

COPYRIGHT NOTICE:

Andrei Marmor: Social Conventions

is published by Princeton University Press and copyrighted, © 2009, by Princeton University Press. All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the publisher, except for reading and browsing via the World Wide Web. Users are not permitted to mount this file on any network servers.

Follow links for Class Use and other Permissions. For more information send email to: permissions@press.princeton.edu

CHAPTER ONE

A First Look at the Nature of Conventions

I WANT to begin with an attempt to define what social conventions are. I will start with some intuitive ideas on what seems special about conventional norms, and try to define those features as precisely as possible. If this tack leads us to a single explanation of the point, or function, of conventions in our lives, so be it. But we should not assume in advance that a single explanation is available, and we should certainly not predetermine what it is.

First, conventional rules are, in a specific sense, *arbitrary*. Roughly, if a rule is a convention, we should be able to point to an alternative rule we could have followed to achieve basically the same purpose. Second, conventional rules normally lose their point if they are not *actually followed* in the relevant community. The reasons for following a rule that is conventional are tied to the fact that others (in the relevant population) follow it too. To give one familiar example, consider the convention of saying “Hello” when responding to a telephone call. Both features are manifest in this example. The purpose of the convention is to have a recognizable expression that indicates to the caller that someone has answered the phone. But of course, using the particular expression “Hello” is arbitrary; any other similar expression would serve just as well—that is, as long as the expression I use is one that others use too. If the point of the convention is to have an expression that can

be quickly recognized, then people have a reason to follow the norm, that is, use the expression, that others in the community follow as well. If, for some reason, people no longer use this expression, I no longer have a reason to use it either.

In fact, both of these intuitive features of conventional rules derive from a single, though complex, feature that I will call “conventionality,” defined as follows:

A rule, *R*, is conventional, if and only if all the following conditions obtain:

1. There is a group of people, a population, *P*, that normally follow *R* in circumstances *C*.
2. There is a reason, or a combination of reasons, call it *A*, for members of *P* to follow *R* in circumstances *C*.
3. There is at least one other potential rule, *S*, that if members of *P* had actually followed in circumstances *C*, then *A* would have been a sufficient reason for members of *P* to follow *S* instead of *R* in circumstances *C*, and at least partly because *S* is the rule generally followed instead of *R*. The rules *R* and *S* are such that it is impossible (or pointless) to comply with both of them concomitantly in circumstances *C*.

In this chapter I will explain and defend this definition, showing how it applies to the variety of conventions we are familiar with. I will then present David Lewis’s theory of social conventions, arguing that it successfully explains some cases, but that it fails in many others. Finally, I will consider Margaret Gilbert’s alternative account of conventions, arguing that it raises more questions than it answers. An alternative to Lewis’s theory or, more precisely, a supplement to it, will be presented in the next chapter.

DEFINITION

Let us take up the details of the definition.

1. There is a group of people, a population, *P*, that normally follow *R* in circumstances *C*.

This condition indicates that conventions are *social rules*:¹ A convention is a rule that is, by and large, followed by a population. Not all rules have to meet this condition. People can formulate rules of action and regard them as binding even if nobody is actually following those rules. Conventions, however, must be practiced, that is, actually followed, by a population in order to exist. Furthermore, I use the term “followed” advisedly. In many contexts people’s behavior can conform to a rule without the rule being followed, as such. A rule is followed when it is regarded as binding. Conventions, I assume, must be regarded as binding by the relevant population. To say that a certain behavior is conventional is to assume that, at least upon reflection, people would say that they behave in a certain way because the relevant conduct is required by the convention.

Margaret Gilbert raised some doubts about this, relying on the following counterexample: suppose that it has been a convention in a certain community that people should send thank-you notes after being invited to dinner parties. As it happens, conformity with this convention has dwindled and most people no longer abide by it. Either they tend to express their gratitude in some other way, or not at all. “Does *this* mean that there is no longer a convention [to send thank-you notes]?” Gilbert asks. “The present author has no such clear sense,” she replies.²

¹ A note on usage: I assume here that “rule” is the content of a linguistic form and thus that rules can be valid or correct irrespective of practice. Paradigmatic example of rules would be rules of conduct, such as “In circumstances C do X.” Later we will have to include more complex rules, including those that determine how to create or modify other rules. The word “norm” I will use to indicate a rule that is followed by a population or, at least, is regarded as binding by a population. Thus the terms “social rule” and “norm” will be used interchangeably.

² See Gilbert, *On Social Facts*, 347. Notably, at other places in the same chapter (e.g., 345), Gilbert is more explicit in claiming that conventions can exist without conformity of behavior. See also Millikan, “Language Conventions Made Simple,” 170, who seems to share that view. Millikan’s examples, however, are somewhat ambiguous. (“Few actually hand out cigars at the birth of a boy, nor does everyone wear green on St. Patrick’s Day, or decorate with red and green on Christmas” [170].) Either they are like Gilbert’s, namely, conventions that have, by and large, ceased to be practiced, or else, they concern

The lack of a clear sense about such cases is understandable. It often happens that a conventional rule gradually ceases to be practiced, and at some point it might become impossible to determine whether the convention still exists or not. In other words, the idea of a practice is rather vague, and borderline cases are not uncommon. As with most distinctions, however, the vagueness of the concepts constituting the distinction does not entail that the distinction itself is problematic. The importance of this first condition pertains to the unique structure of reasons for following rules that are conventional. It is a unique feature of such reasons that they are closely tied to the fact that others generally comply with the rule. This will become clearer as we complete the explanation. Suffice it to say at this point that Gilbert's example is not a counterexample to condition 1; what we face here is precisely what is described, namely, a convention that has dwindled. There used to be such a conventional practice, but now it is no longer clear that there is one. The essential point remains that conventional rules are social rules and that there must be a community that by and large follows the rule for it to be conventional.

Similar considerations apply to the question of what constitutes a community that practices a certain convention. Some conventions are practiced almost universally; others are much more parochial. However, a rule followed by just a few people is typically not a convention, even if the other conditions obtain. As social rules, conventions must be practiced by some significant number of people. Numbers matter here because in small numbers the relevant agents can create, modify, or abolish the rules at will, by simple agreements between them. As David Lewis was right to observe, conventions typically emerge as an alternative to agreements, precisely in those cases where agreements are difficult to obtain because of the large number of agents involved.

cases in which only a small part of the general population actually conforms to the convention. Naturally, it all depends on how one identifies the relevant population. Either the convention, e.g., to hand out cigars at the birth of a boy, is one that has dwindled and no longer practiced, and therefore, is no longer a convention, or else it is a convention that is practiced by a small subset of the general population, as the case may be.

