

Normativity in Lewis' and Bicchieri's Accounts of Conventions and Norms



MARTINA VALKOVIĆ

Abstract

Lewis [3] argues that, generally, we ought to conform to conventions because that answers (1) our own preferences, and (2) the preferences of others. While (1) is based on instrumental rationality, (2) is based on a moral principle or norm: other things being equal, we should do what answers others' preferences. Bicchieri [1] claims there is a third kind of normativity, neither rational nor moral, that applies to social norms. I argue that conventions draw their normativity from instrumental rationality and other independent moral principles or norms, and that it is unclear what further normativity could there be.

Keywords: *convention, social norm, normativity, coordination, game theory*

Introduction

In her influential book *The Grammar of Society*, Cristina Bicchieri puts forward a theory of social norms according to which social norms give rise to normative expectations, which distinguishes them from the so-called descriptive norms and conventions, which generate only empirical expectations. Bicchieri's account of conventions differs from the one proposed by David Lewis, who claimed conventions are a species of norms, drawing their normativity from the normativity of instrumental rationality and normativity of "external" norms and moral principles. Both Lewis and Bicchieri claim that conventions arise as solutions to coordination problems, but Bicchieri also argues that a social norm never arises as a solution to a coordination problem. Instead, a social norm transforms an existing mixed-motive game to a coordination problem by changing the agents' payoffs. These payoffs can be changed by introducing negative and positive sanctions for players, but also by agents recognising "others' normative expectations as well-founded" [1, p.23]. This, according

Kriterion – Journal of Philosophy, 2018, 32(2): 47–60.

<http://www.kriterion-journal-of-philosophy.org>

© 2018 The author

to Bicchieri, gives them a reason to conform to the social norm, even if it goes against their self-interest, and even if they are not privately committed to a norm. In other words, there exists what Francesco Guala [2] termed “intrinsic normativity” of norms, a normativity which is neither based on rationality nor moral (or other-regarding) considerations.

The aim of this paper is to compare Bicchieri’s and Lewis’ conceptions of conventions and social norms and the expectations they generate, together with the concepts of normativity they employ. One of the reasons to investigate this is the wish for conceptual clarity – to see how similar Lewis’ and Bicchieri’s terms are, and in which exact aspects they differ. However, what makes the topic especially interesting is that it may shed some light on the origin of normative expectations, or provide an insight into how they come to exist, and how we should understand the normativity in question. I will argue with Lewis that conventions are a type of norms that draw their normativity from instrumental rationality and other independent other-regarding or moral¹ principles or norms, such as not to harm other people, and that it is not clear what further normativity (i.e. that which is not rational nor moral) there could be. Among other fields, the topic of this paper is of interest for game theory and evolutionary game theory as both of these approaches operate with concepts of conventions, norms and normativity.

An important preliminary point is that I will limit my discussion only to Lewis’ and Bicchieri’s theories. Needless to say, there are many other authors who have written on social norms and conventions from various perspectives. Of those close to Lewis and Bicchieri, in that they embrace game-theoretic and individualistic approach, one could point out Edna Ullmann-Margalit, Robert Sugden, Russell Hardin, Peyton Young, Peter Vanderschraaf and others. Then, there are those accounts that start to diverge from this approach, such as Ruth Garrett Millikan’s theory of convention as sustained due to the weight of a precedent, which requires no rationality. Another example is Seumas Miller, who writes about collective ends, but remains close to Lewis in that his account is reducible to individualism. Finally, some authors completely abandon the individualist picture in favour of an irreducibly collective one. Probably the most influential account here is that of Margaret Gilbert, who introduces the notion of plural subjects. This broader discussion on social norms and conventions, however, will not be addressed here as it is considerably beyond the scope of this paper.

I will start by presenting Bicchieri’s account of social norms and conventions, followed by presenting Lewis’ account. Then I will compare the

two, pointing out their similarities and differences, especially regarding the normativity of conventions. Finally, I will argue in favour of Lewis' theory of normativity of conventions, and against Bicchieri's third kind of normativity.

