

## Conventions and Status Functions

Conventions and status functions are central features of social life. How are they related? We argue that

- a) there is a variety of convention that has not been adequately identified in the literature,
- b) status functions constitutively involve this variety of convention, and
- c) what is special about it explains the central feature of status functions, namely, that objects with status functions can perform their functions only insofar as they have been collectively accepted as having them.

We will call this variety of convention *effective coordinating agreement* (ECA). It need not involve explicit agreement, but it is the kind of state sincere explicit agreement leads its parties to. Its purpose is to solve coordination problems. Its key feature is that it does so through a structure of interrelated, generalized, conditional intentions directed at a collective action plan.

We believe that ECA is an important explanatory social kind, so far unrecognized or underappreciated, and that it is a kind of convention. Here we are primarily interested in arguing for the first thesis. We advance considerations in favor of regarding ECA as a subtype of convention, but we do not wish to become mired in debates over common usage.<sup>1</sup> ECA shares salient features with practices widely recognized as conventions

---

<sup>1</sup> There are different strands in the ordinary notion of convention. There are conventions (like driving on the right) that solve coordination problems, which David Lewis took as

and in virtue of special features it plays a role in assigning status functions that practices identified by other analyses of convention cannot.

In section I, we explain what an ECA is. In section II, we compare it with two influential notions of convention—Lewisian conventions<sup>2</sup> and Gilbertian conventions.<sup>3</sup> We argue that ECAs share key features with these familiar kinds of convention. They are thus well-suited to play the role that convention plays in the social world. In section III, we show that they play this role by showing that status functions, given how they are defined by constitutive rules governing social transactions in which they figure, constitutively involve ECAs in their assignment. What is distinctive about ECAs explains why they are suited for this role where other kinds of convention are not, and it gives precise content to the claim that objects with status functions can perform their functions only insofar as they have been collectively accepted as having them. In section IV, we

---

his subject in *Convention* (Cambridge: Harvard University Press, 1969). There are practices like sending thank you cards after a dinner party or eating with forks or chopsticks (Margaret Gilbert, *On Social Facts* (London: Routledge, 1989)). Some of these may be sustained by the weight of precedent (Ruth Garrett Millikan, *Language: A Biological Model* (Oxford: Clarendon Press, 2005)). Our interest lies in the strand in which conventions are seen as solutions to coordination problems. Here, we will claim, something important has been overlooked, or maybe overshadowed, by the Lewisian precedent in particular.

<sup>2</sup> Lewis, *Convention*, *op. cit.*

<sup>3</sup> Margaret Gilbert, "Social Convention Revisited," *Topoi*, xxvii, 1-2 (2008).

show that this requires a sharp distinction between constitutive rules and conventions. We summarize and conclude in section v.

#### I. EFFECTIVE COORDINATING AGREEMENT AS A KIND OF CONVENTION

An ECA is a shared understanding among members of some group concerning how *they* are to perform a collective intentional action that could be performed in several ways. Take an example: a group of friends decide to form a book club that meets at a café every Thursday at 8 pm. Initially, some of them arrive at 7:45 pm so that the discussion can start at 8 pm, while others arrive at 8 pm, taking a further 15 minutes to settle in. Each member would prefer that no one have to wait for the book discussion to start more than necessary, and it makes no difference to them whether they arrive or start at 8 pm. Eventually, either through explicit agreement or by simply settling into a pattern, they coordinate together on arriving at 8 pm. If asked about the timing of the meetings, each member would say: *we* arrive at 8 pm.

The book club has an ECA to arrive at the meetings at 8 pm. ECAs have two important characteristics: they are structures of *intentions* of the parties, and they solve coordination problems.

The structure of intentions characteristic of ECA is akin to the effect of a sincere explicit agreement to do something together. The book club has an ECA to arrive at the meetings at 8 pm only when each member intends to participate in *their* arriving at the café at 8 pm *whenever* they have a meeting. These are generalized conditional *we*-intentions.

A *we-intention* is the sort of intention that is directed at doing something with others when participating with them in collective intentional action.<sup>4</sup> We take an individual member of a group *G* who intends to participate in *G*'s performance of a collective intentional *J*-ing to have a *we-intention directed towards the group's J-ing*. We say that this individual *we-intends* that *G J-s*. Since she is a part of *G*, she intends to participate in and contribute to *G*'s *J*-ing. More specifically, as we understand it, a *we-intention* is a commitment to act with others in accordance with a collective action plan *P* (at the time of action), and consequently an individual member of *G* who *we-intends* that *G J-s* thereby *we-intends* that *G J-s* in accordance with their collective plan *P* for *J*-ing. Since intending that *G J* in accordance with a collective action plan requires for its success that each member of *G* intentionally act in accordance with the plan as their plan, it follows that that no member's *we-intention* can be fulfilled unless the *we-intentions* of every member of the group are fulfilled. When each member of *G* *we-intends* that the group *J*, the members of *G* *share an intention to J* in accordance with some *P*.<sup>5</sup> The group *J*'s jointly intentionally when they successfully execute their shared intention, which is to say, each of them successfully executes her *we-intention* that they *J*.

---

<sup>4</sup> See Raimo Tuomela and Kaarlo Miller, "We-Intentions," *Philosophical Studies*, LIII, 3 (1988); Kirk Ludwig, *From Individual to Plural Agency: Collective Action 1*, 2 vols. (Oxford: Oxford University Press, 2016).

<sup>5</sup> For further discussion of the shared plan component see Ludwig, *From Individual to Plural Agency: Collective Action 1*, *op. cit.*, §14.3.

A *conditional* intention is an intention to do something on a certain contingency obtaining.<sup>6</sup> A *conditional we-intention* is a we-intention to do something *with others* on a certain contingency obtaining. For example, Henrietta's intention to go to the prom with Harold if he asks is a conditional we-intention. A *generalized conditional we-intention* is general with respect to occasion of execution, that is, it is a conditional we-intention to do something with others *whenever* a certain contingency obtains. Henrietta's intention to go out with Harold *whenever* he asks is a generalized conditional we-intention.

The book club's ECA comes to exist when each member forms a stable *generalized conditional we-intention* directed at contributing to their arriving at the café at 8pm for their meetings (that is, when each member *we-intends* that *they* arrive at 8 pm *whenever* they have a meeting). These are the kinds of intentions they would have if they explicitly and sincerely agreed to arrive at 8 pm whenever they have a meeting, but they may arrive at them in other ways, for example, by settling into a pattern. What is characteristic of ECAs are the interrelated intentions of the parties, not how they arrive at them.

