Legal Theory, 8 (2002), 387–441. Printed in the United States of America
Copyright © Cambridge University Press 1352–3252/02 $9.50

LAW, PLANS, AND PRACTICAL REASON*

Scott J. Shapiro
Benjamin N. Cardozo School of Law, Yeshiva University

I. MORALITY OR CONVENTION?

American lawyers are audacious by nature and have been known, on more than one occasion, to make outrageous arguments on behalf of their clients. Yet there is one claim that even the most shameless of lawyers would not put forward in an American court: No one would ever challenge the idea that the United States Constitution is law in the United States.[1] For if there is one proposition of law that enjoys universal acceptance by the bench and is held with unshakable conviction, it is that the United States Constitution determines the authority structure of the federal government and is legally binding on all officials.

It will not be my purpose in this paper to challenge this consensus either. Quite obviously, the United States Constitution is good law. Rather, my aim will be to figure out *why* this proposition is true: What makes it the case that the United States Constitution is law in the United States? What justifies the conviction of the courts?

This question is philosophically interesting because it does not admit of an ordinary legal answer. That is, the legality of the Constitution cannot be established by pointing to some other rule authorizing its creation. The Constitution is ultimate law and as such, its existence as law does not depend on the existence of any other law. To the contrary, the Constitution sets out the authority structure of the federal government and hence plays the role as the foundational norm for the federal system. It validates; it is not itself validated.

This question is quite clearly not unique to the United States legal system

*This paper was presented at the Law, Participation, and Democracy Conference held at the Yale Law School, the University of Pennsylvania Law School legal theory workshop, and the University of Toronto Law school legal theory workshop, where I benefited from comments by Bruce Ackerman, Matt Adler, Michael Bratman, Jules Coleman, Margaret Gilbert, Mark Greenberg, Chris Kutz, Stephen Perry, Philip Petit, Arthur Ripstein, Ori Simchen, Martin Stone, and Nicos Stravropolous.

1. Imagine a lawyer saying to a judge: "Yes, your honor, I concede that my client purposefully engaged in multiple acts of counterfeiting. But, unfortunately for the prosecution, Article I of the Constitution is not good law, and hence Congress lacks the authority to criminalize counterfeiting. The court has no choice but to set my client free."

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

and can be asked about the ultimate laws of other systems as well. What makes it the case, for example, that the British or French or Swedish or Iranian Constitutions are the law of their respective countries? Or, put in completely general terms: What ultimately makes it the case that a particular legal system has a particular authority structure?

One can imagine two sorts of answers, each corresponding to one of the two major jurisprudential schools of thought. The "natural law" approach would claim that moral facts ultimately determine authority structures. Persons have legal authority in a particular system just in case it is morally appropriate that they have such legal power. In the case of the United States Constitution, for example, one might argue that the fact that the Constitution was ratified by a supermajority of state legislatures invests it with moral authority and hence legal authority.

The legal-positivist answer, on the other hand, would point not to a moral fact, but rather to a social fact. The positivist would argue that it is simply true by legal convention that the United States Constitution is the law of the United States. The fact that legal officials agree to treat the Constitution as setting out the authority structure of the federal government makes it the case that the federal government has the authority structure prescribed by the Constitution.

This paper will focus almost exclusively on the legal-positivist answer. I will begin by setting out, and then criticizing, the argument that legal positivists have used to demonstrate the primacy of convention in the determination of legal authority. It will take me most of the remainder of the paper to develop an alternative argument. I will end by showing why, irrespective of the particular argument I develop in favor of legal positivism, the natural-law position is implausible as an account of the determination of legal authority.

II. COORDINATION, CONVENTION, AND AUTHORITY

A. The Right Kind of Social Fact

If one were to ask a lawyer why the United States Constitution, as opposed to, say, the French Constitution, is authoritative in the United States, he or she would probably base his or her answer on the fact that courts in the United States look to the United States Constitution, and not the French Constitution, as setting out the authority structure of their system. Or he or she might say something very similar, such as that the United States Constitution has been ratified in accordance with the requirements accepted as necessary by U.S. courts, and the French Constitution has not. Vice versa for the case of France—the French Constitution, and not the United States Constitution, is authoritative in France, because French courts consult, or accept ratification criteria satisfied by, the former and not the latter. From

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

a lawyer's point of view, social facts determine the authority structure of a legal system—the principles of morality do not enter into the analysis.

Legal positivists have spent much time and effort attempting to justify this lawyerly intuition. Some attempts, however, have been more successful than others. John Austin, for example, thought that legal authority is determined by habits of obedience.[2] The sovereign is the sovereign, on Austin's view, because others habitually obey him and he does not habitually obey anyone else.

Yet Austin's account will not do. As Hart pointed out, there is no sense in which habits can create obligations or establish standards of conduct or thought.[3] The mere fact that I habitually pace when I lecture places me under no obligation, or even gives me a reason, to continue pacing when I lecture.

In order to demonstrate that social facts determine authority structures, therefore, it is necessary to show that the putative facts are the sorts of things that are capable of creating obligations. In the next two sections, I will examine an argument that legal positivists have developed that purports to do just that. These positivists argue that coordination conventions determine authority in every legal system.[2] Since coordination conventions are capable of creating obligations, these positivists conclude that the existence of a convention is the right kind of fact necessary for fixing authority relations.

B. Choice of Authority as Second-Order Coordination Problem

The legal-positivist argument begins by noting that one of the most vital functions authorities serve is to solve recurring coordination problems. Roughly speaking, a coordination problem is a game in which every player strictly prefers that every player act in a uniform manner.[3] A recurring coordination problem is a repeated game in which each round is a coordination problem. For example, motorists want to drive on the same side of the road as all other motorists. However, because both the left and

2. *See, e.g.*, Gerald Postema, *Coordination and Convention at the Foundation of Law*, 11 J. LEGAL STUDIES 165; Jules Coleman, *Negative and Positive Positivism*, *in* MARKETS, MORALS AND THE LAW (1985); *Incorporationism, Conventionality and the Practical Difference Thesis*, *in* HART'S POSTSCRIPT (2001). *See also*, Eerik Lagerspetz, THE OPPOSITE MIRRORS: AN ESSAY ON THE CONVENTIONALIST THEORY OF INSTITUTIONS, Ch. 7 (1995). *See* John Austin, *The Province of Jurisprudence Determined*, Lecture vi.

3. More specifically, a coordination problem exists just in case two conditions are met. First, the interests of the players must predominately coincide. Second, there must exist at least two proper coordination equilibria. A combination of choices S is a proper coordination equilibrium just in case everyone strictly prefers that everyone choose S over a situation where everyone but one chooses S. To require that a coordination problem contain two coordination equilibria, therefore, is tantamount to requiring the existence of two combination of choices S and S′ such that (1) everyone strictly prefers that everyone play their part in S, just in case almost everyone plays their part in S; and (2) everyone strictly prefers that everyone play their part in S′, just in case almost everyone plays their part in S′. H.L.A. Hart, *The Concept of Law*, 60 (2nd ed., Penelope A. Bullock and Joseph Roz, eds., 1994)

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

right sides are equally good choices, they will have problems knowing on which side of the street others will drive and, hence, on which side they should drive. Authorities are able to solve coordination problems because they can designate one of the strategies as the choice for all to follow. In marking one of the combinations as binding on all, everyone's expectations are focused on that combination and the problems of incomplete information are overcome. Motorists know that they should drive on the right side of the road because the law has selected this side by imposing a rule requiring it. A motorist should drive on the right side because he knows that others will drive on the right, given that they expect him to do likewise.

The ability of authorities to solve coordination problems does not stem from any expertise. When the law designates the right side of the road as the proper way to drive, it does not do so because the right side is better than the left—the law was needed precisely because the right side was as good as the left.

Given that expertise is not necessary to solve a coordination problem, there will exist many different sets of authorities that are able to perform such a task. The desirability of authority will be determined in large part by whether members of society already look to them as solvers of coordination problems. It appears, therefore, that the choice of an authority structure is itself a recurring coordination problem. Members of society will prefer that every member follow the same set of authorities just in case almost everyone else does. On this account, everyone in the United States prefers that everyone always heed the United States Constitution just in case almost everyone does. And everyone in the United States also prefers that everyone always heed, say, the French Constitution just in case almost everyone consulted it instead.[4]

These legal positivists do not, however, go so far as to claim that everyone prefers that everyone follow the same authority structure *no matter the structure*. Some people need not prefer heeding the Constitution of the Third Reich even if every one else did. Rather, the claim is that there will always be a set containing more than one authority structure such that everyone would prefer universal conformity to any member of the set over unilateral defection. While reasonable people would not prefer heeding the Constitution of the Third Reich under any circumstances, they would be happy listening to the United States or the French Constitutions, depending on what everyone else did.

4. These positivists do not claim that the selection of an authority structure must be a pure coordination problem. They readily admit that the players' preferences may conflict. Some players might prefer that one authority structure be selected over another, while others have the opposite preferences. What matters is that everyone prefers universal agreement on a structure given that almost every one else has chosen to follow that structure. Coordination games of partial conflict are often known as "The Battle of the Sexes" in the game-theory literature.

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

C. Looking to Lewis

The next step in the legal positivists' argument borrows a page from David Lewis. As Lewis famously argued, conventions are commonly used to solve recurring coordination problems.[5] When a convention exists, the fact that everyone acts according to joint strategy selected by the convention renders that strategy "salient," thereby giving everyone a reason to continue acting according to that strategy. The fact that drivers generally ride on the right side of the road gives drivers reason to believe that the other drivers will continue to ride on the right side, thereby making it rational for them to continue to ride on the right side.

On Lewis' model, a convention is defined as follows:[6]

> A regularity R in the behavior of members of population P in a recurring situation S is a convention if and only if, in any instance of S:
>
> (1) everyone conforms to R;
>
> (2) everyone expects everyone else to conform to R;
>
> (3) everyone prefers to conform to R on condition that the others do, since S is a coordination problem and uniform conformity to R is a coordination equilibrium in S.

Notice that Lewis builds into his definition that conventions solve recurring coordination problems. A regularity that solved a noncoordination problem would not be a convention.

Lewis's model of conventions provides the positivist with a powerful argument in favor of the claim that in every legal system, legal officials follow conventions when determining which authority structure to heed. Because (1) the choice of an authority structure is a recurring coordination problem; (2) legal officials manage to solve these problems; and (3) conventions are common solutions to such problems, it is plausible to infer that legal officials solve their recurring coordination problems *via* conventions.

To illustrate how this argument works, let us assume that most individuals in the United States prefer that everyone heed the United States Constitution just in case every other one heeds the United States Constitution and prefer that everyone heed the French Constitution just in case every other one heeds the French Constitution. By definition, most individuals in the United States face a recurring coordination problem.

It is a fact that most Americans regularly heed the United States Constitution by obeying the authority structure it sets out. This means that these individuals have managed to solve their recurring coordination problem. Hence it is plausible to suppose that they solve their recurring coordination

5. See David Lewis, CONVENTION (1969).

6. *Id.* at 58. The formulation set out in the text is not Lewis's canonical definition but one that he offers earlier in his discussion. The virtue of the earlier formulation is that it makes explicit the role that coordination problems play in Lewisian conventions.

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

problem in the way that many other recurring coordination problems are solved, that is, by convention. The fact that almost everyone generally heeds the United States Constitution gives almost everyone reason to expect that almost everyone else will continue to heed the United States Constitution, thereby giving them sufficient reason to continue heeding the United States Constitution.

The positivist argument concludes with the attempted demonstration that coordination conventions are able to create obligations.[7] As we have seen, when a convention exists, general conformity to it generates expectations that similar behavior will continue. These expectations will normally induce parties to rely on them and to continue to act according to the convention. Behavior that departs from the convention will, therefore, thwart such expectations and harm those who relied to their detriment on general conformity. Such considerations, it is claimed, are of significant practical weight—so weighty, in fact, that they generate obligations to conform.

On this account, conventions not only guide legal officials to the appropriate authorities, but they would impose obligations upon them to do so. In this sense, conventions determine the authority structure of a legal system. For someone has authority over others in a legal system just in case the latter are under obligations to heed the pronouncements of the former.

III. CONVENTIONS AND COORDINATION PROBLEMS

The argument set out in the previous section is an ingenious attempt to show that conventions determine legal authority in every legal system. It represents the best argument that legal positivists have deployed to date in their debate with the natural lawyers. Unfortunately, it is also seriously mistaken.

As I will argue, the positivists were wrong to claim that the selection of an authority structure is necessarily a recurring coordination problem. Coordination problems are defined in terms of the players' preferences, and it is implausible to suppose that the preferences of legal actors necessarily have the structure required of such games. Once it is realized that the selection of an authority structure is not necessarily a recurring coordination problem, positivists will no longer be able to use Lewis's model in order to establish the primacy of conventions.

A. The Arbitrariness of Coordinative Conventions

As many have noted, solutions to coordination problems are arbitrary.[8] They are arbitrary in the sense that the players' preferences that all act on a certain solution can always be "flipped" to some other solution by simply

7. See Postema, *supra* note 3, at 196–197; Coleman, *Incorporationism, Conventionality and the Practical Difference Thesis*, at 119–120.

8. Lewis, supra note 6, at 16–24; Andrei Marmor, *Legal Conventionalism, in* HART'S POSTSCRIPT 201 (2001). Postema refers to this property as "ambiguity." Postema, *supra* note 3, at 174.

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

changing the background behavioral assumption: If everyone prefers that everyone act on some combination of choices (because almost everyone acts this way), everyone would prefer that everyone act on some different combination under the supposition that almost everyone acted in this different way instead. For example, preferences for everyone riding on the right can always be changed to the left if it were supposed that almost everyone rides on the left instead.