Bicchieri on norms and conventions

Bicchieri [1] puts forward a theory of social norms and conventions in which they are explained in terms of expectations and conditional preferences of its followers. These social constructs exist because a sufficient number of people believe they exist and expect that a sufficient number of people will follow them, and these expectations give them conditional preferences for obeying the norms and conventions in question. These preferences and expectations then result in conforming behaviour, which gives the followers a further proof of the existence of the norm [1, p.2].

What distinguishes social norms from the so-called descriptive norms and conventions, Bicchieri claims, is that the latter never go against self-interest of the agents involved, while the former often do [1, p.2]. Descriptive norms and conventions, such as dress codes and signalling systems, can be modelled as solutions to coordination problems, i.e. situations in which there are multiple possible equilibria, and the agents prefer to coordinate with others on an equilibrium, whatever that might be. Thus, even though we might have a preference for one particular equilibrium, that preference is trumped by our preference for coordination. A typical example of a coordination problem is deciding on which side of the road we should drive. We do not especially care whether we drive on the left or the right, but we do care that we all drive on the same side, because otherwise we are likely to crash. Our problem is solved when a convention has emerged to drive on, say, the right side of the road. That could happen if, say, a bit more than half of the population followed the strategy of driving on the right. This would make them more successful than those driving on the left, which would, in turn, cause more people to choose the strategy of driving on the right, until everyone started doing so. In that case, it would be in our interest to drive on the right. But, if initially there were more people choosing to drive on the left, the opposite would happen – the convention of driving on the left would have evolved, and it would have been equally in our interest to drive on that side.

Bicchieri argues that, in contrast, social norms, such as cooperation, reciprocity or fairness, are *never* a solution to an original coordination

game, but to mixed-motive games, where there exist both some shared interest and a conflict of interests [1, p.3]. Social norms transfer these mixed-motive games into coordination games, by changing the agents' payoffs. As an example of a social norm we can take refraining from taking other people's belongings. The shared interest element here is that, presumably, we all prefer to have an exclusive right to decide what happens to our belongings, including who uses them. The element of conflict is that we would like to use other people's belongings when we please, without having to ask for the permission. Thus, we might be tempted to simply take other people's possessions, although that might easily result in others doing the same to us. Here the social norm comes into play. If we know that the norm of respecting other people's possessions is established, we know that other people expect us to conform to it. According to Bicchieri, for some individuals recognising that these expectations are legitimate is enough to give them a reason to obey the norm. Other people might have other reasons to obey, such the expectation of punishment if they disobey, or a desire to please others.

In more formal terms, Bicchieri lists the following as conditions for a social norm to exist [1, p.11]:

Let R be a *behavioral rule* for situations of type S , where S can be represented as a mixed-motive game. We say that R is a social norm in a population P if there exists a sufficiently large subset $P_{cf} \subseteq P$ such that, for each individual $i \in P_{cf}$:

1. *Contingency*: i knows that a rule R exists and applies to situations of type S ;
2. *Conditional preference*: i prefers to conform to R in situations of type S on the condition that:
 - (a) *Empirical expectations*: i believes that a sufficiently large subset of P conforms to R in situations of type S ; and either
 - (b) *Normative expectations*: i believes that a sufficiently large subset of P expects i to conform to R in situations of type S ;
 - or
 - (b') *Normative expectations with sanctions*: i believes that a sufficiently large subset of P expects i to conform to R in situations of type S , prefers i to conform, and may sanction behavior.

[1, p.11]

These four conditions – contingency, conditional preference, empirical expectations, normative expectations – are individually necessary and jointly sufficient conditions for a social norm to *exist* in a population [1, p.7]. For a social norm to be followed by a population, however, the following must also be true [1, p.11]:

A social norm R is *followed* by population P if there exists a sufficiently large subset $P_f \subseteq P_{cf}$ such that, for each individual $i \in P_f$, conditions 2(a) and either 2(b) or 2(b') are met for i and, as a result, i prefers to conform to R in situations of type S .

