commitments are subject to the *sedan-chair principle*. Suppose Susan and Tom are two porters carrying Veronica in a sedan chair. They cannot pick the chair up, or set it down, without doing it together. If one of them tries, they risk not only spilling Veronica onto the street, but injuring each other. Likewise, in assembling the TV stand, Ann and Burton cannot start, or stop, without doing so together. Acting alone risks causing harm. If Ann suddenly stops holding a sidepiece while Burton is screwing in a screw, she may damage the sidepiece, hurt herself, or hurt Burton.

Projective Pairs

It is one thing to characterize joint commitments, but quite another to say how they get established. One way is with projective pairs.

After Ann and Burton attached the crosspiece to the two sidepieces, they were at a choice point: What to do next. They needed to establish a joint commitment to a course of action. They could not count on such a commitment arising spontaneously and simultaneously. They had to *make* it happen, and they did it this way:

    (3) Burton   ((Now let's do this one)) [picking up the top-piece]
         Ann   Okay

Burton proposed to Ann that the two of them ("let's") assemble the top piece ("do this one") next ("now"). Ann took up his proposal by agreeing to it ("Okay"). In just two turns, they established a joint commitment to assemble the top piece next. They specified the ensemble ("us") and goal ("do this one") as well as the commitments to do their parts in reaching the goal.

This pair of turns is what Schegloff and Sacks (1973; see also Schegloff this volume) called an *adjacency pair*. In such pairs, one person produces the first part, and another person, the second part. The first part is of a type for which it is conditionally relevant for the second part to be of a type *projected* by the first part. Burton produced a *suggestion*;

that projected her *consent* to go ahead as the second part; and Ann immediately gave her consent with "Okay" (see Bangerter and Clark 2003)

What is needed here, however, is the more general notion of *projective pair* (Clark 2004). In Schegloff and Sacks's account of adjacency pairs, both parts must be turns at talk, yet in many situations, one or both parts of analogous pairs are gestural. Later in assembling the TV stand, Ann and Burton produce this sequence of actions:

>   (4)   Ann   [Extends hand with screw] So you want to stick
>   the screws in?
>   Burton   [Extends hand to take screw]

In line 1, Ann proposes that Burton stick the screws in. In line 2, Burton could take her up with "Okay." Instead, he extends his hand to take the screw. She construes that move as signaling consent roughly as if he had said "Okay." I will use the term *projective pair* to cover adjacency pairs as well as analogous pairs with gestures.

A projective pair, then, is a *proposal* plus an *uptake*. By *proposal*, I mean any signal that raises the possibility, at any strength, of a joint action or position by the initiator and addressees. By uptake, I mean any action that addresses that possibility. So when Ann makes her proposal to Burton, she is simply initiating a process. Burton has the options of accepting, altering, rejecting, or even disregarding her proposal. Here are examples of the four options.

Full acceptance of proposal. In 4, Ann proposes, "So you want to stick the screws in [extending her hand with a screw]." Burton could have replied, "No, you do it" or "Hold on. I've got a better idea," but instead he takes hold of the screw and thereby accepts her proposal *in full*. Now, they are jointly committed to transferring the screw from her to him.

Altered acceptance of proposal. In 5, Burton asks a question that projects "Yes, it is" or "No, it isn't" as uptake:

>   (5)   Burton   Is this one big enough?
>   Ann   Oh ((xxx)) I guess cause like there's no other side
>   for it to come out.
>   Burton   M-hm.

In line 2, instead of saying yes or no, Ann accepts an *altered* version of Burton's proposal, and he accepts her alteration, "m-hm."

Rejection of proposal. In 6, from another corpus (Svartvik and Quirk 1980), we find yet another pattern:

(6) Betty   what happens if anybody breaks in and steals it, —
                 are are is are we covered or .
      Cathy   Um — I don't know quite honestly .

Betty's question projects an explanation for what happens if anybody breaks in. Cathy is not able to provide that explanation, so she turns down, or *rejects*, the proposal that she do so. Not only does Cathy reject the proposal, but she gives a reason why she is rejecting it.

Disregard of proposal. In 7, Ann asks Burton a question, which projects a yes or no in agreement:

(7) Ann   They snap in? [said as she snaps the rollers in]
      Burton   [Silence, no visible gesture or response]

Although Burton presumably has heard Ann, he appears to *disregard* her proposal by going on without addressing it. He appears unwilling to let her engage him in the potential projective pair. He simply opts out.

Projective pairs are efficient ways of creating joint commitments. When Burton realizes that he and Ann need to plan their next joint action, he initiates a projective pair, "Now let's do this one," and Ann completes it, "Okay." With projective pairs people are really *negotiating* joint commitments. No matter what the first person proposes, the second person has options in taking it up. The joint commitments that emerge are shaped by them both.

