



Article

Judith Martens^a and Luke Roelofs^{a,*}

Implicit Coordination: Acting Quasi-Jointly on Implicit Shared Intentions

Winner of the 2018 Essay Competition of the International Social Ontology Society

<https://doi.org/10.1515/jso-2018-0018>

Abstract: We identify a social phenomenon in which large numbers of people seem to work towards a shared goal without explicitly trying to do so. We argue that this phenomenon – implicit coordination – is best understood as a form of joint agency differing from the forms most commonly discussed in the literature in the same way that individual actions driven by “explicit” intentions (those available for reflection and report) differ from individual actions driven by “implicit” intentions (those not thus available). More precisely, implicit coordination is both analogous to wholly implicit individual intentions, and constituted by the partly implicit intentions of participants. We discuss the significance of this category for action theory, social ontology, and social criticism.

Keywords: Joint action; Implicit intention; Social ontology; Collective agency; Collective intentionality.

1 Introducing Implicit Coordination

People in society do things together, but there are important differences among forms of “together”-ness. Sometimes people work together based on an explicit shared intention; sometimes they operate as members of an organisation with explicit rules for determining collective policies. But both

^aThe authors are funded by the Volkswagen Foundation project on Situated Cognition.

*Corresponding author: **Luke Roelofs**, Ruhr-Universität Bochum Institut für Philosophie, Institut für Philosophie II, Bochum, Germany, e-mail: luke.roelofs@utoronto.ca

Judith Martens: Ruhr-Universität Bochum Institut für Philosophie, Institut für Philosophie II, Bochum, Germany

of these are demanding to establish and maintain, and much of social life is not really covered by either model. By contrast, sometimes people “do things together” simply in the sense that the individual things they do add up to a total effect, without anybody intending this. In the former cases (“joint action” and “group action”) it makes sense to speak of collective *agency*; in the latter case (“aggregate action”) the individual actions sum together in a non-agential way.

But there is also an intermediate layer, where people’s actions seem to coordinate and interlock without an explicit plan, but with something like an implicit goal. This sort of case (which we call “implicit coordination”) is not well-covered by the existing literature: it is not quite joint action (as discussed by, e.g. Gilbert 1990, 2009; Bratman 1993, 2014), and not quite group action (as discussed by, e.g. List and Pettit 2011); but it is more agential than merely aggregate action (Chant 2007). Moreover, although there has been discussion of “emergent coordination” and others sorts of less-explicit coordination (e.g. Lakin et al. 2003; Richardson et al. 2007; Knoblich et al. 2011; Pacherie 2013) this has focused on highly automatic mechanisms of coordination that usually operate in small-scale, face-to-face contexts, not on large-scale social phenomena. In this paper we take the first steps towards rectifying this deficiency.

For an example of implicit coordination, consider a sudden decline in purchases of a particular brand over a period of a few weeks. This might be completely uncoordinated: perhaps the weather changed and now that brand’s products are less useful. On the other hand, this decline might represent a deliberate boycott campaign, whose participants talk explicitly about reducing the sales of that brand and coordinate their efforts to do so. This would plausibly be a joint action, at least on the part of the organizers of the campaign. But it might also be something in between. Perhaps some new piece of information about the company’s disreputable business practices came out, and even though plenty of such information had come out before, this time was different. Only a few people actually changed their purchasing behaviour *just* because of the new information, but they, combined with all the talk about it, gave many other people the impression that those around them were deliberately avoiding that brand. These people, out of some mixture of wanting to be good and wanting to look good, then started scaling back their purchases where possible, further contributing to the growing sense that buying that brand is mildly disreputable. This chain reaction spread a lot like a successful boycott campaign might, but without anyone needing to explicitly call for a boycott. This collective turning away from that brand is implicit coordination: it has a measure of the alignment of goals and interdependence of actions characteristic of a joint action, but not enough to fall squarely into that category.

Or consider the many ways in which people coordinate their production and consumption of symbolic artefacts – words, songs, stories, and so on. Suppose several authors have written stories in the same genre dealing with the same themes. This might reflect a joint action: perhaps they got together with that as their explicit aim, perhaps even under a shared name (“the So-and-So Writers’ Group”). Or it might be mere aggregate action: perhaps they have never heard of each other and just wrote similar stories entirely independently. But many cases lie in between these two extremes: often several authors write similarly because they have been inspired by reading each other’s work, or have been inspired by the same older sources, or have encouraged each other, guiding each other thereby towards the creation of a collective body of work with certain features, but never explicitly making that their aim. Moreover, their readers may play a substantial role in coordinating their efforts – encouraging certain directions more than others, comparing notes and recommending things to each other. Each writer wrote their own stories, but the *corpus* was created by implicit coordination.

Arguably, implicit coordination plays a major role in the tissue of social life in general – in customs, norms, (sub-)cultures, ways of life, social milieus. And the critical evaluation of society and social events often seems to imply something like implicit coordination. Suppose someone says: “Together, we’ve built a vibrant, thriving literature on this topic, motivated by our conviction of its importance.” Or, for a more negative example, suppose someone says: “Those with power in this industry have actively, collectively, perpetuated a fertile climate for sexual harassers, motivated by their desire to preserve their power and status.”

Both claims identify something done collectively (individuals can not build a literature or perpetuate a culture all by themselves), and allege that it was done actively and for a reason. But in both cases, it may be that the individuals involved can credibly claim not to have aimed explicitly at the outcome identified, or coordinated explicitly with the others involved. So this is not a joint action of the sort standardly analysed, but it is more than a mere aggregation of individual effects. Claims like examples A and B are best read, we think, as alleging implicit coordination. In this paper we provide some conceptual tools for thinking about this grey area in between the merely aggregate and the fully joint.

In Section 2 we discuss the way that individuals may act less-than-fully-reflectively, distinguishing both between implicit and explicit intentions, and between manifest and operative intentions. Section 3 translates these idea to the collective case, arguing that individuals whose intentions only implicitly make reference to each other often act in a “quasi-joint” manner. Section 4 contrasts our construct with other sorts of large-scale social phenomenon in terms of two dimensions – aim-sharing and interdependence. Section 5 briefly discusses some of the implications of our proposal for reactive emotions and ethical appraisal.

2 Implicit Intentions

We believe that a clear understanding of implicit coordination requires recognising the importance of *implicitly intentional* action. In this section we explain what we mean by this.

2.1 Examples of Implicit Intentions

Consider first actions which seem to happen “automatically”, but have identifiable intentions when “the same” action is done in the past or future (see Mele 1992; Roughley 2016; Fridland 2015). I walk into a room, and flick on the light without thinking about it; I step out my door, and turn left (my usual route to work) without thinking about it; I sit down at a seat where a pint of beer happens to be placed, and take a sip without thinking about it. When I first learned to do these things, they were explicitly intentional, and they still seem to be intentional actions in a certain weak sense: they have identifiable aims (e.g. illuminating the room), and target those aims in light of a representation of reality (of the room’s present illumination, the switch’s position, and the general power of light switches to operate lights). But they operate implicitly, without offering up those goals and assumptions for reflection or for integration with the agent’s other mental states. Their implicitness is manifested both in our often being unaware of doing them at the time, or not remembering doing them a second later, and in the risk of “absent-minded” mistakes, where we act out of habit despite having a belief or desire that ought to interfere (e.g. the room is dark because someone is sleeping; I need to go in the opposite direction today; it is not *my* beer).

Consider also skillful action (See Papineau 2013; Fridland 2014; Christensen et al. 2016). I want to serve the tennis ball to a particular spot, and all sorts of postural and muscular sub-actions are carried out to bring my racquet into the position it needs to be in; I want to drink some coffee and my arm, eyes, neck, lips, tongue, and throat all do their parts. These *subordinate* actions aim at goals in light of representations, and some but not all could be reported with a bit of thought, but they are not explicitly represented at the moment of action.¹ More careful scrutiny might let us articulate more: we might be able to identify the

¹ We do not think that skills are entirely mindless. Recently a number of philosophers and psychologists have reacted against this kind of view: Cognition makes an important contribution to skilled action (see Montero 2010; Sutton et al. 2011; Papineau 2013; Stanley and Krakauer 2013; Fridland 2014).

steps in our action “serve” by mentally or physically rehearsing it with attention focused on particular body parts, thereby making the implicit explicit.