(More on this below.) But again, the concept of a population or a community is vague, and borderline cases are bound to exist.

2. There is a reason, or a combination of reasons, call it A, for members of P to follow R in circumstances C.

The second condition requires three clarifications. First, what is a *reason* for following a rule? I take it that reasons for action are facts that count in favor (or against) doing (or not doing) something. Thus reasons for action are closely tied with the idea of value. A reason to ϕ typically derives from the fact that ϕ -ing is good, valuable, or serves some worthy end. Some philosophers claim that it is the other way around: to say of something that it is valuable is to say that it has certain properties that provide reasons for action.³ Either way, reasons are closely related to values or goodness. A reason to follow a rule necessarily assumes (or, is suggested by the assumption) that following the rule is valuable under the circumstances, that it serves some purpose or point, that it is good in some respect (not necessarily moral, of course).⁴

Second, it is not part of this condition of conventionality that members of P must be aware of the reason, A, to follow R. People may follow conventional rules for various misconceived reasons or, in fact, for no reason that is apparent to them at all. The conventionality of a rule does not depend on the subjective conception of the reasons for following the rule by those who follow it. As an example, consider this: there are some Orthodox Jewish communities who believe that Hebrew is a holy language descended directly from God. In fact, they only speak Hebrew in religious contexts, and use Yiddish for

³ Roughly, this is Scanlon's view, called the "buck passing" account of values. See Scanlon, *What We Owe to Each Other*, 96–97. The details of Scanlon's view are controversial; see, for example, Heuer, "Explaining Reasons."

⁴ This formulation assumes that there is a fact of the matter about reasons for action, as about some matters of value. (For an argument that the latter does not necessarily assume *realism* about values, see my *Positive Law and Objective Values*, chap. 6.) This objectivist assumption is not necessary for the rest of the argument in this book. Expressivists can follow the argument on their own terms.

everyday life. Surely Hebrew remains conventional (to the extent that it is, of course), even when spoken by those Orthodox Jews. The fact that they misconceive the reasons for following the rules does not prove the contrary. Reasons for following a convention do not have to be transparent.

Similar considerations apply to the question of whether people have to know that the rule they follow is a convention or not. Lewis suggested that at least in one crucial sense, the answer is yes.⁵ He claimed that, for the rule to be conventional, the arbitrary nature of the rule must be *common knowledge* in the relevant population. Tyler Burge has rightly argued that this is wrong and no such condition should form part of the definition. People can be simply mistaken about the conventional nature of the rules they follow. To mention one of Burge's examples, he asks us to imagine a small, completely isolated linguistic community, none of whose members ever heard anyone speaking a different language. "Such a community would not know—or perhaps even have reason to believe—that there are humanly possible alternatives to speaking their language. . . . Yet we have no inclination to deny that their language is conventional. They are simply wrong about the nature of their activities."⁶ Notably, such mistakes can go both ways; for example, some people believe that all moral norms are social conventions. They may be quite wrong.

There is, really, nothing surprising about the fact that conventionality is often opaque. The conventionality of various domains, or types of norms, could not be controversial had it been the case that conventionality is necessarily transparent. However, the fact that people need not be aware of the correct reasons for following a convention, or even of the fact that it is a conventional rule they follow, does not entail that there are no epistemic constraints that apply to conventionality of norms. The idea of following a rule, in itself, is a rather complex condition. It normally entails that the agent regards the rule as binding under the circumstances, which would normally entail that the agent must be aware of the

⁵ Lewis, *Convention*, 58.

⁶ Burge, "On Knowledge and Convention," 250.

fact that he or she is following a rule. It may be tempting to think that people can act in ways which conform to a convention without it being the case that in this they follow a rule. According to this line of thought, then, conformity to a convention is not always a matter of following a rule. There is, I think, some truth in this, but only in a very limited sense. We often follow norms without being self-consciously aware of the fact that we do so. In reading these words you follow numerous norms of the English language (notation of symbols, spelling, meaning, syntax, etc.). It is not something you do in a self-conscious way; you don't tell yourself that the symbol "A" stands for a particular sound, that "B" stands for a different sound, and so on. Following a rule does not require that the agent be self-consciously aware of the fact that he or she follows the rule. But in using these symbols in reading or writing, we do follow rules, and, generally speaking, we know that we do. We normally become aware of the fact that we follow a rule when our attention is drawn to it by some particular need, say, when there is a question about what the rule is, or how to interpret it under some doubtful circumstances. None of this proves, of course, that conventions are necessarily rules. I will deal with this question shortly. The point here is that to the extent that conventions are rules, and to the extent that conformity to a convention is an instance of following a norm, there is always the *potential* of awareness that in complying with a convention one follows a rule. But again, what the reasons for the rule are, or what kind of rule it is, is not something that the agents must be aware of.

The assumption that there must be a reason for following a rule that is a convention requires another clarification. On the one hand, it is normally the case that people follow rules for reasons. On the other hand, we must make room for the possibility that a convention is silly or just plain wrong. There would seem to be two ways to account for this. In many cases it would be appropriate to say that there is a reason to follow a convention, R, but there are also reasons not to follow R, and perhaps the latter ought to prevail. A reason to follow a rule does not entail an all-things-considered judgment that one ought to follow it. Second, it might be the case that the reason to follow the

convention is just not a good reason. But this is problematic: according to a plausible view about the nature of reasons, there is no such thing as a bad reason (just as there is no such thing as a bad value). Either there is a reason, or there isn't. If this is correct (and I think that it is), then we cannot say that there is a reason to follow a convention that is just a bad reason; we would have to say that there is a reason, perhaps a very weak one, that is somehow immediately defeated by countervailing considerations.

The third condition explains the sense in which conventional rules are *arbitrary* and, as we shall see, *compliance dependent*:

3. There is at least one other potential rule, S, that if members of P had actually followed in circumstances C, then A would have been a sufficient reason for members of P to follow S instead of R in circumstances C, and at least partly because S is the rule generally followed instead of R. The rules R and S are such that it is impossible (or pointless) to comply with both of them concomitantly in circumstances C.