Joint Actions and Joint Positions

What do people make joint commitments about? Recall that the first pair part of an adjacency pair is an action of a particular type (e.g., a suggestion), which projects an action of a second type (e.g., a consent) as the second pair part. A question projects an answer, a greeting a greeting, a request a promise, and so on. Action types like these have long been studied, from quite a different perspective, as illocutionary acts (e.g., Austin 1962; Bach and Harnish 1979; Searle 1969, 1975). Although the analysis of illocutionary acts has its problems, it still offers useful insights.

Illocutionary acts have usually been treated as autonomous. For Searle (1969), a question "counts as an attempt to elicit [certain information] from [the hearer] H." But surely, questions are more than that. In 5, when Burton asks, "Is this one big enough?" he is not trying simply to elicit "yes" or "no." He is proposing that Ann join him in establishing

whether or not a particular peg is big enough. In her uptake, she offers useful information, but without answering yes or no. Together, they establish a joint commitment to the proposition, roughly stated, that "the peg is probably big enough because there's no other side for it to come out." This proposition is not either Ann's or Burton's alone. It is their *joint position*—an amalgam of contributions by them both.

Projective pairs can also be used to establish *joint courses of action*. Recall this exchange from the TV stand assembly:

(3) Burton   ((Now let's do this one)) [picking up the top-piece]
     Ann   Okay

Burton suggests a course of action with one illocutionary act, and Ann consents to it with another. The result is a joint commitment to a course of action.

The idea, then, is that illocutionary acts are better viewed as participatory acts. A question is a question because it can be the first part of a projective pair in which the expectable second part is an assertion that "answers" it. The projective pair establishes a joint position. Likewise, a suggestion is a suggestion because it can be the first part of a projective pair in which the expectable second part is consent. The projective pair establishes a joint course of action. According to Searle (1975; see also Bach and Harnish 1979), there are four main types of illocutionary acts: assertives, directives (which include questions), commissives, and expressives.[4] In this system, assertives, questions, and expressives are used for establishing joint positions, whereas directives (other than questions) and commissives are used for establishing joint courses of action. All are used for establishing joint commitments.

The picture so far is this. Ann and Burton need to coordinate their actions and their positions if they are to assemble the TV stand together. They do that largely with projective pairs in which one of them proposes a next joint step, and the other takes up that proposal—accepting, altering, rejecting, or disregarding it. In this way, they negotiate joint commitments that are mutually satisfactory.

Emergence of Joint Commitments

No matter what the joint action, agreement must be reached, explicitly or implicitly, on at least five elements:

Participants. Who are to take part in the joint action?
Roles. In what roles?
Content. What actions are they to perform, or what positions are they to adopt?
Timing. When are the actions to take place, or the positions to be in effect?
Location. And where?

Let me call these *joint elements*. Reaching agreement on these elements tends to be incremental and hierarchical, leading to the gradual emergence of joint activities.

Incremental Commitments

Joint elements tend to get fully specified piecemeal. When Ann and Burton arrived at the lab room, they were asked to participate in a psychology experiment. When they agreed, all they were committed to was "doing some activity together in this room for the next hour." It was only after further instructions that this commitment got narrowed to "assembling a TV stand together." It got narrowed further to "doing the top-piece together" with this adjacency pair:

(3) Burton   ((Now let's do this one)) [picking up the top-piece]
     Ann   Okay

With this exchange, Ann and Burton agreed on the content ("doing the top-piece"), timing ("now"), and roles (Burton in control, Ann helping) of their next joint action.

  For these elements to be part of a joint commitment, they must be taken as common ground (Clark 1996). There are many ways to establish them as common ground (see Clark and Marshall 1981; Clark and Schaefer 1989; Clark et al. 1983; Enfield this volume; Goodwin this volume; Hutchins this volume; Lewis 1969; Schelling 1960):

Explicit commitments. Burton and Ann used all the talk in 1 to commit explicitly to the roles, content, and timing of their next joint actions.
Joint salience. In 3, Burton picked up the top piece as he spoke, making it obvious that he would affix the top piece to the piece Ann was holding. That helped fix their roles and the content of their action. Indeed, they *presupposed* that as they carried out the next joint action.
Precedent. Ann and Burton often established who would do what on

the basis of what they had just done. Once, for example, when Ann had just inserted one peg, the two of them presupposed that she would insert the second one too.

Conventional practice. Several couples assembling the TV stand (although not Ann and Burton) presupposed that the man was in charge, and the woman was the assistant: Building furniture was a man's job. This presupposition was apparently based on their idea of conventional practice.