A third example might be actions based on implementation intentions, where one states a very explicit intention including a lot of situational details earlier on to heighten the chance to act accordingly later on. Such if-then plans produce automatic action control by “intentionally delegating the control of one’s goal-directed thoughts, feelings, and behaviors to specific situational cues” (Gollwitzer et al. 2005, p. 485). Here, as with habits and skilled actions, what was once explicit can be implicit when it is repeated at a later time (Gollwitzer 1999; Wieber et al. 2015).

In all these cases, there is the potential to articulate an explicit intention. When asked “why did you do that?”, we can often answer, either readily or after some reflection, by identifying the goals and assumptions implicit in the action. Such an answer typically comes with a feeling of familiarity, a sense that we are articulating what was already there.

Habits, often, were once explicitly intentional but now are not; (the automatic elements of) skilful actions can be made explicitly intentional by being scrutinised. But there is no reason to think that these historical facts are constitutive of the action being intentional at all. If a habit-like automatic action was never explicitly intentional to begin with, and never received the scrutiny needed to become so, it can still be meaningfully described as intentional, as aiming at something in light of a representation of reality.² Maybe as a baby I learnt to turn my head a certain way when talking to people, or adjust my speech patterns in certain ways, and learnt this because it made communication and conversation go more smoothly. Maybe I never explicitly identified that I was doing this, either at the time or subsequently. And maybe I will die without making it explicit. Nevertheless it seems right that it involves a certain picture of the environment (that I am trying to communicate with someone) and a certain aim (to facilitate communication). It would make sense to say that in some cases I might automatically do this, and unintentionally bring about some other effect, even though its typical effect is not explicitly intended either. Since the distinction between doing things unintentionally and intentionally is a key feature of intentions, we take this to be further support for accepting implicitly-intentional actions. But to substantiate this claim we ought to clarify what we mean by “intentional”.

² Maybe intentions need to be *potentially* explicit: maybe any action-guiding state that can be called an intention of mine must be *the kind of thing* that sufficient attention and scrutiny could articulate explicitly. But that is compatible with recognizing that this potential often goes unrealized, and that given the difficulty of actualising it we have no way of knowing in advance all the different implicit intentions we may be acting on.

2.2 Implicit Intentions: A Definition

By “intentional action” we mean simply any sort of action that can be explained in terms of a goal it aims at, and a representation of the present situation, such that the action is “rationalised” by the fact that it would accomplish the goal if the representation were accurate. This “aiming at goals” and “representing the situation” might be propositional or not, conceptual or not, and the terms “desire” and “belief” might be used to cover all their forms, or might be reserved for a particular, propositionally structured, subset. What is key is that a distinction can be made between effects which were intended and those which were not, and between reality being the way the agent takes it to be in acting, and it not being that way.

We use “intention” for the mental state that guides a given instance of this kind of action, without attempting to provide a comprehensive theory of the nature of intentions. Looking to the intention behind an action allows us to divide its consequences into three categories: what is intended (the aim), what is unintended but still done intentionally (necessary first steps and collateral damage – what lies in the intention’s “motivational potential”), and what is unintentional (the product of failure, mistakes, or unforeseen circumstances) (Bratman, 1984).

This sense of “intentional” is grounded in the intuitive distinction between the things that merely happen to people – the events they undergo, even those which involve endogenous movements of their body – and what they do. Sometimes “what they do” is associated specifically with “reflective” intentions (or “planning intentions”), those which are metacognitively represented by the agent and endorsed as such, and fit into a coherent plan with that agent’s other intentions in the past, present, and future. But there are cases that are neither reflectively intentional, nor “mere bodily movements”. Frankfurt (1978) points out that animal behaviour is often purposeful without being reflective: when a spider walks across the table, the spider is in control of its legs, and aims to get from one location to another. Similarly, the idle, unnoticed movements of my fingers may have the goal of releasing the candy wrapper from my grasp.

Other philosophers have also made a three-fold distinction among mere bodily movement, mere purposeful activity, and reflective intentional action (O’Shaughnessy 1980, 2008; Frankfurt 1988; Velleman 2000; Bratman 2014). Our use of “intentional” aims to cover both the latter two, which differ in their relation to reflection but are both distinguished from mere bodily movement in that the agent is in some sense “in control”.

We divide “intentions”, understood in this deliberately broad way, into “implicit” and “explicit”. “Implicit” here means “not explicit”, and “explicit” means “available for reflection”, where “reflection” is the attentive, introspective, second-order processing characteristic of mature human agency. An explicit

intention is one which is accessible to me so that I can know that I have it and can ask why I have it, whether I endorse it, whether it coheres with my other intentions, etc. More precisely, an explicit intention is one for which my ability to know about does not depend on taking time or effort to discern it, or inferring it from other evidence. An implicit intention is one for which this is not the case – if it can be known about at all, this requires time and effort, or must be inferred from considering other evidence. It may still tend to cohere with other intentions, but its coherence cannot be actively considered by the subject. Such active consideration “makes it explicit” in the sense that subsequent reflection can then be easy and fast. But equally, a repeated action may start out explicit and become implicit (as often happens with habits) in the sense that both noticing what one is doing, and identifying its aims, become gradually harder and slower.

We are deliberately open in our understanding of the term “available”, in “available for reflection”.³ Different states might be unavailable for reflection in completely different ways, and consequently different implicit states may be implicit for different reasons. They may be completely inaccessible just by their nature (much of early perceptual processing, for instance), or they may be part of a process so rapid as to defy scrutiny (as in the many steps in a smooth, skilful, action, which could only be attended to by slowing the action down so much as to disrupt it) (Hohwy 2011). Alternatively, they may simply be hard to detect, either from their subtlety and faintness, or from their unpalatability: that is, our failure to notice everything going on in our minds may be sometimes due simply to the difficulty of noticing dim or convoluted processes, and sometimes due to a motivated refusal to face an unpalatable truth (McKay and Dennett 2009). These different states may have nothing positive in common: in calling them implicit we mean merely to identify something they lack.

2.3 Operative and Manifest Intentions

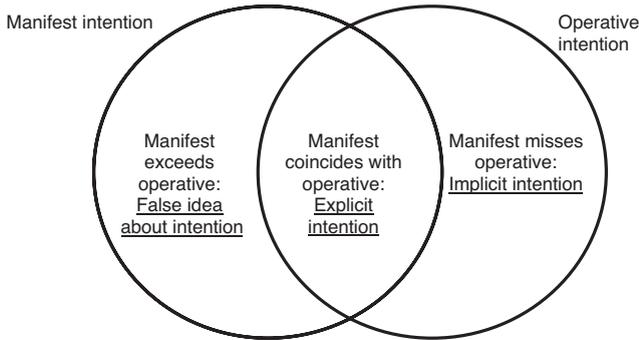
Not only can there be implicit intentions, there can be cases where implicit and explicit intentions work together to guide the same action, so as to function as a single complex intention with implicit and explicit parts. Suppose, for instance, that I am debating with someone. I’m getting annoyed with them. There’s a piece of information I could bring up that is only slightly related to the point they just made,

³ “Available for reflection” is not quite the same as “accessible” or “access-conscious” in the sense used by Block 1995 (cf. Tye 2000), which is defined in terms of a wide range of uses – to be access-conscious is to be available for control of action, for formation of memory, for future planning, and also for metacognitive reflection. An implicit state might be substantially accessible, if it flexibly informed our actions, memories, and plans despite not being available to reflection.

but which would make them very upset. What I keep explicitly before my mind is the intention to rebut their point by bringing up that piece of information, simply because of its relevance to that point. But if I was not aware “at the back of my mind” that they would be hurt by it, I would have probably picked a different piece of information, that undermines their point more directly. The awareness that these words would hurt them was there, and it made me choose them. Now suppose I bring it up and they visibly crumble; I might tell myself that was not my intention, because my “official”, explicit, intention certainly was not that. But in fact my desire to hurt them determined what I said. If bringing up this piece of information ends up hurting the feelings of another person present, it is a meaningful psychological question whether (1) I intended to hurt their feelings (did my awareness of their sensitivity encourage my impulse to bring it up?), (2) did not intend to but did so intentionally (I foresaw their hurt but accepted it as collateral damage), or (3) did so unintentionally (I did not know they were there, did not realise they would be hurt, etc.). Since we can make that distinction, it seems appropriate to think of there being an implicit part of the intention to raise that point.