When it comes to descriptive norms, agents have no normative expectations, and the conditions for the norm's existence are as follows [1, p.31-32]:

Let R be a *behavioral rule* for situations of type S , where S is a coordination game. We say that R is a descriptive norm in a population P if there exists a sufficiently large subset $P_{cf} \subseteq P$ such that, for each individual $i \in P_{cf}$,

1. *Contingency*: i knows that a rule R exists and applies to situations of type S ;
2. *Conditional preference*: i prefers to conform to R in situations of type S on the condition that:
 - (a) *Empirical expectations*: i believes that a sufficiently large subset of P conforms to R in situations of type S .

As with social norms, the condition for a descriptive norm to be followed in a population is [1, p.32]:

A descriptive norm is followed by population P if there exists a sufficiently large subset $P_f \subseteq P_{cf}$ such that, for all $i \in P_f$, Condition 2(a) is met for i and as a result i prefers to conform to R in situations of type S .

Conditions for existence of conventions, which are for Bicchieri only those descriptive norms which have “endured the test of time”, are as follows [1, p.58]:

A descriptive norm is a convention if there exists a sufficiently large subset $P_f \subseteq P$ such that, for each individual $i \in P_f$, the following conditions hold:

1. *Empirical expectations*: i believes that a sufficiently large subset of P conforms to R in situations of type S and
2. S is a coordination game without nonstrict Nash equilibria.

Since empirical expectations condition always has to be met in order for a descriptive norm to be followed, this means that a convention is always a followed descriptive norm [1, p.58].

Lewis' account of conventions and their normativity

Similar to Bicchieri, Lewis [3] takes conventions to be solutions to recurrent coordination problems, situations in which two or more agents' interests predominantly coincide and, as the outcomes of every agent depend not only on his but also on other agents' actions, they have to coordinate to achieve the desired outcome [3, p.8]. Combinations of agents' actions in which "no one would have been better off had *any one* agent alone acted otherwise" are called *coordination equilibria* [3, p.14]. A *proper* equilibrium, on the other hand, is that which each agent likes "*better than* any other combination he could have reached, given the others' choices" [3, p.22]. According to Lewis, a coordination problem must have at least two proper coordination equilibria, otherwise it would be trivial – with only one proper coordination equilibrium and without a meaningful conflict of interest, the equilibrium will be reached if all the agents understand the structure of the situation [3, p.16]. This is not the case with coordination problems.

Since in coordination problems each agent depends on the other agent's action for achieving her desired outcome, she has to choose her course of action according to what she expects others will do. These expectations are formed by trying to replicate the other person's practical reasoning, using knowledge of what that person believes about the matters of fact, her preferences, and assuming some level of practical rationality [3, p.27]. To take one of Lewis' examples, assume that both of us want to meet each other and that it does not matter to us where exactly we will meet. So, we will want to go where the other person will go.

To figure out where this place is, we will try to replicate the other person's reasoning, using our knowledge about the places she knows, thinks are suitable for meeting people, likes to go to, are easy to access, etc. Since we know that the other person is trying to replicate our reasoning, we will also try to replicate that person's reasoning about our reasoning, and her attempt to replicate our attempt of replicating her reasoning, and so on. In other words, we want to know where she expects us to go, and where she expects us to expect her to go, and similar with other higher-order expectations.

Convention arises when we start to follow a certain precedent which at one time lead to successful coordination. We notice this past regularity of behaviour and we extrapolate to future cases, expecting that we will continue with this behaviour so that we would coordinate. This expectation gives us a reason to follow the precedent and, in this way, we create "a metastable self-perpetuating system of preferences, expectations and actions capable of persisting indefinitely" [3, p. 42]. Thus, Lewis' final definition of a convention goes as follows [3, p.78]:

A regularity R in the behaviour 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 that, and it is common knowledge in P that, in almost any instance of S among members of P ,

- (1) almost everyone conforms to R ;
- (2) almost everyone expects everyone else to conform to R ;
- (3) almost everyone has approximately the same preferences regarding all possible combinations of actions;
- (4) almost everyone prefers that any one more conform to R , on condition that almost everyone conform to R ;
- (5) almost everyone would prefer that any one more conform to R' , on condition that almost everyone conform to R' .

where R' is some possible regularity in the behavior of members of P in S , such that almost no one in almost any instance of S among members of P could conform both to R' and R .

