

LEWIS ON CONVENTION

Lewis' theory of convention¹ is in terms of, firstly, regularities in action rather than rules, and secondly, certain types of psychological attitude, namely, conditional preferences for action and expectations about the actions of others. Conventions according to Lewis are the solutions to co-ordination problems. Here are samples of some such problems. They are Lewis'.²

Meeting at A/B: We wish to meet. There are two possible places A and B. Neither of us cares much whether we meet at A or B, so long as we meet. If we both go to A the convention is to go to A, the conventional alternative to go to B, and we've solved the co-ordination problem. If one went to A and the other of us to B we would have failed to solve it.

Driving on the left/right: We wish to avoid auto collisions and we do so by all driving on the left, (or all on the right, it doesn't matter so long as we stick to the same side of the road).

Stag/rabbit: We all need to eat and we can do so by individually hunting rabbits - each can hunt and kill rabbits on his own - or by collectively hunting the stag. The fact that in this case the alternatives differ in that one, rabbit hunting, involves individual, the other collective, action, makes this example crucially different from our other paradigms of conventions, as we shall see.

Clothes case: Each of us prefers to wear a white tie providing most of the others do, and would prefer to wear a red tie if most of the others wore red ties. However in this example each prefers a minority not to wear a white tie when the majority do so that he can laugh at that minority. This latter small difference makes the clothes case a non-convention on Lewis' analysis of convention. This is a defect in that analysis.

Co-ordination problems are given by Lewis a technical definition in game-theoretic terms.³ However the intuitive idea involves:

- (1) the existence of a common interest
- (2) there being at least two mutually exclusive types of action, either one of which solves the co-ordination problem, e.g. driving on the left. The problem is to co-ordinate our actions so that one or other of the types of action is performed.
- (3) the choice between alternative being largely arbitrary
- (4) the performing of any conventional action being rational in the sense that the agent acts in accordance with his preferences in the light of the likely actions of other agents. His action of, say, driving on the left is rational because he prefers to drive on the left given that others drive on the left, and he expects them to drive on the left. Clearly he requires a basis

for this expectation and this Lewis provides in terms of common knowledge⁴ of past conformity to the convention. I won't in this paper be concerned with the notion of common knowledge or with the expectations component⁵ in Lewis' account of convention.

So much for the sketch of Lewis.

This paper will concern itself with the kind of interdependence of action involved in conventions, and hence with the conditional preference component of Lewis' account. It is assumed that conventions do involve interdependent actions.

Lewis' basic intuition is that conventions are in essence solutions to co-ordination problems. I think that while many conventions are solutions to co-ordination problems many are not. If I am right it would seem to follow that the structure of interdependence of actions underlying all conventions is not whatever structure is peculiar to those conventions which solve co-ordination problems.

I set myself three tasks:

- (1) To show that Lewis' account fails to adequately characterise the type of interdependency of action involved in conventions
- (2) Present an alternative account of the interdependence of action involved in conventions
- (3) On the basis of (1) and (2), as well as by counter-example, show that Lewis wrongly makes it definitive of a convention that it have a conventional alternative, whereas this is a feature only of some conventions, namely those that solve co-ordination problems.

The relevant clauses in Lewis' definitions for our purposes are:

- (1) everyone prefers that everyone conform to R, on condition that at least all but one conform to R;
- (2) everyone would prefer that everyone conform to R1, on condition that at least all but one conform to R1, where R1 is some possible regularity in the behaviour of members of the population P in the recurring situation S, such that no-one in any instance⁷ of S among members of P could conform both to R1 and to R;
- (3) almost everyone prefers that any one more conform to R, on condition that almost everyone conform to R.