To claim that the choice of an authority structure is a recurring coordination problem commits one to holding that the players will see the solutions to the game as arbitrary in the sense just described. But is this assumption plausible? Is it true that in every legal system, everyone's preference for heeding some set of authorities would change if it were supposed that almost everyone acted differently? This, I think, is rather doubtful.[9] Consider, for example, the ancient Hebrews. They believed of the Old Testament that it was the word of God. Surely many, if not most, would not have preferred to follow a different text just in case everyone else did. Indeed, this thought would have been heretical!

In fact, I am not even sure that most Americans would view the United States Constitution as an arbitrary solution to a recurring coordination problem. My guess is that many would believe that they had a moral obligation to heed a text that had been ratified by the representatives of the people of the United States, regardless of what everyone else did. Many officials would probably also share such an attitude. If most officials suddenly abandoned the United States Constitution, this would not lead all others to similar action. Some would resign in protest, while others would continue applying the rules validly enacted under the United States Constitution. To be sure, if the French Constitution had been ratified instead, these judges would prefer following the French, instead of the United States, Constitution. But since the United States Constitution was ratified, they prefer to follow it and it alone.

I am not denying that many individuals in many legal systems regard their authority structure as arbitrary in some way and hence prefer agreement on such structures to disagreement. The point is that it does not seem to be conceptually necessary that such an attitude be universally held. And this constitutes a fatal problem for the positivists' argument. For in those (possible or actual) legal systems where individuals do not view selected authority structures as arbitrary, the selection of an authority structure will not constitute a recurring coordination problem. But if no recurring coordination problem exists in those legal systems, the positivist supplies us no reason to believe that they solve such (noncoordination) problems using conventions. Indeed, it actually follows from Lewis's model that conventions do not exist in such systems because, as we saw, Lewisian conventions solve only recurring coordination problems.

9. Marmor makes a similar point, although in a somewhat different manner. Marmor, supra note 8, at 201–2.

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

B. From Coordination to Cooperation

Although I have argued that the legal positivists' argument fails, I think their basic strategy is a sound one. Positivists will have presumptively established that conventions exist in every legal system if and when they can demonstrate that the selection of an authority structure is a type of decision problem that conventions are especially able to solve. The positivists whose argument we examined had the bad luck of picking the wrong game and the wrong model of convention. As I tried to show, the selection of an authority structure is not necessarily a recurring coordination problem, and legal conventions, if they exist, are not Lewisian ones.

In the remainder of this paper I will present what I believe to be a better implementation of this strategy. As I will argue, legal practice ought to be understood as a "joint intentional activity," that is, as a series of actions performed by members of a group with the intention that their actions will contribute to some collective outcome. Examples of joint intentional activity range from the relatively trivial, such as taking a walk and baking a cake, to the incredibly complex, such as fighting a war and building a skyscraper. The chief practical problem faced by participants in joint intentional activities is how to organize their behavior so that they may successfully achieve the collective outcome.[10] The function of legal conventions, I will suggest, is to organize this behavior so that the participants are more likely to bring about the collective outcome by acting jointly than if they acted separately. Conventions don't just enable coordination—understood as working in tandem—but also enable cooperation—understood as working together.

In the next section, I will set out the account of joint intentional activity recently developed by Michael Bratman. Using Bratman's model, I will try to show how law can be understood as a joint intentional activity and legal conventions as the plans that enable participants to engage in such activities. Unfortunately, certain modifications to Bratman's accounts will be necessary, as he self-consciously ignores cases involving authority structures, and law clearly employs such structures. Once these modifications are in place, it will be relatively easy to show not only that these models fit the case of law rather nicely, but also why conventions are necessary to the existence of such a practice.

IV. BRATMAN ON JOINT INTENTIONAL AND SHARED COOPERATIVE ACTIVITY

A. Individual Action versus Collective Action

Action theorists have long noted that collective action cannot be analyzed as the simple mereological sum of individual action. Let us suppose that you

10. As we shall see, participants need not desire the achievement of the collective outcome in order for a cooperation problem to exist. They may simply desire that they contribute to the achievement of the collective outcome.

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

and I cooked dinner together. While it is true that *we* cooked dinner, it is not true that *you* cooked dinner and *I* cooked dinner. You may have cooked only the vegetables and I may have cooked only the chicken. Or let us suppose that you are going to Chicago and I am going to Chicago. It does not follow that we are going to Chicago together. We may not know each other. Even if we know each other, we may not care about our going to Chicago together. And even if we cared about our going to Chicago together, each of us may not have lifted a finger to make sure that we went together (our being on the same plane might simply be fortuitous).

In a series of important articles, Michael Bratman has argued that collective action of the sort I have been describing, what he calls "joint intentional activity" (hereinafter "JIA"), arises when two general conditions are met.[11] They are:

> (1) Mutual Responsiveness: Each participating agent attempts to be responsive to the intentions and actions of the other, knowing that the other is attempting to be similarly responsive. Each seeks to guide his behavior with an eye to the behavior of the other, knowing that the other seeks to do likewise;[12]
> (2) Commitment to the Joint Activity: The participants each have an appropriate commitment (though perhaps for different reasons) to the joint activity, and their mutual responsiveness is in the pursuit of this commitment.[13]

To illustrate, consider our cooking together. First, when we cook together, we are responsive to each other's culinary intentions and actions. If you are planning on using the stockpot, I will plan on using some other pot or try to persuade you to use some other pot, and if I am using the mixer, you will wait until I finish before using it or try to persuade me to let you use the mixer instead. Second, our responsiveness to each other's intentions and actions are not merely expressions of etiquette but are in service of our joint activity. You do not wrestle the mixer away from me because my use of the mixer is in pursuit of our goal of cooking dinner.

All JIAs involve some form of cooperation. We cannot cook together unless we cooperate with one another in planning and action. But, as Bratman argues, some JIAs are more cooperative than others.[14] These JIAs satisfy the additional following principle:

> (3) Commitment to Mutual Support: Each agent is committed to supporting the efforts of the other to play her role in the joint activity. If I believe that you need my help to find your note (or paintbrush), I am prepared to provide such help; and you are similarly prepared to support me in my role. These

11. *See, e.g.*, Michael Bratman, *Shared Cooperative Activity* and *Shared Intention*, *in* FACES OF INTENTION (1998). *See also Shared Intention and Mutual Obligation* and *I intend that we* J, *in* FACES OF INTENTION. I will cite to *Shared Cooperative Activity*, although much of the material is repeated in the other articles as well.
12. *Id.* at 94.
13. *Id.* at 94–5.
14. *Id.* at 103–04.

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

> commitments to support each other put us in a position to perform the joint activity successfully even if we each need help in certain ways.[15]

Bratman calls JIAs that satisfy this third principle "shared cooperative activities" (hereinafter "SCA").[16] In the cooking example, if each of us is prepared to help the other do his or her part in preparing the meal, then our cooking dinner is an SCA.

B. From Individual Intentional Activity to Joint Intentional Activity

Having set out these general principles, Bratman develops a precise characterization of JIA and SCA. He begins with an explication of the commitment to the joint activity.

According to Bratman, individuals are committed to a joint activity when each of the participants forms an intention in favor of the joint activity.[17] In the cooking example, each of us must intend that we cook dinner. Bratman allows that each participant may form these intentions for very different reasons.[18] I may want to cook dinner because I am hungry. You may want to cook dinner to brush up on your culinary skills. We each share an intention to cook, even though it is for different reasons.

As Bratman points out, certain care must be taken when characterizing the joint activities that are the objects of these intentions.[19] It would be circular if the definition of cooperative activity appealed to intentions in favor of acting together *cooperatively*. Luckily, there is a clear sense in which we can talk about joint acts in a cooperatively neutral manner. If I told you that Jones and I were on the same plane to Chicago, you would not know whether our being together on the plane was part of a cooperative venture. "Our traveling on the same plane to Chicago" is a cooperatively neutral description of a joint act because it is consistent with instances of both cooperative and noncooperative behavior. Hence, when characterizing the commitment in favor of the joint activity, the object of the intentions must be understood under a cooperatively neutral description.

Bratman proceeds to argue that JIA cannot be characterized purely by the fact that participants in a group form an intention in favor of some joint activity.[20] Assume that we each decide to paint a house together, but I want to paint the house completely red, you want to paint it completely blue, and neither of us is willing to compromise. Even if we manage to paint the house some combination of red and blue, our activity has not been cooperative at

15. *Id.* at 95.
16. *Id.* at 94.
17. *Id.* at 96.
18. *Id.*
19. *Id.* at 96–7.
20. *Id.* at 98.

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

all. Although our overall plans of painting the house agree, our individual subplans disagree. You intended to paint the house by painting it blue all over, I intended to paint the house by painting it red all over, and a house that is blue all over cannot be red all over.

Although the subplans carried out must not disagree in a JIA, Bratman does not require that they in fact agree.[21] Suppose that you intend to paint the house with an inexpensive paint and I intend to paint the house using paint bought from the corner hardware store. Our subplans neither agree nor disagree. In Bratman's terminology, these subplans "mesh."[22] Subplans in favor of a joint activity mesh just in case there exists a way to engage in the joint activity that satisfies each of the individual subplans. Your subplan of buying an inexpensive paint meshes with my subplan of buying paint from the corner hardware store, for it is possible to paint the house using an inexpensive paint bought from the corner hardware store.

In JIA, therefore, participants don't just form an intention in favor of a joint activity. Their intention is to perform the joint activity *in accordance with meshing subplans*. It should be noted that Bratman builds the "meshing subplan" condition into the content of the participants' intentions.[23] Even if individual subplans happen to mesh, the participants must also be committed to this meshing. If, midway through our painting, you decide that you want to buy the rest of the paint from the hardware store across town, I must be prepared to work with this modified subplan. I must be committed to either changing my subplan regarding the corner hardware store, convincing you that you are making a mistake, or bargaining with you by offering to alter another of my subplans so that you revert to your original subplan.

Participants can form intentions to perform the group activity in accordance with meshing subplans even though they have not identified the relevant subplans that are supposed to mesh. Their intentions simply amount to a commitment to finding and implementing such subplans. Moreover, as the last example demonstrated, participants can act cooperatively even if their subplans do not happen to mesh. Should conflict arise, the participants will be committed to deliberate, negotiate, and/or bargain in order to overcome the impasse. This is not say, of course, that this commitment is an absolute one. It may be that each party cannot accept the subplans of the other, in which case the attempt at cooperation will fail.

Commitment to the joint activity demands that the participants each form an intention to act in accordance with meshing subplans. But, as Bratman points out, still more is required.[24] Suppose that you and I decide to drive to Dubuque together. You want us to get there by my driving the car and I want us to get there by my locking you in the trunk. Although our

21. *Id.* at 98.
22. *Id.* at 99.
23. *Id.*
24. *Id.* at 100.

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

subplans mesh, my intentions are not cooperative in nature. Participants act together cooperatively only when they treat each other as intentional agents and intend that the joint activity be achieved through intentional agency. Our going together to Dubuque will not be a joint venture because your agency will play no role in our getting there. It is not enough that all participants intend to perform the joint activity in *accordance with* meshing subplans—they must also intend to perform the activity *because of* those meshing subplans.

The intentions that form the core of JIA are, therefore, strongly interdependent. In Bratman's helpful terminology, they are "interlocking."[25] The intentions interlock in the sense that they favor both a consistency of subplans as well as the efficacy of those subplans. Each participant is committed to meshing his or her subplans with the others, not only in order to ensure his or her own successful completion of the task, but also to ensure the completion of each other's contributions.

Yet interlocking intentions, while being necessary to JIA, are not sufficient. Consider cases where we each share intentions that satisfy the conditions set out above, but we do not reveal these intentions to each other. Bratman plausibly argues that activities cannot be jointly intentional unless the existence of interlocking intentions is common knowledge between the participants.[26] Summing up, these considerations provide the following characterization of JIA:

> For a cooperatively neutral J, our J-ing is a JIA only if:
>
> (1)(a)(i) I intend that we J.
>
> (1)(a)(ii) I intend that we J in accordance with and because of meshing subplans of (1)(a)(i) and (1)(b)(i).
>
> (1)(b)(i) You intend that we J.
>
> (1)(b)(ii) You intend that we J in accordance with and because of meshing subplans of (1)(a)(i) and (1)(b)(i).
>
> (2) It is common knowledge between us that (1).[27]

C. From Joint Intentional to Shared Cooperative Activity

The conditions laid out in (1) and (2) are necessary and sufficient to yield jointly intentional behavior. But they do not yet capture the higher level of cooperation characteristic of shared cooperative activity. Let us suppose that we travel to Chicago together, but only because I threaten to harm you if you do not accompany me on the plane. I will intend that we travel

25. *Id*. at 102.
26. *Id.*
27. *See id*. at 100–102.

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

together, that we travel together in accordance with and because of our meshing subplans, but our traveling together is not truly cooperative.

Even if our interlocking intentions are not coerced, JIA is not necessarily an SCA. To demonstrate this, Bratman presents the case of the unhelpful singers.[28] The unhelpful singers set out to sing a duet together. Each fully expects the other to hit his or her notes and thus intends to coordinate his or her own singing to the singing of the other. However, neither of the singers is disposed to help the other should he or she need the help. Indeed, they prefer that the other stumble rather than the duet succeeds. Because they each don't expect the other to stumble, their intentions interlock. Yet the duet is not a shared cooperative activity because they are not committed to helping the other do his or her part in the event that he or she needs the help.