A second feature of ECAs is that the members' conditional we-intentions solve a coordination problem for the group. Settling on when to arrive is a coordination problem, in the sense that the group needs to coordinate on one of several possible ways of achieving their goal of no one having to wait. Thus, their ECA to arrive at 8 pm whenever they have a meeting *solves* a coordination problem for the club. More

---

<sup>6</sup> Kirk Ludwig, "What Are Conditional Intentions?," *Method: Analytic Perspectives*, IV, 6 (2015).

generally, a group's shared intention to act in accordance with a plan of action  $P$  in certain circumstances  $C$  solves a coordination problem when, in acting in accordance with  $P$ , members of  $G$  we-intend to achieve an end that they could have achieved roughly equally well by acting in accordance with an alternative collective action plan  $P'$ .

Putting this together, we analyze *effective coordinating agreement* (ECA) as follows.

Members of a group  $G$  have an *effective coordinating agreement* (an ECA) just in case there is a collective action plan  $P$  and an end  $E$  such that

- (i) each member of  $G$  we-intends that they act in accordance with  $P$  in order to achieve end  $E$ , whenever she is in a circumstance of type  $C$ ;<sup>7</sup>
- (ii) there is an alternative,  $P'$ , to  $P$  in  $C$  by which the members of  $G$  could collectively achieve  $E$  in any instance of  $C$  by acting together intentionally in accordance with  $P'$ .

---

<sup>7</sup> Miller characterizes a convention as a set of conditional practices in a group that satisfy a collective end the agents mutually believe one another to have (Seumas Miller, *Social Action: A Teleological Account* (Cambridge: Cambridge University Press, 2001), ch. 3, esp. pp. 115-22). For Miller, a *collective end* is an end several agents have that cannot be realized by any of them alone (*ibid.*, pp. 56-72). Miller's collective ends can be satisfied even when agents who have them do not act together intentionally. Since it is insufficient for shared intention, this account does not capture the sort of conceptually central social convention to which we aim to draw attention.

When these conditions are satisfied, *G* has an effective coordinating agreement to act in accordance with *P* in *C*.

We highlight four important features of ECAs that are standardly associated with conventions that solve coordination problems.

(1) They are *arbitrary* in the roughly Lewisian sense<sup>8</sup> that effective agreement on an alternative action plan *P'* would have enabled the parties to achieve *E* roughly equally well.

(2) They are *social* both in the sense that they involve a plan of action that is collectively accepted by the group and in the sense that the parties' acting in accordance with an ECA is collective intentional action.

(3) They are *stable* due to the stability of intentions, which tend to resist reconsideration once formed.

(4) They are *reciprocal* in that the participants see their involvement and that of others in the same way.

The fourth feature, *reciprocity*, requires more explanation.

---

<sup>8</sup> Lewis, *Convention*, *op. cit.*

Lewis observed that not all social practices that solve coordination problems are conventions.<sup>9</sup> To see why, suppose that the book club establishes the practice of arriving at 8 pm in the following way. Each member prefers that no one wait unnecessarily to arriving either at 8 pm or at 7:45 pm, but does not think the others do. Rather, each has the false belief that the others will arrive at 8 pm out of habit, for no reason and regardless of what they expect anyone else to do. In such a case, a rational club member will arrive at 8 pm. Since there is no incentive for anyone to deviate, the club will have a stable practice of arriving at 8 pm that “solves” a coordination problem. But, as Lewis notes, we do not see such a practice as a convention. Moreover, requiring that no one think that others will arrive at 8 pm for no reason does not ensure that the practice is a convention. For suppose that every member had the following false second order belief: that everyone else expects her to arrive at 8 pm out of habit. The group would again have a practice of arriving at 8 pm. But, Lewis thinks, and we agree, it would not be a convention. And so on—the cases can be iterated infinitely.

We take *reciprocity* to be whatever feature of convention makes it incompatible with cases like these, where there is broadly speaking an “asymmetry between the way the parties view themselves and the way they view [the other parties].”<sup>10</sup> This asymmetry is reflected in the way each member sees herself as *taking advantage of the others’ behavior*. In the initial case, each member thinks that everyone else arrives at 8 pm simply out of habit, while she takes advantage of their expected behavior in

---

<sup>9</sup> *Ibid.*, p. 59

<sup>10</sup> Gerald J. Postema, “Salience Reasoning,” *Topoi*, xxvii, 1-2 (2008): at p. 44.



choosing to arrive at 8 pm. A principled consideration in favor of requiring reciprocity is Hume's observation that convention of the variety of interest is motivated by "a general sense of common interest; which sense all the members of the society express to one another."<sup>11</sup> The agents in the examples above do not seem motivated by a general sense of common interest, as each thinks that others either lack this motivation (that is, that each of the others will arrive at 8 pm no matter what she does), or think that she lacks it, and so on. We take reciprocity to express in part this Humean requirement.

Lewis took non-reciprocal regularities of the kind described above to show that the book club's practice of arriving at 8 pm is a convention only if there is *common knowledge* in the group that everyone conforms, expects the others to conform, and would prefer to take part in an alternative practice if others did.<sup>12</sup> In the examples above,

---

<sup>11</sup> David Hume, *A Treatise of Human Nature*, 2d ed. (New York: Clarendon Press, 1978), at p. 490.

<sup>12</sup> According to Lewis, *Convention, op. cit.*, p. 56, the proposition that *p* is common knowledge in the group *G* iff a state of affairs *A* obtains such that

1. Everyone in *G* has reason to believe that *A* holds;
2. *A* indicates to everyone in *G* that everyone in *G* has reason to believe that *A* holds;
3. *A* indicates to everyone in *G* that *p*.

Clauses (1)-(3), along with suitable assumptions about shared inductive standards and background information (*ibid.*, pp. 52-56), entail an infinite series of propositions about what agents have a reason to believe, so that when it is common knowledge in *G* that *p*,

each member has a false belief about the others' reasons for conforming. In the first iteration, each thinks that everyone else arrives at 8 pm simply out of habit. In the second, each thinks that everyone else thinks that everyone else arrives at 8 pm simply out of habit. And so on. If it were common knowledge among the members that everyone arrives and expects others to arrive at 8 pm, given their public preferences, they would have a reason to think these beliefs are false. (For everyone would have a reason to think that: everyone expects everyone else to arrive at 8 pm, and so does not arrive at 8 pm simply out of habit; everyone has a reason to think that everyone expects everyone else to arrive at 8 pm and so does not expect others to arrive at 8 pm simply out of habit; and so on.) Thus, unless they believe (or think that others believe, or that others believes that others believe, and so on) what they have a reason to think is false, they do not commonly know that everyone arrives and expects others to arrive at 8 pm, and so they do not have a Lewisian convention to arrive at 8 pm.<sup>13</sup>

ECA's also rule out non-reciprocal regularities. The agents in these cases do not, provided they are rational, we-intend to arrive at 8 pm when they have a meeting. In the first iteration, each agent believes that the others will arrive at 8 pm out of habit and so believes that the others *do not* we-intend to arrive at 8 pm when they have a meeting. If Alex believes that Farah does not intend that they *J*, it is not rational for him to form a we-intention that they *J*, given that they cannot *J* intentionally without Farah's intentional

---

then everyone in *G* has a reason to think that *p*, everyone has a reason to think that everyone has a reason to think that *p*, and so on.