Hierarchies of Commitment

Most joint activities, as I noted earlier, can be viewed as hierarchies of joint positions and actions. These hierarchies, too, emerge bit by bit, and so, therefore, do the joint commitments that coordinate them. Consider the assembly of the TV stand first by Peter working alone and then by Ann and Burton working together.[5]

Peter assembles the TV stand more or less according to a standard *means–end analysis* (Newell and Simon 1972). He begins with the problem, "How to assemble the TV stand from its parts," which he then decomposes recursively into subproblems. He does the decomposition one piece at a time. What emerges is a hierarchy of self-commitments that can be represented as a standard outline:

1. Build TV stand
  1.1. Arrange parts
    1.1.1. Put sides in pile
    1.1.2. Put screws, pegs in pile
    1.1.3. Put wheels in pile
  1.2. Assemble parts
    1.2.1. Attach top-piece to side 1
      1.2.1.1. Insert pegs
      1.2.1.2. Affix top piece to pegs
    1.2.2. Attach side 2 to top-piece
  Etc.

Peter first decomposes the entire task into "arranging the parts" and "assembling the parts." He then decomposes "arranging the parts" into "gathering the sides into a pile" plus "gathering the screws and pegs" plus "gathering the wheels." And so on. Each line represents a self-commitment to a state or action.

Ann and Burton, too, do a means–end analysis (more or less), but with a crucial difference: They do it together. They establish a hierarchy of *joint* commitments, not *self*-commitments, which looks something like this:

1. Build TV stand
    1.1. Attach cross-piece to side-piece
        1.1.1. Stick pegs into side-piece
            1.1.1.1. Find pegs
            1.1.1.2. Insert pegs into side-piece
        1.1.2. Affix cross-piece to side piece
    1.2. Attach top-piece to side-piece
    Etc.

Ann and Burton establish most of these joint commitments by negotiation. They agree on 1.1, for example, by means of the adjacency pair in (2), repeated here:

(2)   Ann   Should we put this in, this, this little like kinda cross bar, like the T? like the I bar?
      Burton   Yeah ((we can do that))

They next agree on 1.1.1, but that takes eight more turns. And so it goes. Although Ann and Burton assembled the same TV stand as Peter did, they had to coordinate *two* people's ideas, *two* people's commitments, *two* people's actions.

Stacking and Persistence of Joint Commitments

Joint commitments are complicated, therefore, not just because they require two decision makers—recall the sedan-chair principle—but because their emergence is incremental and hierarchical. This leads to two properties that I will call *stacking* and *persistence*.

People's commitments to each other accumulate, or *stack up*, the further they get into any joint activity.[6] At the beginning of their task, Ann and Burton's joint commitment was simple: "1 build TV stand." Their next move added a commitment to the stack: "1.1 attach cross-piece to side-piece." Then they added: "1.1.1 stick pegs into side-piece." If we think of each commitment as written on a sheet of paper, then Ann and Burton stacked up more and more sheets the further they got

into the hierarchy. And to complete their task, they had to discharge all of these commitments from the top of the stack down.

Joint commitments get added to a stack in two ways. One is vertical. For Ann and Burton to establish "1.1.1 stick pegs into side-piece," they had first to establish, or presuppose, all the joint commitments in the stack below it—"1 build TV stand" and "1.1 attach cross-piece to side-piece." The other way is horizontal. Once Ann and Burton were committed to, "1.1.1.1 find the pegs," they were also committed to the *next* steps at that level, here "1.1.1.2 insert the pegs into the side-piece." That is, joint commitments 1.1.1.1 and 1.1.1.2 were placed next to each other on top of the stack so far. Ann and Burton had to discharge both of them before considering the joint commitment below it (1.1.1) to be complete.

Joint commitments at the bottom of the stack *persist* even when those on top of them are renegotiated or reneged on. For the TV stand, Ann and Burton committed themselves first to "1 build a TV stand" and then to "1.1 attach cross-piece to side-piece." Next they negotiated on 1.1.1 and 1.1.2. When they negotiated the next level, they retained all of the commitments in the stack below it—both vertically and horizontally. If Ann had objected to sticking in the pegs and had got Burton to do it, that would not have changed their commitment to "1.1 attach cross-piece to side-piece" and "1 build a TV stand."

Entanglements of Joint Activities

Why are stacking and persistence so important? Because they help explain how hard it is to extricate oneself from joint activities. Look at Ann just as she begins negotiating with Burton on "1.1.1 stick pegs into side-piece." She is committed to being part of 1, 1.1, and 1.2, and she is about to add 1.1.1 and 1.1.2. If she reneges unilaterally on 1.1.1, she is letting Burton down not only on 1.1.1 but on 1.1.2—a double injury. Even if she reneges on 1.1.1, she is still committed to 1, 1.1, and 1.2. Plainly, once you get into a joint activity, it is hard to take unilateral actions.