As David Lewis explained in his account of social conventions, it is crucial to note that arbitrariness (thus defined) should not be confused with indifference.⁷ This condition does not entail that people who follow the convention ought to be indifferent about the choice between R and S. The rule is arbitrary, in the requisite sense, even if people do have a reason to prefer one over the other, but only as long as the reason to prefer one of the potential rules is not stronger than the reason to follow the rule that is actually followed by others. A typical game theory model of this is the so-called battle of the sexes game (which Lewis calls an imperfect coordination problem). Assume two agents, X and Y, both having a dominant preference to act in concert with the other; however, X prefers option P over Q, and Y prefers option Q over P. As long as their dominant preference is to act in concert with each other (namely, if Y opts for Q, X would rather Q as well, even though he would have preferred P; and same goes for Y), the condition of arbitrariness

⁷ See Lewis, *Convention*, 76–80.

as defined above is satisfied. Consider, for example, the case of a greeting convention. Suppose that there is a reason to greet acquaintances in some conventional manner. Now let us assume, for the sake of simplicity, that there are only two possible alternatives: we can either greet each other by shaking hands, or else by just nodding our head. Presumably, some people would prefer the hand-shaking option, while others (perhaps because they are not so keen on physical contact) would prefer nodding. The point is that we need not assume indifference. As long as the reason to act in concert with others is stronger than people's preference for one of the options, whichever rule evolves as the common practice is likely to be followed. And it would be arbitrary in the sense defined here.⁸

In a sense, then, arbitrariness admits of degrees. We could say that a rule is completely arbitrary if the reason to follow it entails complete indifference between the rule, R, that people do follow, and its alternative(s), S, that they could have followed instead, achieving the same purpose. Then a rule becomes less and less arbitrary as we move away from complete indifference, up to the point at which the reason to follow the rule that is actually followed by others is just slightly stronger than the reason to prefer a different alternative.

Arbitrariness is an essential, defining feature, of conventional rules.⁹ This is actually a twofold condition. First, a rule is arbitrary if it has a conceivable alternative. If a rule does not have an alternative that could have been followed instead without a significant loss in its function or purpose, then it is not a convention. Norms of rationality, and basic moral norms, for instance, are not conventions; properly defined and qualified, they do not

⁸ As we shall see later, the rationale of following arbitrary rules is not confined to coordination situations. For a much more sophisticated game-theoretical treatment of these issues, see, for example, Sugden, *Economics of Rights*, and Vanderschraaf, "Convention as Correlated Equilibrium."

⁹ This is undoubtedly one of the most important insights of Lewis's account, an insight that is shared by most of his critics. (See, for example, Millikan, "Conventions of Language Made Simple.") A notable exception is Miller, "Conventions, Interdependence of Action, and Collective Ends." Gilbert also disagrees, and I will consider her alternative account at the end of this chapter.

admit of alternatives (in the sense defined above).¹⁰ Admittedly, it is not easy to define what a relevant alternative to a rule might be. Surely not every imaginable alternative to a rule would satisfy this condition. First, it has to be a rule that the same population could have followed in the same circumstances. Second, it has to be an alternative rule that is supported by the same reasons or functions that the original rule serves for the relevant population. Third, in some loose sense that I cannot define here, the alternative rule has to be one that the relevant population can actually follow so that the cost of following it would not exceed the rule's benefits. Finally, the alternative rule has to be a genuine *alternative*, and not just an additional rule that people could follow in the same circumstances as well.

The second aspect of arbitrariness concerns the nature of reasons for following a convention: The reason for following a rule that is a convention depends on the fact that others follow it too. But we have to be more precise here. There are two possible ways in which the reasons for following a rule are affected by general compliance (in the relevant population), and only one of them is a defining feature of social conventions. The reason for following a rule may sometimes be lost, or seriously compromised, if too many people in the relevant community infringe the rule or otherwise fail to follow it. This, in itself, does not indicate that the rule is arbitrary. Consider, for example, a rule that prohibits smoking in public places. This is not a convention but, let us assume, a rule that is required by the reasons to avoid causing harm to others. Nevertheless, if in a certain place nobody follows the rule and the vicinity is filled with smoke anyway, the reason for following the no-smoking rule is lost. More accurately, we should say that the reason is still there, but it is not an operative reason under the circumstances.¹¹ Practice-

¹⁰ I realize that this is not uncontroversial. For example, Bruno Verbeek in a recent article ("Conventions and Moral Norms") argues that moral norms are conventional. So perhaps no example is free of controversies.

¹¹ I am assuming here that the room is so smoky that the marginal harm of any additional smoker is basically zero. If you prefer a different example, consider the case of pollution: assume it is wrong to pollute the river, but if the river happens to be so polluted anyway that it makes no difference whether

dependence is a defining feature of conventions, however, only when the fact that others actually follow the rule *initially* forms part of the reasons or the rationale for following it. This is not the case with the no-smoking rule: people should not smoke in public places because it causes harm to others. A certain level of compliance with the rule may be required to ensure its success in solving the problem that the rule is there to solve. But it is not the case that in order to explain what is the point of the no-smoking rule we must appeal to the fact that it happens to be the rule most people follow in the relevant circumstances. Note that I am assuming here that even one smoker in the vicinity of others is causing some harm and that this in itself is a reason to refrain from smoking in the presence of others—unless, that is, the vicinity is so full of smoke anyway that an additional smoker makes no difference. In the case of conventional rules, however, there is a crucial sense in which it is uniquely appropriate to say that we follow the rule partly *because* others follow it too. (We drive on the right side of the road because others drive on it too; or we wear a suit and tie to this party because others will come similarly dressed, etc.)¹²

Let me call this the condition *compliance-dependent reasons*. A reason for following a rule R is compliance dependent if and only if, for a population P in circumstances C,

1. there is a reason for having R, which is also a reason for having at least one other alternative rule, S, and,
2. part of the reason to follow R instead of S (in circumstances C) consists in the fact that R is the rule actually followed by most members of P in circumstances C. In other words, there is a reason for following R if R is generally complied with, and the same reason is a reason for an alternative rule if that alternative is the rule generally complied with.

Arbitrariness, and therefore conventionality, assumes that the relevant reasons for following the rule are compliance dependent

we add some or not, then the reason not to pollute is not an operative reason under the circumstances.

¹² Hume suggested a very similar observation; see *Treatise*, 490.

in this sense. A rule is conventional if and only if there is at least one other potential rule that the relevant community could have followed instead, achieving the same purpose, as it were. The reason for following a convention partly depends on the fact that it just happens to be the rule that people in the relevant community actually follow. Had they followed an alternative rule, the same reason, A, would provide a sufficient reason to follow the alternative rule, namely, the one that people actually follow.