<sup>13</sup> See Lewis, *Convention, op. cit.*, p. 59, for a more detailed explanation.

participation.<sup>14</sup> In the second iteration, each agent believes that the others believe him to not we-intend to arrive at 8 pm, and so believes that they also will not we-intend to arrive at 8 pm (again, provided they are rational). Given that, it would not be rational for him to we-intend to arrive at 8 pm. And so on. Therefore, it will be irrational for book club members who have the beliefs described in these examples to have the we-intentions characteristic of ECA.

Thinking of a convention as an ECA provides a robust explanation of reciprocity. When a group acts to solve a coordination problem on the basis of an ECA, each member intends that the group act in accordance with a shared plan and intends the others to act on the same intention. Each member of the book club we-intends that *the group* arrive at 8 pm, and as a part of this, intends that they do this in part because of

---

<sup>14</sup> If Farrah we-intends that she and Alex *J*, then if her intention is satisfied, they *J* intentionally. For them to *J* intentionally (in the way required for the satisfaction of their we-intentions), their *J*-ing must come about as a result of their each successfully executing their we-intentions, for they each aim at their intentionally *J*-ing. So one could not be satisfied without the other being satisfied. Therefore, for Farrah's we-intention to be successful, they must *J* in part because of the we-intentions of each of them that they *J*. Supposing that the content of a state is given by its satisfaction conditions, it follows that the content of Farrah's we-intention requires they *J* in part because of Alex's we-intention that they *J*. See for example, Michael Bratman, "Shared Cooperative Activity," *The Philosophical Review*, CI, 2 (1992); Michael Bratman, *Shared Agency: A Planning Theory of Acting Together* (Oxford: Oxford University Press, 2014).

the others' we-intentions that the group arrive at 8 pm. So each of them may say: "I will arrive at 8 pm because this is my part in what we all do when we have meetings." And each *intends* (rather than merely expects) the others to act on the basis of the same intention. So when members act on the basis of their ECA, no one thinks of herself as taking advantage of the others' behavior but rather as *acting with them*. And not seeing oneself as taking advantage of others' behavior is what the requirement of reciprocity essentially comes to.

## II. EFFECTIVE COORDINATING AGREEMENT AND OTHER ACCOUNTS OF CONVENTION

ECAs are arbitrary, social, stable, and reciprocal. They are therefore well suited to play the role of convention. They are also a so far unrecognized kind of practice that can play this role. We highlight this by comparing ECAs with L-conventions (practices that satisfy the conditions of Lewis's 1969 account<sup>15</sup>) and G-conventions (practices that satisfy the conditions of Gilbert's account<sup>16</sup>).

*II.1 ECAs and L-conventions.* According to Lewis:

A regularity  $R$  in the behavior of members of a population  $P$  when they are agents in a recurrent situation  $S$  is a convention if and only if it is true and it is common knowledge in  $P$  that, in any instance of  $S$  among members of  $P$ ,

1. everyone conforms to  $R$ ;

---

<sup>15</sup> Lewis, *Convention*, *op. cit.*

<sup>16</sup> Gilbert, *On Social Facts*, *op. cit.*; Gilbert, "Social Convention Revisited," *op. cit.*

2. everyone expects everyone else to conform to  $R$ ;
3. everyone has approximately the same preferences regarding all possible combinations of actions;
4. everyone prefers everyone else to conform to  $R$ , on condition that at least all but one conform to  $R$ ;
5. everyone would prefer that everyone conform to  $R'$ , on condition that at least all but one conform to  $R'$ ;

where  $R'$  is some regularity in the behavior of members of  $P$  in  $S$ , such that no one in any instance of  $S$  among members of  $P$  could conform both to  $R$  and  $R'$ .<sup>17</sup>

Like ECAs, L-conventions are (i) *social* in the sense that everyone conforms and expects others to; (ii) *arbitrary* in the sense that they solve coordination problems as defined by (3)-(5); (iii) *stable* in the sense that agents with these preferences and expectations do not have an incentive to deviate from the established regularity; and (iv) *reciprocal* because no one thinks of herself as merely taking advantage of the others' behavior as everyone conforms because everyone else conforms, everyone has a reason to think that everyone conforms because everyone else conforms, and so on.

The main difference between L-conventions and ECAs is that the former is a structure of the parties' expectations and preferences while the latter is a structure of their intentions. L-convention is neither necessary nor sufficient for an ECA. It is not sufficient since agents with Lewisian expectations and preferences need not we-intend

---

<sup>17</sup> Lewis, *Convention, op. cit.*, p. 76

to conform to *R*. For example, each driver could individually intend to drive on the right in order to not collide with others, given that he expects the others to do the same.<sup>18</sup> It is not necessary since an ECA can exist between (even rational) agents in the absence of *common knowledge* of conformity and the others' preference orderings. For example, Alex may rationally intend that he and Farah shake hands when they meet even if it is *not* the case that Alex believes or has any reason to believe that Farah has the corresponding we-intention. It may be very important for Alex that they shake hands, or (as is more likely in this case) his extending a hand may not be very costly, and this can make the attempt worthwhile even if he is in doubt of Farah's reciprocating.<sup>19</sup>

---

<sup>18</sup> The point that expectations and preferences that establish an equilibrium are not sufficient for shared intention is well established in the literature. See, for example, Bratman, "Shared Cooperative Activity," *op. cit.*; Gilbert, *On Social Facts*, *op. cit.*; Margaret Gilbert, "Walking Together: A Paradigmatic Social Phenomenon," *Midwest Studies in Philosophy*, xv (1990); John Searle, "Collective Intentions and Actions," in Philip R. Cohen, Jerry Morgan, and Martha E. Pollack, eds., *Intentions in Communication* (Cambridge, Mass.: MIT Press, 1990); Raimo Tuomela and Kaarlo Miller, "We-Intentions and Social Action," *Analyse & Kritik*, vii, 1 (1985); Tuomela and Miller, "We-Intentions," *op. cit.*

<sup>19</sup> See Ludwig, *From Individual to Plural Agency: Collective Action 1*, *op. cit.*, §14.5; Olle Blomberg, "Common Knowledge and Reductionism About Shared Agency," *Australasian Journal of Philosophy*, xciv, 2 (2016). We invoked a belief requirement on we-intention to explain why non-reciprocal regularities described above are not ECAs.