It is useful here to borrow some terminology from Haslanger’s discussion of social construction (2006). Haslanger observes that often the way that a person applies a concept may diverge from how they describe their application of it. In her example, someone employs a concept of “misogynistic music”, but in practice applies it only to music of a certain genre, or only to sexually graphic music. Haslanger thus distinguishes between the “manifest” concept (what the subject thinks they are applying) and the “operative concept” (what actually best fits the pattern of discriminations they make). In a similar spirit, we suggest distinguishing “manifest intentions” (the intentions that agents take themselves to act on) from “operative intentions” (the intentions that best capture their actual behaviour). The “manifest” is about metacognition, while the “operative” is, as Haslanger puts it, what “best captures the distinction that I in practice draw” (2006, p. 96). But whereas concepts are characterised by the distinctions they make possible – by their dividing things into instances and non-instances – intentions must be characterised differently, in terms of what we will call their “aim” and their “motivations.” An intention is associated with both a distinction between the circumstances in which we would take it to have been fulfilled, and those in which we would take it not to have been (its conditions of fulfilment, or “aim”), and a distinction between the preceding circumstances in light of which we formed it, and the circumstances in which we would not have (its reasons for formation, or “motivations”). Then the operative intention is what in fact “best captures” the conditions under which we would feel satisfied, and the conditions under which we would have formed and acted on that intention, while the manifest intention is our own idea about those conditions.

Clearly the manifest-operative distinction has some connection to the implicit-explicit distinction, but the distinctions do not line up exactly, as illustrated in the below diagram:



An “operative” intention might be explicit or implicit, as long as it makes a difference to actual practice. Whereas a “manifest intention” is, it seems, necessarily *explicit*, since it involves metacognition, it might not for all that be an explicit *intention*. When the manifest intention departs from the operative intention – i.e. when we think we’re acting on things that we’re not – what is explicit is a *false idea about our intention*, rather than the intention we actually act on. An explicit intention corresponds to the overlap of manifest and operative: an actual action-driving intention that is also reflected in the subject’s metacognition (however exactly we want to think of the intention and the belief about it as connected). An implicit intention is operative but not manifest: it is what’s left of the action-driving psychology when we subtract what is available to reflection.⁴

3 Shared Implicit Intentions

3.1 Social Aspects of Intentions

Sometimes our intentions are complex, with an explicit part which mentions only individual sorts of factors, and an implicit part which is heavily social in content.

⁴ The reason this relationship is so complicated is that the distinction “manifest intention” vs. “operative intention” is between two different *senses* of the term “intention”, while the distinction “explicit intention” vs. “implicit intention” is between two varieties within a single sense of “intention”.

To put it in the terms of the last section, sometimes our manifest intentions leave out the social factors that enter into our operative intentions.

Recall an example from Section 1, where a new revelation about brand X's unethical business practices causes a sudden shift in consumer behaviour, even though similar unethical practices by that company were known about beforehand. Suppose one of the people no longer buying brand X is disposed to say, if asked about their choice, "I am trying to not buy things that are produced unethically, because I think it's wrong to support practices I consider unethical." This manifest intention mentions individual aims ("that I not buy unethically produced products") and individual motivations ("that I consider the practices involved unethical.") But the operative intention might be quite different. It might be that if this person had considered the practices involved in producing brand X products unethical, but had *not* thought that others in their community felt the same way, they would have continued buying those products. Indeed, it might be that they would not have even formed a definite opinion about those practices, had they not picked up on disapproval from those around them. In this case their motivations – the factors necessary to get them to act – include social factors that are left out of their explicit intention. Likewise, it might be that if they had stopped buying brand X, but been alone in doing so, feeling like an eccentric in their community, they would have felt a sense of disappointment, of surprise, of an aim unfulfilled. In this case their aim – the conditions they would count as fulfilling their intention – is less "that I not buy unethically produced products", and more "that I and my peers not buy unethically consumed products". That is, their actual feelings and behaviour do not fit "I want to bring my shopping habits into line with my ethical judgements", but rather "I want us to bring our shopping habits into line with our ethical judgements".

It seems likely that cases like this may be quite widespread. What we like, what we value, what we consider good or bad, desirable or undesirable, is often influenced by the attitudes of others around us (indeed, some theorists define the category of "social norms" by the constitutive role of expectations about the judgements of others, e.g. Bicchieri 2006, 2016). We are systematically driven by social motivation: "a set of psychological dispositions and biological mechanisms biasing the individual to preferentially orient to the social world [...] to seek and take pleasure in social interactions [...] and to work to foster and maintain social bonds" (Chevallier et al. 2012, p. 231).

This pervasive social orientation is manifest, for instance, in the tendency to conform with the opinion of those around one, as demonstrated in the famous Asch conformity experiments (Asch 1956; Bond and Smith 1996), in the tendency to unconsciously synchronise our movements to those of others (Bargh et al. 1996;

Chartrand and Bargh 1999), and to be more favourably disposed to those with whom we thus synchronise (Kirschner and Tomasello 2010; Reddish et al. 2013; for a review, see Knoblich et al. 2011).

In game theory Bacharach (2006), Bicchieri (2006), and Sugden (2001, 2007) have argued that pro-social behaviors influence the choices made in social dilemmas such as the prisoners' dilemma and the Hi-Lo game. Zawidzki (2013) argues that we shape each other's minds, allowing for better social understanding and smoother interactions without explicit communication.

Nevertheless, often our explicit ideas of what we feel and why we feel it leave this out, either as an active denial of our social embeddedness or simply by passively declining to consider it. After all, it would be impractical to mention all the factors that fed into a certain decision, so in a social environment that is taken for granted it makes sense to mention the individual factors without referring to the social organization and background that shapes them.

3.2 Shared Implicit Intentions

Suppose we have a large number of people doing individual actions which together add up to some sort of overall outcome. For this to be a joint action as classically conceived, the individuals should explicitly intend the outcome and perform their individual acts because they know that the others are performing theirs (on Bratman's analysis, e.g. 1993, 2014), explicitly intend to do their part while believing that the outcome is feasible and that everyone knows this (on Tuomela and Miller's 1988 analysis), take themselves to be jointly committed to bringing about that outcome (on Gilbert's, e.g. 1990, 2009), or similar. But human sociality is often less explicit than that. In implicit coordination, we suggest that the individuals are more like the implicitly-socially-influenced people described in the previous subsection: their operative intentions, but not their manifest intentions, feature the total outcome of the many individual acts as a component of what they aim at, and their motivations are sensitive to their awareness, implicit or explicit, of the other's actions. If we ask them about their aims and motivations, the story they provide might be entirely individualistic: they describe the things that made them act, and the goals they aimed at, without mentioning anything about other people. But the other participants are in fact relevant to their action – whether they perform the action, and whether they feel that their action has attained its aim, is determined in part by their implicit awareness of social facts.