The notion of “sufficient reason” calls for a clarification. Some philosophers may believe that all reasons for action, as such, are sufficient reasons, in that other things being equal, one’s failure to act on a reason that applies to the circumstances is, *ipso facto*, wrong. I do not think that this is a correct view of reasons for action. A person may have a reason to play chess, for instance, but it would not be wrong in any sense if she decides not to play chess for no reason at all. In any case, it will not be assumed here that failure to act on a reason is, *ipso facto*, wrong. I do mean, however, that if A is a reason for playing the game according to rule R, R is arbitrary if and only if there is at least one other rule, S, so that if the game was actually played in compliance with S, A would provide a sufficient reason (for the same agents, under the same circumstances) to follow S instead of R.¹³ All this indicates that *conventionality is relative to reasons*. Given that A is a reason for R, if R is generally complied with, and A would have been a sufficient reason for an alternative rule if that alternative is generally complied with, then R is conventional relative to A. Note that R need not be conventional altogether, as there may be some other reasons to comply with R that are not convention-type reasons. A norm would be purely conventional, however, if there are no such other reasons.

You may think that all this talk about arbitrariness of conventional norms has not yet been proved; I have not provided any argument to support the assumption here that arbitrariness, as defined above, is an essential feature of conventions. I doubt, however, that any straightforward argument is available.

¹³ There is a sense, of course, in which it might not be the same game. I will elaborate on some of the difficulties about identity of practices in chapter 2.

My assumption here is that the way we characterized conventional norms captures some basic intuitions we have about such norms, most importantly, the intuition that if a norm is conventional, we could have had a different, alternative, norm that would have served us just as well (in the same circumstances, of course). I have tried to offer a detailed formulation of this intuition; but the formulation cannot be vindicated by an argument. Its validity depends on how theoretically fruitful the formulation is, on how it is put to work in the rest of the argument. In particular, we will have to see how arbitrariness (and compliance dependence) help to clarify not just whether certain types of norm are conventional or not, but also why would it matter that they are, and how could it be philosophically controversial. I believe that defining conventionality in terms that are relative to reasons is going to be helpful in these respects. But of course, whether it is really helpful and to what extent, depends on of the arguments that will be deployed in the rest of this book.

As a final clarification of the definition of conventionality, we need to say something about the idea of rules. Conventions, I have claimed, are social rules. Rules should be distinguished from both regularities of behavior and from generally recognized reasons. Not everything we do regularly we do as an instance of *following a rule*. People regularly eat breakfast, but eating breakfast is not an instance of following a rule. It is just something we tend to do regularly. Rules are normative; the validity of a given rule that applies to the circumstances is a factor in practical reasoning. I will assume here that a rule of conduct exists when there is a certain population¹⁴ that regards the rule as binding. (Needless to say, not every rule that is *regarded* as binding is binding or valid; it is, if there is an adequate reason that supports it.) Without an attempt to define what rules are, we can say at least this: The basic function of rules of conduct is to *replace* (at least some of the) first-order reasons for action. When we make it a rule to ϕ under circumstances C , it is as if we have made a decision that under circumstances C there is no

¹⁴In some cases, the relevant population can be a single-member set. People can make rules for themselves.

need to deliberate, or to try to figure out, whether to ϕ or not (that is, up to a point, of course, and under certain conditions, etc.).¹⁵ The rule replaces the need for separate deliberation in each and every context of its application. We take the fact that there is a rule to count in favor of doing ϕ . In this sense, we take the rule to be binding, namely, we take it as a *pro tanto* reason for action. A regularity of behavior, as such, does not have such normative significance. Consider John who tends to skip breakfast: perhaps it is not wise, and he may be criticized for not being responsive to reasons (assuming there are reasons to eat breakfast), but John could not be sensibly criticized for breaking a rule. There is simply no such a rule.¹⁶ There are many things we have reason to do with some regularity because the circumstances that give rise to the reasons appear in some regular fashion. This, in itself, does not make our conformity with those reasons instances of following a rule. The normativity of rules consists in the fact that the rule, as such, forms part of one's practical reasoning; if the rule is sound, or valid, its application to the circumstances is a fact that counts in favor (or against) doing (or not doing) something.¹⁷

Rules should also be distinguished from generally recognized reasons. Consider, for example, a game, like chess. It is constituted by a set of rules, some of them very explicit and others, perhaps, less so. In playing the game, players follow the rules of chess. And then, additionally, there may also be some strategies that are widely recognized among players as sound strategies. Now those strategies are not necessarily rules. They might be reasons that apply to some aspect of playing the game, and they may well be widely recognized as such.¹⁸ Admittedly, such strategies can be formulated as rules. When instructing a novice player how to play, one can say, "Don't ever move the king when . . ."; and this sounds very much like a rule. But the

¹⁵ See Raz, *Practical Reasons and Norms*.

¹⁶ Unless, of course, John made it a rule to himself to eat breakfast. I am assuming here that this is not the case.

¹⁷ See Hart, *The Concept of Law*, chap. 5.

¹⁸ The distinction, as well as the game example, has been proposed by Warnock, *The Object of Morality*, 45–46.

formulation is potentially misleading. When you point to a rule of the game, you cite the existence of the rule as the reason for action. (“You may only move the bishop diagonally”; “Why?” “*Because* this is the rule.”) When you point out a sound strategy, you are not entitled to point to a rule as the relevant reason for action. You just sum up the reasons that apply independently, in a rule-like formulation. Once again, failing to abide by a sound strategy can be criticized as foolish or wrong, but not as breaking a rule. To miss this point is to miss the normative significance of conventions. Conventions are rules of conduct, and they are normatively significant as such.

At one point Lewis raised some doubts about this, arguing that conventions need not necessarily be rules. In many games, he claims, players normally develop a set of tacit and informal understandings about what they are entitled to do in circumstances that are not covered by the rules of the game. These conventions, he contends, are ones left open by the “listed rules” of the game. Lewis concedes that “we might call these understandings rules—unwritten rules, informal rules—if we like.” But, he claims, “we would also be inclined to emphasize their differences from the listed rules by saying that they are not rules, but only conventions.”¹⁹

These cryptic remarks are very misleading. For one, they seem to suggest that our concept of a rule ties rules to some sort of formality, as if rules are only those things one would find in rule-books, enlisted and codified systematically, as it were. There is no good reason to hold such a formalistic conception of rules. It is a convention of many games, for instance, that the participants ought to make their moves within a reasonable period of time. As procrastination and lengthy delays would defeat the purpose or the enjoyment of such games, these conventions are often taken for granted and hence left unstated, as it were, by the listed rules of the game. They would normally surface once there is some tendency to deviate from them, and then they might get enlisted and codified like any other rule

¹⁹ Lewis, *Convention*, 104–5; a similar view is expressed by Searle, *Construction of Social Reality*, 28.

of the game. Whether codified or not, however, such conventions are certainly part of the rules of the game. In other words, whatever rules are, they are not necessarily “formal,” written somewhere, or explicitly promulgated as such.