Just how hard it is, is illustrated by the closing of telephone conversations (Schegloff and Sacks 1973). Once two parties think they have finished a conversation, they do not just hang up. They first reach agreement that they have completed the last topic and then open up a closing section. Consider the end of a telephone call in 8 (from Svartvik and Quirk 1980, S.7.2p.1397):

(8)	1	Topic talk	Ned	I don't know whose car it'll [be,
	2		Molly	[uhuh,
	3		Ned	I don't think it'll be Chris's,
	4		Molly	m,
	5		Ned	but . uh I'll be there directing traffic,
	6	*Pre-closing statement*	Molly	okay? . right,
	7	*Response*	Ned	okay,
	8	Future plans	Molly	great, — yeah, . well, — see you then, .
	9		Ned	all right, [see you Friday,
	10		Molly	[that's wonderful,
	11	Leave taking	Ned	right, bye now,
	12		Molly	bye,
	13	Terminating contact	Both	[*hang up*]

Once Ned and Molly finish the topic about cars in line 6, Molly offers to start closing the conversation with "right," and Ned agrees, "okay." With that exchange, they begin the actual closing, in which they make future plans, take leave (with "bye now" and "bye"), then hang up. Closing a conversation is a joint decision, but once it is made, the two parties still have work to do. Even routine telephone calls, like calls to directory enquiries, have closing sections, although the closings are briefer, reflecting the less intimate activity just completed (Clark and French 1981; Clark and Schaefer 1987).

Risks in Joint Commitments

To enter a joint commitment is to give up a bit of one's autonomy. When I join Helen in juggling six pins between us, I lose some of my options. I must work closely with her or risk hurting one or both of us. And when I drive out into the street, I must coordinate with all the other drivers or risk collision or injury. It is not just that I give up a bit of my autonomy. I cede to Helen, and to the other drivers, partial control over what I do. Normally we have good reasons for ceding control like that. I enjoy juggling with Helen, and I want to get to my destination safely (see Boyd and Richerson this volume, for the costs and benefits of cooperation).

Sharing control in joint activities, however, carries risks. One risk is *exploitation*. Partners are tempted to exploit the partial control they

have over us. Helen might draw me into a juggling routine that I do not know or do not want to do. Another driver might cut in front of me and force me to brake suddenly. Just as I could injure Helen or other drivers if I do not cooperate, they can injure me by exploiting my cooperation. Being drawn into unwanted, unforeseen, or regrettable actions has its moral and emotional consequences. I may get angry at the reckless driver and feel he was wrong. I may feel embarrassed at not knowing the juggling routine and hold Helen responsible.

Another risk is *overcommitment*. Joint commitments, once negotiated, are generally difficult to *re*negotiate, forcing the parties to honor their original commitments. Some joint commitments are impossible to renegotiate. Once I am on the road with a reckless driver, or in a difficult routine with Helen, it is too late to change or back out. I must make the best of the situation, however much I may resent it.

To illustrate these risks, I turn to one of the most famous studies in social psychology in the last half century—Stanley Milgram's experiments on "obedience to authority." These experiments have caused a great stir because they are taken as evidence that people will obey authority blindly even when that causes harm to others. The experiments may indeed show that. But for us, they are excellent examples of the risks of joint commitments.

The Milgram Experiments

In 1962, Milgram advertised in the newspaper for paid male volunteers to come either to "the elegant Yale Interactional Laboratory," or to a modest, unaffiliated "Research Associates of Bridgeport [Connecticut]," for a "study of memory." The subjects ranged from factory workers to professors. When a subject arrived at the laboratory, he and another subject drew straws to see who would be the "teacher" and who the "learner" in a learning experiment. The second subject was a confederate of Milgram's, and always became the learner. He was played by a 47-year-old accountant. The "experimenter" was played by a 31-year-old biology teacher dressed in a gray technician's coat; "his manner was impassive and his appearance somewhat stern."

The learner's job was to memorize a list of word pairs, and the teacher's job was to punish him for each wrong response. The learner was strapped into a chair with electrodes, and the teacher sat in front of an impressive "shock generator." The generator had 30 switches labeled 15 to 450 volts and further labeled (in sets of four): Slight Shock; Moderate Shock, Strong Shock, Very Strong Shock, Intense Shock, Extreme Intensity

Shock, Danger: Severe Shock, and XX. The teacher was instructed to "move one level higher on the shock generator each time the learner gives a wrong answer" and to announce the voltage level before each shock.