Although Lewis asserts that the term "convention" itself is not a normative term under his analysis, he argues that conventions are a species of norms, i.e. regulations to which we believe one should conform [3, p.97]. This is so because of the probable consequences of the fact that

a certain action would conform to a convention, which are, according to our shared opinion, presumptive reasons why that action ought to be performed. Namely, in most cases in which people are deciding whether they should conform to a convention, the following conditions must hold [3, p.97]:

- (1) Most other members of *P* involved with me in situation *S* will conform to *R*.
- (2) I prefer that, if most other members of *P* involved with me in *S* will conform, then I conform also.
- (3) Most other members of *P* involved with me in *S* expect, with reason, that I will conform.
- (4) Most other members of *P* involved with me in *S* prefer that, if most of them conform, I conform also.
- (5) I have reason to believe that (1)–(4) hold.

Furthermore, if conditions (1)–(5) hold, then also these two conditions must hold [3, p.98]:

- (6) I have reason to believe that my conforming would answer to my own preferences.
- (7) I have reason to believe that my conforming would answer to the preferences of most other members of *P* involved with me in *S*; and that they have reason to expect me to conform.

And, Lewis claims, these two conditions, if true, are presumptive reasons why we ought to conform to the convention. The reason is that we assume, other things being equal, that we should do what answers our own preferences, and what answers other people's preferences, especially when they have reasonable expectations that we will do so. Although the conditions (6) and (7) are presumptive reasons why we should conform to a convention, there might always be other reasons why the conforming action should or should not be done. Lewis' claim is not that these are particularly forceful reasons why one should conform. Rather, he claims that the importance of these reasons is that they are general – they hold for any convention whatsoever.

Lewis further argues that conventions are socially enforced norms, that is, that we are expected to conform and that our failure to do so will likely evoke negative responses from others [3, p.99]. By the definition of convention, if *R* is a convention in population *P* in situation *S*, and if I am a member of *P* in *S*, then it is probably true that [3, p.99]:

- (8) Most other members of P involved with me in S expect me to conform.
- (9) Most other members of P involved with me in S have reason to believe that conditions (1)–(5) hold.

Whenever (9) is true, so is the following:

- (11) Most members of P involved with me in S have reason to believe that my conforming would answer to my own preferences.
- (12) Most members of P involved with me in S have reason to believe that I have reason to believe both that my conforming would answer to their preferences and that they have reason to expect me to conform.

This means that, whenever others observe my failure to conform to a convention, I have gone against their expectations, and they are likely to be able to infer that I have knowingly acted contrary to my own and their preferences and reasonable expectations. They will tend to form poor opinions of me and respond unfavourably, with punishment or distrust, and this is likely to reflect negatively on how they treat me in the future. Our interest in avoiding all these negative consequences further reinforce our conditional preference for conforming to a convention [3, p.99-100].

Contrasting the two accounts

Both Lewis and Bicchieri construct their definitions of conventions around beliefs, expectations and conditional preferences of members of a certain population. For both of them, conventions are always followed, that is, it does not make sense to talk of a convention which exists in a population, but it is not the case that almost everyone follows it. While both think that conventions arise as solutions to coordination problems with multiple equilibria, Bicchieri claims that a social norm can never arise as an answer to an original coordination problem, but as a solution to a mixed-motive game. This would mark a difference from Lewis, who considers conventions to be a type of social norms. However, Bicchieri allows for conventions to become social norms, for example when breaking them causes negative externalities:

“When breaking a convention creates negative externalities, people prefer not just that the convention is in place, but

also that everyone follows it. Such violations are usually legally sanctioned, but, even more importantly, they are also informally sanctioned by society. A reckless driver is blamed as irresponsible: We think he should have observed traffic rules. When breaking a coordination mechanism produces negative externalities, we may expect conventions to become full social norms.” [1, p.39]

One might ask what exactly is the difference between preferring the convention to “be in place”, and preferring that everyone follows it. After all, from the definition of a coordination problem and the fact that a certain convention exists in a population, it follows that everyone prefers everyone else to follow the established convention. Leaving this aside, Bicchieri also allows conventions to become social norms even if they do not involve any externalities, but “become so well entrenched that people start attaching value to them”, citing examples such as the dining etiquette and foot-binding in China [1, p.40]. Then, the real distinction between conventions and social norms, on Bicchieri’s account, is that the latter are “accompanied by what are perceived as legitimate expectations of compliance” [1, p.42].