How strong does the commitment to mutual support have to be in order for the activity to count as truly cooperative? According to Bratman, this commitment must entail that there is one possible scenario in which you need my help to perform your task, I can help you without undermining my own contribution to the joint activity, I have acquired no new reason to help you (e.g., you have not offered me money for the help), all of this is common knowledge, and I am willing to help.[29] If my intentions satisfy this condition, then these intentions are "minimally cooperatively stable." The intentions in an SCA must be minimally cooperatively stable.

Summing up, we get:

> For a cooperatively neutral J, our J-ing is an SCA only if:
>
> (1)(a)(i) I intend that we J.
>
> (1)(a)(ii) I intend that we J in accordance with and because of meshing subplans of (1)(a)(i) and (1)(b)(i).
>
> (1)(b)(i) You intend that we J.
>
> (1)(b)(ii) You intend that we J in accordance with and because of meshing subplans of (1)(a)(i) and (1)(b)(i).
>
> (1)(c) The intentions in (1)(a) and (1)(b) are not coerced by the other participant.
>
> (1)(d) The intentions in (1)(a) and (1)(b) are minimally cooperatively stable.
>
> (2) It is common knowledge between us that (1).[30]

D. Mutual Responsiveness

Having explicated the commitments to the joint activity and to mutual support, Bratman concludes by discussing the condition of mutual respon-

28. *Id.* at 103–4.
29. *Id.* at 104–5.
30. *Id.* at 105.

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

siveness.[31] In order for JIA or SCA to take place, it is not enough that the participants perform the joint activity with the attitudes specified in (1) and (2). Those attitudes must be appropriately connected to the performance of the joint activity.

Bratman specifies two ways in which the attitudes in question must be so connected. First, there must be mutual responsiveness of intention. During the planning of the joint activity, each participant must be attuned to the subplans of the other participants.[32] Each must be committed to adjusting his or her intentions in light of the intentions of the other members of the group or at least be committed to convincing the others to adjust their intentions so as to achieve a mesh. Second, there must be mutual responsiveness in action.[33] During the execution of the joint activity, each participant must be attuned to the behavior of the others and adjust his or her behavior accordingly so as to achieve the intended results. When we sing a duet together, I guide my behavior by listening carefully to which parts of the song you are currently singing. Bratman contrasts mutual responsiveness in action to "prepackaged cooperation."[34] In prepackaged cooperation, each participant acts according to, and only according to, the previously set plan. In a prepackaged duet, I might gauge my parts by looking at my watch, rather than my partner, and coming in at the predesignated time. In such a case, the planning was a JIA or SCA, but the singing is not.

In conclusion, a JIA results whenever we J, we share the appropriate attitudes, and these attitudes are connected to our J-ing.

> For a cooperatively neutral J, our J-ing is a JIA if and only if:
>
> (A) We J.
>
> (B) We have the attitudes specified in (1)(a) and (1)(b) and (2) [restricted to (1)(a) and (1)(b)] and
>
> (C) (B) leads to (A) by way of mutual responsiveness (in the pursuit of J-ing) in intention and action.[35]

An SCA is a JIA where the second condition is modified to:

> (B) We have the attitudes specified in (1) and (2).

E. Levels of Mesh

As Bratman notes, there are many types of activities that combine aspects of cooperation and competition.[36] Consider a game of chess. We cannot

31. *Id.* at 106–7.
32. *Id.* at 106.
33. *Id.*
34. *Id.*
35. *Id.*
36. *Id.* at 107.

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

play chess together unless we cooperate a great deal. We must each intend to abide by, and abide by, the rules and practices of chess. However, because the game is competitive, I will not try to mesh all of my subplans with yours, and vice versa. Nor will I be committed to the efficacy of all of your intentions, given that one of your principal intentions is to win the game.

Bratman suggests that we model this type of mixed situation using the idea of "levels of mesh."[37] When we agree to play chess, we do not commit ourselves to act in accordance with subplans that mesh *all the way down*. The mesh stops at the level of the rules and practices of chess. We each intend to abide by the rules and practices and are supportive of each other's attempt to abide by such commitments. To this extent, playing the game is a joint intentional (or shared cooperative) activity. However, past the point of rules and practices, it is purely competitive. Our intentions to play the game must interlock, but our respective intentions to win the game must not. A game where each player guides his behavior in order that the other may win is not a game.

V. COOPERATION AND PLANNING

A. The Importance of Plans

Like all nineteenth-century novels, all JIAs have plots. Participants always engage in their joint activity in accordance with a set of joint plans. These plans represent the steps that the participants are required or empowered to take so that they may achieve the collective end.

Why do plans play a central role in such activities? Here is Bratman's answer:

> Why do we need stable intentions and plans to support coordination? Why don't we simply figure out, at each moment of action, what would then be best given our predictions about what we and others will do in the future if we act in certain ways in the present? One answer echoes the work of Herbert Simon: We are agents with significant limits on the resources of time and attention we can reasonably devote to reasoning and calculation. Given these resource limits, a strategy of constantly starting from scratch—of never treating prior decisions as settling a practical question—would run into obvious difficulties. A second, related answer is that coordination requires predictability, and the actions of planning agents are more easily predicted.
>
> If we are to meet in Philadelphia next week, my friend needs to predict that I will be there then. The fact that I am planning to be there helps support such a prediction. His prediction that I will be in Philadelphia next week need

37. *Id.*

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

not depend on his detailed knowledge of my deepest values or on a prediction of complex calculations I will make just before leaving.[38]

According to Bratman, plans serve two main functions. First, they compensate for our cognitive limitations. For imperfectly rational agents such as us, deliberation is time-consuming, energy-depleting, and error-prone. By formulating plans in advance, agents can avoid the need down the line to deliberate about their next move. They may consult their plans in *medius res* and thereby conserve precious cognitive resources and increase the reliability of their decision-making.

Second, plans are useful in securing coordination because they render our actions more predictable. Although my actions might be reasonably foreseeable if you knew my beliefs and preferences, this information may not be available to you and, even if it were, it might be costly for you to figure out what I will do given my cognitive and conative attitudes. Plans obviate the need for this detailed information and calculation. Plans go to the bottom line of practical reasoning: By knowing what I plan to do, you can reasonably predict that this is what I will end up doing.

While Bratman is correct to single out these two functions as important, I would like to argue that plans play several additional roles in practical reasoning. In what follows, I will suggest not only that plans compensate for cognitive limitations and increase predictability, but also that they guide the uniformed, assure the uncertain and uncommitted, and solve coordination problems.

B. Guidance and Assurance

Compensating for cognitive limitations is one way in which plans guide agents. But it is not the only way. Suppose I had no idea how to paint a house. In my state of ignorance, I did not know that one had to scrape the old paint off before new paint could be applied. Being told by my more informed painting partner that scraping is needed before painting would be quite helpful. This plan would guide my behavior and enable me to contribute to our joint project of painting the house.

In contrast to the mode of guidance Bratman discusses, this form of guidance has nothing to do with bounded rationality. My ignorance of how to paint a house is not a cognitive defect. The plan guides my behavior by informing me about what is needed in the circumstances. It acts much like expert advice—by relying on it, I do better than if I tried to go it alone.

Because plans are capable of guiding participants, they are capable of assuring them as well. When my painting partner tells me that I am supposed to scrape before he paints, and I adopt this plan, he will be assured that I know what to do in order for us to achieve our goals. This assurance

38. *Id.* at 36–7.

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

may then serve as sufficient inducement for my partner to continue investing in the project and carrying out previously adopted subplans.

Assurance is important even when every one knows what to do. My partner and I may have painted many houses before. Yet I may not know whether he wants to paint this house blue or red. I might care so much about painting this house blue that I would not have agreed to paint the house at all if we had to paint it red. My partner's adoption of the subplan to paint this house blue would tell me something important about his preferences and this might then induce me to continue with the project.

The modes of assurance just described work in almost the exact opposite manner to that mentioned by Bratman. According to Bratman, plans facilitate coordination by enabling agents to predict the behavior of others without knowing their beliefs and desires. Here, plans help organize behavior precisely by allowing agents to infer the beliefs and/or desires of the planners and then to form a prediction based on such information.

Each type of assurance has its advantages. Bratman's method allows agents to coordinate without having to infer beliefs and desires and without having to perform any calculations. The method I mentioned may be more cognitively demanding, but it may also allow agents to form more accurate predictions. For example, I may not have any information about how resolute my new painting partner is. But if I know that he knows how to paint and really wants to paint, then I may become very confident that he will end up helping me paint the house.

It is important to note that plans may play a role in assuring participants even when everyone knows what to do and what each prefers, and when all of this is common knowledge. Suppose I hate painting (I am allergic to the fumes) and you hate scraping (you can't stand the noise). It will be rational for me to scrape if and only if you paint, and it will be rational for you to paint if and only if I scrape. It will not, therefore, be rational for either of us to commit to perform our preferred jobs unless we know that the other will perform his or her preferred job. One solution to such a problem is for one of us to commit to a plan and to manifest that commitment to the other person. If I show my painting partner that I intend to scrape, he or she will be induced to form an intention to paint. Through the unilateral commitments to plans, assurances are generated and stalemates are ended.

C. "Cooperation" Problems

Plans serve another important function that has not yet been mentioned: They solve coordination problems. I may not care about scraping or painting, and neither may you. Yet we have to decide who is going to paint and who is going to scrape. A plan that allocates scraping to me and painting to you will serve to mark off one of the two arbitrary solutions. Both will have

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

a reason to act on this plan, given that they believe that the other will similarly act on this plan and this is common knowledge.

As we have seen, however, plans might be useful even when no coordination problems exist. I may refuse to paint and you may refuse to scrape. Still, plans are useful in such cases, for I may not yet know that scraping is necessary for painting. Or I may tend to forget this crucial piece of information during the long hours of our working together. Or I may need to assure you that I will scrape before, and not after, you paint.

Plans, therefore, serve three main functions: they guide, assure, and solve coordination problems. Because it is helpful to have one word to describe these three roles, we will say that plans "organize" the behavior of the participants. We will also say that plans solve "cooperation" problems. A cooperation problem is a game where each player has an interest in working together in order to achieve some joint outcome and the players increase the likelihood of achieving that outcome through the organization of their behavior. Plans help solve cooperation problems because they are capable of organizing through their guidance, assurance, and coordinative functions and hence increase the likelihood that the sought after outcome will be achieved.

VI. JOINT INTENTIONAL ACTIVITY WITH AUTHORITY

In the previous two sections, I set out Bratman's account of JIA and SCA as well as his theory of plans. Eventually, I would like to argue that legal practice is a JIA and always holds itself out as an SCA. This will then allow me to show that legal conventions are plans that organize such activities and enable participants to achieve the sought-after outcomes.

However, more work needs to be done before this demonstration can be performed. As Bratman himself notes, he developed his theory by focusing on joint activities that did not involve authority structures.[39] In this section, I would like to show how his models of JIA and SCA can be extended to cover these sorts of cases as well.

A. Can Bratman's Model Be Extended?

As a general matter, when adults cook dinner together, sing a duet together, or fly to another city together, neither has authority over the other. If while preparing dinner I order you to "fire" the chicken, our cooking together would probably end then and there. There are many cases, however, where one member of the group has authority over others and yet it is entirely appropriate to consider their joint activity a cooperative venture.

Imagine that a group of friends decides to sail from New York to Nova

39. *Id.* at 94. See also *id.* at 110.

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

Scotia. Having sailed together once before, they remember the problems that arose the last time they set out. They recall that during storms, despite their best efforts, they were unable to coordinate their behavior effectively, and as a result, everyone's safety was seriously put at risk. In addition, the friends were generally indifferent between stopping in Maine and proceeding directly to Nova Scotia, and this collective indifference made the decision difficult to make, causing substantial delays. In order to prevent the same problems from arising again, they resolve to appoint a group member as the captain for the trip. They figure that the trip will be faster, safer, and more pleasant if one is designated as having authority over the others.

Let us assume that their first trip was an SCA. I think it is reasonable to assume that the mere appointment of an authority would not preclude their second trip from being an SCA as well. That the new captain has authority over her crew is not inconsistent with their cooperating with one another in order to achieve a common goal. Indeed, they set up the authority structure precisely because they thought that they could not achieve their goal, or achieve it as well, without hierarchy.

At first glance, however, it would seem that this second trip could not be an SCA or even a JIA. Consider the commitment to a joint activity. As we saw, participants are committed to a joint activity only when each intends to act in accordance with subplans that mesh. However, this condition appears to fail in joint activities involving authority. The captain of the boat does not intend to mesh her subplans with those of her crew—rather, they are supposed to mesh their plans with her. Or consider the commitment to mutual support. Let us assume that the captain assigns to one of the crew the task of rigging the sail. If the captain believes that the crew member has done a poor job, she might order another to take the sail down and re-rig it. Rather than helping the first crew member do his job, the captain will have undermined his contribution to the joint project. Finally, even mutual responsiveness does not seem to be present in the second voyage. While the crew must be responsive to the captain, the captain does not have to be responsive to the crew. Such is the nature of hierarchy: Responsiveness to intention and action need not be mutual. It must go from top to bottom, not bottom to top.

These incompatibilities, though, are only minor or nonexistent. As I will argue, the conditions that Bratman imposes on joint intentional activity and shared cooperative activity are either consistent with the existence of joint intentional activities involving authority or shared cooperative activities involving authority (hereinafter "JIAA" and "SCAA"), respectively, or must be slightly modified in order to achieve the requisite fit. In addition to showing the adaptability of Bratman's general model, I will set out further conditions that will distinguish an ordinary JIA and SCA from one where authority relations are also present. This will pave the way for my proposal that the law is an instance of a JIAA and holds itself out as an SCAA.