Have these people acted jointly? Have they done something together, based on a shared intention? We think they have: the implicit social content of their

individual intentions, and the ways this content interacts across different individuals, allow the complex of those intentions to serve the essential function as a shared intention. Bratman argues that the essential functions of shared intentions are: (i) to coordinate people's actions towards a goal; (ii) to likewise coordinate their plans towards that goal; and (iii) to provide a background framework that structures bargaining (1993, p. 99; 2014, p. 132). Clearly, the phenomenon we are describing can play that first role: indeed, that is what draws our attention to it, that many people can act in ways tailored to bringing about a certain aim, even though none of them individually avows that aim. It is less clear how such implicit states coordinate planning; indeed, it is in general unclear how implicit states interact with planning, since forming and following long-term plans typically relies on precisely the sort of high-level reflective processes that implicit states are not accessible to (though see, e.g., Seligman et al. 2016). But for this very reason, we do not consider the coordination of plans essential to the role of intentions generally, but only to the role of reflective intentions. Thirdly, connected implicit intentions can and do provide a framework to structure bargaining: the way we approach practical disputes about social affairs frequently draws on unexamined, unconscious assumptions we had made about “what people do”, about what is normal and expected and supports “what we want” (cf. Dworkin 1986, p. 195–205).

The role of shared implicit intentions as a background for negotiation comes out particularly in those cases where, as with individual habits, the same activity goes on implicitly at one time and explicitly at another time. Picture a family that is preparing for dinner. At one point they might have decided on a particular time for everyone to come to the table, and a particular seat for each person. At first each has to deliberately remember the assigned time and seat, and regards their coming at the right time and taking the right seat as “doing their part” in the joint activity. But over time this becomes so habitual that each person tends to get hungry and wander towards the kitchen just before the assigned time, without explicitly intending to, and sit in their assigned seat automatically. Their awareness that they are coordinating has faded into the background, but the coordination they are engaged in has remained. Another example might be the way people in different countries interact in shops. When traveling from Germany to Australia, the first reaction when the cashier asks “how are you?” is to semi-ignore the question: “he cannot be truly interested in how I am doing” one might think. In Germany they will wish you a “nice day”, but usually with no eye contact, and an uninterested intonation. And who could blame them? Most customers react in the exact same way. Participating in such a conversational culture is a conscious choice for the new arrival at first, but can eventually become habitual. But conversely, when confronted with the newcomers' surprise, the natives may

find themselves coming to more explicitly articulate their practices, even while continuing to engage in them.

3.3 Joint Action and Quasi-Joint Action

Consider two possible objections to our proposal. First, a critic might worry that the lack of explicitness in the social aspects of the individual intentions will transfer to their shared intention, which will be entirely implicit. And they might deny that a wholly implicit intention makes sense: an intention must be at least partly explicit, or else, not being available to reflection, it will not be a real intention. We agree with the first part of this objection: insofar as there is a joint intention only because of the individual intentions' implicit content, it is reasonable for us to say that the joint intention is implicit, and that the participants act only jointly-implicitly-intentionally, despite each individual acting explicitly-intentionally. But we disagree that intentions cannot be wholly implicit. As argued in the previous section, an action can be intentional in our sense – guided both by a way the world is taken to be and by a way the world is to be made – even if none of what guides it is available for rational reflection.⁵

Second, a critic might worry that the implicit social components to the individual intentions are not central enough to the actions to bind them into an intentional joint action. Even if aspects of the states driving your action are left out of the explicit description one gives, however, it seems a reasonable assumption that the explicit intentions capture the core motivation and aim. We should be reluctant, after all, to attribute to people a complete and utter ignorance of their own intentions. Even if we suspect that someone who stops purchasing brand X just as it stops being popular, citing ethical concerns, may be leaving out the social aspects of their intention, we should stop short of suggesting that they have no actual ethical opinion about brand X, or no desire to act in line with their ethical opinions. Consequently it might seem that the implicit factors can only play a supplementary role in their behaviour.

We think this second objection has some force, which is why we do not claim that implicit coordination is jointly intentional *in the very same manner* as classical examples of joint action. Rather, they typically have *looser forms* of the same

⁵ It may be that, paralleling what we said in footnote 2, any shared implicit intention must be at least the right sort of thing to potentially become explicit – might be brought to reflective awareness by the right combination of individual introspection and cooperative discussion and sharing of views. Indeed, a shared implicit intention's possibility of becoming explicit might be essential to its ability to play this third role, of providing a background for negotiation.

structure found more fully in classical joint actions. Consider, for example, the seven conditions for a joint action laid out by Bratman (2009): that each (i) intend that the joint activity proceed (ii) by way of the other's similar intention, and (iii) intend to coordinate their sub-plans with the other to this end and (iv) help the other if needed, and (v) persists in intending all this *because* of the other's so intending, in (vi) mutual responsiveness to the other's activities, with (vii) common knowledge that all these conditions obtain.

We think it very likely that not all these conditions will be met by many cases of implicit coordination: in particular, even if individual's operative intentions satisfy (i), (iii), and (iv), there will usually not be enough communication and interaction among the participants for them to know much about each other's actions as individuals, violating (vi) and (vii), or for their actions to be dependent on the specific actions of specific other individuals, violating (ii) and (v).

But these conditions are presented by Bratman as "robust sufficient conditions" (2009, p. 159), rather than necessary conditions, and other authors have convincingly argued that they cannot all be necessary. Pacherie (2013), for instance, argues that small children are capable of joint action despite lacking the cognitive skills needed for thinking about the intentions of others as such (conditions ii and vii); Shapiro (2014, p. 273–274) argues that large-scale social agency typically cannot involve any strong form of interdependence (conditions v and vi). So we should not assume that any phenomenon which does not meet all seven conditions cannot be joint agency; it is simply a different way to meet the fundamental functional requirements of joint agency.

We might put this, perhaps, by speaking of "quasi-joint action", when individual actions are bound together in a looser version of the sort of bonds found in joint action. In a joint action, individuals all share the same goal; in a quasi-joint action their goals overlap more or less fully. In a joint action, individuals act because of their awareness of the others' actions; in a quasi-joint action they are made more likely to act by their awareness of the others' actions.

The category of quasi-joint action is not coextensive with that of shared implicit intention: one concerns, so to speak, how tightly individual actions are bound together, while the other concerns whether these bonds are available to reflection. But in practice it is common for looser bonds to be unavailable to reflection, and we think of implicit coordination as fitting both categories. Implicit coordination would then be defined as a quasi-joint action coordinated by a shared implicit intention. That is, it differs from classically-conceived joint action both in that the component actions show a looser form of coordination, and in that the mental states that accomplish this coordination are unavailable to reflection.

3.4 Operative and Manifest Intentions in Joint Action

A key implication of the implicitness of the shared intention is that it may be resistant to reflective change or reform. An individual whose intentions are wholly or partly implicit may act “against their better judgement”, because their intentions are not subjected to the process of rational reflection, whether that implicitness comes simply from speed and automaticity, or from a more robust failure of self-knowledge. We often perform habitual actions even in circumstances where they are inappropriate, because our knowledge of those special circumstances fails to connect with the habitual intention before it is executed. And often someone who cannot admit their own aims to themselves cannot pursue those disavowed aims efficiently because they cannot thoughtfully plan how to bring them about.

A similar sort of failure of reflection is a natural risk with implicit shared intentions, though in this case it may be a failure of communication (“collective reflection”) rather than individual reflection. If everyone involved self-ascribes a purely individual intention, they might expect their behaviour to change in response to information about, e.g. better options for accomplishing their supposed aims. But because their social awareness influences their action in ways that ignore their explicit reasoning, this may prove ineffective. Changing the collective behaviour requires not just changing individual minds, but changing what individuals believe about other individuals, changing what is “common knowledge” – failure to do so may leave “pluralistic ignorance”, where each individual has changed their mind but continues to assume others have not (see, e.g. Miller and McFarland 1991). Moreover, the same group which explicitly works together for one goal, may at the same time be implicitly working together in ways that directly undermine that goal; avoiding such failures requires not just changes in individual intentions, but communication that makes such changes commonly known.

For example, consider our running example where brand X becomes reputed as “unethical”, and each person’s explicit intention not to buy from unethical companies is bolstered by their implicit intention not to buy things that others in their peer group are avoiding and think are unethical. You might persuade one individual that actually, brand Y is significantly worse than brand X, but be unable to change their buying habits because they remain implicitly aware that they face social penalties for buying brand X but not for buying brand Y. Indeed, you might persuade many, even *all* the individuals, that brand Y is worse, and yet find them each continuing to focus on avoiding brand X because all the others are (cf. Bicchieri and Mercier 2012; Potter 2018; Sunstein Ms.).