If this is right, then ECAs are neither necessary nor sufficient for L-conventions. However, there is a straightforward relation between them. The psychological conditions that Lewis identifies provide a rational underpinning for the formation of ECAs. If those conditions are in place, then it makes sense for members of the group to take an extra step to form an ECA directed toward a collective action plan  $P$  that represents one of the ways of their achieving a shared end  $E$ . This makes sense not only because of the naturalness of our joining together intentionally in projects that are clearly of mutual benefit but also because it adds stability to the project, which then does not have to rest on the maintenance of common knowledge.

ECAs are present in many examples Lewis uses to fix the notion of convention his account targets. In the call-back convention, when the line gets disconnected, the original caller is to call back. When there is such a convention, each party will we-intend—whenever a phone-call is dropped—to do her part in the joint action of reestablishing the connection, specifically by acting in accordance with an action plan that requires her to call back if she is the original caller and wait otherwise.

*II.2 ECAs and G-conventions.* According to Gilbert, social convention is a jointly accepted simple fiat.

---

But what is required for someone to rationally we-intend that her group  $J$  is that she not have this belief: the others *do not* we-intend that they  $J$ . An individual can satisfy this requirement while lacking the belief (or even any reason to believe) that the other members of her group intend that they  $J$ .

A population  $P$  has a convention of conformity to some regularity in behavior  $R$  in situations of type  $S$  if and only if the members of  $P$  are jointly committed<sup>20</sup> to accept as a body, with respect to themselves, the fiat:  $R$  is to be conformed to [in  $S$ ].<sup>21</sup>

The fiat in question is *simple* in the sense that “no particular rationale for it is presupposed”<sup>22</sup> or is “regarded as holding in the absence of any special justification which may be available.”<sup>23</sup>

---

<sup>20</sup> For more on joint commitment, see for example Margaret Gilbert, *A Theory of Political Obligation* (Cambridge: Cambridge University Press, 2006), ch. 7; Margaret Gilbert, “Joint Commitment,” in Marija Jankovic and Kirk Ludwig, eds., *The Routledge Handbook of Collective Intentionality* (New York: Routledge, 2018). For our purposes, the significant features are that (i) it is a commitment of two or more people, typically created by each expressing, in conditions of common knowledge, readiness to be jointly committed to doing something, accepting a certain proposition, and so on; (ii) it entails normative constraints on all parties in that each is obligated to the others to do her part, and the concurrence of all parties is needed to rescind the joint commitment; consequently, (iii) joint commitment can exist in the absence of corresponding personal commitments or intentions.

<sup>21</sup> Gilbert, “Social Convention Revisited,” *op. cit.*, p. 12

<sup>22</sup> Gilbert, “Social Convention Revisited,” *op. cit.*

<sup>23</sup> Gilbert, *On Social Facts*, *op. cit.*, p. 373.



Examples that Gilbert mentions that satisfy these conditions include sleeping on mats versus mattresses, wearing skirts versus pants,<sup>24</sup> brushing one's teeth before going to bed,<sup>25</sup> sending thank you notes to hosts of dinner parties,<sup>26</sup> as well as the call-back convention and driving on the right.

G-conventions are social, stable, and reciprocal, in at least the sense that agents who have a G-convention will be able to refer to it as "our" convention.<sup>27</sup> Like ECAs, they essentially involve attitudes (that are said to be) responsible for collective action.<sup>28</sup> A major difference between G-conventions on one side and ECAs and L-conventions on the other is that G-conventions are not necessarily solutions to coordination problems, and so are not arbitrary in the operative sense (though they are perhaps arbitrary in another sense connected to the fiats in question being simple). This is something Gilbert explicitly aims for, since she thinks that, for example, a department's decision to dress formally for meetings just because it would be nice (and not because everyone

---

<sup>24</sup> *Ibid.*, p. 343

<sup>25</sup> Margaret Gilbert, "Agreements, Conventions, and Language," *Synthese*, LIV (1983): at p. 234.

<sup>26</sup> Gilbert, *On Social Facts*, *op. cit.*, p. 316.

<sup>27</sup> See Gilbert, "Social Convention Revisited," *op. cit.*, p. 6.

<sup>28</sup> Though Gilbert's joint commitment to a principle of action is not the same thing as a shared intention in favor of a joint action, as we have understood it. See, for example, Margaret Gilbert, "Shared Intention and Personal Intentions," *Philosophical Studies*, CXLIV, 1 (2009).

wants to be dressed at similar levels of formality) counts as a convention. Relatedly, G-conventions can concern individual actions that are not intended as contributions to a joint action, such as sleeping on mats, using forks, and sending thank you notes.

Though G-conventions are not necessarily solutions to coordination problems, it is clear that Gilbert intends the account to apply to practices such as driving on the right, meeting at a particular time, or calling back if one is the original caller.<sup>29</sup> So it is useful to compare a subgroup of G-convention, namely, those that solve a coordination problem, with ECAs. Are they the same thing? Despite significant affinities, the answer is no.

Consider a G-convention to the effect that regularity  $R$  (solving a coordination problem) is to be conformed to in situation  $S$  and an ECA to conform to  $R$  in  $S$ . Gilbert emphasizes that joint commitment to a principle of action can exist in the absence of corresponding personal commitments. Suppose, for example, that some of the parties who have jointly committed to the fiat that  $R$  is to be conformed to in  $S$  stop we-

---

<sup>29</sup> See for example, Gilbert, "Agreements, Conventions, and Language," *op. cit.*, p. 230. It is not clear that Gilbert's account actually applies to practices that solve coordination problems, for it is not clear that fiats governing them hold "in the absence of any special justification" (*ibid.*). Such fiats are accepted *because* they solve the coordination problem (for example, we settle on driving on the right because we want to avoid collisions when there is two-way traffic). What does not have special justification is the choice between alternative solutions (for example, driving on the right vs. left). But, since it is clear that Gilbert intends the account to apply to such cases, we will take it to as well.

intending that they *R* in *S*. When this happens, there is no longer an ECA to conform to *R* in *S* in that group. But Gilbert holds that no party can terminate a joint commitment without concurrence of the others.<sup>30</sup> Suppose that there is no such concurrence. Then the agents who fall off the ECA are still a part of a joint commitment to the fiat that *R* is to be conformed to in *S*, and so are still parties to a G-convention. So ECA is not necessary for a G-convention.

ECA is also not sufficient for G-convention. Gilbert distinguishes agreeing to do something on a regular basis with jointly accepting a norm “commanding such regular action.”<sup>31</sup> For example, if we agree to have lunch once a week, settling on Tuesdays, this would be an ECA among us. But we may establish this ECA without accepting a fiat to the effect that having lunch on Tuesdays *ought to* be conformed to.