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

B. Orders and Meshing Subplans

Let us begin with the commitment to joint activity. Upon reflection, it should be apparent that the captain of the boat is committed to acting in accordance with meshing subplans. Given that she wants everyone's behavior, including her own, to be organized in a certain way, she must ensure that everyone's subplans, including her own, mesh. If the captain believes it preferable to rig the sail herself, she must find out whether anyone else intends to rig the sail as well. More importantly, if the captain has decided that the boat will dock in Maine, she must guarantee that this decision meshes with everyone else's plans. The captain courts disaster if she plans for the crew to dock the boat in Maine but they plan to dock it in Nova Scotia.

That the captain is committed to meshing subplans does not mean, however, that she is committed to revising *her* plans in case of conflict. To be sure, the captain may adjust her plans to mesh with those of her crew; she may also negotiate and bargain with them in order to persuade them to revise their subplans, rather than vice versa. What is special about the captain's position, though, is that she possesses a tool for achieving a mesh that her crew does not have: She may *order* her crew to act as she intends. If the captain exercises her authority, her crew will be committed to revising their subplans so that they mesh with the captain's. Once orders are given, it will then be inappropriate for the crew to negotiate and bargain with the captain—to do so would be a slight to the captain's authority.

In a JIAA, therefore, every participant is committed to acting in accordance with meshing subplans. Authorities and subjects are distinguished, however, by what those commitments require them to do, given their differing roles. Whereas authorities have the option of revising their subplans or issuing orders, their subjects do not have that luxury. If they are ordered to act by someone having authority over them, they must revise their subplans just in case their intentions do not yet mesh.

As this discussion suggests, there is a significant difference between an authority's orders and its intentions. We might say that orders are devices that an authority uses in order to enforce its intentions.[40] When someone in authority issues an order, that person intends the order to be a reason for the subject (1) to adopt the content of the order as his or her subplan; and (2) to revise his or her other subplans so that they mesh with the order. If the captain orders the crew to head towards Maine, the captain intends every crew member to treat the order as a reason to adopt her plan as their plan, as well as to revise their other plans so that they are consistent with this newly adopted plan.

40. There are other differences between an order and a mere intention. For example, an order must be expressed in order for it to be an order. An intention, on the other hand, is a mental act that need never be expressed publicly.

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

The commitment to the joint activity in an JIAA can now be characterized as follows:

> For a cooperatively neutral J, our J-ing is a JIAA only if:
>
> (1) J is a JIA.
>
> (2) If one of us has authority over the other in J, then:
>
> (a) The authority intends his or her orders to be reasons for the subject to adopt the content of the order as a subplan as well as reasons to revise the subject's subplans so that they mesh with the orders.
>
> (b) The subject intends the authority's orders to be reasons to adopt the content of the orders as subplans as well as reasons to revise his or her subplans so that they mesh with the orders.
>
> (c) (a) and (b) are common knowledge.
>
> (3) Either I have authority over you in J or you have authority over me in J.

This formulation explicitly relativizes the authority relation to the joint activity, i.e., someone has authority over another *in J*. This relativization is necessary because someone can have authority in one practice but fail to have it in another. Maritime law may, for example, recognize the most experienced sailor as the captain. The most experienced sailor will have authority in law but not in the friends' sailing venture.

One limitation of this formulation, however, is that it takes the authority relation as primitive. It does not show us how to determine whether someone has authority over another, only that certain inferences can be generated from such a determination. In principle, though, it should be possible to construct a definitive test for establishing whether someone has authority in a joint intentional activity. After all, authority relations in a JIAA are created by that very activity. The captain became the captain of the boat because the friends decided, during the planning phase of their trip, that one will have the power to regulate certain activities on the boat. The test for determining who has authority over whom should track the process by which authority relations are generated.

How, then, are authority relations generated from JIA? Here's a hypothesis: Someone comes to have authority over another in a JIA just in case he intends his directives to be conclusive reasons for the others to adopt and revise certain subplans and the others intend to respond to these directives in the manner intended. Authority arises, in other words, just in case the intentions of the participants "vertically" interlock. We can capture this idea by converting the "if . . . then" in condition (2) to an "if and only if" and treating it as an explicative definition of "authority in J."

> (D) If our J-ing is a JIA, then one of us has authority over the other in J if and only if:

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

(a)(i) I intend my orders to be reasons for you to adopt the content of the order as a subplan as well as reasons to revise your subplans so that they mesh with my orders.

(a)(ii) You intend my orders to be reasons to adopt the content of my orders as subplans as well as reasons to revise your subplans so that they mesh with my orders.

(a)(iii) You comply with my orders most of the time.

OR

(b)(i) You intend your orders to be reasons for me to adopt the content of the order as subplans as well as reasons to revise my subplans so that they mesh with your orders.

(b)(ii) I intend your orders to be reasons to adopt the content of your orders as subplans as well as reasons to revise my subplans so that they mesh with your orders.

(b)(iii) I comply with your orders most of the time.

(c) Either (a) or (b) is common knowledge.

For a cooperatively neutral J, our J-ing is a JIAA only if:

(1) J is a JIA.

(2) Either I have authority over you in J or you have authority over me in J (as defined by (D)).

C. Commitment to Mutual Compensation

At the beginning of this section, I questioned whether the commitment to mutual support must be present in every SCAA. I set out the case of the captain who initially ordered one of the crew to rig the sails and, after the task was performed incorrectly, ordered another member of the crew to re-rig the sails. In this example, not only did the authority not help the first crew member perform his task, but she prevented him from ever being able to complete it.

I believe that this counterexample is effective in showing that participants in an SCAA need not be committed, even in a minimal way, to helping each other complete their tasks. However, I would like to reformulate the commitment to mutual support in such a way that it is plausible to believe that a similar, but more general, principle is present in every SCAA.

Recall why Bratman originally introduced the commitment to mutual support. He presented the case of the unhelpful singers who believe that each will perform their parts in the duet effectively but are not disposed to help each other should one need the other's help. Bratman persuasively argued that the unhelpful singers would be engaged in a joint, but not a shared cooperative, activity.

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

Similarly, we can imagine the case of the resentful captain. This captain has full confidence in his crew. He believes that when he gives an order, not only will his crew obey, but also they will carry out the directive perfectly. However, this captain has a bad case of *ressentiment*: He deeply resents the excellence of his crew and prefers to see them fail. If it ever turned out that one of the crew bungled his task, the captain would take no corrective action. He would neither try to help the crew member to complete his task nor reassign other crew to undo the damage.

The voyage of the resentful captain and his crew is not a shared cooperative activity. It is not a SCAA because the captain has no commitment to ensuring that the tasks assigned are performed correctly. This suggests the following general principle applicable to all SCAA:

> Commitment to Mutual Compensation: Each agent, unless otherwise directed, is committed to compensating for the lack of information or ability of the other in playing her role in the joint activity. This compensation may take the form of helping others complete their tasks or, in the case of an authority, reassigning tasks to others who are able to perform the tasks adequately.

The commitment to mutual compensation is more general than the commitment to mutual support. If I help you perform your task in the joint enterprise, then *ipso facto* I am doing something in order to compensate for your deficits. However, if I, as an authority, reassign someone else to perform the task, then I am compensating for your deficits but I am not thereby helping you perform your task. In fact, by reassigning your task to someone else, I guarantee that you will have failed to do your job.

In the context of an SCAA, the greater generality of the commitment to mutual compensation is desirable because, as the first counterexample showed, there is no reason to believe that authorities must be committed to helping their subjects fulfill their tasks. Authorities are generally responsible for the success of the joint enterprise, or some portion of it. They may judge it extremely inefficient to cure the problems brought about by one participant through directly helping that one participant. Reassignment of tasks may be thought to be the best way to ensure the success of the project.

The commitment to mutual compensation admits of one exception. Members of the group may be prohibited from aiding fellow participants by those in authority. The captain, for example, might forbid the crew from helping each other in the belief that it is best for the crew to stay focused on their jobs and not worry about the progress of their fellow sailors. In such a case, the commitment to the joint activity would preclude the crew from helping each other. Hence the commitment to mutual compensation cannot require compensatory behavior.

If I am right that the commitment to mutual compensation is necessary to distinguish joint activities with authority from SCAA, then Bratman's definition of minimal cooperatively stable intentions must be modified. Recall that, according to Bratman, intentions are minimally cooperatively stable if there is one possible scenario in which you need my help to

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

perform your task, I can help you without undermining my own contribution to the joint activity, I have acquired no new reason to help you (e.g., you have not offered me money for the help), all of this is common knowledge, and I am willing to help. In light of the commitment to mutual compensation, this characterization should be changed to:

> (M) P_i's intention is minimally cooperatively stable if and only if for each participant P_j (where $i \neq j$) there is one possible scenario in which:
>
> (i) P_j lacks some information or ability to perform his task.
>
> (ii) P_i is able to compensate for this lack so that there will be a greater probability that the participants' task will be successfully completed (possibly by some other participant).
>
> (iii) P_i can compensate for this lack without undermining his contribution to the joint activity.
>
> (iv) P_i has acquired no new reason to compensate for this lack.
>
> (v) P_i has not been directed not to compensate for this lack.
>
> (vi) (i)–(v) are common knowledge.
>
> (vii) P_i is prepared to compensate.

D. Coercion

In Bratman's model, a JIA cannot qualify as an SCA if some participants successfully coerce others. Bratman imposed this condition because he believed that coercion is inimical to true cooperation. This requirement, however, is too strong, at least as it applies to situations involving authority. For it is often the case that authorities use the threat of coercion precisely in order to facilitate cooperation.

For example, during a recent faculty meeting at my school, the faculty voted overwhelmingly to grant the dean the authority to impose a series of fines on individuals who handed in their grades late. During the meeting, several reasons were set out to support the measure. Many endorsed the measure because they thought it would serve as an effective precommitment strategy: By giving the dean authority to impose sanctions, these faculty members sought to bind their future selves so that they would comply with the deadlines. Others wanted to send a message to those who routinely and brazenly flout the rules. In condemning such irresponsible behavior, they tried to promote a greater ethic of professionalism among the faculty. Still others wished to support those who did follow the rules, by assuring them that the others would not free-ride on their diligence.

It would be unfair to the faculty to see the grant of penal authority to the dean, or his subsequent exercise of it, as being an essentially uncooperative act. If anything, these acts were intended to foster a spirit of cooperation. Everyone recognized that he or she had an obligation as a faculty member

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

to hand in his or her grades on time. The fines were designed to ensure that the faculty lived up to its acknowledged obligations.

Authoritative coercion is consistent with cooperation, I would submit, when it is treated as a mere backup strategy. Members of the group must recognize that coercion follows only when they have failed to live up to their previous uncoerced commitments. This suggests that condition 1(c) of Bratman's model must be amended to the following:

> (1)(c) The intentions in (1)(a) and (1)(b) are not coerced by the other participant, unless coercion is understood by the participants to be a method for enforcing uncoerced intentions.

E. Mutual Responsiveness

It should be clear by now that authorities are as responsive to the intentions and actions of their subjects as subjects are responsive to the intentions and actions of authorities. The function of authorities in JIAA and SCAA is to ensure that the participants' actions are organized so that their common goals can be achieved. This will require that authorities be attuned to the intentions and actions of the participants, modifying their own subplans in some cases and ordering the revision of their subjects' subplans in others. Moreover, they must be sensitive to the success or failure of their subjects' actions so that compensatory action may be taken if judged appropriate.

There is one important sense, however, in which the participants in JIAA and SCAA are not mutually responsive to one another. When an authority exists within a group, subjects may look towards the authority in order to coordinate their behavior rather than taking their cues from each other. Moreover, when orders are issued, the directives may remove discretion from subjects on how to proceed. Depending on the scope of authoritative regulation, aspects of a JIAA or SCAA may start to resemble prepackaged cooperation, insofar as the subjects will be forced to respond to the rules laid down rather than to the each other's intentions and actions.

F. Generalizing the Definition

The definitions of JIAA and SCAA we have been constructing apply only to the simple case of a two-person JIA. If they are to apply to cases of legal authority, it will be necessary both to generalize and to relax the definitions in several ways.

First, the JIA definition should be expanded so that it can accommodate *n*-person JIAs. Second, the definition should allow for "impersonal" authority relations. That is, instead of participants committing themselves to following the orders of a particular person, they should be able to commit themselves to the directives of a person who satisfies certain qualifications

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

(e.g., the friends might have a rule that confers captain status on the most experienced sailor).[41]

Third, the definition should relax the common-knowledge requirement. Such a strict condition is problematic, for as the size of a JIAA grows, it becomes increasingly unlikely that everyone will know the contents of everyone else's intentions. Accordingly, it will be best simply to require that the contents of each other's intentions be "publicly accessible," that is, that the contents of each others' intentions should be accessible in some way to each of the participants. This condition weeds out cases where the participants refuse to disclose to each other their intentions or are intentionally deceitful as to their content.

Fourth, a generalized definition should not require commitment to the joint activity, or mutual responsiveness, from all participants. Only JIAs involving the smallest groups could pass a test of universal commitment, for it is inevitable that some participants will be apathetic, lazy, misguided, rebellious, heroic, evil, or on the take and will therefore not be committed to the joint activity. The only requirement that should be imposed is that most participants are so committed. The "most" is intentionally vague, as the concepts being explicated are vague in just this way.