This brings us to the crucial difference between Lewis’ and Bicchieri’s accounts. Bicchieri claims that in the case of conventions we have only empirical expectations, i.e. expectations about what other people do in this sort of situations. She argues that, for social norms, we need a further type of expectations, normative expectations with or without sanctions. Normative expectations tell us that other people expect us to conform, and normative expectations with sanctions tell us that they expect us to conform, prefer us to conform, and may sanction behaviour. In Lewis’ account, on the other hand, the fact that everyone conforms to the convention (his first condition for the existence of conventions; empirical expectations in Bicchieri’s terminology) gives rise to the expectation that everyone will expect everyone to conform to the convention (Lewis’ second condition; Bicchieri’s normative expectations). This follows from the fact that the agents find themselves in a coordination problem, where their interests are tied in such a way that they all lose if they choose different actions. Since a convention already exists, the safe way to achieve coordination is to act in accordance with it, rather than deviate. It seems that this fact is enough to account for expectances of conformity, without bringing any additional normative expectations into the picture.

In turn, this leads us to our main question, which is the normativity

of conventions. As Lewis points out, we ought to conform to established conventions because that would (1) answer our own preferences, and also (2) answer the preferences of others. While (1) is based on instrumental rationality, (2) is based on a certain moral principle or an other-regarding norm: other things being equal, we should do what answers others' preferences, especially when they have a reason to expect us so, this reason stemming from the fact that they know that the convention exists. Bicchieri accepts that some people could conform to a norm only because of the negative or positive sanctions [1, p.23], or because they are personally committed to the norm in question, meaning that they "attribute some value to the norm in question" and obey "out of a deep conviction of the norm's merits", as is often the case with rules such as reciprocity [1, p.43]. In the first case, we have compliance for the reasons of rationality, and in the second case, compliance for the reasons of morality or "personal norms", when we would feel guilty, ashamed or uncomfortable with ourselves for disobeying. However, Bicchieri also claims that there is a third kind of reason for obeying the norm, which cannot be reduced to either rational or moral reasons, and that it arises when "individuals recognize the legitimacy of others' expectations and feel an obligation to fulfill them" [1, p.23] and thus have a reason to "fulfill others' normative expectations" [1, p.29]. In what follows, I will argue that we cannot intuitively make sense of this idea unless we interpret "reason to fulfill others' normative expectations" either in terms of rationality or morality. Then, I will show how Lewis can account for the normativity of conventions and social norms without appealing to any third kind of normativity.

It is simply unclear what a reason to fulfil others' normative expectations would mean, unless we understand it as a rational or moral reason. If you expect me to conform to a certain convention or a social norm, e.g. to divide a pie we are sharing into two equal pieces, one reason to conform is to avoid you punishing me. This punishment could consist in you throwing your significantly smaller piece of pie in my face, or avoiding future interactions with me, or just raising your eyebrow a little. Or, I might do it to get on your good side, because I want you to like me, to think of me as a fair person or say nice things about me to others. In all of these cases, the reason to conform is based in rationality – it is in our self-interest to avoid punishment and, if we value the opinion and affection of others, to seek their approval. Another reason I could have for conforming is that I do not want to go against your expectations and preferences, because that might be harmful or unpleasant to

you. If I give you a smaller piece of pie, you might end up hungry or sad because you interpret my behaviour as an act of disrespect. I might want to avoid this not because of the potential negative consequences for myself, but because of a certain level of care or respect for you. In this case, the reason to conform would be based in morality or a kind of other-regarding considerations. It seems that these two kinds of reasons exhaust the potential meanings of “a reason to fulfill others’ normative expectations”, and that we cannot make sense of the idea as Bicchieri would like us to.