More generally, ECA is a psychological category, a structure of intentions among the parties, while G-convention is a normative one, a structure of obligations among the parties. Once agents express to one another their readiness to jointly commit to a fiat, the commitment and associated obligations exist even when parties no longer intend to act in accordance with it.

---

<sup>30</sup> See Gilbert, *A Theory of Political Obligation*, *op. cit.*, pp. 141-44, on the concurrence condition. She notes that there are cases in which there are background agreements or conventions that in effect conditionally grant concurrence to exit the joint commitment should one party want to. But absent such background understanding the concurrence of all parties is required.

<sup>31</sup> Gilbert, *On Social Facts*, *op. cit.*, p. 384.

*II.3 ECAs as ordinary conventions.* We end the comparison of ECAs with other accounts of conventions with two points in favor of construing (at least some) ordinary conventions as ECAs rather than L-conventions or G-conventions.

First, an ECA does not require common knowledge of conformity (unlike L-convention) or of any individual's readiness to jointly commit to a principle of action (unlike G-convention). It is implausible that common knowledge of either conformity to a regularity or readiness to jointly commit to a principle of action is a feature of every instance of ordinary convention.

Common knowledge of conformity plays a crucial role in underpinning the reciprocity of L-conventions. Lewis's account explains how conventions can be reciprocal only among game-theoretically rational agents who share inductive standards.<sup>32</sup> However, it is hardly clear that ordinary agents who participate in conventions are always or even typically game-theoretically rational or share inductive standards. More significantly, Lederman has argued that people can in fact *never* have common knowledge if this requires them to satisfy the conditions of Lewis's definition.<sup>33</sup>

---

<sup>32</sup> Lewis, *Convention, op. cit.*, pp. 52-57. See note 12 above.

<sup>33</sup> Harvey Lederman, "Uncommon Knowledge," *Mind*, cxxvii, 508 (2018). Lewis's notion of common knowledge is different from the one Lederman discusses, but the argument applies to it too—as well as to Gilbert's definition of common knowledge in *On Social Facts, op. cit.*, pp. 194-95—as Lederman notes (Lederman, "Uncommon Knowledge," *op. cit.*, n. 3). See also Ken Binmore, "Do Conventions Need to Be Common

Putting aside technical objections to specific definitions of common knowledge, any requirement that parties to a convention know, or even have more reason to think than not, that there are others participating in it appears too strong for many ordinary cases. One may participate in a convention without even being confident that it is still in place. One might hear about the call-back convention but have reason to think it has died out, but not be certain, and then give it a try, and so participate in it without knowledge that it is in place.

In contrast, there is no obstacle to others adding themselves to a group that embraces a community-wide ECA as long as they all conceive of the group involved as being those who sign on in the community. They can intend to participate even without any confidence that the convention is embraced by others. ECA is thus a more plausible candidate of the sort of thing that goes on in ordinary, less than perfectly rational, cognitively limited agents implementing everyday social conventions of the sort we focus on here.

Second, conventions are often said to involve obligations in the sense that someone violating a convention in a community is open to rebuke for failing to conform to it. Lewis's account is sometimes criticized for not being able to accommodate this.<sup>34</sup> Gilbert's account is designed to accommodate this, since Gilbertain conventions involve

---

Knowledge?" *Topoi*, xxvii, 1-2 (2008), who suggests that it is difficult for common knowledge in Lewis's sense to obtain in large groups though many conventions do.

<sup>34</sup> Gilbert, *On Social Facts*, *op. cit.*; Andrei Marmor, *Social Conventions* (Princeton: Princeton University Press, 2009)

joint commitment to a fiat whose form is: we ought to do *R*. But this seems to go to the opposite extreme. Not all ordinary conventions can be described in terms of their parties accepting a principle of the form: we ought to follow *R* in *C*. It would be odd to describe our book club members as accepting the principle that they *ought* to arrive at 8 pm. It is simply something they do, and something that can be described, in an ordinary sense, as their convention. Their recognizing that they have an ECA of course can give rise to expectations that generate derived obligations to conform or to provide some notification if they expect not to be doing so, since failure to do so may cause inconvenience. We may say here that they have obligations to conform, but this is just to say each of them has a pro tanto obligation to the degree to which they have set up expectations in others about their performance that the others rely on. There is no principle they have taken as theirs which is thereby the source of its authority over them.<sup>35</sup>

ECA steers a course between the Lewisian Scylla of self-interested strategic reasoners and the Gilbertian Charybdis of joint commitment to a norm of action. Seeing at least some sorts of conventions as ECAs provides some rationale for the thought that conventions generate obligations without having to accept that all conventions are acceptances of a normative principle of action.

---

<sup>35</sup> Cf. Michael Bratman, "Shared Intention and Mutual Obligation," in *Faces of Intention: Selected Essays on Intention and Agency* (Cambridge: Cambridge University Press, 1999); Bratman, *Shared Agency: A Planning Theory of Acting Together*, *op. cit.*, pp. 110ff.

### III. STATUS FUNCTIONS REQUIRE ECAs

A status function is a function an object has in a social activity in virtue of the status it has in a group which engages in that type of social activity. Examples are being a royal seal, a twenty-dollar bill, or a driver's license. Searle introduced the notion in the following passage.<sup>36</sup>

The radical movement that gets us from such simple social facts as that we are sitting on a bench together or having a fistfight to such institutional facts as money, property, and marriage is the collective imposition of function on entities, which—unlike levers, benches, and cars—cannot perform the functions solely by virtue of their physical structure. ... The key element in the move from the collective imposition of function to the creation of institutional facts is the imposition of a collectively recognized status to which a function is attached. Since this is a special category of agentive functions, I will label these *status functions*.

An agentive function is a function "assigned relative to practical interests of conscious agents."<sup>37</sup> A status function is an agentive function whose special feature is that objects can perform it only in virtue of the collective acceptance that they are to have that

---

<sup>36</sup> John Searle, *The Construction of Social Reality* (New York: Free Press, 1995), at p. 41.

<sup>37</sup> *Ibid.*, p. 20

function.<sup>38</sup> Being a twenty-dollar bill is a status function because nothing can function as a twenty-dollar bill in monetary transactions unless those involved collectively accept that it has that function.

*Why* do status functions have this special feature? Searle never explains it. For Searle what marks out status functions as special is just that for objects to have them those objects must be collectively accepted as having them. But why is something like collective acceptance necessary? The answer we give is that, given how the relevant functions are defined, imposing them on objects is a coordination problem, one that has to be solved jointly intentionally if it is to be solved at all. Thus, the problem has to be solved by the adoption of a shared intention directed at one of the possible ways members of the relevant group can coordinate in the relevant circumstances. If this is right, then what Searle calls collective acceptance turns out to be what we have characterized as an ECA.