Finally, the definition should be weakened so as to accommodate the existence of alienated but cooperative participants. As Chris Kutz has argued, Bratman-like accounts of JIA are inapplicable to cases of large-scale cooperation, for they require that participants intend that their group bring about the collective outcome.[42] However, it is often the case that participants will intend to contribute to the group effort but not care at all about whether their group is successful. Suppose I am hired by Microsoft to work on a new version of its operating system. Because I get paid only if I contribute to the programming effort, I intend to contribute. However, I may not care at all if the software group is successful, given that I am paid regardless of whether the group is successful. To use Kutz's terminology, as an alienated worker, I have a "participatory" intention to contribute to the collective project but I don't have the "group" intention that such a project be successful.

The generalized definitions of JIAA and SCAA are set out below:

For a cooperatively neutral J and a group G, the J-ing of G is a JIA if and only if

(1) G J's.

(2)(a)(i) Most participants intend to contribute to G's J-ing.

(2)(a)(ii) Most participants intend to contribute to G's J-ing in accordance

41. Unfortunately, the generalization I shall offer will not be as general a definition as one might wish. It will not allow for the characterization of the authority that corporate bodies such as boards of directors and legislatures might have over other corporate bodies such as business associations and governments. To do so would require treating these corporate entities as JIAAs and SCAAs and defining "authority in J" in these terms. Although an interesting exercise, I shall not carry it out as it would be somewhat complicated and would add little to the analysis that follows.

42. *See* Chris Kutz, COMPLICITY 92 (2000).

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

with and because of meshing subplans of those participants that similarly intend to contribute.

(3) (2) is publicly accessible to most participants.

(4) The attitudes in (2) lead most participants to contribute to J by way of mutual responsiveness in action.

For a cooperatively neutral J and a group G, the J-ing of G in an SCA if and only if

(1) G J's.

(2)(a)(i) Most participants intend to contribute to G's J-ing.

(2)(a)(ii) Most participants intend to contribute to G's J-ing in accordance with and because of meshing subplans of those participants that similarly intend to contribute.

(2)(b) The intentions in (2)(a) are not coerced by any of the other participants, unless coercion is understood by most participants to be a method for enforcing uncoerced intentions.

(2)(c) The intentions in (2)(a) are minimally cooperatively stable, as defined by (M).

(3) (2) is publicly accessible to most participants.

(4) The attitudes in (2) lead most participants to contribute to J by way of mutual responsiveness in action.

(E) For a cooperatively neutral J and group G of participants P, if G's J-ing is a JIA or an SCA, the P_i has authority over P_j in J if and only if

(1)(a) Most participants who satisfy qualifications $QA_1, \ldots, QA_m$ intend their directives to be reasons for participants who satisfy qualifications $QS_1, \ldots, QS_n$ to adopt the content of their directives as subplans as well as reasons to revise their subplans so that they mesh with the directives.

(1)(b) Most participants who satisfy qualifications $QS_1, \ldots, QS_n$ intend the directives of participants who satisfy the qualifications $QA_1, \ldots, QA_m$ to be reasons to adopt the content of their directives as subplans as well as reasons to revise their subplans so that they mesh with the directives.

(2) (1) is publicly accessible to most participants.

(3) P_i satisfies $QA_1, \ldots, QA_m$.

(4) P_j satisfies $QS_1, \ldots, QS_n$.

(5) Most participants who satisfy qualifications $QS_1, \ldots, QS_n$ comply most of the time with the directives of those participants who satisfy qualifications $QA_1, \ldots, QA_m$.

For a cooperatively neutral J and a group G, the J-ing of G is a JIAA (SCAA) if and only if

(1) The J-ing of G is a JIA (SCA).

(2) Some participant has authority over another participant in J, as defined by (E).

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

VII. THE NATURE OF THE COOPERATIVE RELATIONSHIP

A. Intentions, Beliefs, and Desires

Like Bratman's accounts of JIA and SCA, the model presented in the previous section characterizes JIAA and SCAA in terms of the intentions of the participants. Due to the fact that JIAA is a type of JIA, the participants must intend to contribute to the collective project and these intentions must lead to a mutual responsiveness of planning and action. And because SCAA is a type of SCA, the participants must also intend to compensate for each other's weaknesses. Furthermore, JIAA is a type of JIA and SCAA is a type of SCA in which certain individuals intend to provide reasons for other individuals by issuing directives and the latter intend to treat the formers' directives as reasons. A practice counts as a JIAA or SCAA, therefore, only when the intentions of the participants interlock both horizontally and vertically.

By focusing solely on shared intentions, this modified Bratmanian account abstracts away from the particular beliefs and desires of the participants. Such a blinkered approach has considerable merit, for it neatly captures the fact that cooperation is possible even in the face of considerable disagreement about facts and values. Members of a group need not agree on every matter related to the sought-after outcome in order to work together in order to achieve it. Each participant can intend to engage in a joint activity but do so for very different reasons.

Although JIAA and SCAA are defined solely by reference to a practice's intentional structure, it would be a mistake to assume that the beliefs and desires of the participants are irrelevant to a proper understanding of these joint activities. To the contrary, not only do these attitudes affect the future direction of cooperation, for example, how conflicts will be resolved, which new subplans will be adopted, and so on, but they also determine the way in which participants understand their present cooperation. If, for example, members of a group are united in their beliefs and desires, then they will see their venture to be a strongly cooperative one. However, if their attitudes are highly divergent with respect to means and ends, then they will inevitably interpret their joint activity to be weakly cooperative, as a mere compromise sought for mutual advantage. Some cooperative relationships, it is fair to say, are more "cooperative" than others.

To some extent, of course, crossing the line between JIAA and SCAA signifies that there has been an improvement in the cooperative relationship. Participants who are committed to compensating for each other's weakness are more cooperative than those who are not. Yet, even within these categories, the nature of the cooperative relationship can vary considerably. As I will attempt to show, every cooperative relationship has its own unique character, and this character is determined by certain of the attitudes of the participants. In particular, I will claim that the character of the cooperative relationship is a function of the degree of competition that

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

exists between the participants and the level of trust that they share. I will then go on to suggest that this character actually plays an essential role in the design of JIAAs and SCAAs.

B. Competition

Instead of friends forming a crew and setting sail, let us assume that a group of strangers assembles in order to charter a boat from New York to Nova Scotia. As they start planning the trip, though, they realize that their preferences conflict a great deal. Some of them desperately want to stop over in Maine, while others strongly desire to go directly to Nova Scotia. Some prefer a two-hour lookout schedule, others prefer a four-hour one. Some refuse to cook, while the others insist on doing no more than an equal share of cooking. After much squabbling, it becomes clear to all not only that they will be unable to resolve these issues now, but also that other disagreements will arise during the trip and that resolving them at that point will prove to be even more intractable. Accordingly, they decide to appoint one member of the group as the captain and entrust him with authority to resolve these disputes.

Assuming that a captain is appointed and that the members of the group are mutually responsive and committed to the joint activity, mutual compensation, and the authority structure, the charter trip will be an SCAA. However, this SCAA will be significantly less cooperative than the one involving the friends. In the friends scenario, their preferences were highly convergent and, in those situations where their wants diverged, they were able to resolve their disagreements through negotiation and bargaining. The appointment of their captain was not motivated by the need to overcome conflict but rather to coordinate behavior in difficult situations and to break ties. By contrast, the strangers require authoritative intervention precisely because they are unable to resolve their disagreements through negotiation and bargaining. The function of the captain is to compel a mesh that would not come about otherwise.

The difference in the cooperative relationship between the two scenarios is not usefully captured by reference to differing "levels of mesh." In both scenarios, the participants are committed to meshing their subplans "all the way down." The contrast, though, resides in the method that each group is willing to use in order to achieve this mesh. The friends are inclined to negotiate and compromise with another, whereas the strangers are not similarly disposed, but instead are committed to deferring to the captain so as to achieve the requisite mesh.

The strangers case reveals one way in which JIAAs and SCAAs may be highly competitive: that is, where participants disagree about the ends to pursue. Conflict may also arise with regard to the means that participants wish to employ to achieve agreed-upon ends. These forms of competition may play themselves out either through lobbying authorities to resolve disagreements in a certain direction or, in practices where authorities are elected, through voting for candidates who possess the desired preferences or beliefs.

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

Not only may subjects compete with one another with respect to means and ends, but authorities may also be similarly competitive with one another. Authorities of the same rank may jockey among themselves for greater power or influence. Authorities may try to be the first to set policy in the direction that they favor and, where permitted, they may seek to distinguish or overturn precedent and past decisions.

The lines of competition, though, are not limited to the horizontal dimension but may extend vertically as well. Not only may subjects be competitive with one another, but also they may be competitive with authorities as well. Authorities and subjects may have differing preferences and/or beliefs about which subplans are desirable. Authorities with different amounts of power may also be competitive with each other and may struggle with one another with respect to the direction that the group takes in order to achieve their goals.

C. Trust

It is a truism that cooperative ventures cannot proceed unless the members of the group trust each other, at least to a minimal extent. In order for cooperation to be rational, participants must have a certain degree of confidence in each other's competence to complete their assigned tasks. Moreover, they must have a fair degree of conviction in the goodwill of their fellow members to shoulder their burdens. This truism is even truer with respect to cooperative ventures involving authority. Subjects will comply only when they have some faith that those in authority will use their power responsibly; likewise, authorities will issue directives only when they trust subjects to heed their directives. Without trust between authority and subject, the intentions of the participants will never vertically interlock; if trust is lost, their intentions will decouple.

Although some minimal level of trust is necessary for their existence, JIAAs and SCAAs can exhibit a great deal of variation on how trusting their participants are towards each other. Subjects and authorities may trust each other implicitly or they may be rather suspicious of each other. Subjects may trust one another but distrust those in authority. Or they may trust the competence of those in authority but distrust their goodwill. Authorities may have confidence in some of their subjects but not in others. Higher-ranking authorities may trust each other but have an unfavorable opinion of those in the lower ranks. Subjects and authorities may trust each other on certain matters but not on others. The set of possible distributions of trust and distrust is infinite.

Like the level of competition, the degree of trust present in a group will profoundly affect how the participants understand their cooperative relationship. A practice that involves a high degree of trust will be viewed as an activity among equals or betters that is marked by excellence and goodwill. On the other hand, if the practice is marked by a high degree of distrust,

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

the association will be seen as a necessary evil, one that is tolerated for lack of a better alternative.

Often there is a direct connection between the level of competition and the degree of trust in a group. Strong competition can breed distrust, while convergence of interest may similarly generate much trust. However, these two factors also operate as independent variables. It is possible, for example, to distrust those with whom one shares much in common. The friends may want to sail to Nova Scotia together, but the captain and crew may be inexperienced and everyone may know this. Conversely, highly competitive individuals may trust each other to act in ways that further their collective self-interest.

D. Reading the Cooperative Relationship

As I have argued, a description of a JIAA or SCAA that focused exclusively on shared intentions would be incomplete. By omitting references to the beliefs and desires of the participants, the report would neglect the very factors that determine the specific character of the cooperative relationship for that joint activity.

Such an omission, however, is not as consequential as it might first appear. In general, it will be possible to reconstruct the nature of the cooperative relationship from a description of the participants' shared intentions. This is so because participants will generally take their cooperative relationship into account when structuring their JIAA or SCAA.

Consider, for example, the level of competition in a group. If the members of the group are very competitive, they may find it highly desirable to give those in authority broad power to decide disputes. To be told, then, that members of a group have invested broad arbitration powers in those in authority would provide one with evidence that their cooperative relationship is competitive. To be sure, this would be weak evidence, as there are explanations for this authority structure that are compatible with a highly cooperative relationship. But if the level of competition is indeed high, we would expect to see other manifestations of this fact in other plans that the members adopt.

VIII. LEGAL PARTICIPANTS AND LEGAL ACTIVITY

In the last several sections I argued that certain practices involving authority should be regarded as cooperative in nature. When authorities and subjects are committed to working together so as to achieve a certain outcome and are prepared to adjust their behavior in light of each other's behavior, then it is appropriate to consider such practices as JIAs. The modified Bratmanian models of JIAA and SCAA I developed are meant to isolate these important types of joint activity.

In the next several sections, I will argue that legal practice is always a JIAA and that it always holds itself out as an SCAA. In order to establish this claim,

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

I will attempt to demonstrate that the participants in every legal system generally conform to the model of JIAA and hold themselves out as though they conformed to the model of SCAA. More specifically, I will try to show that legal participants are committed to the legal activity by way of meshing subplans and are mutually responsive to each other's intentions and action and that they claim that they are uncoerced and are committed to mutual compensatory behavior.

But before any of this can be done, I must address a few preliminary matters. First, I need to identify the individuals who constitute the class of legal participants. Second, I must describe in some detail the joint activity to which these participants intend to contribute.

A. The Participants

If the claim that legal practice is a JIAA has any chance of being true, then the class of participants will have to be restricted to legal officials. This is so because there are legal systems where large segments of the population conform to the law out of nothing more than fear of punishment and evade the law if and when they can. While such individuals are responsive to the directives of those in authority, it would be a cruel joke to interpret their behavior as motivated by a commitment to the joint activity.

To be sure, many ordinary citizens consider the law to be legitimate and willingly obey those in authority. However, even in these cases, it strains ordinary usage to describe these individuals as "participants in the legal system." The business of law is the creation and exercise of authority, and its participants are the officials who operate the levers of legal power. Compliant citizens are best regarded as cooperative participants in the social life of the community, and if these citizens are numerous enough, the social life of that community will count as a JIAA.[43]

Although ordinary citizens are not, strictly speaking, participants in legal practice, it is the conceit of the law that they are its intended beneficiaries. Governors generally claim that they create and administer law in order to serve the governed. This does not, of course, mean that citizens are the intended beneficiaries, only that this is what they are told. And it should be noted that some segments of the population may be so powerless that the authorities do not even bother to lie to them.