Second, Lewis’s remarks here seem to conflate unlisted rules with generally recognized reasons. “Tacit understandings” that evolve in playing games can be generally recognized strategies that apply, not necessarily rules. To count as rules or conventions, they would have to attain a normative significance. We should be able to say that by not making such and such a move, the player has broken a rule and *thus* did something impermissible. Whether this necessarily calls for a sanction is, of course a separate question. But it is worth noting that typically breaking a rule, as such, can be cited as an adequate justification for a sanction, whereas a failure to abide by a sound strategy cannot.

Lewis is quite right to complain that it is very difficult to define what rules, in the relevant sense, are, and that “the class of so-called rules is a miscellany, with many debatable members.”²⁰ However, it is not a precise definition of rules that we need here. What we need is just to avoid a confusion. And the confusion is to conflate instances of following rules, with mere regularities of behavior or with cases of abiding by a generally recognized reason. It is important to avoid this confusion in order to be able to appreciate the normative significance of social conventions. Conventions, *qua* social rules, are normatively significant. They are cited as reasons for action, and it is precisely the challenge of a theory of conventions to explain what kind of reasons they are.²¹ Furthermore, by downplaying this normative significance of conventions as a species of rules, Lewis has opened himself to the criticism that his own account of conventions, which heavily relies on the fact that conventions involve compliance-dependent reasons, is wrong. If we allow generally recognized reasons that are not rules to be examples

²⁰ Lewis, *Convention*, 105.

²¹ It seems that in “Languages and Language” Lewis modified his views about this, and basically acknowledged that conventions are essentially normative.

of conventions, it is all too easy to show that conventions need not be practiced in order to exist. If R' is a reason that applies to circumstances C , and is widely recognized as such, R' can be thought to be a convention even if most people in the relevant community fail to follow R' in circumstances C .²² But this is misguided. It is true of generally recognized reasons, but not of conventions, that they exist regardless of practice. Nothing in the logic of generally recognized reasons requires those reasons to be compliance dependent. One way to block this misguided criticism, then, is by keeping in mind that generally recognized reasons are not conventions.

A possible objection needs to be answered. It might be thought that *beliefs* can be conventional as well, not just rules. For example, people in a certain community may widely share the belief that women should not wear pants; or they may share the belief that it is better to sleep in beds than in hammocks. Countless such beliefs are conventional, according to this line of thought, simply in the sense that they are widely shared in a certain community, and held as the beliefs they are, mostly because they are just widely shared. Now, if beliefs can be conventional, then a widely shared belief in the reasons that apply to certain actions might be conventional as well. And then it is no longer true that conventions can be distinguished from generally recognized reasons.

The problem with this argument lies in its assumption: beliefs are not appropriate candidates for conventionality. Whereas it makes perfect sense, under certain circumstances, to act in a certain way only because others act in the same way (e.g., on which side of the road to drive), it makes no sense to believe that P only because others believe it too. We cannot provide any sensible rationale for conventionality of beliefs, as opposed

²² Indeed, many of Gilbert's counterexamples to Lewis's analysis involve precisely this mistake. The examples she gives are cases of generally recognized reasons, not rules. Hence it is all too easy for her to argue that their existence does not depend on general compliance; see *On Social Facts*, 344–55. My remarks here are also meant to apply to Millikan's view (see her "Conventions of Language Made Simple," 173–74, and, more generally, her *Language: A Biological Model*).

to rules that guide action. It is true, of course, that we often come to have beliefs, that is, we acquire them, because we get a sense that others in our vicinity have those beliefs and perhaps hold them strongly. This may be a fact about our psychology, but it is beside the point. In justifying a belief he holds, a person would make a fool of himself by saying, "I believe that P because everybody else does."²³ Beliefs are appraised by their truth or falsehood, not by their compliance with others' beliefs. In other words, whereas we may have compliance-dependent reasons (for action), something like compliance-dependent belief is not a coherent idea. There is no case in which it is true that (1) X believes that P and (2) had it been the case that everybody else believed that not-P, X would have had a *sufficient reason* to believe that not-P. So what about the examples? Is there no sense in which we can properly speak about conventionally held beliefs? Perhaps there is, in a derivative or metaphorical sense. We can talk about conventional beliefs when we want to indicate that the belief in question is not warranted by its truth, but is widely held nevertheless. Or perhaps when we intend to indicate that the belief represents a traditional way of thought, reflecting a social custom, or such. Either way, the notion of convention is used imperfectly here, and does not withstand philosophical analysis. Conventions are social rules that purport to guide action; they are not beliefs.

A final question before we turn to Lewis's theory. One may wonder whether there are *social* norms that are not conventions. Is the category of social norms wider than that of social conventions? Now of course, this very much depends on how we understand the idea of social norms. But let us assume a simple, intuitive understanding here, whereby social norms are simply those norms that are widely followed in a certain community, and they do not originate in any institutional enact-

²³ Unless one has independent reasons for thinking that a certain population is more likely to have correct beliefs on a particular issue. I can justify my beliefs on certain scientific matters by citing the fact that those in a position to know such things hold those beliefs. In this sense, I use others' beliefs as a kind of indication or evidence for truth. But none of this would show that beliefs can be conventional.

ment (like legal norms or regulations of a college, etc.). So the question is, are all social norms necessarily conventions? I think that the answer is no, because many of the social norms we follow are not arbitrary in the sense defined above. Some social norms do not have genuine alternatives that we could have followed instead without a significant loss in their purpose or function. Consider, for example, cultures in which it is a widely accepted social norm that younger people should care for the elderly. Presumably, this is a morally sound norm that is supported by good reasons, and those reasons are not compliance dependent. The fact that the norm happens to be socially accepted and followed in the community does not make it conventional. And of course, similar considerations apply to countless other social norms that instantiate sound moral principles. True, it often happens that such norms are actually manifest in a variety of specific norms that are conventional. For example, conventions may determine ways in which care and respect for the elderly is actually instantiated in the relevant community. But the underlying social norm that requires such care and respect remains nonconventional, that is, even when there is a variety of norms that conventionally determine its application. In short, I don't think that we are entitled to assume that all social norms are conventional; each case needs to be examined on its own terms.