4 What Is Implicit Coordination and What Is Not

4.1 Contrasting Cases

Human social life has diverse forms, and we have no wish to flatten this diversity into one sprawling category of “quasi-joint action”. So it is worth emphasizing some contrasts between implicit coordination and other social phenomena. First, implicit coordination is not structured by things like explicit rules or authority relations: Shapiro (2014) argues that many forms of shared agency, and especially those where large groups are involved, are unlikely to succeed in the absence of authority and explicit centralization of control. Our interest, by contrast, remains fixed on the kind of unstructured, spontaneous hanging-together of individual actions which forms the general social fabric out of which these structures can emerge.

Second, implicit coordination needs to be more than an aggregate action, and consequently goes beyond the kind of “collective” actions studied in the literature on collective harm (Kagan 2011; Nefsky 2011, 2018), where the only combination among the individuals involved happens post-action, in the accumulation of their small effects. The puzzle of how to assign responsibility for harms caused in this cumulative way is an interesting one, but it is not our target here.

The difference between implicit coordination and mere aggregate action, however, is better thought of in degree-terms than in category-terms. To a first approximation, they differ along two key dimensions: (1) the degree to which individuals’ aims coincide or overlap (aim-sharing), and (2) the degree to which each individual acts because of the others acting (interdependence). To illustrate these dimensions, let us consider an example of an aggregate action, and see what would need to be added to get implicit coordination.

Start with the following merely aggregate action (adapted from Chant’s “Two Irresponsible Shepherds”, 2007, p. 251). If I walk a certain route across a lawn, and very slightly erode the soil there, and then you do the same later that day, and then someone else, and then someone else, it may in the end come out that “we have worn away the grass and made a dirt path”. But this is just an aggregate action, not a collective action in any interesting psychological or sociological sense. For us to have done it intentionally – even implicitly intentionally – more must be added.

4.2 Degrees of Aim-Sharing

Let us first focus on how the individuals think of the effects they collectively bring about. When we make a path across a lawn, we typically just want to get to our destination, and do not think about the erosion of the ground. If we did, we might

feel mildly regretful about it, but usually would not care that much. If it were a full-fledged joint action, by contrast, we would be aiming at it: for instance, if a group of us actively conspired to wear away the grass, we would each be regarding that collective erosion as our aim, and the small contribution we each made as “doing our bit” towards that aim. Here the accumulation of our individual effects is intentionally aimed at by the overlap of our individual intentions: our guiding mental states each represent the same goal and this is common knowledge.

But as we argued in the last two sections, people’s intentions can often be complex, confused, indeterminate, and implicit: there may be cases where someone is partly motivated by their awareness of an effect, without admitting it to themselves, or without it being enough to motivate their action by itself. Perhaps I dislike the guy who manages this lawn, and this colours my background awareness that he would not want me walking on his grass with a faint positive emotion, a little glimmer of defiance at his imagined disapproval. Maybe this faint attraction to the prospect affects my choice of route right at the margin, just by making it seem more attractive to walk across the lawn when I would otherwise be indifferent. Then, while it would not be true to say that “I walked across his lawn in order to erode the soil”, it would not be accurate to deny it completely, either.

Enriching an aggregate action with this sort of partially-motivating attitude towards a goal, even if it is shared by all the participants (e.g. a hundred people walk the same route across the lawn, wearing away a path, animated in part by their shared dislike of the gardener), might move it a little closer to being a joint action, but the shared aim is only one relevant factor. Even the highest degree of aim-sharing would not render an aggregate action joint. Consider, for example, a hundred people across the world, each struggling in their own way to promote the same set of ideas in their societies. Their efforts might in fact add up to an aggregate effect (the increased global prominence of those ideas), and they might each explicitly think of themselves as aiming at this goal, and think of their individual actions as contributing to this goal, as “doing their part” in this great work. But if they acted quite independently, in ignorance of each other’s efforts, with no coordination and no way to encourage each other, they would not be “working together”, even if they worked towards the same goal.

One might argue there is a *qualitative* difference, rather than a *quantitative* difference, between cases where there is aim-sharing and awareness of aim-sharing and cases where there is not such an awareness.⁶ Although we think that awareness of aim-sharing can indeed influence the degree to which the aim

⁶ We would like to thank an anonymous reviewer for pointing at this potential issue.

is shared, we argue this can come in degrees (and hence is quantitative). Awareness (and the possibility to develop shared intentions and common knowledge of such intentions based on such awareness) will help to align aims. This, however, does not commit us to the idea that there is a qualitative difference between aim-sharing with and without awareness. Approval and reinforcement of a (shared) aim can come in many forms, of which (explicit) common knowledge is only one.

4.3 Degrees of Interdependence

So we must add a different sort of factor: the actual interdependence of each participants' actions on the others'. We imagined each person's probability of walking across the lawn as independent of whether anyone else does so; but when they actively conspire to make a path, they each act *because* they think the others have acted or will act. In between complete independence (in merely aggregate action) and complete dependence (in fully joint action), there is a spectrum of milder sorts of influence, where each one's action affects the odds of another choosing the act, but without making it a sure thing.

Of course, the actual formation of paths may often involve a sort of non-social interdependence: once there is a path partly formed, it will be a more salient or more attractive route, and I am more likely to walk along it. But this sort of interdependence is not quite what is relevant: this is acting because of the effects of another's action, not acting because of my awareness of another's action. A slight degree of a more social sort of interdependence can also be present in everyday cases of walking along a path: when I see the beginnings of a path, I see it as "where others have walked", and may well be more inclined to walk there just because of this feeling that it is "where people walk" (even if I am not at all aware that this motivates my choice). Indeed, there will probably be all sorts of examples: as argued in the last two sections, we are very prone to automatically pick up on what those around us do, or expect, and to align our actions. And we can imagine a higher degree of (social) interdependence being present: perhaps because I dislike the gardener, I feel pleased when I see signs that someone else has been violating his lawn, and feel inspired to reinforce this action by walking there myself.⁷

⁷ Awareness of interdependence can, over time, influence interdependence just as awareness of aim-sharing can influence aim-sharing. Again, this is compatible with thinking that interdependence with awareness is not qualitatively different from interdependence without.

Interdependence by itself is not enough for jointness, any more than aim-sharing by itself. Indeed, we are familiar with a particular type of social phenomenon which shows extremely high interdependence: the “trend”. Suddenly everyone is buying onesies: why? For each individual person, it is probably largely because they suddenly became aware that others were doing so, and in buying one for themselves they spread this influence on to others. In general, all sorts of “fashions” (not just sartorial fashions but academic fashions, lifestyle fashions, etc.) involve waves of individual actions produced by awareness of others performing similar actions. Should we therefore say that the participants in the trend or fashion are acting jointly to bring about the cumulative result of these individual actions, namely the sudden buying (wearing, using, etc.) of large quantities of the now-fashionable item? Usually no. Trends typically lack the aim-sharing characteristic of joint actions: when I follow the crowd in what I buy, wear, research, use, etc., I am not even implicitly attracted to the outcome I contribute to, namely that the trendy thing should become (more) trendy. My motives are typically entirely self-regarding. Indeed, I might actively dislike that outcome – I might hate the trend, even as I “bow to the crowd” and go along with it. Perhaps my ideal outcome might be for me to be among the only few to follow the trend (I want to be “ahead of the curve”), or else for the trend not to happen at all, so I do not feel the need to keep up (I fear being “behind the curve”).

In some cases, it might even be that many or most participants in the trend feel, on-balance, negatively about the trend, but nevertheless go along with it because they do not want to be left out. Cases like this would, in fact, be a form of another familiar sort of social phenomenon: the “collective action problem”, whose essential game-theoretic structure is exhibited in the “prisoners’ dilemma” (Tucker and Kuhn 1950). Sometimes everyone would prefer that nobody ϕ s (“defects”), but each prefers the outcome where *they* ϕ along with everyone else (or where they ϕ and others do not) to the outcome where everyone else ϕ s but they alone do not. Thus it can easily happen that everybody ψ s despite their dislike of ψ ing, together bringing about an outcome none of them wants. A wide range of social problems, from climate change to consumerism, have been modelled in these terms (e.g. Heath 2000; Cole 2008).