---

<sup>38</sup> As Searle puts it in a more recent book, “for the status functions to actually work, there must be collective acceptance or recognition of the object or person as having that status” (John Searle, *Making the Social World: The Structure of Human Civilization* (Oxford: Oxford University Press, 2010), at p. 7). We give a different account of what ‘collective acceptance’ comes to.



To establish this, we take a detour through the topic of constitutive rules.<sup>39</sup> Constitutive rules are rules the intentional following of which partially constitutes the type of activity they govern. We do not insist that there is only one use of the phrase 'constitutive rule'. But there is a straightforward way of explaining what a (type of) constitutive rule is that gets at the central idea and makes constitutive rules unmysterious.

We start with the idea that there are action types that are essentially intentional and which involve patterns of activity. Essentially intentional action types have to be performed intentionally (unlike, for example, falling). Essentially intentional *collective* action types have to be performed by more than one agent jointly intentionally. Certain types of essentially intentional action involve patterns of activity, for example, playing chess or tic-tac-toe, or solitaire, or a baseball game, and so on. As the examples show, the pattern can involve one or more people. The pattern of activity can be separated from the requirement that it be instantiated intentionally. Two people could go through the motions involved in the play of a game of tic-tac-toe or chess without doing anything together intentionally. One can extract from the rules of chess or tic-tac-toe a neutral description of an activity pattern type. We can think of the activity pattern as a higher-level type that subsumes all the variations of movements by two agents that are in accordance with the rules. We get the concept of chess or tic-tac-toe, and so on, when

---

<sup>39</sup> See Searle, *The Construction of Social Reality*, *op. cit.*, §2.5; Kirk Ludwig, *From Plural to Institutional Agency: Collective Action 2*, 2 vols. (Oxford: Oxford University Press, 2017), ch. 7.

we add to the description of the activity pattern the requirement that it be instantiated by the agents acting jointly intentionally in accordance with the rules, together with the additional requirement (for these cases specifically) that each intend to achieve a certain final state which only one can achieve, what each game defines as a winning position.<sup>40</sup>

Constitutive rules define patterns of activity. They are constitutive *relative to an action type* that requires that the rule-specified pattern of activity be instantiated intentionally—jointly intentionally for joint activity patterns. To see this, notice that any regulative rule can be constitutive relative to a further activity type. Robert's Rules of Order, RRO, are given as a typical example of a set of regulative rules, for they govern an activity type, a meeting, that can exist without being governed by the rules. However, if we define a *parliamentary meeting* as a meeting conducted in accordance with RRO, then RRO are constitutive rules for parliamentary meetings.<sup>41</sup>

---

<sup>40</sup> Whether playing chess to win is a constitutive requirement or regulative ideal will not be relevant to anything that follows. Concepts that invoke constitutive rules can accommodate deviations, intentional (cheating) or unintentional (mistakes), from the canonical type. Activities defined by constitutive rules can embed regulative rules. Rules against fouling in basketball are regulative rules; the constitutive rules of the game specify what to do when they are violated.

<sup>41</sup> Our account has the same analysandum as Searle's, but there are important differences. Searle says that constitutive rules canonically have the form 'X counts as Y in C' (John Searle, *Speech Acts: An Essay in the Philosophy of Language* (London:

With this as background, the key points in thinking about what is special about status functions are the following. The collective activities in which status functions figure are essentially intentional joint action types that are governed by constitutive rules. The constitutive rules that introduce status functions are ones that define activities involving items that play a role in them without, however, specifying which particular things are to fill those roles.

---

Cambridge University Press, 1969), §2.5; Searle, *The Construction of Social Reality*, *op. cit.*, pp. 43-8; Searle, *Making the Social World: The Structure of Human Civilization*, *op. cit.*, pp. 97-98). This is a mistake. Placing an opposing player's king in check while affording him no move that would not put his king in check *counts as* checkmating your opponent *in a context* in which two people are playing chess. This is not a rule the following of which helps to constitute the activity of playing chess. It defines an expression in relation to a particular moment in an activity governed by constitutive rules. The expression 'X counts as Y in C', though underwritten by constitutive rules in these cases, does not express a constitutive rule itself.

Constitutive rules are not norms either. They describe, for a collective activity, a sequence of action types by more than one agent. The description can be transformed into a set of rules if agents wish to instantiate the activity type. You can generate hypothetical practical norms from this as for any bit of means-end reasoning: if you want to play chess and you are the white player, then you must .... Equally, if you want to start your car, you must turn the key.

This presents a coordination problem for those who intend to engage in the activity. For if the rules that constitute the collective intentional action type mention roles for things without uniquely specifying the role fillers, those who want to perform that action type have to choose among various different ways of doing it, that is, they have to choose which items are to be used in the roles, and there will typically be, as in the case of units of monetary exchange and chess pieces, many things that can fill the relevant role. If the participants anticipate performing the activity defined by the rules repeatedly, then they may solve the coordination problem by agreeing (explicitly or tacitly) to use the same item or items of the same type on these repeated occasions. When they reach this agreement, they come to have generalized conditional we-intentions in favor of engaging in the collective intentional action type in question, when it comes up, by using a certain item or a type of item. Therefore, the imposition of a status function on an object or type of object for use in recurring circumstances involves an ECA in the relevant group.

Consider the concept of a pawn in chess. For something to be a pawn is for it to have a certain status function. Nothing can function as a pawn in the play of chess unless it is collectively accepted as having that role. The property of being a pawn is *functionally defined* by the rules of chess. To be a pawn is simply to be an item that starts from thus-and-such a position on the board, that is subject to thus-and-such movement restrictions in thus-and-such a joint intentional action, and so on. In short, to be a pawn is to be an item that is subject to pawn-rules in a certain collective intentional action type. This functional specification does not uniquely determine what is to play the pawn role. This explains why it is possible for people to invent endlessly novel chess

sets, and why two travelers can play chess without a standard set by using adventitiously whatever items are at hand, such as bottle tops, or annotated slips of paper.

What makes a particular item  $X$  a pawn in a game of chess? Since to be a pawn is to be subject to pawn-rules, the players must treat  $X$  as subject to those rules. And, since the rules in question do not uniquely determine which item is to be treated as subject to pawn-rules, the players must agree or settle on treating  $X$  as subject to pawn-rules. Being a pawn, then, is a function an item can perform only if the participants collectively accept, in this sense, that it is to perform that function.

Jointly intentionally coordinating on treating  $X$  as a pawn for the purpose of playing chess is an ECA provided that the parties intend to use  $X$  as a pawn in several games. For the members of the group (say two travelers using bottle caps for pieces while on a long train trip), there is an end  $E$ , to play a game of chess, and a collective action plan  $P$ , to use certain items in the roles of pieces in accordance with the rules, such that (i) each of them we-intends that they act in accordance with  $P$  in order to achieve end  $E$ , whenever they feel like playing chess on the train; and (ii) there is an alternative,  $P'$ , to  $P$  (differing in assignment of piece roles across items), by which they could achieve  $E$  in the same circumstances by acting together intentionally in accordance with  $P'$ .