DAVID LEWIS ON CONVENTIONS

David Lewis provided an ingenious account of social conventions that, at least in its core, has been widely accepted ever since.²⁴ Lewis claimed that conventions are social rules that emerge as practical solutions to wide-scale, recurrent, coordination problems. Interestingly, Lewis's account of social conventions was aimed at answering Quine's doubts about the conventionality of language. Quine argued that language cannot be conventional because conventions are essentially

²⁴Lewis, *Convention*.

agreements, and of course, we have never agreed with each other to abide by the rules of language. More importantly, we would have needed at least some rudimentary language in which to express the first agreements. Lewis replied that Quine's assumption was simply wrong: conventions do not result from agreements. On the contrary, conventions tend to emerge precisely in those cases where an agreement is very difficult or impossible to reach (because of the large number of people involved, for example), and a solution to a recurrent coordination problem is needed.

A coordination problem arises when several agents have a particular structure of preferences with respect to their mutual modes of conduct: namely, that between several alternatives of conduct open to them in a given set of circumstances, each and every agent has a stronger preference to act in concert with the other agents, than his own preference for acting upon any one of the particular alternatives.²⁵ Most coordination problems in our lives are easily solved by simple agreements between the agents to act upon one more or less arbitrarily chosen alternative, thus securing concerted action among them. However, when a particular coordination problem is recurrent in a given set of circumstances, and agreement is difficult to obtain (mostly because of the large number of agents involved), a social rule is very likely to emerge, and this rule is a convention. Conventions, in other words, emerge as solutions to large-scale recurrent coordination problems, not as a result of an agreement, but as an alternative to such an agreement, precisely in those cases where agreements are difficult or impossible to obtain.

Lewis's analysis of conventions in terms of solutions to recurrent coordination problems embodies remarkable advantages. Quite apart from the fact that it is capable of explaining the emergence of conventions without relying on the need for agreement, it also explains the essential conditions of conventionality.

²⁵ Lewis defined coordination problems in terms of a standard game-theoretical model. The details of such models have been substantially revised since, and the models game theorists work with have become much more sophisticated. See, for example, Sugden, *Economics of Rights*; Vanderschraaf, "Convention as Correlated Equilibrium"; Bacharach, *Beyond Individual Choice*; Bicchieri, *The Grammar of Society*.

If the whole point of a conventional rule is to secure concerted action among a number of agents in accordance with their own dominant preferences, then it is quite clear why the reasons for following a convention are compliance-dependent reasons. Solving a coordination problem is a paradigmatic example of compliance-dependent reasons: I have a reason to do P (and not Q or R) if, and only if, I have reason to assume that others will do P as well (and not Q or R). And then it also becomes quite clear why it is the case that conventional rules are arbitrary in the requisite sense. Once again, it is the very structure of a coordination problem that there are at least two alternatives of conduct open before the agents in question, and that they have a stronger preference to act in concert than to act upon any other alternative they may (subjectively) prefer. If a given rule is a solution to a coordination problem, then it is already assumed that there is an alternative rule that the agents could have followed instead, solving the relevant coordination problem they had faced.

So far so good. Lewis, however, added another layer to his explanation, concerning the ways in which conventional rules tend to emerge. Following standard game-theoretical analysis (as known at the time), Lewis suggested that when there is a recurrent coordination problem and there are many agents involved, so that agreements are difficult, if not impossible, to reach, people would tend to opt for the option that happens to be salient, assuming that others would opt for the salient option as well, thus securing the relevant concerted action. I am sure that this often happens. I doubt, however, that it is the business of a philosophical analysis to explain how conventions emerge, as a matter of a historical account. I would guess that many of the conventions we follow emerged almost by accident, and many have emerged for various obscure reasons that we can hardly even trace back to their historical origin. (I have heard that the convention of a salute in the army evolved from medieval knights' practice of greeting their opponents by raising the visor of their helmet before battle. Whether this is a true story, I don't know; the point is that there are many like it.)²⁶

²⁶ Here's another example: I have heard that the convention to say "hello" when responding to a telephone call comes from Hungarian; the Hungarians

In any case, as long as we understand that large-scale recurrent coordination problems can be solved without the necessity of an agreement between the relevant agents, Lewis's theory proves its point. The suggestion that we tend to follow salient options, and thus solve coordination problems, is one plausible speculation, but not more than this.²⁷

Be this as it may, the main problem with Lewis's analysis concerns its scope. It is difficult to deny that many conventions are, indeed, normative solutions to large-scale recurrent coordination problems. Clear examples would be notational conventions, like the sound-sense conventions of languages, early conventions of the road, such as on which side of the road to drive (that is, before such conventions were replaced by legal regulation), and many others. But the mistake here is to generalize from some cases to all. There are important types of social conventions that do not fit this analysis. Generally, the problem is this: Lewis's analysis assumes that *first* there is a recurrent coordination problem in a given set of circumstances, and then a social rule evolves that solves the problem for the relevant agents. For many types of familiar conventions, however, this story does not make sense. Antecedently to the emergence of the convention, there is no coordination problem that we can identify, at least not without already assuming that the conventional practice is in place.²⁸

Consider, for example, a structured game, like chess. Presumably, the rules constituting the game of chess are (or at

were, they claim, among the first to invent the contraption, concomitantly with Bell and Gray, and the word "hello" is a slight distortion of the Hungarian word "hallod," which literally means "do you hear." Others credit Edison with the introduction of this convention. Who knows? Either way, salience seems to have very little to do with it.

²⁷ For a much more empirically based account, see, for example, Bicchieri, *The Grammar of Society*, chap. 2.