It might be objected that the exclusion of nonlegal officials from the class of legal participants makes it impossible for legal practice to be a JIAA. According to (E), participants have authority just in case other participants recognize that authority. Authorities require subjects in order to count as authorities, and by restricting the class of legal participants to officials, it would seem as though I am effacing the very individuals who are responsible for the existence of legal authority.

43. The line between ordinary citizen and legal participant may not always be so crisp. In democracies, for example, citizens are generally entitled to vote for officials and hence they do participate in the operations of the legal system.

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

This objection fails, however, because it neglects the possibility that legal officials can be subject to each other's authority. Judges, for example, are subject to the authority of the legislature, for they are required to apply the rules enacted by them. Indeed, the legislature has the authority it has precisely because judges are committed to treating their directives as reasons for action and modifying their intentions and actions accordingly.

B. The Joint Activity of Law

Having identified the class of legal participants, it behooves me to say a bit about the kind of activity that is legal practice. If most participants in a JIA intend to contribute to their group's J-ing, where J is some joint activity, what is the "J" of law?

One might be tempted to say that the joint activity characteristic of law is the creation of a legal system if one does not yet exist, or the maintenance of a legal system if one already exists. This will not do, however, as it is circular. We cannot say that the practice of law creation and administration is a type of JIAA and then define this type of JIAA using the concepts of law creation and administration.[44]

A better approach is to identify the joint activity as the creation or maintenance of a system of rules. This breaks the circle, as the description of the activity does not use any legal concepts. However, this description is not yet adequate because it does not specify the kind of system of rules that participants create or maintain. After all, there are many types of normative systems in the world (morality, etiquette, games, religion, etc.), and law is just one of them. To understand the nature of legal practice, therefore, we must identify the specific type of normative system that legal participants bring about.

Unfortunately, detailing the essential features of a legal system is no trivial matter. Whole books have been written on this subject.[45] Rather than enter into this difficult subject, let me mention two properties instantiated by every legal system and which help distinguish them from many other types of normative systems.

As Hart argued in *The Concept of Law*, legal systems are not simply collections of rules for the guidance of conduct in a certain community. Legal systems possess a characteristic unity—every rule in the system must pass the same tests for admission, the so-called "criteria of legal validity." In the United States, for example, all legal rules must satisfy the conditions set out in the United States Constitution. A rule that satisfies such conditions is a member of the system, and those that do not are not members of the system.

Hart, of course, made a stronger claim. He argued that the criteria of legal validity are set out in a conventional rule that he called the "rule of

44. A legal system is a system of rules where law is created and administered.
45. *See, e.g.*, Joseph Raz, The Concept of a Legal System (2nd ed. 1980).

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

recognition."[46] I don't want to assume the existence of such a convention, however, given that this is the proposition that I am trying to prove. I will only assume the weaker claim that legal participants intend that the group create or maintain a "unified" system of rules. Whether they look to a conventional rule or to the principles of morality in order to determine which rules to apply is, for now, an open question.

How exactly do officials create or maintain a unified system of rules? First, I will assume that, in order to achieve this goal, officials must look to the same tests for validity and apply those rules that pass such tests. The commitment to do one's part in the creation or maintenance of a unified system of rules, therefore, entails the commitment to look to the same tests as other officials do and to apply the rules that pass those tests. The resulting system will be unified in two senses. First, the rules of the system will pass the same tests for admission. Second, the officials of the system will look to the same rules when guiding their own conduct and evaluating the conduct of others.

However, the mere appeal to the same set of rules does not yet suffice to create *one* system of rules—it may create *many* systems of rules. If French legal officials suddenly decided to follow the United States Constitution, but did not defer to United States courts in constitutional interpretation, the French would still possess their own legal system. Like identical twins, the French and American systems would have the same genetic makeup, and might appear exactly the same, but would be two separate entities.

In order for legal officials to be part of the same legal system, and for them to create and maintain one system of rules, it is not enough that they look to the same set of rules. They must also be committed to resolving disputes that might arise concerning the membership of this set or the proper application of any of its elements. The French system is separate from the American one not only because they look to different sets of rules, but also because French courts are not committed to resolving their disputes with American courts regarding American law.

The commitment to resolve disputes is usually discharged by the commitment on behalf of legal participants to accept a strongly hierarchical authority structure. An authority structure is strongly (weakly) hierarchical if, for virtually any dispute (relatively few disputes) that might arise between two officials regarding the rules that they are under a duty to apply, either one official will have authority over another to resolve the dispute, a third agent will have the authority to resolve the dispute, or some agent will have the authority to confer authority on another to resolve the dispute. Hierarchy is the glue that holds officials together and makes it the case that they are officials of the same system.

Although they are sufficient, strongly hierarchical authority structures are not necessary for the unification of a legal system. Legal systems with weakly hierarchical structures are unified just in case legal officials are committed to negotiating with one another until consensus is reached. In

46. *See* H.L.A. Hart, THE CONCEPT OF LAW, 92.

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

the limit case, hierarchy is dispensable when such commitments to negotiate are pervasive. To be sure, systems with weak or no hierarchy would probably be very inefficient and unstable. Yet I don't see why we should not countenance them as conceptual possibilities.

In order to do their part in creating or maintaining a unified system of rules, therefore, most legal participants must discharge two separate subtasks. First, they must look to the same tests of validity as other officials do and must apply the rules that pass those tests. Second, they must commit themselves to resolving their disputes regarding which rules pass such tests and how to apply such rules.

C. Institutional Design and Substantive Goals

In what follows, I will assume that most legal participants in every legal system are committed to doing their part in the joint activity of law. That is, they are committed to doing their part in the creation or maintenance of a unified system of rules. I hope that this assumption is sufficiently uncontroversial. The contribution to the creation or maintenance of a unified system of rules can be common to all legal participants in all legal systems because it is a purely formal goal. It makes no mention of political or moral objectives and hence can be shared by individuals committed to very different ideologies.

By claiming that in every legal system officials are committed to doing their part in the creation or maintenance of *a* unified system of rules I am not presupposing that they are committed to creating or maintaining *any* such system. In fact, they may be committed to creating only one particular system. "Doing their part in the creation or maintenance of a unified system of rules" is the *weakest* description that must be true of their commitments in order for it to be the case that they are committed to the joint legal activity. Stronger descriptions of their commitments may also be apt, as will typically be the case. But stronger descriptions necessarily entail the weaker ones. Members who intend to contribute to the creation and maintenance of a particular unified system of rules *a fortiori* intend to contribute to the creation or maintenance of a unified system of rules.

In addition to the formal goal of creating or maintaining a unified system of rules, legal participants often engage in the practice with the hopes of achieving certain substantive moral or political goals. This is especially true of those who design legal systems. Unless an individual cares only for the stability that a system of rules tends to engender, legal participants will usually prefer certain political or moral objectives to others and will structure their institutions so as to further their preferred goals.[47] The framers

47. We should not discount the possibility that the designers of a legal system only pretend to care about certain political goals and have structured legal institutions so as to enrich themselves. Even in such cases, however, they claim that the design of the legal system actually furthers the putative political goals rather than simply furthering their own self-interest. I ignore this complication in what follows.

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

of the United States Constitution were not, for example, indifferent between democratic, communist, socialist, theocratic, and monarchic forms of government. They were committed to a form of democratic self-rule, and this commitment explains why they designed the authority structure of the United States in the way that they did.

The design of a legal system will, therefore, be a function of the substantive goals that designers intend to advance. The relation between authority structure and substantive goals may be either intrinsic or instrumental. It is intrinsic when the goals themselves demand that certain individuals be empowered to act in certain ways. The promotion of democratic self-rule, for example, necessarily requires that the citizenry have some say in determining the rules of the system. A system designed to further such a goal will, therefore, make provisions such as elections, representation, voting, and protection of free speech so that individuals are empowered to determine the terms and direction of social interaction.

The relation between structure and goal may also be highly instrumental. Federalism, for example, is often justified by the effect that partial decentralization has on the protection of individual liberty. In theory at least, depriving the centralized authority of certain power diminishes its ability to oppress the citizenry. Moreover, by removing barriers to internal immigration, the units that make up the federal system are forced to compete for residents, thereby giving them an incentive not to enact abusive regulations. Finally, by dividing authority between sets of rivalrous political actors, the participants in each set serve as checks to the encroachment of the others.

Because participants generally design their system so as to achieve certain moral or political goals, they are usually sensitive to the nature of their cooperative relationship. As we saw in the voyage of the strangers, a highly competitive environment may significantly impede the ability of individuals to achieve their goals. The strangers were so conflicted and competitive that they could not agree on how to proceed with their trip and had to appoint a captain with broad authority to decide such questions. Competition between legal participants might, therefore, necessitate the establishment of extensive adjudicatory institutions so as to resolve the conflicts that will inevitably arise. The existence of competition, though, need not be a hindrance. As previously indicated above, federalism attempts to harness the rivalry between political actors and use it for the protection of individual liberty.

Likewise, designers of legal systems are also attuned to the degree of trust that they believe it is appropriate to accord to various participants. A high degree of distrust, for example, would result in strict constraints on the use of power in the system. Constitutional democracies show evidence of this distrust, for they place limits on the ability of the people's representatives to enact legislation. To be sure, the fact that these representatives have power at all also indicates that the designers have reposed in them a certain confidence that they will act in their constituents' interests.

When legal systems are designed to achieve certain moral or political

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

goals, it is usually possible to recover the goals of a system by a close examination of its institutional structure. For example, a system that made provisions for voting, representation, elections, and some protection for public deliberation would be one in which democratic self-rule was prized. By contrast, an institutional structure that empowered clerics to decide matters of principle and policy and minimized the degree to which secular forces can affect the direction of the law would be a system in which religious values were being promoted.

Reverse engineering, however, is not the only method available for ascertaining the goals of a system. Usually, legal systems are proud of their objectives and proudly proclaim them to the world. Consider for example the preamble to the United States Constitution.

> We the People of the United States, in order to form a more perfect union, establish justice, insure domestic tranquility, provide for the common defense, promote the general welfare, and secure the blessing of liberty to ourselves and our posterity, do ordain and establish the Constitution of the United States of America.[48]

The preamble sets out, either explicitly or implicitly, the following goals of the federal system: (1) democratic self-rule; (2) political and economic integration; (3) personal security against domestic and foreign threats; (4) promotion of the general welfare; (5) protection of the individual from governmental tyranny. These goals are served by the institutional arrangements set out in the Constitution. For example, the goal of democratic self-rule is furthered by the creation of a dual system of political representation, one based on population (for the House of Representatives), the other based on statehood (for the Senate), the mandate of periodic elections for congressional representatives and for the president, the institution of majority voting to enact legislation (in the absence of a presidential veto), supermajority voting to ratify and amend the Constitution, the protection of freedom of speech and the press, and so on. Political and economic integration is furthered by the creation of a national Congress comprised of representatives from every state, the grant to Congress of the power to regulate foreign and interstate commerce, the denial of power to the states to close their borders to other states or impose on them import duties, the grant of power to the Senate and the president to enter into treaties with other countries and the denial such power to the states, and so on.

The diversity of goals that legal systems are designed to achieve is actually quite striking. Consider the following excerpts from the preambles of some prominent constitutions:

> *Soviet Union*: The supreme goal of the Soviet state is the building of a classless communist society in which there will be public, communist self-government.

48. U.S. CONST. preamble.

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

> The main aims of the people's socialist state are: to lay the material and technical foundation of communism, to perfect socialist social relations and transform them into communist relations, to mould the citizen of communist society, to raise the people's living and cultural standards, to safeguard the country's security, and to further the consolidation of peace and development of international cooperation.[49]

> *Iran*: The Constitution of the Islamic Republic of Iran advances the cultural, social, political, and economic institutions of Iranian society based on Islamic principles and norms, which represent an honest aspiration of the Islamic Ummah.[50]

> *Sweden*: Article I — (1) All public power in Sweden proceeds from the people. (2) Swedish democracy is founded on freedom of opinion and on universal and equal suffrage. It shall be realized through a representative and parliamentary polity and through local self-government. (3) Public power shall be exercised under the law.

> Article II — (1) Public power shall be exercised with respect for the equal worth of all and for the freedom and dignity of the individual. (2) The personal, economic and cultural welfare of the individual shall be fundamental aims of public activity. In particular, it shall be incumbent upon the public administration to secure the right to work, housing and education, and to promote social care and social security and a good living environment. (3) The public administration shall promote the ideals of democracy as guidelines in all sectors of society. The public administration shall guarantee equal rights to men and women and protect the private and family lives of the individual. (4) Opportunities should be promoted for ethnic, linguistic and religious minorities to preserve and develop a cultural and social life of their own.[51]

These preambles set out political objectives not only of differing *content*, but also of differing *depth*. Compare the goals of Iranian versus American law. The goals of the Iranian system are politically quite deep, in that they can be accepted from only one comprehensive worldview, namely, an Islamic one. By contrast, the goals of American law can be accepted from many different comprehensive conceptions of the Good. One can be a Kantian, utilitarian, Aristotelian, libertarian, or Rawlsian and still believe in the moral necessity of promoting democratic self-rule and protecting individual rights.

Legal officials can, therefore, be committed to certain substantive political or moral goals without agreeing on comprehensive worldviews. This, as we have seen, is typical of JIA: Participants may converge on a set of common goals despite differing on beliefs and values.