Finally, let us show how Lewis does not need any elusive third kind of normativity to account for the normativity of conventions. As we have seen, since others know that a convention is in place and that I know that, too, they have legitimate *empirical* expectations of my conformity to the convention, but also *normative*, rational and other-regarding, expectations, since they know that conformity answers both mine and their preferences. As failure to conform to an established convention is likely to result in miscoordination and negative consequences for the agents involved, a possibility of sanctions also arises. It is easy to see that on our traffic convention example. In a society in which driving on the right side of the road is the established convention, people know this, know that others know it, and know that they all prefer to drive on the same side of the road. For this reason, they will all expect others to stick to the established convention, since this will result in successful coordination. If, under these conditions, I suddenly decide to switch to driving on the left side of the road, it is likely that I will cause a collision or at least some disturbance both to myself and to others. As we do not commonly think that hurting ourselves or putting ourselves in danger is in our self-interest, this behaviour is likely to be interpreted as irrational. Also, as we know and expect others to know that suddenly changing the side on which we are driving puts other people in danger, my behaviour is equally likely to be interpreted as immoral or inconsiderate. Both of these interpretations are likely to cause negative responses, ranging from people forming negative opinions of me, to punishing my behaviour in some more palpable way. My behaviour could result in loss of reputation and various unpleasant future consequences. We ought to obey established conventions in order to avoid all of this. Thus, we can explain the normativity of conventions without appealing to anything else but the normativity of rationality and other-regarding norms.

It is important to note, as Lewis does [3, p.98], that these other-regarding considerations, just as rational considerations, give us only a

presumptive reason to conform. In other words, there are many reasons that enter into consideration when we are deciding whether to conform to a convention or not, and they can cancel each other out. So, if I live in a very corrupt society where bribery is the norm, it is likely that, when I meet a public official, she will expect me and prefer me to bribe her. Although her expectations and preferences give me *a* reason to consider bribing her, they do not give me a decisive reason. This reason can easily be overpowered by, say, my deep moral conviction that bribery is wrong.

Conclusion

Having compared Bicchieri's and Lewis' accounts of conventions and social norms, we saw that the main difference between the two accounts lies in the conception of normative expectations, expectations that other people have about our conforming to a certain regularity. In Bicchieri's theory, normative expectations mark the distinction between conventions and social norms. This neat distinction fails if, as we have seen in Lewis, we understand normative expectations as arising from empirical ones and the fact that the agents are in a coordination problem situation. We have also looked into concepts of normativity, and seen that the idea of a third kind of normativity, which is neither rational nor moral, is quite an elusive and problematic concept, and the question is why we should assume that it exists. It is far from clear, intuitively, what this third kind of normativity would be, and it does not seem we need it to explain the normativity of conventions and social norms, since it appears that their normativity can be based on only the normativity of rationality and morality.

Acknowledgments

I would like to thank Bart Geurts, Hannah Rubin, and an anonymous reviewer of this journal for their helpful comments on earlier versions of this paper.

Notes

- ¹ Throughout the text, I will use “other-regarding” and “moral” interchangeably. This is not to say that moral principles must always be understood as prohibitions to harm others or prescriptions to help them. However, since this paper is about normativity of social norms and conventions, and *not* normativity and morality in general, presumably it is other-regarding principles that matter in this context.

Martina Valković
Radboud University Nijmegen
Faculty of Philosophy
Nijmegen, Netherlands
<mvalkovic@gmx.com>

References

- [1] Cristina Bicchieri. *The Grammar of Society: The nature and dynamics of social norms*. New York: Cambridge University Press, 2006.
- [2] Francesco Guala. “The Normativity of Lewis Conventions”. In: *Synthese*, 190 (2013), pp. 3107-3122.
- [3] David Lewis. *Convention: A philosophical study*. Oxford: Blackwell Publishers, 2002.