²⁸ Several writers have criticized Lewis's generalization in this respect, though mostly on game-theoretical grounds. See, for example, Miller, "Rationalizing Conventions"; Sugden, *Economics of Rights*; and Vanderschraaf, "Convention as Correlated Equilibrium," and see also Davis, *Meaning, Expression, and Thought*. Gilbert and, to some extent, Millikan, reject Lewis's model on other, more general grounds. More on this below.

least were, before institutionally codified) conventions.²⁹ Does it make sense to suggest that the rules of chess are there to solve a coordination problem between potential chess players? Have these rules evolved as a solution to some large-scale recurrent coordination problem that we could have specified before the emergence of the game itself? I think it is extremely unlikely that playing by the rules of chess solves a coordination problem between the players; as if there had been a coordination problem between potential chess players before chess was invented, as it were, and now they play by the rules to solve that problem. This seems implausible. Of course you can structure a very vague and highly general coordination problem, say, a desire to play some intellectual board game. (Note that even for this general problem to be specified, we would need a fairly elaborate conception of what board games are, which in itself is conventionally determined. But let's ignore this complication for now.) So allegedly, you can say that there was this general coordination problem: there we are, wanting to play some structured board game, and then rules have evolved to solve *that* problem. The obvious difficulty is that such a coordination problem would be too abstract and underspecified. If you allow for almost any concerted action between a number of agents to count as a solution to a coordination problem in the relevant sense, then you will end up with the conclusion that all social interactions are solutions to coordination problems. That seems to be very unlikely and philosophically unhelpful.

More importantly, there is something seriously amiss about the suggestion to characterize the rationale of playing chess as a solution to a coordination problem. When asked, for example, why I drive on the right side of the road, it makes perfect sense to reply that I do it because I need to coordinate my

²⁹ You may have some doubts about the appropriateness of the example. In chapter 2 I explain in greater detail why the rules of chess are conventional, and I also discuss the ways in which conventions tend to be codified and how codification affects the conventionality of the relevant practice. For now, chess is really just one possible example; if you have doubts about the conventionality of chess, think about other cases, like practices of etiquette, various social rituals, conventions of artistic genres (discussed below), etc.

driving with others. But when asked why I play chess right now, it would be perplexing to reply that I do it because I need to coordinate my behavior with my fellow players. There are two related points that I want to make here: first, even if we could tell a story that would explain the *emergence* of chess as a solution to some vague coordination problem, such a story would be quite irrelevant to our present concerns when playing chess. To solve a coordination problem is not why players indulge in this game. In the case of coordination conventions, the reasons for the emergence of the convention and the reasons for complying with the convention in each and every instance, are basically the same: to solve the relevant coordination problem. In the case of a game, like chess, however, this is not necessarily, or even typically, true. The reasons for the emergence of the relevant practice need not be the same as the reasons for complying with its norms on specific occasions. Once the game is there and can be played, it may give reasons for the agents to follow its rules that are quite independent of the story of why and how the game has emerged. (More on this in chapter 2.)

Second, and more important, the reasons people normally have for playing chess have very little to do with solving a coordination problem. Of course, once the game is there and it is being played, it may give rise to various coordination problems that might then get solved by additional rules or conventions. But the essential rationale of the game is ill explained in terms of a solution to a coordination problem. It is true, of course, that in playing any structured game, part of the reason we have to stick to the rules, and be sure that we all know what the rules are, is to make sure that our actions are well coordinated. So yes, undoubtedly there is a coordinative function to any such rule-guided activity, like playing a structured competitive game. But this is only one aspect of playing chess, not its main rationale. And it is an aspect that is present in any rule-governed activity, whether conventional or not.

Consider a different example, from the realm of arts. Most forms of art are at least partly conventional. Conventions determine, for example, what theater is, how it is staged, what one is to expect from such a performance, and so on. Would it make

any sense to suggest that conventions constituting such artistic forms or genres are also solutions to coordination problems? Does it make sense to assume that before theater evolved as a specific form of art, there was some recurrent coordination problem that needed to be solved, and then the conventions of theater evolved to solve it? What would that coordination problem be, and for whom? So I hope you see where I am heading here: either we construe the idea of a coordination problem in such abstract and general terms as to be completely unhelpful, or else we must give up the idea that all social conventions are explicable in terms of solutions to recurrent coordination problems. The latter option should be easy to follow. Social conventions evolve as responses to numerous kinds of social needs, they serve a wide variety of social functions, and we have no reason to assume that all those needs are reducible to coordination problems. This idea will be further developed in the next chapter.

AN ALTERNATIVE ACCOUNT?

The details of Lewis's theory of conventions have been criticized by many philosophers, but few have suggested an entirely different approach.³⁰ Margaret Gilbert's work is a notable exception: she has argued for an alternative account of what social conventions are, one that is quite different from Lewis's. The gist of Gilbert's characterization of social conventions is the following:

³⁰ See note 28 above. How to classify, in this respect, Millikan's account of conventions as patterns of behavior sustained by weight of precedent, is something that I find very difficult to determine. On the one hand, she seems to share some of Lewis's basic insights about the nature of conventionality; in particular, she quite explicitly endorses Lewis's account of arbitrariness (see her "Conventions of Language Made Simple," 167). On the other hand, her account deviates from Lewis's in many crucial points, some of which are at odds with my suggestions here (e.g., that conventions are necessarily rules/norms), while others not necessarily. Be this as it may, the main difference between the account I offer here and Millikan's view consists in the relation of conventions to reasons. On my view, conventionality is relative to reasons; on her view, conventions have basically nothing to do with reasons.

[O]ur everyday concept of a social convention is that of a jointly accepted principle of action, a group fiat with respect to how one is to act in certain situations. . . . conventions on this account are essentially collectivity-involving: a population that develops a convention in this sense becomes by that very fact a collectivity. Further, each party to the convention will accept that each one personally ought to conform, other things being equal, where the “ought” is understood to be based on the fact that together they jointly accept the principle.³¹

Before we try to untangle some of these concepts, a few points have to be mentioned. First, as we have already noted, Gilbert explicitly rejects the assumption that conventions have to be practiced by a certain community in order to exist as social conventions. Second, and more problematically, Gilbert explicitly rejects the idea that arbitrariness is an essential characterization of conventional norms. The problem is, however, that Gilbert explicitly assumes that Lewis’s definition of arbitrariness is equivalent to indifference.³² Since we have already seen in some detail that indifference is not what is entailed by an appropriate definition of the arbitrariness of conventions, Gilbert’s criticism misses its target here. Furthermore, failing to realize the essentially arbitrary nature of conventions results in the failure to see that the reasons for following conventions are compliance-dependent reasons. And then it becomes very doubtful that it is really “our everyday concept of a social convention” that Gilbert is explicating here.

Be this as it may, there are two main concepts in her characterization of conventions worth looking at. First, Gilbert defines conventions in terms of a “group fiat” or of a “jointly accepted principle of action.” Second, she claims that conventions are “essentially collectivity-involving: a population that develops a convention in this sense becomes by that very fact a collectivity.” I think that both concepts are misguided. The latter is really just a generalization from some (very few, actually)

³¹ *On Social Facts*, 377.