To summarize, what makes status functions special among agentic functions is that items with them cannot perform their function unless they are collectively accepted as having it. This is explained by the fact that the relevant functions are defined by constitutive rules for a kind of social transaction which specify a kind of role for one or

more items in the transaction without specifying any particular role fillers. Then when a group of agents want to engage in that type of activity, they face a coordination problem. If they settle on items or types of items to play the relevant roles for repeated occasions on which they want to engage in the activity, they solve the coordination problem by way of an ECA, which we now see is a very central sort of convention.

That conventions, conceived of as ECAs, are required for assigning status functions explains why something like collective acceptance is necessary in order for things with status functions to perform their functions.<sup>42</sup> Cashing this out in terms of an

---

<sup>42</sup> A referee raised the question why ECAs are *necessary* for collective acceptance, especially after status functions are initially assigned in the context of an ECA. The basic answer is that status functions are defined in terms of roles of things in essentially intentional collective action types: intentional social transactions. Thus, the parties must coordinate on what things fill the roles, but they must do so jointly intentionally given that it is a role in an essentially intentional collective action type. That requires an ECA. A second worry may be that once a core group has established a certain convention with respect to what things will play certain roles in intentional social transactions, others may feel pressure to conform even if they would rather not. One might say that *they* do not accept the convention. But if they do participate in it, they we-intend the relevant things to fill the roles in conformity with the ECA, that is, they are party to the collective acceptance by which those things fill those roles, in the sense we have described—they just wish that another convention had been adopted. They are parties to the ECA under protest, as it were.

ECA also gives precise content to the idea of collective acceptance. Collective acceptance is neither *sui generis* nor captured by the idea that everyone in a group believes (or that it is common knowledge) that everyone in the group believes something has a certain function. For if they are not willing to act as if it had that function, it cannot perform it. It is rather captured by their sharing an intention to treat it in accordance with a functional role in a social transaction. It will function in that way if the participants share an intention to so treat it whether or not they all believe that they all believe it has the function.

In contrast, L-conventions cannot assign status functions to objects. Because status functions play a role in essentially intentional collective action types, their assignment to objects or types of objects requires that the parties to it share intentions with respect to its use in that role. L-conventions do not secure that the participants have appropriate generalized conditional we-intentions toward the use of the items in a type of joint intentional action. They say nothing about the concept of shared intention. Their satisfaction, though compatible with shared intention, does not require it. In this respect, the concept of an ECA is essentially richer than the game-theoretic notion that Lewis constructs. But although in this respect it is essentially richer, it is also possible for an ECA to be realized though not all of Lewis's conditions are, since (among other things) an ECA does not require, as Lewis's account does, common knowledge to secure reciprocity.

#### IV. THE DISTINCTION BETWEEN CONSTITUTIVE RULES AND CONVENTIONS

An upshot of this discussion is a sharp distinction between constitutive rules and conventions.<sup>43</sup> Constitutive rules have sometimes been treated as conventions. Davidson, for example, treated the rules of games as conventions.<sup>44</sup> Goldman's account of the conventional generation of actions presupposes that constitutive rules are conventions.<sup>45</sup> Marmor more recently has called them *constitutive conventions*.<sup>46</sup> This is a mistake, if we are right. This has been pointed out before. Searle draws the distinction between constitutive rules and their conventional realization as far back as *Speech Acts*,<sup>47</sup> and later grounds this in the observation that constitutive rules are not arbitrary

---

<sup>43</sup> Weirich distinguishes rule (Gilbert) and practice (Lewis) conceptions of convention (Paul Weirich, "Conventions and Social Institutions," *Southern Journal of Philosophy*, xxvii (1989)). However, he has in mind a rule in force in a community for the former. The thesis that rules in force in a community are conventions is not the thesis that constitutive rules are as such conventions.

<sup>44</sup> Donald Davidson, "Communication and Convention," in *Inquiries into Truth and Interpretation* (New York: Clarendon Press, 2001), at p. 265; Donald Davidson, *The Structure of Truth: The Locke Lectures 1970* (Oxford: Oxford University Press, 2020), at p. 23.

<sup>45</sup> Alvin I. Goldman, *A Theory of Human Action* (Englewood Cliffs: Prentice-Hall, 1970), at p. 25.

<sup>46</sup> Marmor, *Social Conventions*, *op. cit.*, ch. 2.

<sup>47</sup> Searle, *Speech Acts: An Essay in the Philosophy of Language*, *op. cit.*, pp. 38-40.



whereas conventions are.<sup>48</sup> Williamson has also argued against classifying constitutive rules as conventions on the grounds that conventions are arbitrary and contingent and constitutive rules are not.<sup>49</sup> What our account provides is a more precise account of that difference. Constitutive rules define activity types. When they define essentially intentional collective activity types that express patterns while leaving open aspects of the realization that require participants to coordinate on the same things, they set the stage for the introduction of conventions, in the sense we have identified.

Are constitutive rules just a different *type* of convention? No. Constitutive rules provide directions for action and are constitutive relative to an action type that requires the pattern that they define be instantiated intentionally. So understood, they bear none of the standard marks of conventions.

First, they are not in any relevant sense arbitrary. They are not arbitrary *relative to the action types they govern* (this is the guiding idea behind Williamson's argument<sup>50</sup> and also Searle's 1969 remarks in *Speech Acts*<sup>51</sup>). While it might be said that it is arbitrary which among a range of activities governed by constitutive rules we choose to engage in to satisfy some more general interest, this is true for most things we do whether governed by constitutive rules or not, even when the choice is influenced by past practice. For example, walking together in a certain way (staying within five feet of

---

<sup>48</sup> Searle, *The Construction of Social Reality*, *op. cit.*, p. 28.

<sup>49</sup> Timothy Williamson, "Knowing and Asserting," *Philosophical Review*, cv, 4 (1996).

<sup>50</sup> *Ibid.*

<sup>51</sup> Searle, *Speech Acts: An Essay in the Philosophy of Language*, *op. cit.*

each other) satisfies an interest that walking together slightly differently (staying within four and a half feet of each other) would satisfy as well. But this does not make walking together arbitrary in any relevant sense. An ECA in particular is arbitrary in the sense that following it is not only one of several ways to satisfy some interest (which is true of almost anything we do) but also *construed* by its parties as one of the several ways of satisfying some interest. This distinguishes following constitutive rules and acting in general from the way in which ECAs are arbitrary.<sup>52</sup>

---

<sup>52</sup> Marmor classifies constitutive rules like those of chess as conventional on the grounds that they are arbitrary (Marmor, *Social Conventions*, *op. cit.*, p. 42). His idea is that playing chess answers a more general human need that could be met by engaging in alternative practices governed by constitutive rules (chess-like games, checkers, go, mahjong) and that we play chess because of an established practice. It is not clear why this would make the *rules* conventional. But even so, is the practice of playing chess conventional for the reasons given here?