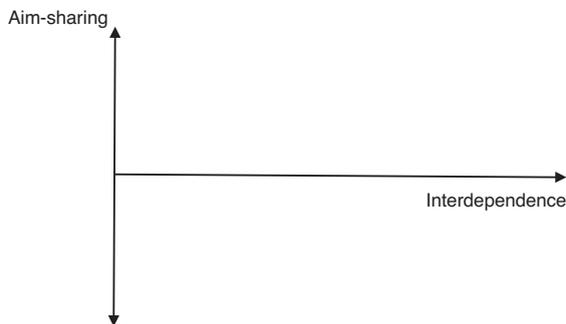
A collective action problem is not a matter of joint action, but of its absence: what is needed to avoid the bad outcome is more coordination among individuals in their striving for the goals they all share. Collective actions problems, however, *are* characterised by interdependence among individual actions: each person acts the way they do because of how they see others act or expect them to act. So interdependence all by itself is insufficient for jointness. Similar remarks apply to trends: they may display high interdependence among

individual actions, but that is not by itself enough to make them anything like a joint action.

4.4 Aim-Sharing and Interdependence Together

What if we add enough of both shared goals *and* interdependence to an aggregate action? Will that be enough to get us a joint action? More or less, we think “yes”. If we all want the same outcome, and each do our bit towards it because we expect the others to do theirs, it seems reasonable to treat the result as one we brought about jointly.

This suggests the possibility of plotting a graph with aim-sharing on one axis and interdependence on the other:

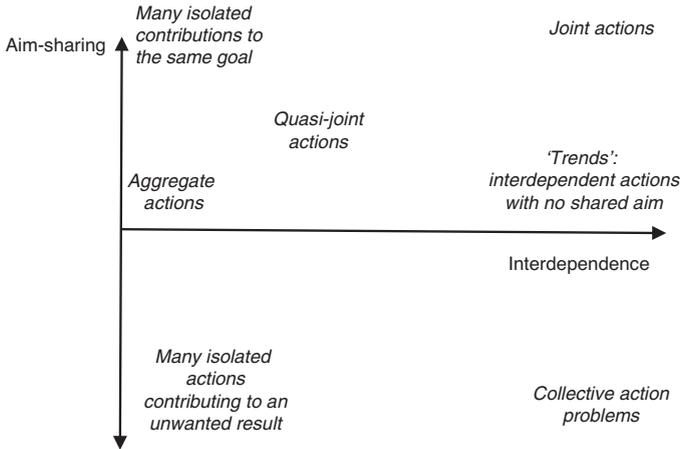


Suppose we plot on this graph various cases where many individuals act, and their actions add up to some overall result. In the middle-left (zero aim-sharing, zero interdependence) are mere aggregate actions, where people act independently and are indifferent to what they bring about together (like those who walk across the lawn and do not care about eroding the soil). At the top-left (high aim-sharing but zero interdependence) would be our scattered crusaders, each labouring independently for their shared cause. At the bottom-left (negative aim-sharing but zero interdependence) would be the opposite: people who actively do not want the cumulative outcome of their individual actions, but do not coordinate in any way (we might consider people who together wear away a path but are saddened by that result).

Over on the right are cases where each person acts because others do; at the top-right (high aim-sharing, high interdependence) in a joint action, working together for a shared goal; at the bottom-right (negative aim-sharing, high interdependence) in a collective action problem, spiralling together towards a hated outcome; in the middle-right (zero aim-sharing, high interdependence), in

something like a trend, indifferent to the overall result but nevertheless sensitive to what others do.

Implicit coordination, as a form of quasi-joint action, is located in the upper-centre of this graph:



That is, we are interested in cases where people act with enough overlap of their aims to distinguish them from participants in a trend or a collective action problem, and with enough interdependence to distinguish them from many isolated crusaders, but without enough of either to qualify them as engaged in a fully joint action in the usual sense.⁸

Thinking of quasi-joint actions, and in particular implicit coordination, as lying somewhere in a two-dimensional space with fully joint actions at one corner and merely aggregate actions at the opposite corner provides us with tools for characterising them more precisely than just saying they are “like joint actions but not quite”. We can distinguish those which are more “cooperative” in the sense of exhibiting greater sharing of aims than others, and those which are more “coordinated” in the sense that participants’ actions are more interdependent.

We could also use this framework to understand the diachronic development of different cases: a movement or a tradition might change on these dimensions over time, or a fully joint action might evolve slowly from a quasi-joint action by activities (like conferences and debates) which steadily increased its interdependence and aligned the participants’ aims with one another.

⁸ Perhaps sometimes people’s actions have high enough aim-sharing and interdependence to count as a full-fledged joint action, all in virtue of a shared implicit intention. But this seems likely to be rare.

Return to our example where everyone stops buying brand X. Perhaps at first this causes only a mild downturn in sales, but then many individual activists who had agitated against the brand on previous occasions decide to try and “ride” the public mood, and reach out to make contact with each other. Their explicit calls for a boycott get more traction now that they are reinforced by listener’s vague awareness that “people are starting to avoid brand X”, and as the momentum builds the interdependence of the process increases. Perhaps there are “pledges” not to invest in brand X that organisations sign, or are asked to sign, in direct response to their competitors signing them. But perhaps also increasingly many people or groups start to dissociate themselves from brand X not because of any, even socially influenced, sense of its being unethical: maybe they start to feel compelled by public pressure even though they dislike the whole phenomenon and privately think brand X did nothing wrong – and in bowing to this pressure they strengthen the pressure on others. Over the course of these events, the movement against brand X has gone from the “scattered crusaders” corner (top left) to the “trend” corner (mid-right), as interdependence has increased and aim-sharing decreased.

5 Applications

We set out to analyse implicit coordination, and to argue that it constitutes a sort of quasi-joint action, guided by a shared implicit intention which is constituted by the complex of interdependent overlapping partly-implicit intentions of individuals. We believe that understanding implicit coordination as quasi-joint action, lying between emergent coordination and explicit coordination, can be helpful not just in enriching social ontology and action theory, but in interpreting and understanding the targets of social criticism.

Recall our examples from Section 1: a positive and negative claim about what a group has done collectively, that seem like they could be true even if the individuals involved never explicitly aimed at doing it.

“Together, we’ve built a vibrant, thriving literature on this topic, motivated by our conviction of its importance.”

“Those with power in this industry have actively, collectively, perpetuated a fertile climate for sexual harassers, motivated by their desire to preserve their power and status.”

Looking at claims like this as asserting implicit coordination suggests a very plausible set of truth-conditions for them. The first claim would mean that lots of people did things that contributed to there being a thriving literature (writing, reading, citing, organising conferences, raising the topic in other discussions,

etc.); that they were encouraged to do so by a background awareness of others doing similar things; and that their intentions, though differing in explicit aim (e.g. winning an argument, getting a publication, padding a CV, etc.), often had the implicit aim of building a literature. This implicit aim might be manifest just as a tendency to be pleased by the topic being discussed or taken seriously, and to be more attracted to actions associated with that.

Similarly, the second claim would mean that many people did things that contributed to protecting sexual harassers (looking the other way, normalising what should be warning signs, being cold with accusers and supportive of anyone accused, etc.); that they were encouraged to do so by a background awareness of others doing similar things; and that their intentions, though differing in explicit aim (e.g. dodging difficult work, avoiding bad publicity, cultivating a useful friendship, etc.), often had an implicit aim of protecting the power held by senior over junior people in the industry. This implicit aim might be manifest just as a sense of anxiety at any suggestion of junior people being able to ignore the wishes of, or affect the interests of, senior people, and a tendency to react negatively to such things.