³² *Ibid.*, 341.

cases to all, and the former confuses conventions with generally recognized reasons or, alternatively, agreements. Let me take up these points in this reversed order.

It is true that the acceptance of certain social norms is sometimes constitutive of a group identity. One becomes a football fan, for example, by, *inter alia*, accepting the norms that govern this practice. The group's adherence to such norms is partly what defines the identity of the group. Are such norms that constitute group identity necessarily conventions? Probably not. To take an extreme example, some philosophers suggest that our adherence to norms of rationality and practical reasoning is what defines or constitutes our humanity. Surely this is not a matter of convention. Or think about the common suggestion that American (that is, U.S.) identity is partly constituted by accepting certain principles of political morality, like cherishing individual freedom, free enterprise, constitutional protection of civil rights, and so on. In short, it seems that "essentially collectivity-involving" norms need not be conventions. More importantly, however, conventions need not be collectivity-involving. Some conventions are, but many conventions have no such role to play. There are numerous conventions we follow that have nothing to do with group identity. Consider, for example, something like the notational convention of the arrow sign. As far as I know, in most cultures the sign of an arrow is a conventional means of pointing to a certain direction in space. Does it make any sense to suggest that by following the arrow convention all of us become a collectivity? What collectivity would that be? And do we become a collectivity by following conventions of the road, or notational conventions of a language, and so on? The connection between conventions and group identity is contingent, at best. This collectivity involving feature is not a defining feature of social conventions.

The concept of "group fiat" or "jointly accepted principle of action" is also very problematic. To begin with, it might be overinclusive. A principle of action that is jointly accepted can be an instance of a generally recognized reason for action, or some other type of social norm, not necessarily a convention. Suppose, for example, that it is a "jointly accepted principle of

action” in my department that we do not burden our junior faculty with committee assignments. This is not necessarily a convention. More plausibly, it is a principle of action we adhere to because we recognize it to be a good practice. We generally recognize that it is better for the junior faculty to be able to focus on their research and teaching until they acquire more experience, and we try to implement this principle by excluding them from administrative work. Perhaps we regard it as a “group fiat” or a “jointly accepted principle,” but this would not make it a convention. Generally speaking, how people see themselves committed to a given norm does not entail anything about the kind of norm it is.

Gilbert would reply that I have missed a crucial point here about the idea of a joint acceptance. Presumably, what she means is some notion of *conditional* acceptance based on reciprocity: I commit myself to N if and only if you commit yourself too, and you commit yourself to N if and only if I do. If something like this is what Gilbert has in mind, then perhaps the above example is not a counterexample. But then two other problems emerge: first, this notion of joint acceptance as conditional on reciprocity would seem to be inconsistent with some of Gilbert’s critical remarks about Lewis’s account. Second, and more important, it is very difficult to make sense of the idea of conditional acceptance without relying on the idea of agreement. Let me explain.

One of Gilbert’s main criticisms of Lewis’s account of conventions consists in her claim that conventions are not based on what I have called compliance-dependent reasons. As she put it: “social conventions can exist in the absence of expectations of conformity.”³³ But then, if you do not need to expect general conformity, what could you mean by “joint acceptance” as a reciprocal condition? I accept N only if you accept it too, but I do not need to expect you to comply with N? That is a strange form of reasoning. In other words, if conventions do not have to involve compliance-dependent reasons, then the idea of reciprocity cannot be rationalized either.

³³ *Ibid.*, at 348.

More importantly, however, the idea of joint acceptance construed in this conditional way would seem to require some notion of agreement to give it any hold in reality. Conditional acceptance, in this sense, is precisely the rationale of a contract: I accept the terms only if you do, and vice versa, you would accept if I do. Gilbert is aware of this. Conventions, she says, are norms of *quasi*-agreement. Now the need for the “quasi” qualifier is clear enough: as Quine and Lewis rightly observed, there is no hope for any account of the conventionality of language if we have to assume that conventions emerge from agreements. So if we are back to the agreement requirement, it better not be anything like an explicit agreement. Hence the suggestion that conventions are *quasi*-agreements. The difficulty, of course, is to provide this “quasi” qualifier with any concrete substance in this context. I don’t see how it can be done.

In numerous contexts and for different reasons, philosophers have tried to rely on the idea that there are cases that are like an agreement, but without the explicit speech acts that would normally constitute an agreement. For example, we can have agreement by some other speech act equivalent, such as a body movement or an inference from some particular behavior. This is still an explicit form of agreement and one that is parasitic on the standard speech act cases. If this is what is meant by “quasi-agreement,” Quine’s problem remains.

Gilbert seems to agree, since she characterizes quasi-agreement as a situation in which no agreement of any kind had taken place, but the situation is one in which “it is as if they have agreed.”³⁴ What are we to make of this? There are two possibilities to make sense of such a locution, but none of them would be helpful. On one possible reading, something may look as if it resulted from an agreement if it is an instance of an invisible-hand mechanism. Those are cases in which people’s behavior over time *looks like* conduct that has emerged from concerted, purposeful, action, when in fact, it has not. This familiar invisible-hand mechanism would not help Gilbert’s account, however, since it explicitly denies precisely what Gilbert

³⁴ Ibid., at 369.

needs here, namely, the idea of some reciprocity, or conditional acceptance. (Note that it is a crucial aspect of an invisible-hand mechanism that the agents act independently of each other, pursuing her own individual ends.)

The second way in which we can make sense of the idea that there are situations that we can treat *as if* there was an agreement is a moral construal. Sometimes we think that people are under an obligation to ϕ even in the absence of an agreement (or consent) on their part to ϕ , because we are entitled to treat them as if they had agreed to ϕ . What those situations are, and how to explain this kind of justification, is a familiar and daunting challenge to many moral philosophers, but we need not go into this here. What Gilbert needs in order to substantiate the idea of a quasi-agreement is not this kind of moral justification. The question we face is not how to justify, morally, a requirement of compliance with norms that do not emerge from agreements. We need an account of what such norms *are*. So here is what Gilbert's account needs: how to explain, without relying on the idea of an agreement, a situation in which the basic rationale of a given social norm is that I will do ϕ only if you ϕ , and vice versa, you will ϕ only if I do. We know one way in which it works: this is a standard coordination problem. So it seems that we have not avoided Lewis's analysis after all. At least not on Gilbert's account. What we need is a fresh start. We need an account of the rationale of those conventions that does not consist in a solution to coordination problems. I propose such an account in the next chapter.