49. U.S.S.R. CONST. preamble.
50. IRAN CONST. preamble.
51. SWED. CONST. art. I, art. II.

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

Let me add two caveats to the preceding discussion. In speaking of the "design" of a legal system, I do not mean to imply that every legal system has been totally designed in advance during some constitutional convention. Although this is true of most modern systems, legal systems can, and do, evolve without the participants following some master plan. As long as the various alterations in institutional structure are either brought about and/or maintained in order to further certain substantive goals, it will be proper to consider the system designed for these goals.

Moreover, I don't want to rule out the possibility that the design of a system can change over time. Institutional arrangements are rarely static and often shift based on the changing goals of those in power. These changes are of two kinds: authorized and unauthorized. In an authorized change, the system provides a method for its own revision. Most constitutions, for example, contain amendment provisions allowing participants to change the constitutional structure in order to achieve different objectives. By contrast, an unauthorized change is one that occurs outside the authorized channels of revision. Legal history is replete with cases where legal officials either abruptly or gradually change the fundamental structure of the law in a manner not recognized by law.

D. Substantive Goals in Stable Systems

As I have argued, in every legal system most legal officials are committed to the formal goal of contributing to the creation or maintenance of a unified system of rules. In addition, I contended that designers of every legal system are usually committed to furthering certain substantive moral or political goals and that they design their systems so as to promote these goals in light of the cooperative relationship that exists between the parties.

Please note that I did not claim that participants in existing legal systems must be committed to these intended substantive goals. It is possible that in certain legal systems many legal officials are so apathetic or alienated that they care not a wit for the ideals upon which their system was founded. It is necessary, though, that they be committed to seeing that the rules of the system are carried out and to resolving their disputes regarding these rules; if they do not have this as their goal, then a legal system will cease to exist.

In stable legal systems, though, most legal officials will intend that the group bring about the goals for which their system was designed. This is not to say that they will be happy with every aspect of the system within which they participate. An official may deeply regret that the legal system has certain features and may try hard to reorient the law to his moral liking. But an unqualified endorsement of a particular institutional arrangement is not required in order to ground a commitment to the joint activity. In order for such a commitment to be rational, the official must only prefer to participate, given everyone else's participation, over not participating. Thus, even

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

though certain officials might wish to pursue other goals in different circumstances, they commit themselves to pursuing the goals of the system simply because, in the present circumstances, others are following them as well.

While most officials in stable legal systems will be committed to pursuing the goals that their systems are designed to achieve, they rarely pursue them directly. If a legal system is geared toward, say, ensuring the equality of resources, officials need not always allocate goods or services on a strictly equal basis. Such behavior is likely to be extremely inefficient and/or counterproductive. In a legal system, the goals of the system are generally achieved indirectly through the rules that determine the powers and duties of particular officials. Officials pursue the basic goals of the system by following the rules that are designed to achieve such goals.

IX. LEGAL PRACTICE AS JIAA AND SCAA

Having identified the class of legal participants and the nature of the joint activity constitutive of legal practice, I would now like to argue that legal practice is indeed a JIAA. To show this, I will assume that in every legal system, legal participants are indeed committed to doing their part in creating or maintaining a unified system of rules. I will then attempt to demonstrate that legal participants are thereby committed to acting in accordance with and because of meshing subplans, their intentions are publicly accessible, and they are mutual responsive in planning and action.

After arguing that law is necessarily a JIAA, I will then go on to show that legal practice is not necessarily an SCAA. Although it is often true that legal participants are uncoerced and committed to mutual compensatory behavior, it is not necessarily true that they have such intentions. I will end the section by arguing that legal participants necessarily hold themselves out as though they were engaged in SCAA even if this amounts to mere pretense in some cases.

A. From Vertical to Horizontal Interlock

It might be thought that the commitment to act in accordance with subplans that mesh would trivially follow from the commitment of officials to look to the same tests of validity and to apply those rules that pass such tests. Because this latter commitment puts pressure on officials to heed a common set of rules, their subplans will mesh simply in virtue of the fact that they will intend to apply rules that belong to the same set.

This inference, however, is fallacious. First, it does not follow from the fact that participants are generally committed to applying the same set of rules that the rules that they eventually end up intending to apply will always pass the same tests. For it is possible that some participant will make a

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

mistake and intend to apply a rule that is not valid according to the tests that all others take for granted. The inference assumes that every participant will perfectly execute his or her general plans, and this assumption is not necessarily valid. Second, even if every participant ends up intending to apply the same rules, this would show only that their subplans meshed *in fact*. It would not indicate the existence of a *commitment* to such a mesh.

Although the commitment to mesh does not follow from the shared commitment to apply the same set of rules, it can nevertheless be derived from the second component of the commitment to the joint legal activity. As we saw in the previous section, legal officials are committed not only to appealing to a common set of rules but also to resolving their disagreements about the membership of such set or the application of any of its members. The commitment to resolving their disagreements puts the pressure on officials to mesh their subplans with respect to the creation or maintenance of a legal system.

Here's an example from the U.S. legal system. Suppose Trial Judge A applied a rule that Trial Judge B thinks is unconstitutional. Suppose further that the parties who lost in front of Judge A head to Judge B seeking relief. In order to show that B is committed to acting in accordance with meshing subplans, we must ask whether he would be prepared to do anything in order to seek a mesh with A.

If B is committed to the joint legal activity, the answer is clearly "Yes." A's decision as applied to the parties before her court would be *res judicata*, which is to say that the party who lost in A's courtroom would be precluded from pressing his or her suit in any other courtroom in A's jurisdiction.[52] Because B accepts the rules of the jurisdiction, which includes the rule of *res judicata*, B will be committed to not adopting any subplan to rule in the losing party's favor and to adopting subplans to dismiss the losing party's case. B's subplans will, therefore, seek a mesh with A's insofar as he will not adopt any subplan that conflicts with the subplan pursued by A and will revise any that do.

B might still apply a different rule to different parties who came before his court. In this sense, his subplan will not mesh with A's subplan because they will be applying rules that pass different tests. But B would be prepared to mesh his subplans with A if an appellate court ruled that A's rule was correct. Once overruled by the appellate court, B would be committed, under the rule of *stare decisis*, to revise his subplans so that he will be committed to applying the rule A applied. Or the legislature in B's jurisdiction might amend its constitution, making it clear that it believes that A's rule is constitutional. This would again force B into revising his subplans so that they mesh with A.

Res judicata, *stare decisis*, and amending provisions of constitutions are all

52. Under the Full Faith and Credit Clause of the United States Constitution, the parties would be barred under most circumstances from pursuing their claim in any other U.S. jurisdiction as well.

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

examples of "mesh-creating" mechanisms. When disputes between officials break out with respect to the proper way to maintain the system, these mechanisms create a mesh between the subplans of the participants by demanding that both sides accept the authoritatively imposed solution. A and B are committed to acting according to meshing subplans because they accept the authority of these mesh-creating mechanisms. Should some authority decide in a way contrary to their present subplans, they will be committed to revise those subplans so that they mesh with their fellow officials.

Mesh-creating mechanisms are not unique to the legal systems in the United States. All actual legal systems possess them. This can be seen from the fact that existing authority structures are all hierarchical in nature: each step of the command structure has the power to resolve disputes between officials of the lower rung. In particular, these hierarchical structures tend to exhibit the familiar pyramidal form, where larger sets of participants are subject to the authority of smaller sets. And it is obvious why legal systems have this structure—because conflict between officials with respect to the rules of the system is potentially destabilizing, there is a need to force a mesh when one does not yet exist. By accepting the authority of this structure, officials are thus leaving open the possibility that they will change their subplans so that they mesh with their fellow officials whenever their superiors demand it.

It turns out, then, that the horizontal interlock between the intentions of legal participants is guaranteed by the vertical interlocking that exists between the intentions of superior and subordinate. By accepting the authority of the law to settle disputes, legal participants are committed to meshing their subplans with one another just in case the law determines a resolution to such disputes. To be sure, legal participants need not rely solely on authorities in order to establish a mesh. They may also bargain, negotiate, and deliberate with one another in order to achieve interpersonal consistency of plans. The point is that they need not establish the horizontal interlock directly if they are committed to the authority structure of their legal system and that structure is sufficiently hierarchical in order to create mesh out of conflict. Such vertical interlocking leads to the requisite horizontal interlock of a JIAA.

B. Intentional Agency, Public Accessibility, and Mutual Responsiveness

Having established that legal officials intend to act in accordance with meshing subplans, we must now show that they are also committed to act because of each other's meshing subplans. This is apparent from the fact that the law guides conduct by attempting to provide reasons for action. When legislators want judges to act in certain ways, they do not walk over to

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

the courthouse and physically force judges to utter certain words and sign certain pieces of paper. Officials could not possibly hope to achieve their ends through actual puppeteering. Rather, legal authorities attempt to guide conduct by affecting the intentional agency of their subordinates. This is accomplished by issuing rules that impose duties and confer powers and by threatening sanctions if the rules are not obeyed. Subordinates, in turn, are committed to treating the rules as reasons for action and to acting on such reasons.

It can also be shown that the public-accessibility condition is satisfied in every legal system. Legal officials do not hide the fact that they are acting in the name of the law. Given that most of their activities result in some denial of a good to some individual or group, they must cite the rules that they are applying in order to justify their actions. Such attempts at justification manifest their intention to maintain the system of rules by applying legally valid rules. Moreover, the justification of such denials might require officials to apply rules that they believe are invalid but that have been forced upon them through authoritative intervention or through negotiation with fellow officials. By referencing these mesh-creating mechanisms, they are acknowledging that they are committed to acting in accordance with meshing subplans.

Finally, it is clear that most authorities and their subordinates are responsive to one another, for a legal system exists only if its officials apply a unified set of rules. This implies, first, that subordinates must be responsive to the demands of the rules created by their superiors and must act in accordance with them. Second, it implies that their superiors must be responsive to conflict generated by the actions of their subordinates with regard to the proper membership or application of the rules and must act to resolve such disputes so that the system remains sufficiently unified.

C. Coercion and Mutual Compensation

Having established that law is a JIAA,[53] we might ask whether the law is also an SCAA. If legal participants are committed to creating or maintaining a unified system of rules and they are mutual responsive to one another, is it also true that they are uncoerced and committed to mutual compensation?

It is tempting to answer this question in the affirmative. Consider the noncoercion condition. In most legal systems, legal officials are not threat-

53. I have not argued that whenever individuals intend to create a unified system of rules and their intentions satisfy the conditions for a JIAA, they will have succeeded in creating a legal system. While all legal practice is an instance of JIAA, one needs more than a JIAA to engage in legal practice. As Hart argued in THE CONCEPT OF LAW, a legal system exists only when, in addition to commitment by officials to the law, the bulk of ordinary citizens obey the rules at least most of the time. See Hart, *supra* note 47, at 109–14. We can say, though, that legal practice exists just in case it is a JIAA and legal officials and ordinary citizens are mutual responsive in intention and action.

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

ened with punishment for nonfeasance (as opposed to malfeasance such as accepting bribes or manufacturing evidence). A judge who refuses to apply the law is generally impeached, not imprisoned.[54] Even in those systems where officials are subject to coercive measures for dereliction of duty, the noncoercion requirement may still be satisfied. Recall that this condition allows participants to coerce each other just in case the participants see the coercion as a method for enforcing uncoerced commitments. As long as the officials were not coerced into accepting their positions, the threat of coercion will not violate the noncoercion condition.

Next, consider the mutual-compensation condition. Legal systems always provide for certain compensatory resources, as evidenced by their pyramidal structure. The most obvious example of such a resource is an appellate court. Appellate courts exist in order to correct the mistakes made by lower courts. Along the same lines, the purpose of judicial review is to ensure that legislative and executive action conforms to law. The ability of a legislature to repeal past legislation may also be seen as a compensatory resource. Subsequent legislatures are empowered to undo the errors made by earlier legislatures and are committed to using that power when the goals of the systems are not otherwise being met.

While the noncoercion and mutual-compensation conditions will be met in many legal systems, there is no guarantee that they will be met in all legal systems. It is possible, for example, that in some legal system a small cadre of well-armed individuals decide to coerce a much larger group into accepting official positions within their regime. The noncoercion condition would fail in this case, for the coercion would be used to enforce coerced commitments. Admittedly, this case is a bit far-fetched, but it is not inconceivable.

It is similarly possible to imagine a failure of the mutual-compensation condition. Recall Bratman's unhelpful singers. Each singer believes that the other would always perform competently and hence is committed to acting in accordance with, and because of, meshing subplans. However, each wishes for the other to fail and, hence, is not prepared to provide assistance should the other need it. As this case demonstrates, the commitment to the joint activity can exist independently of the commitment to mutual compensation in those situations where participants are sufficiently confident of the abilities of their fellow participants.

We can similarly imagine a legal system of unhelpful subordinates. Each official of the same rank believes of each other that they will competently carry out their duties. Although each official prefers to perform his or her duties competently, given that others perform their duties competently, each prefers that his or her fellow officials fail at their task. Consequently, none is prepared to assist in any way should his or her fellow officials need

54. Because impeachment is based on a finding of unfitness for office, it is best not understood to be a punishment for wrongs committed.

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

the help. In such a system, the condition of mutual compensation would fail because officials of the same rank would not be committed to any compensatory behavior. The existence of compensatory mechanisms cited above only evinces a commitment of *unilateral*, not mutual, compensation. Superiors may be committed to correcting the errors of their subordinates, but the subordinates need not be committed to compensating for each other's failures or for the failures of their superiors.