Milgram's interest was in how far subjects would go before opting out of the experiment. In one experiment, the learner was in a second room, but could be heard making an escalating series of protests as the voltage was increased. The experimenter, sitting at a table behind the subject, responded to the subject's objections with prods, "using as many as necessary to bring the subject into line."

Prod 1: Please continue, *or*, please go on.
Prod 2: The experiment requires that you continue.
Prod 3: It is absolutely essential that you continue
Prod 4: You have no other choice, you *must* go on.

When Milgram described this experiment to various groups of psychiatrists, college students, and middle-class adults, each group predicted that no one would reach the maximum of 450 volts ("XX"), and that the average subject would max out at 135 volts ("Strong Shock"). Their predictions were quite wrong. In fact, 62 percent of the subjects went all the way to 450 volts, and the average subject maxed out at 368 volts ("Extreme Intensity Shock"). We are poor in imagining how we would behave in this experiment.

Milgram carried out 18 experiments, each with 40 participants. The experiments varied on such features as where the subject, learner, and experimenter sat, how they communicated, whether the laboratory was at Yale or in Bridgeport, and who gave the orders. Only one experiment had women as subjects, and they complied as often as the men. In most of the experiments, a majority of subjects continued the shocks to the maximum. Milgram described them as "obedient to authority":

With numbing regularity good people were seen to knuckle under to the demands of authority and perform actions that were callous and severe. Men who are in everyday life responsible and decent were seduced by the trappings of authority, by the control of their perceptions, and by the uncritical acceptance of the experimenter's definition of the situation into performing harsh acts. [Milgram 1974:1.23]

But did these people really "knuckle under to the demands of authority"? Did they show "uncritical acceptance of the experimenter's definition of the situation"?

The Milgram Experiment as two Joint Activities

The psychology experiment, Martin Orne (1962) argued, is "a very special form of social interaction" (p. 782). The demands it places on experimenters and subjects are like the demands placed on the participants in any social interaction—in any joint activity. The Milgram experiment is an example par excellence of Orne's argument.

The Milgram experiment, as viewed by the subject, is really two joint activities, one embedded within another:

Memory task. This joint activity has two participants, whose roles are teacher and learner. The goal is for the teacher to teach the learner a list of word pairs. The basic activity is a series of cycles of joint action: the teacher gets the learner to learn the word pairs one by one. On each cycle, the two of them coordinate through scripted projective pairs: the teacher presents a test word, and the learner responds with the paired word; the teacher does or doesn't shock the learner, and the learner does or doesn't groan or yell.

Psychology experiment. The larger joint activity has three participants, whose roles are experimenter and subjects. Their goal is to carry out a psychology experiment. They also coordinate through projective pairs. These include the experimenter's instructions as well as long exchanges between subject and experimenter.

The subjects believed that the joint activity of interest was the memory task. But the *real* activity of interest was the psychology experiment: At what point would they opt out of it?

In Milgram's book *Obedience to Authority*, the first description of his experiments was this: "A person comes into the psychological laboratory and is told to carry out a series of acts that come increasingly into conflict with conscience" (p. 3). But this is a quite misleading characterization of what went on—and a good example of what Orne was speaking of. The extensive dialogues between experimenter and subject, quoted by Milgram, reveal something very different. The subject was not simply "told to carry out a series of acts." He *negotiated* with the experimenter on almost every act and position he took. These negotiations were often prolonged and intense, shaping what the subject did.

Mitigation

The experimenter relied on a range of negotiating tactics. One was *mitigation*. Although the entire experiment hinged on the harm the subjects thought their shocks were causing, what constituted harm was negotiated by the experimenter and subject:

If the subject asked if the learner was liable to suffer permanent physical injury, the experimenter said: "Although the shocks may be painful, there is no permanent tissue damage, so please go on." (Followed by Prods 2, 3, and 4, if necessary.) If the subject said that the learner did not want to go on, the experimenter replied: "Whether the learner likes it or not, you must go on until he has learned all the word pairs correctly. So please go on." (Followed by Prods 2, 3, and 4, if necessary.) [Milgram 1974:21–22]

Many subjects were skeptical of the experimenter's reassurances, which were belied by the labels on the shock generator, and that led to extended—sometimes heated—negotiations. These negotiations go to the heart of the study. If the shocks were not genuinely harmful, the subject had less reason to abort the experiment.