Clause (1) characterises what in game-theoretic terms is a proper co-ordination equilibrium. For the sake of brevity I shall refer to this clause as the pce clause and the combinations it refers to as pce's. Clause (2) is the clause requiring that each convention have an alternative convention. Clause (3) is merely a reformulation of clause (1) and

is framed by Lewis to cope with exceptions to the rule. (Note that clause (1) occurs within clause (2))

The combination of all hunting rabbits in the stag/rabbit example is different from Lewis' other 'conventions' in that there is not the same sort of dependence. Given stag hunting was not possible, the combination in which each hunts rabbits would involve no conditional preferences at all - no agent's action would be dependent on others' actions. That is, the combination of all hunting rabbits would not involve interdependent actions and would not therefore be, or be part of, a co-ordination problem. But in other conventional combinations e.g. all hunting stag, all driving on the left, all meeting at A etc., there is dependence of each on every other, even if the conventional alternative insisted upon by Lewis is ruled out. For example, if in the meeting at A/B case the only alternative meeting place to A, namely B, was ruled out, that each preferred to go to A would still be dependent on the others going to A.

All hunting rabbits is in fact of the same logical type in terms of interdependence of action as those combinations of actions which are not alternative conventions but rather non-conventional alternatives to conventions. Thus if the three agents in the meeting at A/B were to give up trying to meet and instead stayed at home, the resulting combination would not be a convention though it would be of the same logical type in terms of the structure of the interdependence of action as the combination in which all hunt rabbits. For the all staying home and the all hunting rabbits combinations don't involve interdependence, at least of the same sort as the other regularities i.e. all hunting stag, all meeting at A, do. Thus if in the meeting at A/B case the possibilities of meeting are ruled out then each stays at home completely independently of what the others do, in the same way as each hunts rabbit independently of the others given that stag hunting is ruled out.

Like the all hunting rabbit situation, however, all staying home involves a certain type of interdependence when a conventional alternative is introduced as a possibility. In particular in the all meet at A/stay home and all hunt stag/rabbit cases, each prefers to stay at home (or in the latter case hunt rabbits) conditionally on the others not meeting, (or in the other case, not hunting stag). That is, the non-conventional alternatives to conventions (which include it seems the all hunting rabbits combination) are such that each prefers to perform some action of type x only on the condition that the others do not perform actions of type y, where to y is one, but not the only, way of not x-ing. Here the regularity of all y-ing would be a standard conventional regularity.

However, a necessary (though not as we shall see a sufficient) condition of a combination being a standard convention is that each prefers to x on condition the others x. Thus all prefer to meet at A on condition the others meet at A and all to hunt stag if the rest are hunting stag and so on. All x-ing on condition the others x is not however a feature of the non-conventional alternatives to conventions, or of the all hunting rabbit combination. It follows that the all hunting rabbit combination and the non-conventional alternative to conventions are logically alike, having the same structure of interdependence of action.

On some construals, but only on some construals, of the stag/rabbit case, all hunting rabbits and the non-conventional alternatives are, in terms of Lewis' clause (1), logically alike and different from the standard conventions. They are logically alike in that they violate, on these construals, his clause (1). Take the circumstance in which a majority do not hunt stag in the stag/rabbit case. We can assume that each doesn't care whether anyone else is hunting rabbits just so long as he is. If all are in fact hunting rabbit, all hunting rabbit is nevertheless not a pce. For if some were not hunting rabbit the rest would not prefer that they did - the rest, on our assumption, couldn't care less. The same would hold of the staying at home combination in the meeting at A case, if we made the assumption that no-one cares less what the others were doing so long as they weren't at place A or B.

Thus the all hunting rabbits combination is, on at least one construal, not a convention according to Lewis' clause (1). It is not a convention by our account on any of the construals we will consider. But on the above construal, both according to Lewis and ourselves, whereas all hunting rabbits is not a convention, hunting the stag is. But now the stag/rabbit paradigm of conventionhood involves only one conventional regularity. Since there is no conventional alternative there is no co-ordination problem. Thus conventions are not always solutions to co-ordination problems. The stag/rabbit case, on at least one construal of it, is a counter-example to Lewis' thesis that conventions are solutions to co-ordination problems; and a counter-example in terms not only of our account but of his own account of convention. All hunting rabbit violates clause (1) and clause (2). The violation of clause (2) entails the non-existence of the conventional alternative R1.