Consider also the idea of “microaggressions” (Pierce 1970), small individual acts which subtly disparage a certain group, and which are sometimes thought to play a key causal role in perpetuating social inequality. Huebner (2016) describes a certain category of microaggression that adds up particularly clearly to a social-level phenomenon, namely the use of the racial slur “Redskins” for an American football team. In the face of repeated requests by indigenous Americans to change the name, its defenders claim that their use of it carries no hostile intention. In Huebner’s analysis, “Microaggressions [...] are constituents of spatially and temporally distributed macroaggressions, which are carried out by multiple people in multiple places; put differently, aggregated behaviors can yield aggressive acts, even if individuals do not, and perhaps cannot perceive the aggressiveness of their individual actions” (2016). Individual utterances of the name add up to the cumulative effect of enshrining a slur in everyday discourse, and in doing so alienating and disparaging indigenous Americans. But is this a mere aggregate fact, or something *done*: is it a quasi-joint action on the part of those who use this name? Is it implicit coordination?

Whether this is implicit coordination depends on what people’s operative intentions are, and that is not always easy to determine. Certainly it is not something that we can realistically demonstrate in this paper. But the possibility which our analysis casts some light on – and which is suggested by the anger felt toward this phenomenon – would be that some large number of individual microaggressions display both aim-sharing and interdependence. Their interdependence would come through the way that each person’s use of the term normalises the next, makes it easier and more comfortable to use it while making it more effortful to avoid it. The aim-sharing would come through implicitly intending a

state of affairs which might be explicitly described as “maintenance of a culture where whiteness is the default and other ethnicities feature not as fully-realised subjects but as colourful caricatures to entertain the white spectator.”

Of course that explicit description would not likely be admitted to by many people: it would not typically be part of the manifest intention. But perhaps the implicit, inarticulate attraction towards that sort of outcome were partly responsible for motivating each individual action – perhaps, in particular, they are disposed to experience a sense of frustration, confusion, or even betrayal at the prospect of finding themselves in a culture *not* fitting that description (perhaps such a culture would prompt them to say “this doesn’t feel like America anymore”, or “my country has been taken away from me”). If so, then it would not be inappropriate to say that they implicitly ‘use the name “Redskins” *in order to* maintain such a white-centric culture. If their uses were also substantially interdependent, then we could further say that they do so quasi-jointly. If that were the situation, then this case would be an instance of implicit coordination.

Why does it matter? Does it affect our evaluations of the phenomenon or of the participants? When evaluations focus on results, and the morality or immorality of bringing about those results, implicit coordination may not matter. However, our evaluations are not just sensitive to results, but also to intent, and here the possibility of implicit coordination is important. Many of our emotional reactions to things only really make sense if they are intentional actions, if some sort of motive can be discerned behind them.

Sometimes large-scale social phenomena lack any guiding intention, and so are evaluatively rather like a weather event: just as I may be upset when the weather inconveniences me, but could not reasonably feel resentment, indignation, or a desire for retribution, I may be upset that my institutions are too weak, or that economic fluctuations have taken away my job, or that something I used to like is no longer cool, but cannot sensibly feel the kind of emotions towards these phenomena that would attribute them motives. Of course sometimes there are individuals who can be blamed for their specific actions, or for their failure to better regulate matters, but the social phenomena themselves cannot sensibly be blamed for what they do.

But when the social phenomenon is a joint action, the presence of a shared intention makes reactive emotions appropriate. Moreover, when there is an intention behind what happens, we can make more fine-grained evaluations based on its particular content. And injury that is caused by an intentional act might be excused by your not knowing it would happen, or justified by the need to prevent a greater harm; it might qualify as “callous” if you foresaw it but were indifferent, or as “cruel” if you sought to inflict it for its own sake.

Indeed, people who want to be able to evaluate social phenomena this way sometimes lapse into, or say things to suggest, thinking of those phenomena

as explicit joint actions, as a sort of “grand conspiracy” by capitalists/movie executives/liberal academics/someone else. Part of the usefulness of analysing implicit coordination is to undercut the visceral appeal of this sort of distortion.

With implicit coordination, we suggest, the presence of a shared implicit intention also makes such judgements appropriate, at least to some extent. The more that, for instance, the use of racial slurs or other microaggressions flows from a network of implicitly connected intentions – the more it displays high aim-sharing and high interdependence – the more appropriate it is to judge it in intentional terms, as disrespectful, excusable, vindictive, lazy, oblivious, or so on, at a collective level rather than simply as the aggregate effect of individual acts. Much social criticism has this sort of character: identifying large-scale phenomena which nobody has explicitly planned or coordinated but which still display a sort of systematic hanging-together. If our argument succeeds, then it is often justifiable to evaluate such phenomena in the kinds of terms, and with the kind of emotions, appropriate to an intentional act.

Acknowledgement: The authors are funded by the Volkswagen Foundation project on Situated Cognition.

Bibliography

- Asch, Solomon (1956): “Studies of Independence and Conformity: I. A Minority of One Against a Unanimous Majority”. In: *Psychological Monographs: General and Applied* 70. No. 9, p. 1–70.
- Bacharach, Michael (2006): *Beyond Individual Choice: Teams and Frames in Game Theory*. Princeton: Princeton University Press.
- Bargh, John, Mark Chen and Lara Burrows (1996): “Automaticity of Social Behavior: Direct Effects of Trait Construct and Stereotype Activation on Action”. In: *Journal of Personality and Social Psychology* 71. No. 2, p. 230–244.
- Bicchieri, Cristina and Hugo Mercier (2014): “Norms and Beliefs: How Change Occurs”. In: *The Jerusalem Quarterly* 63, p. 60–82.
- Bicchieri, Cristina (2006): *The Grammar of Society: The Nature and Dynamics of Social Norms*. New York: Cambridge University Press.
- Bicchieri, Cristina (2016): *Norms in the Wild: How to Diagnose, Measure, and Change Social Norms*. Oxford: Oxford University Press.
- Bicchieri, Cristina and Hugo Mercier (2012): “Norms and Beliefs: The Dynamics of Change.” In: B. Edmonds (Ed.): *The Dynamic View of Norms*. Cambridge University Press.
- Block, Ned (1995): “On a Confusion about a Function of Consciousness”. In: *Behavioral and Brain Sciences* 18. No. 2, p. 227–287.
- Bond, Rod and Peter Smith (1996): “Culture and Conformity: A Meta-Analysis of Studies Using Asch’s (1952b, 1956) Line Judgment Task”. In: *Psychological Bulletin* 119. No. 1, p. 111–137.