Another negotiating tactic by the experimenter was to accept responsibility for any harm done. An exchange with a subject called Prozi went as follows:

Subject: I mean who's going to take the responsibility if anything happens to that gentleman?
Experimenter: I'm responsible for anything that happens to him. Continue please.
[36 turns intervening]
Experimenter: Continue. Go on.
Subject: You accept all responsibility?
Experimenter: The responsibility is mine. Correct. Please go on. [p. 74–76][7]

As Milgram said about Prozi, "Once the experimenter has reassured the subject that he is not responsible for his actions, there is a perceptible reduction in strain" (p. 160). Another subject, called Rensaleer, was interviewed after the experiment, "When asked who was responsible for shocking the learner against his will, he said, 'I would put it on myself entirely'" (p. 51).

Negotiations of responsibility (as with Prozi) also go to the heart of the study. If the experimenter was fully responsible for harm done, then subjects had less reason to call off the experiment. But how many subjects negotiated responsibility? Milgram did not say. Still, Rensaleer called off the experiment midway, whereas Prozi continued his shocks to the maximum.

Risks of Exploitation

Other negotiating tactics by the experimenter were patently exploitative. One was to use disregard in uptake, as described earlier. Here is an illustration:

Subject: I can't stand it. I'm not going to kill that man in there. You hear him hollering?

Experimenter: As I told you before, the shocks may be painful, but—

Subject: But he's hollering. He can't stand it. What's going to happen to him?

Experimenter: (*his voice is patient, matter-of-fact*): The experiment requires that you continue, Teacher. [p. 73]

In the first exchange, the subject suggests that he may "kill that man in there" and asks "You hear him hollering?" Although the suggestion and question are serious, the experimenter disregards both by simply repeating a point he had made before. In the second exchange, he disregards all three of the subject's proposals.

To disregard a proposal is to imply that it is not worthy of consideration. It may be unimportant. It may be irrelevant. It may be misconceived. It may be too obvious to deal with. So when the experimenter disregards "You hear him hollering?" and "What's going to happen to him?" the subject can take him as implying that the questions are misconceived or irrelevant. These interpretations are reinforced when the experimenter speaks "with detached calm." Exchanges like this were common in the transcripts quoted in Milgram's book and sometimes lasted for 20 to 30 turns.

As for the learner, the experimenter disregarded everything he said, even when he screamed, "Let me out of here, you have no right to keep me here. Let me out of here, let me out, my heart's bothering me, let me out!" What could the subject conclude except that the learner's demands were of no importance or relevance?

Subjects' decisions were, indeed, influenced by how they negotiated. In an analysis of the complete unpublished transcripts of one of Milgram's experiments, Modigliani and Rochat (1995) found that subjects who raised questions or made objections early in the session were significantly more likely to abort the experiment early. "Evidently," Modigliani and Rochat argued, "certain forms of verbal resistance can alter the dynamics of interaction sufficiently to change its future course

and facilitate escape" (p. 1.19). In another experiment, the experimenter and subject communicated only by telephone, so their negotiations were presumably fewer, briefer, and less intense. In the standard experiment (in Bridgeport), 65 percent of the subjects continued the shocks to the maximum. In the telephone experiment, only 20 percent did.

Risks of Overcommitment

The participants in the Milgram experiments created tall stacks of joint commitments. When a volunteer arrived at the lab, he agreed first to be in the psychology experiment, then to be in the memory study, and then to enter each part of the memory study. Consider a hypothetical subject named Sam who is just about to shock the learner for failing on word pair 14. From Sam's perspective, the hierarchy of joint commitments at that moment looks something like this (with the critical line in italics):

    1. Enter experiment with others at Yale laboratory
    1.1 Arrange roles of teacher, learner, experimenter for memory task
    1.2 Establish procedure for memory task
    1.3 Enter memory task proper
            1.3.1 Instruction on word pair 1
            1.3.2 Instruction on word pair 2
            1.3.3 Instruction on word pair 3
            …
            1.3.14 Instruction on word pair 14
                    1.3.14.1 Exchange word pair 14
                    1.3.14.2 Exchange feedback on word pair 14
                            1.3.14.2.1 Teacher gives learner feedback, e.g. a major shock
                            [1.3.14.2.2 Learner responds to feedback
            [1.3.15 Instruction on word pair 15
            …
    [1.4 Exit memory task
    …
    [2. Exit experiment with others at Yale laboratory

By line 1.3.14.2.1, Sam has entered into and acted on 99 joint commitments (those up to 1.3.14.2.1), and he has entered into others to be acted on later (1.3.15 on), as marked by left square brackets. And the

current stack is five joint commitments high. So, by this moment, Sam and the experimenter have a long record of joint actions achieved and a tall stack of joint commitments yet to be achieved. This leaves Sam at this juncture with four main options:

First, Sam might *refuse* to deliver the shock. But to do that, he would have to renege not just on joint commitment 1.3.14.2.1, but on all of the joint commitments in the stack below it. And to do that unilaterally would destroy everything he and the experimenter had accomplished together. Taking this option, then, has costs, and few subjects took it.