However our account of convention does not square with Lewis'. Descriptions in terms of pce's (Lewis' clause (1)) do not mirror descriptions in terms of our principle, 'x on condition others x'. An ad hoc assumption could turn the all hunting rabbit and all staying home combinations into pce's. The assumption that everyone prefers that anyone not hunting rabbit, given most are, actually hunt rabbit, would have this effect. This is not as it should be. Such ad hoc assumptions make no difference to the structure of the interdependency of action. Those assumptions make no difference to that structure because the preferences they assume make no difference to anyone's actual actions, or to what anyone's actions would have been if others had acted differently. The ad hoc assumptions consist of preferences which do not determine actions. If I prefer the minority hunting stag in vain to join myself and the others hunting rabbit, I do nothing other than indicate my capacity for sympathy. This preference makes no difference to my action nor to the actions of those hunting rabbit already. Nor does it make any difference to the actions of those hunting the stag in vain.

Such ad hoc preferences are not part of the underlying structure of the interdependence of action. Variations in these ad hoc assumptions cannot make the difference between conventionhood and non-conventionhood. But on Lewis' account they do. So much the worse for Lewis' account. We must then reject clause (1), Lewis' characterisation of the interdependence of action in conventions.

Can Lewis' reformulation of Clause (1), namely clause (3), avoid the objections here raised? No, once again on some construals combinations are conventions according to Clause (3), on other construals they are not, where the difference in the construals is purely a difference in preferences that do not determine actions. Thus, in the stag/rabbit case, if no-one gives a damn what anyone else does given that he himself is hunting rabbit and most are not hunting stag, then the combination in which all hunt rabbit is not, in terms of clause (3), a convention. Whereas if most do give a damn then that combination is a convention. We must reject clause (3) and for the same reasons we rejected clause (1).

What are the implications of our rejection of Lewis' account and our at least provisional acceptance of our own 'x if the others x' principle? Firstly that all hunting rabbits and the non-conventional alternatives to conventions are alike in not being conventions. Secondly, and this follows from the latter, that conventions do not have to have conventional alternatives. So Lewis' clause (2) must be rejected. An actual case of a convention without a conventional alternative is Lewis' own (Rousseau is the real owner) all hunting stag combination in the stag/rabbit case. But if conventions don't have to have conventional alternatives they can't simply be solutions to co-ordination problems.

It might be argued that all my instances of non-conventional alternatives to conventions are really conventional alternatives. If this is so, it doesn't help Lewis' analysis. Our criticism of his theory was based on the fact that only on some construals of combinations such as the all hunting rabbit and all staying combinations, would those combinations be conventions according to Lewis; only on some construals, where the difference between construals that made the combinations in question conventions and the construals that didn't was not a difference in the structure of the interdependence of action. This difference in construals was rather a difference only in preferences for actions, preferences which made no difference to anyone's actions.

But Lewis aside, are the combinations I have called non-conventional alternatives to conventions nevertheless conventions? We have distinguished between conventions and non-conventional alternatives to conventions in terms of our principles, 'x if others x' and 'x only if others do not y'. I leave it to others to show that this is not the way to characterise the difference between conventions and non-conventional alternatives to conventions, and thus to show that cases like the all hunting rabbits and the all staying at home are really conventional regularities after all.

In rejecting Lewis' clauses (1), (2) and (3) we have completed two of our three tasks. We have, that is, shown that, contra Lewis, conventions don't require conventional alternatives and thus are not necessarily solutions to co-ordination problems; and that Lewis fails to provide an adequate characterisation of the structure of the interdependency of action involved in conventions. Our other task was to provide such an adequate characterisation. The principle 'each x-s on condition the others x' goes some way in this direction but not far enough.