People play chess on a regular basis. It has rules participants must follow (by and large) in order to play it. It answers to a general need—the recreational value of low stakes competition involving the exercise of domain restricted skill. Similar activities could serve the same need. We engage in it in part because it is an established practice. We also take walks together on a regular basis. There are loose rules involved: walk at roughly the same speed and direction and in proximity to one another. This answers a more general need which could be met otherwise, by jogging, biking, skiing, paying chess, and so on. People who do so are influenced partly by there being

Second, they are not essentially social or reciprocal. Following a rule is not as such acting in accordance with a convention. If someone follows a recipe for making a raspberry pie, he is following a set of rules, but not a convention. If someone follows directions for getting to the courthouse, she is following rules but not a convention. These are rules for individual behavior. The social element is not essential to the concept of following a rule. Similarly, the idea of a constitutive rule does not *by itself* give us the idea of a convention because it has nothing necessarily to do with the social. Constitutive rules may govern individual action, as in playing solitaire, just as well as collective action. While there is a general practice of playing solitaire in our communities, this is contingent: anyone can invent a game designed to be played by one person and then play it herself without communicating it to others.<sup>53</sup> For the same reason constitutive rules do not essentially involve reciprocity.

---

a practice of doing so—we learn patterns from others. But the rules we follow in walking together are not conventions, nor is walking together participating in a convention.

Marmor emphasizes that some variations in rules of chess we are willing to classify under the same term and relates this to “the kind of values that the constitutive rules give rise to” (*ibid.*, p. 44). It might be objected the alternatives to walking just cited are not forms of walking. It is unclear why classification under the same term is important to the relevant kind of arbitrariness. In any case, as noted in the text, there can be a variety of modest variations in rules for walking together.

<sup>53</sup> García-Carpintero suggests that constitutive rules, though not conventions, might be thought to be conventionally realized in a community (Manuel García-Carpintero,

A rejoinder is that while constitutive rules governing individual action are not conventions, those governing collective action are. But this could not be simply in virtue of their being constitutive. Nor could it be the virtue of the action being collective. Collective intentional action is not by nature conventional. There is nothing of convention in our intentionally moving a bench together. Even if we thought of *jointly lifting a bench intentionally* as an action brought about by intentionally following rules for its realization, we would not be following a convention. Suppose we define *joint-bench-lifting* as the activity type of two people moving a bench together intentionally by one (Righty) picking up the right end and the other (Lefty) the left and then moving it together to another location and setting it down. Then the rules—Righty picks up the right end, and Lefty the left, and then they move to another location, each then lowering

---

“Conventions and Constitutive Norms,” *Journal of Social Ontology*, v, 1 (2019)). The example cited in the text shows that they can be followed without involving conventions. This leaves open that practices like playing chess might be involved in conventions of one or another sort. It may be a convention in the custom sense that we play chess for entertainment. Playing chess might also in some circumstance be a solution to a coordination problem. When at a diplomatic impasse our leaders play chess, the winner’s position prevailing. There are alternatives: dodgeball, ritual combat, flipping a coin, rock-paper-scissors, and so on. This shows only that activities governed by constitutive rules can figure in conventions, not that constitutive rules, or their being followed, are *by that very fact* conventions.

his end to the ground—are constitutive rules for joint-bench-lifting. Yet instantiating this is not realizing any convention.

Third, and finally, the idea of stability has no straightforward application to constitutive rules. There are constitutive rules relative to any type of activity pattern, individual or collective, conceived as realizable only by the rules being followed intentionally. There are constitutive rules that no one has ever thought of or ever will. They are abstracta. Their existence does not depend on being realized in any community. They exist independently of being followed or even conceived. Conventions do not. The many types of essentially intentional activity patterns, individual and collective, for which there are constitutive rules, are not conventions in any community. People following a rule may realize a convention (of one or another sort) in a community, like following the rule to call back if you initiated a call and the call is dropped. But the rule is not the convention, and a convention is not present whenever a rule is followed. What is stable is not a rule, but a practice or a disposition to engage in one.<sup>54</sup>

Thus, constitutive rules, unlike ECAs, are not arbitrary, social, reciprocal or stable. It is important to understanding how status functions are realized that constitutive rules and ECAs not be conflated. Using one word, ‘convention’, for both obscures their differences and how they are related to one another. This provides good reasons for not extending the word ‘convention’, even with a modifier, to constitutive

---

<sup>54</sup> A similar point is made by Indrek Reiland, “Constitutive Rules: Games, Language, and Assertion,” *Philosophy and Phenomenological Research*, c, 1 (2020): n. 20.

rules. Though we will not enter into it, this has important implications for understanding the distinction between meaning, on the one hand, and conventional meaning, on the other, if particular meanings are status functions. The rules that define meaning roles are not conventions, but words have the defined roles assigned conventionally.

## V. CONCLUSION

There is a kind of convention central to social life, central because it is involved, among other things, in the imposition of status functions on objects, which involves what we have called an *effective coordinating agreement* (ECA). An ECA is realized in a set of generalized conditional we-intentions directed toward a collective action plan that solves a coordination problem for a group anticipating its repeated occurrence. ECA is a subtype of convention that has largely been overlooked in the literature. It is in the same line of business as the kind of conventions Lewis describes in his account. While akin, it requires in some ways more and in some less than what Lewis requires. It requires that its parties share a generalized conditional intention, which involves a distribution of conditional we-intentions across the members of the group, directed at carrying out as needed a collective action plan that solves a coordination problem. Lewis's account in contrast is resolutely individualistic. Both ECAs and Lewisian conventions require a robust form of reciprocity, but to secure it Lewis has to introduce an implausibly strong requirement of common knowledge among members of the group with respect to their beliefs and preferences. In this respect, the sort of convention we identify is less demanding. Most importantly, ECAs are demonstrably central to our understanding of social reality since they are required for the imposition of status functions on their bearers. They are pervasive in the social world. They appear wherever there are status

functions. If to have a linguistic meaning, for example, is to have a status function, then linguistic meaning is conventional in a stronger sense than Lewis's. It is remarkable that this notion of convention has not, up to this point, been cleanly separated from other related notions. This is most likely due to the fact that only recently have all the tools needed to identify it clearly been assembled.

MARIJA JANKOVIC

Davidson College

KIRK LUDWIG

Indiana University, Bloomington