Second, Sam might negotiate with the experimenter on a *joint exit* from the memory task and the experiment. Many subjects tried to do this, but the experimenter refused.

Third, Sam might try to *reframe* the memory task in negotiation with the experimenter: The shocks are not really so harmful, or the shocks are really the experimenter's responsibility. Many subjects tried this option and succeeded.

Fourth, Sam might simply deliver the shock. This way he would continue the long record of achievements in their joint actions, and he would maintain the stack of joint commitments yet to be acted on. As Milgram showed, a majority of subjects took this option to the maximum shock level.

Within a hierarchy of joint commitments, therefore, subjects have sound reasons for continuing—for taking the fourth option. As joint commitments stack up, they become harder and harder to opt out of—even with an uncooperative partner.

Morality and Emotion

Subjects in these experiments often reacted with great emotion. As Milgram describes it:

> Many subjects showed signs of nervousness in the experimental situa-
> tion, and especially upon administering the more powerful shocks....
> Subjects were observed to sweat, tremble, stutter, bite their lips, groan,
> and dig their fingernails into their flesh. These were characteristic rather
> than exceptional responses to the experiment. [1963:375]

But why these reactions and not others? These reactions reflect the subjects' anxiety about hurting the learner and *not* anger at, or dis-appointment with, the experimenter. The subjects apparently took the joint commitments as faits accomplis and focused instead on the harm

they were inflicting on the learner. These reactions might be expected from the stacking and persistence of joint commitments.

Not all subjects reacted this way. Many confronted the experimenter with moral issues, which then became points of negotiation. Different subjects were quoted as saying, "I'm not going to kill that man in there" and "You accept all responsibility" and "Surely you've considered the ethics of this thing. (extremely agitated)" (p. 48). Some subjects were placated in these negotiations, but others were not. When Rensaleer was urged to go on (at 255 volts) by being told "You have no other choice," he responded (p. 51):

> I *do* have a choice. (*Incredulous and indignant*:) Why don't I have a choice? I came here on my own free will. I thought I could help in a research project. But if I have to hurt somebody to do that, or if I was in his place, too, I wouldn't stay there, I can't continue. I'm very sorry. I think I've gone too far already, probably.

Not only did Rensaleer display anger at the experimenter for trying to draw him into this joint commitment, but he offered moral reasons for opting out. And yet Rensaleer apologized for wrecking their session ("I'm sorry") and negotiated a joint exit to the experiment. Despite everything, he took his joint commitments with the experimenter seriously and found a satisfactory way to discharge them.

Conclusions

Sociality is not a mere abstraction. It is a feature of life that gets played out in concrete social actions. These actions depend not only on linguistic acts, as characterized by Schegloff (this volume), but on extralinguistic acts. These range all the way from the pointing gestures in Enfield's (this volume) Laotian women, Goodwin's (this volume) stroke victim, Hutchins's (this volume) ship navigators, Liszkowski's (this volume) infants, and Levinson's (this volume) Rossel Islanders to the head gestures of Gergely's and Csibra's parents and infants, and the manual transfer of screws between Ann and Burton. Social actions also take place in material locations, whether that is a ship's navigation room (Hutchins), a living room (Enfield, Goodwin), a lab room (Clark, Gergely and Csibra), or an outdoor meeting area (Levinson).

Whatever the means and settings, people cannot take social actions—they cannot carry out joint activities—without making commitments to each other. As I argued, entering into joint commitments has

both benefits and risks. The benefits are obvious—the usual reasons for engaging in a joint activity. Working together, Ann and Burton were able to assemble the TV stand quickly and efficiently. But the risks of joint commitments are just as real. Subjects in the Milgram experiment, negotiating with the experimenter, were drawn into actions they did not anticipate, want to do, or approve of. Such is the power of joint commitments—the guiding force inside Levinson's interaction engine.

Acknowledgement

Notes

1. I thank Julie Heiser and Barbara Tversky for use of their video recording of this session.

2. This puzzle has been examined, without resolution, by philosophers (e.g., Bratman 1992; Grice 1989, Harman 1977; Searle 1990; Tuomela 1995), computer scientists (e.g., Cohen and Levesque 1991; Grosz and Sidner 1990), and psychologists (Clark 1996; Clark and Carlson 1982; Tomasello et al. in press). Still, the schema I describe later is close to a consensus solution to the puzzle.