Our principle needs to be altered to 'each x-s only on condition the others x'. This new principle ensures that in conventions the conformity of

others is not simply a sufficient condition for the performance of my conventional part, but also a necessary condition. This cannot be correct as it stands since conventional actions must be such that in general they could have been performed for reasons having nothing to do with the convention and its purposes, i.e. could have been performed outside the convention. However this only requires some sort of relativisation. My suggestion, which I have developed elsewhere and cannot pursue here, is that this relativisation be relativisation to the realisation of some collective end. (Here 'collective end' is a technical term denoting, very roughly, an end individuals appropriate qua members of a group.)

The necessity for the insertion of 'only if' in our principle is shown by the following example. In some community it might be that black ties just happen to be very expensive and that a majority would wear black ties conditionally on a majority wearing other ties which were very cheap, for in so doing the black tie wearers would be flaunting their wealth. Here each prefers to wear a black tie on condition the others do; it's just that each would prefer more strongly that the others wear non-black ties given that he was wearing his black tie. Thus the others wearing black ties is a sufficient condition for me so doing - I don't want to look like a cheapskate - and thus all wearing black ties is a convention by our weaker 'x if others x' principle. However it is not a necessary condition since I would wear my black tie if they wore their cheap ties. (I would not wear my black tie if no-one wore ties.) All wearing black ties is not a convention by our new 'x only if the others x' principle. Clearly all wearing black ties has a strong element of conflict absent in our paradigms of convention. I prefer myself to wear black and you to wear a cheap tie. You prefer the reverse. Nor are these conflicting preferences not determinative of action. If you wore the cheap tie I would wear the black one. This conflict is such as to make us (and certainly Lewis) want to rule such cases out as cases of conventions.

We accept then our new principle - 'each x-s only on condition the others x'. Hunting stag and meeting at A are conventions, according to this principle not simply because each hunts stag or goes to place A on condition the others do, but also because no-one would hunt stag or go to A unless the others did, since there would be no point in so doing.

Our new principle is subject to one final modification in the light of Lewis' clothes case. In Lewis' clothes case each wears a black tie only on condition most of the others do, and would wear a red one only on condition that most of the others did. However each prefers that a minority don't conform to the convention so that they can laugh at that minority. I take such cases to be paradigms of conventions, notwithstanding the small element of conflict. Lewis claims they are not. Certainly they violate his clause (1) and (3), as we shall see. Before doing so let us reformulate our principle to cope with such cases. It now reads 'each (or perhaps most) x-s only if most of the others x'.

This completes our characterisation of the structure of the interdependency of action involved in conventions and thus completes our only remaining task. (There is in fact more to be said about that structure - 'completes as far as this paper is concerned'.)

To return to Lewis and the clothes case. Once again we see that the preferences each has with respect to a minority are not action determining preferences - it makes no difference to my actions, the actions of the minority I wish to laugh at or the actions of the others, whether I have this preference for the minority not to conform. Once again however each having a preference for some minority not to conform makes the combination in which all conform a non-convention by Lewis' account. For if all but one conformed where that one was a despised minority, it would not be the case, as Clause (1) requires, that all would prefer that one to conform.

perhaps the clothes case can be accommodated by Lewis' reformulation, clause (3), designed as it is to cope with exceptions to the rule. It doesn't cope since, in the clothes case, almost everyone wants a minority not to conform, and thus given almost all conform, it is not the case, as clause (3) requires, that almost everyone wants those not conforming i.e. the despised minority, to conform. Most don't want them to conform. Perhaps we could tamper with Lewis' formulation so that it wouldn't rule out such cases. After all, even though most want there to be exceptions, it is only a minority that they wish to be an exception. A major problem with such revising is that that minority to be made an exception of had better not be any old minority. In particular, it had better not be the case that each wants to make an exception of himself. In the clothes case, of course, each wants to make an exception of someone other than himself, but there are cases where this would not be so; e.g. perhaps the social contract. The reason cases where each wants to make an exception of himself are unacceptable is because they are cases of conflict. They endanger Lewis' whole framework of co-ordination problems in the context of common interest. (In fact, it seems many conventions are solutions to conflict problems and our principle under suitable revisions can accommodate this fact. That is another story.) Lewis then will need to find a substantial difference between cases where people want to make an exception of themselves and those where they want to make an exception of other minorities. What is a relevant difference that Lewis can fix on? There is a relevant difference that we can fix on. This difference makes it the case that combinations in which each wants to make an exception of himself are not conventions. (At this stage for us such combinations ought not to be conventions, though other considerations change this, but also change our principle.) The difference we can fix on is that if I want to make an exception of myself this could make a difference to actions, namely my own action in the circumstance that everyone else conforms to the convention. In that circumstance I do not conform. My non-conformity would violate our principle 'each x-s only on condition the others x'.