- Bratman, Michael (1984): “Two Faces of Intention”. In: *Philosophical Review* 99, p. 375–405.
- Bratman, Michael (1993): “Shared Intention”. In: *Ethics* 104. No. 1, p. 97–113.
- Bratman, M. (2009): “Modest Sociality and the Distinctiveness of Intention”. In: *Philosophical Studies* 144, p. 149–165.
- Bratman, M. (2014): *Shared Agency: A Planning Theory of Acting Together*. Oxford: Oxford University Press.
- Chant, Sara (2007): “Unintentional Collective Action”. In: *Philosophical Explorations* 10. No. 3, p. 245–256.
- Chartrand, T. L. and J. A. Bargh (1999): “The Chameleon Effect: The Perception–Behavior Link and Social Interaction”. In: *Journal of Personality and Social Psychology* 76. No. 6, p. 893–910.
- Chevallier, C., G. Kohls, V. Troiani, E. Brodtkin and R. Schultzl (2012): “The Social Motivation Theory of Autism”. In: *Trends in Cognitive Science* 16. No. 4, p. 231–239.
- Christensen, W., J. Sutton and D. McIlwain (2016): “Cognition in Skilled Action: Meshed Control and the Variety of Skill Experience”. In: *Mind and Language* 31. No. 1, p. 37–66.
- Cole, D. (2008): “Climate Change and Collective Action”. In: *Current Legal Problems* 61. No. 1, p. 229–264.
- Dworkin, Ronald (1986): *Law’s Empire*. Cambridge: Belknap Press.
- Frankfurt, Harry (1978): “The Problem of Action”. In: *American Philosophical Quarterly* 15. No. 2, p. 157–162.
- Frankfurt, H. G. (1988): *The Importance of What We Care About: Philosophical Essays*. Cambridge: Cambridge University Press.
- Fridland, E. (2014): “They’ve Lost Control: Reflections on Skill”. In: *Synthese* 191, p. 2729–2750.
- Fridland, E. (2015): “Automatically Minded”. In: *Synthese*. <https://doi.org/10.1007/s11229-014-0617-9>.
- Gilbert, M. (1990): “Walking Together: A Paradigmatic Social Phenomenon”. In: *Midwest Studies in Philosophy* 15, p. 1–14.
- Gilbert, M. (2009): “Shared Intention and Personal Intentions”. In: *Philosophical Studies* 144. No. 1, p. 167–187.
- Gollwitzer, P. M. (1999): “Implementation Intentions: Strong Effects of Simple Plans”. In: *American Psychologist* 54. No. 7, p. 493–503. <https://doi.org/10.1037/0003-066X.54.7.493>.
- Gollwitzer, P., U. C. Bayer and K. C. McCulloch (2005): “The Control of the Unwanted”. In: R. Hassin, J. Uleman and J. A. Bargh (Eds.): *The New Unconscious*. Oxford: Oxford University Press, p. 485–515.
- Haslanger, Sally (2006): “What Good Are Our Intuitions? Philosophical Analysis and Social Kinds”. In: *Aristotelian Society Supplementary Volume* 80. No. 1, p. 89–118.
- Heath, J. (2000): “Ideology, Irrationality and Collectively Self-Defeating Behaviour”. In: *Constellations* 7. No. 3, p. 363–371.
- Hohweg, J. (2011): “Phenomenal Variability and Introspective Reliability”. In: *Mind & Language* 26. No. 3, p. 261–286.
- Huebner, Bryce (2016): “Socialized Attention and Situated Agency”. In: *MindsOnline*, at *Brains-blog.com*, <http://mindsonline.philosophyofbrains.com/2016/2016-2/huebner-socialized-attention-and-situated-agency/>, visited on March 12 2018.
- Kagan, Shelly. (2011): “Do I Make a Difference?” In: *Philosophy & Public Affairs* 39. No. 2, p. 105–141.
- Kirschner, S. and Tomasello, M. (2010). “Joint Music Making Promotes Prosocial Behavior in 4-Year-Old Children”. In: *Evolution and Human Behavior* 31. No. 5, p. 354–364.

- Knoblich, G., S. Butterfill and N. Sebanz (2011): "Psychological Research on Joint Action: Theory and Data". In: *Psychology of Learning and Motivation* 54, p. 59–101.
- Lakin, J., V. Jefferis, C. Cheng and T. Chartrand (2003): "The Chameleon Effect as Social Glue: Evidence for the Evolutionary Significance of Nonconscious Mimicry". In: *Journal of Nonverbal Behavior* 27. No. 3, p. 145–162.
- List, C. and P. Pettit (2011): *Group Agency: The Possibility, Design, and Status of Corporate Agents*. Oxford: Oxford University Press.
- McKay, R. T. and D. C. Dennett (2009): "The Evolution of Misbelief". In: *Behavioral & Brain Sciences* 32. No. 6, p. 493–510.
- Mele, A. R. (1992): *Springs of Action: Understanding Intentional Behavior*. New York; Oxford: Oxford University Press.
- Miller, D. T. and C. McFarland (1991): "When Social Comparison Goes awry: The Case of Pluralistic Ignorance". In: J. Suls and T. A. Wills (Eds.): *Social Comparison: Contemporary Theory and Research*. Hillsdale, NJ: Erlbaum.
- Montero, B. (2010): "Does Bodily Awareness Interfere with Highly Skilled Movement?" In: *Inquiry* 53, p. 105–122.
- Nefsky, J. (2011): "Consequentialism and the Problem of Collective Harm: A Reply to Kagan". In: *Philosophy and Public Affairs* 39. No. 4, p. 364–395.
- Nefsky, Julia (2018): "Consumer Choice and Collective Impact". In: A. Barnhill, M. Budolfson and T. Doggett (Eds.): *The Oxford Handbook of Food Ethics*. Oxford: Oxford University Press, p. 267–286.
- O'Shaughnessy, B. (1980): *The Will: Volume 2, A Dual Aspect Theory*. Cambridge: Cambridge University Press.
- O'Shaughnessy, B. (2008): *The Will: Volume 2, A Dual Aspect Theory*. Cambridge: Cambridge University Press.
- Pacherie, Elizabeth (2013): "Intentional Joint Agency: Shared Intentions Lite". In: *Synthese* 190. No. 10, p. 1817–1839.
- Papineau, D. (2013): "In the Zone". In: *Royal Institute of Philosophy Supplements* 73, p. 175–96.
- Pierce, Chester M. (1970): "Black Psychiatry One Year After Miami". In: *Journal of the National Medical Association* 62, p. 471–473.
- Potter, Andrew (2018): "Is this Time Different? What can we learn from #MeToo and #NeverAgain". At *In Due Course*, <http://induecourse.ca/what-can-we-learn-from-metoo-and-neveragain/>, visited on March 12 2018.
- Richardson, M., K. Marsh, R. Isenhower, J. Goodman and R. Schmidt (2007): "Rocking Together: Dynamics of Intentional and Unintentional Interpersonal Coordination". In: *Human Movement Science* 26, p. 867–891.
- Reddish, P., R. Fischer and J. Bulbulia (2013). "Let's Dance Together: Synchrony, Shared Intentionality and Cooperation". In: *PLOS ONE* 8. No. 8, p. e71182.
- Roughley, N. (2016): "Intentions, Decisional and Nondecisional". In: *Wanting and Intending*. Dordrecht: Springer Netherlands, p. 259–291.
- Shapiro, J. Scott (2014): "Massively Shared Agency". In: M. Vargas and G. Yaffe (Eds.): *Rational and Social Agency: The Philosophy of Michael Bratman*. Oxford: Oxford University Press, p. 257–294.
- Seligman, M., P. Railton, R. Baumeister and C. S. Sripada (2016). *Homo Prospectus*. New York: Oxford University Press.
- Stanley, J. and J. W. Krakauer (2013): "Motor Skill Depends on Knowledge of Facts". In: *Frontiers in Human Neuroscience* 7, p. 503.

- Sugden, R. (2001): "The Evolutionary Turn in Game Theory". In: *Journal of Economic Methodology* 8. No. 1, p. 113–130.
- Sugden, R. (2007): "The Logic of Team Reasoning". In: *Philosophical Explorations* 6. No. 3, p. 165–181.
- Sunstein, Cass (Ms.): "Unleashed". Available at SSRN: <https://ssrn.com/abstract=3025749>, visited on March 12 2018.
- Sutton, J., D. J. F. McIlwain, W. D. Christensen and A. Geeves (2011): "Applying Intelligence to the Re Exes: Embodied Skills and Habits between Dreyfus and Descartes". In: *Journal of the British Society for Phenomenology* 42, p. 78–103.
- Tucker, A. and H. Kuhn (1950): *Contributions to the Theory of Games* (Eds.): Annals of Mathematical Studies. Princeton: Princeton University Press.
- Tuomela, R. and K. Miller (1988): "We-Intentions". In: *Philosophical Studies* 53, p. 367–389.
- Tye, Michael (2000): *Consciousness, Color, and Content*. Cambridge, MA: MIT Press.
- Velleman, D. (2000): *The Possibility of Practical Reason*. Michigan: Michigan Publishing.
- Wieber, F., J. L. Thürmer and P. M. Gollwitzer (2015): "Promoting the Translation of Intentions Into Action by Implementation Intentions: Behavioral Effects and Physiological Correlates". *Frontiers in Human Neuroscience* 9.
- Zawidzki, T. W. (2013): *Mindshaping: A New Framework for Understanding Human Social Cognition*. Cambridge, MA: MIT Press.