3. The term *commitment* is used in game theory (e.g., Schelling 1960) in a sense closest to what I am calling *public self-commitment*.

4. Here I put aside institutionally based illocutionary acts that Searle calls declarations.

5. For Peter, I used a video recording of a lone individual assembling the same TV stand that Ann and Burton assembled. I thank Sandra Lozano and Barbara Tversky for the recording.

6. See *pushdown stacks* in computer programming.

7. There was no mention in the guidelines that the experimenter could negotiate responsibility for harm. Nor is this usually mentioned in discussions of Milgram's findings. Apparently, the experimenter improvised in other unspecified ways, too.

References

Austin, J. L. 1962. *How to do things with words*. Oxford: Oxford University Press.

Bach, K., and R. M. Harnish. 1979. *Linguistic communication and speech acts*. Cambridge, MA: MIT Press.

Bangerter, A., and H. H. Clark. 2003. Navigating joint projects with dialogue. *Cognitive Science* 27(2):32.

Bratman, M. E. 1992. Shared cooperative activity. *Philosophical Review* 101:327–341.

Clark, H. H. 1996. *Using language*. Cambridge: Cambridge University Press.

——. 2004. Pragmatics of language performance. In *Handbook of pragmatics*, edited by L. R. Horn and G. L. Ward, 365–382. Oxford: Blackwell.

Clark, H. H., and T. B. Carlson. 1982. Speech acts and hearers' beliefs. In *Mutual knowledge*, edited by N. V. Smith, 1–36. New York: Academic Press.

Clark, H. H., and J. W. French. 1981. Telephone goodbyes. *Language in Society* 10:1–19.

Clark, H. H., and C. Marshall. 1981. Definite reference and mutual knowledge. In *Elements of discourse understanding*, edited by A K. Joshi, B. L. Webber, and I. A. Sag, 10–63. New York: Cambridge University Press.

Clark, H. H., and E. F. Schaefer. 1987. Collaborating on contributions to conversations. *Language and Cognitive Processes* 2(1):19–41.

——. 1989. Contributing to discourse. *Cognitive Science* 13:259–294.

Clark, H. H., R. Schreuder, and S. Buttrick. 1983. Common ground and the understanding of demonstrative reference. *Journal of Verbal Learning and Verbal Behavior* 22, 245–258.

Cohen, P. R., and H. J. Levesque. 1991. Teamwork. *Nous* 25:487–512.

Grice, H. P. 1989. *Studies in the way of words*. Cambridge, MA: Harvard University Press.

Grosz, B. J., and C. L. Sidner. 1990. Plans for discourse. In *Intentions in communication*, edited by P. R. Cohen, J. Morgan, and M. E. Pollack, 419–444. Cambridge, MA: MIT Press.

Harman, G. 1977. Review of "Linguistic behavior" by Jonathan Bennett. *Language* 53:417–424.

Lewis, D. K. 1969. *Convention: A philosophical study*. Cambridge, MA: Harvard University Press.

Milgram, S. 1963. Behavioral study of obedience. *Journal of Abnormal Psychology* 67:371–378.

——. 1974. *Obedience to authority: An experimental view*. New York: Harper and Row.

Modigliani, A., and F. Rochat. 1995. The role of interaction sequences and the timing of resistance in shaping obedience and defiance to authority. *Journal of Social Issues* 51(3):18.

Newell, A., and H. A. Simon. 1972. *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hall.

Orne, M. T. 1962. On the social psychology of the psychological experiment: With particular reference to demand characteristics and their implications. *American Psychologist* 17(11):776–783.

Schegloff, E. A., and H. Sacks. 1973. Opening up closings. *Semiotica* 8:289–327.

Schelling, T. C. 1960. *The strategy of conflict*. Cambridge, MA: Harvard University Press.

Searle, J. R. 1969. *Speech acts*. Cambridge: Cambridge University Press.

——. 1975. A taxonomy of illocutionary acts. In *Minnesota studies in the philosophy of language*, edited by K. Gunderson, 334–369. Minneapolis: University of Minnesota Press.

——. 1990. Collective intentions and actions. In *Intentions in communication*, edited by P. R. Cohen, J. Morgan, and M. E. Pollack, 401–415. Cambridge, MA: MIT Press.

Svartvik, J., and R. Quirk (eds.). 1980. *A corpus of English conversation*. Lund, Sweden: Gleerup.

Tomasello, M., M. Carpenter, J. Call, T. Behne, and H. Moll. in press. Understanding and sharing intentions: The origins of cultural cognition. *Behavioral and Brain Sciences*.

Tuomela, R. 1995. *The importance of us: A philosophical study of basic social notions*. Stanford: Stanford University Press.