Thus cases where each wants to make an exception of himself are not conventions on our principle, though cases where they want some minority other than themselves not to conform are conventions on our principle. Thus although Lewis' clauses rule out cases where people make an exception of themselves, as they should do, his clauses also rule out cases where people want to except a minority other than themselves - something they oughtn't to. Yet he can find no way, in his own terms, of distinguishing the types of case so as to reformulate his clauses. We can distinguish the types of case in terms of preferences which do and preferences which don't determine actions. This distinction is made room for in our principle

which is why it allows in the clothes case and rules out cases where each wants to make an exception of himself.

We thus return to that major reason we gave earlier for rejecting Lewis' characterisation of the interdependency of action involved in convention and preferring our own.

S.R. Miller

Rhodes University

NOTES:

1. David Lewis "Convention" Harvard University Press, Cambridge Mass., 1969.
2. "Convention" pp 5-7
3. "Convention" pp 14, 24
4. See "Convention" Chapter II
5. See "Convention" pp 107-118
6. "Convention" p 76
7. "Convention" p 78

INDIVIDUATING AGENTS

Introduction - The Problem

Normally we individuate people (agents) by their bodies. We count the number of agents in the room by counting the number of bodies. One person is talking in this room if sounds of a particular type issue in a particular way from one physical object in the room; if sounds of this type emanate thus from two physical objects in the room, then two people are talking. Our assumption is : one body, one agent.

Split brain experiments provide a prima facie case against this assumption. In the late 1940's an operation was developed which cured certain severe cases of epilepsy; it involved severing the higher connections between the two hemispheres of the brain. For some ten years no side-effects were noticed; people who had had the operation seemed to behave completely normally. Later experimenters, however, discovered some startling facts about those who had undergone the operation. Myers and Sperry -- the principal experimenters in the field -- obtained their results by developing techniques for dealing with the hemispheres separately. They presented certain information to people with split brains in such a way that each hemisphere received different information; by, for example, putting an unseen object into one hand only, or by letting only one nostril smell a particular scent. They also set these patients certain tasks which were to be performed by parts of the body controlled by only one hemisphere.

Myers and Sperry discovered that information given to the hemisphere usually in control of speech (the left) -- either by flashing a word on a screen so that only that hemisphere is stimulated, or by putting an object in the hand (usually, the right) that hemisphere controls -- can be reported verbally. If the information is given to the 'non-speaking' (right) hemisphere only, it cannot be reported verbally; but the hand controlled by that hemisphere will point to a relevant object even though the patient denies having been given any information. If two different words, "pencil" and "toothbrush" are flashed to the two hemispheres and the subject is told to get the corresponding object from beneath a screen with both hands, the hands will search independently:

"... the right hand picking up the pencil and discarding it, while the left hand searches for it, and the left hand similarly rejecting the toothbrush which the right hand lights upon with satisfaction."

These, and other surprising experimental results have led to the conjecture that splitting the brain leads to there being two agents in one body. Whether or not this is the best interpretation of the results is a question I am not concerned with here. However, that this supposition has been made at all does show that it is logically possible that there be two agents in one body -- whether or not this is true in the split brain cases. The concept of an agent does not demand that for a single body there be but a single agent; if it did, the two agent conjecture would be incoherent.