The case of the unhelpful subordinates is not as far-fetched as it may first appear. Officials of the former Soviet system might have resembled the unhelpful subordinates. Many Soviet officials believed that the communist system was evil and would have preferred that a different system were in place. Yet the communist system was a fact of life, and participation within it was the only mode of social advancement available. Hence they joined the legal system and hoped to benefit from it by performing their tasks competently. However, they also wished that others would not act competently and that the system would eventually collapse. It is possible that these officials were committed to the joint activity but not to mutual compensation. If an official were at risk of failing in his assigned task, his fellow compatriots might have stood by and watched.

D. Law's Pretense

While I do not believe that legal practice is necessarily an SCAA, I do think that it necessarily holds itself out as an SCAA. By this I mean that legal participants act as though they have the requisite cooperative attitudes, even if they don't really have them. This pretense is manifested by the use of normative language that officials characteristically employ when they carry out their official duties.

Consider the noncoercion condition. By claiming that they are "obligated" to act in certain ways, officials are implying that they are not merely "obliged" to do so. While their motivation in fact might be based solely on coercion, they act as though they are motivated by the legitimacy of the demands that the law places upon them.

Similarly, the use of normative language bespeaks a commitment to mutual compensation. To see this, we should note that the unhelpful singers, resentful captain, and unhelpful subordinates would never sincerely use normative language to describe the behavior of other participants. The use of terms such as "obligation," "permission," and "right" would imply an endorsement of their fellow participants' success, and terms such as "wrong," "bad," and "guilty" would amount to a criticism of their failures. These unhelpful participants prefer each other's failure and hence would not endorse their success and would not look badly upon their failure. Indeed, the opposite is true: They would endorse their failure and criticize their success.

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

By using normative language, therefore, legal participants are holding themselves out as though they care about the goals of the system of which they are a part. They are, at the very least, pretending that they approve of the success of their fellow participants and disapprove of their failures.

As we saw, the commitment to the joint activity rationally leads to the commitment to mutual compensation except in those situations where participants prefer each other's failure. Legal participants always hold themselves out as preferring each other's success. Hence they always hold themselves out as being committed to mutual compensation.

X. CONVENTIONS AND COOPERATION

In the last section, I argued that legal practice is an instance of a JIAA. I would finally like to put this argument to use by showing how it vindicates the legal-positivist position on the determination of legal authority.

Recall that the legal positivists attempted to show that conventions determine legal authority by arguing that the choice of authority structure is a game that conventions are especially able to solve. But in implementing this strategy, the positivists made two crucial mistakes. First, they argued that the choice of an authority structure is always a recurring coordination problem. Second, they assumed that conventions exist only to solve recurring coordination problems. However, once it is acknowledged that the choice of an authority structure is not always a recurring coordination problem, it becomes impossible for the positivists to hold that conventions always select the authority structure of a legal system.

I would like to propose the following fix: Instead of treating the choice of an authority structure as a recurring coordination problem and a legal convention as a Lewisian behavioral regularity, the positivists should approach the choice as a recurring *cooperation* problem and the convention as a Bratmanian *plan*. The defense of the legal-positivist position would then proceed, *mutatis mutandis*, much as the original one did.

It would begin by showing that the choice of an authority structure is a recurring cooperation problem. After noting that legal participants solve their recurring cooperation problems, the positivists would then argue that the only way to explain how participants are able to play such games successfully is by using plans that map out the authority structure of their system. The defense would continue by showing that such plans are conventions and that they come into existence by being practiced, not in virtue of the principles of morality. It is to this modified defense we now turn.

A. Choice of Authority Structure as Cooperation Problem

In a cooperation problem, players have a common interest in working together, and they increase the likelihood that their joint activity will be

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

successful if their behavior is organized in some manner. To show that the choice of an authority structure is a cooperation problem, therefore, the positivist must show that legal participants have a common interest in deferring to the same authority structure and that they increase their likelihood of success by having their behavior organized in a certain way.

As we saw, in every legal system most participants will be committed to the joint legal activity. This entails that participants will be committed to applying the rules that pass the same criteria of legal validity used by all other participants. Because these tests of validity admit a rule into the system just in case some agent with authority has created them, the commitment to apply the same set of rules is nothing more than the commitment to defer to the same set of authorities. Hence, in every legal system, most participants will have a common interest in deferring to the same authority structure.

As we also saw, participants need not see the choice of an authority structure as completely arbitrary. In many legal systems, most participants will have a common interest in promoting a given set of moral or political goals and will see authority as the mechanism by which they can advance such goals. Clearly, by organizing their behavior around a common authority structure, they will be better able to accomplish these goals. Indeed, it is hard to imagine participants even coming close to achieving their collective ends if they do not organize their behavior so that the appropriate people listen to the appropriate people in the appropriate circumstances.

To be sure, the need to organize the behavior of officials around a common authority structure might arise because of the existence of coordination problems. As will often be the case, participants will regard certain aspects of the institutional structure as arbitrary and will be relatively indifferent between the various options. The possibility of cooperation will depend on the ability to break such ties and to focus official attention on the same course of action.

However, as I argued in Section V, a cooperation problem can arise in the absence of a coordination problem. Most importantly, participants might require guidance on how to achieve the goals of the legal system. Many participants may have little or no idea which means are conducive to their ends and will need instructions on which path to take. Or they may be able to figure out what to do, but only if given more time than is available for the task at hand.

Cooperation problems may also arise because of the need for assurance. Certain participants may know what to do but may be unsure whether others are similarly informed. Or certain participants may support a certain course of action, but only if the others do so as well, and they might be uncertain whether the others will commit.

Legal officials, therefore, will need to organize their behavior and the behavior of others, even if they all believe there is a "best" way to construct a legal system and they agree on the political standards by which to assess

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

it. For some participants may not know how to construct such a system.[55] Or they may know but need to be reminded from time to time of how to maintain it. Or they may know but doubt whether their fellow participants know or are willing to go along. It follows, therefore, that in every legal system where most participants are committed to the same substantive goals, the choice of an authority structure will be a recurring cooperation problem.

The need for organization will be even greater for participants who do not share the same substantive goals but are merely committed to the formal goal of contributing to the creation or maintenance of a legal system. These participants will merely be committed to applying the same set of rules issued by the same set of authorities—but which rules and which authorities? With no substantive goals to point the way, the problems of guidance, assurance, and coordination will be quite desperate. Certainly, organization of behavior would greatly increase the likelihood of success. It follows, therefore, that in every legal system—regardless of whether participants share the same substantive goals or not—the choice of an authority structure will be a recurring cooperation problem.

B. The Need for Plans

Obviously, legal officials solve such recurring cooperation problems. They organize themselves into complex authority structures and carry out the tasks that authorities demand. How do they manage to create and maintain these (sometimes incredibly elaborate and complicated) institutions?

A plausible answer is this: Like participants in all JIAs, legal officials manage to cooperate by formulating, adopting, and consulting plans. In particular, legal participants would select their authority structure by first formulating, adopting, and consulting "rules of recognition."[56] These plans would specify those characteristics possession of which by a rule would render it authoritative. Every legal official would be under a duty to follow such plans by according the selected rules authoritative status.

In addition to rules of recognition, legal participants would also formulate, adopt, and consult what Hart termed "rules of change" and "rules of adjudication."[57] A rule of change is a plan that specifies which members of the group have the authority to create or repeal the rules of the system and under which circumstances the exercise of their power is valid. A rule of adjudication is a plan that picks out those members who have the jurisdic-

55. The ancient Hebrews were pretty sure that God would tell them how to build the best legal system—they just need to be told.

56. *Pace* Hart, there is no reason to think that there can only be one rule of recognition in every legal system. Different subgroups can follow different rules of recognition, provided they overlap enough in order to unify the system. *See* Raz, supra note 46, at 200.

57. Hart, *supra* note 47, at 93–4.

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

tion to determine whether the rules have been followed and under which circumstances their decisions are binding.

Not only is it plausible to suppose that shared plans enable legal officials to cooperate in general and select an authority structure in particular, but it is hard to imagine plans not playing such a role. The odds that everyone will listen to the same people, at the same time, and under the same circumstances without previous planning and frequent consultation are effectively zero. The joint activity of law is made possible because legal participants commit themselves to rules of recognition, change, and adjudication and guide their conduct accordingly.[58]

It would seem, therefore, that the legal positivist has an extremely powerful argument in favor of the claim that, in every legal system, legal officials follow general plans when determining which authority structure to heed. Because (1) the choice of an authority structure is a recurring cooperation problem; (2) legal officials solve such problems; and (3) it is extremely unlikely that anything other than general plans can solve recurring cooperation problems in such contexts, it follows that legal officials use general plans to identify the authority structure in every legal system.

C. Plans and Conventions

The defense so far has employed a version of the familiar "inference to the best explanation" strategy. We have argued that plans mapping out the authority structure must exist in every legal system if we are to explain how legal participants are able to solve their recurring cooperation problems. The next step in the defense is to show that such plans are conventions.

In the dictionary-definition sense, most general plans that set out the authority structure of existing legal systems are conventions. *Webster's Third New International Dictionary*, for example, lists for its first definition of "convention": "an agreement between persons or parties." Since most fundamental legal plans are actually adopted pursuant to explicit bargaining and agreement—adopted in the course of a "constitutional convention"—they are conventions in the same way in which the Geneva or Berne Conventions are conventions.

To be sure, some fundamental legal plans are not the objects of explicit agreement. Aspects of the British Constitution come to mind. And it is surely possible to imagine cases where such plans are not even subject to a tacit agreement.[59] But I think it would be a mistake if we took the dictionary definition too seriously and saw agreement as essential for any sense of the

58. The claim, of course, is not that shared plans enable legal participants to achieve the political goal of their system. Plans enable convergence on a given authority structure—whether the authority structure upon which they converge is conducive to the goals of the system is another matter.

59. This would be the case if the fundamental legal conventions were fully Lewisian in nature.

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

term "convention." What makes a convention a convention, I propose, is twofold. First, the convention must be a plan that is shared by the parties to the convention. Second, the acceptability of the plan to each party depends on the general compliance of the other parties.[60]

The benefit of this construal is that it explains why explicit agreements between parties and Lewisian conventions are both considered conventions. The Geneva Convention, as well as the rules-of-the-road convention, are plans that are shared by the parties, and their acceptability depends on the general conformity of others.

If general plans that map out the authority structure of a legal system are conventions, therefore, it must be the case that they are shared by most legal participants and that their acceptability depends on conformity by others. Because I have demonstrated the former in the previous section, I need only show the latter here.

Demonstrating the practical relevance of general conformity proceeds as follows: in every legal system, most participants are committed to the joint legal activity. This entails that most participants are committed to heeding the same authority structure. If most participants do not heed a particular authority structure (call it A), then no participant can satisfy his or her commitment to the joint legal activity by heeding A. Since participants accept the fundamental plans of a legal system in order to fulfill their commitment to the joint legal activity, it follows that participants will not accept plans to heed A when most others do not heed A. Hence general conformity to a particular authority structure must always be part of the reason why legal participants accept the fundamental plans of their legal system.

The fundamental plans of a legal system are thus more strongly conventional than even the rules of the road. General conformity is essential to the acceptance of the rules of the road because it is extremely unlikely that a driver will do better if he drives on a different side of the road from everyone else. However, it is not merely improbable that a legal participant will satisfy his commitment to the joint legal activity if he heeds a different authority structure from everyone else—it is impossible. General conformity to a particular authority structure is a condition of possibility for a legal participant to fulfill his commitment to the joint legal activity. If general conformity to a practice is absent, then there is no possibility that he can contribute to such a practice. Just as one cannot open an open door, one cannot do one's part in a practice if that practice is nonexistent.

The conventionality of legal practice entails that the acceptance by individuals of an authority structure depends in part on the fact that others accept the same structure as well. It does not, however, imply the stronger claim that there exists some other authority structure such that general

60. *See, e.g.*, HART, *supra* note 47, at 243 (a rule is conventional when "general conformity by a group constitutes part of the reason why its individuals accept it.") *See also* Dworkin, *Model of Rules II*, *in* TAKING RIGHTS SERIOUSLY 53 (1975) and Postema, *supra* note 3, at 176.

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

conformity to it would be treated as a reason to accept that structure instead. A plan can thus be conventional without it implying, as Lewis's model does, that the participants see it as an arbitrary solution to their game. This is fortunate, for as we saw, legal participants might not prefer any alternative authority structure. They might prefer their practices and their practices alone. General conformity to another plan would not be seen as an arbitrary solution to a coordination problem but as a moral calamity, one of which they want no part.

XI. TAKING PLANS SERIOUSLY

The defense of legal positivism is nearly complete. Having shown that every legal system contains conventional plans that map out the authority structure of that system, it remains to be shown that such plans come into existence and impose obligations by virtue of being practiced, not because of the principles of morality.

The last step in the argument has a positive and negative part. The positive part aims to show that adoption and practice are sufficient to create the fundamental plans of a legal system. The negative part points out the pernicious theoretical consequences of accepting the natural-law position.

A. The "Internal Point of View" Revisited

The positive part relies on a feature of general plans that, I hope, is uncontroversial: They come into existence in virtue of being adopted and practiced. It is proper to say of me that I have a plan to do 500 sit-ups every night just in case I have adopted such a plan and generally follow it. Smith and Jones have a plan to go to work together every day because they have agreed to such a plan and generally abide by it. The friends have a plan to listen to the captain, given that they have committed themselves to such a plan and ordinarily execute it. The fundamental plans of a legal system would come into existence in much the same way: Legal participants would adopt them as guides to their conduct and guide their conduct accordingly.

Notice that the principles of morality do not enter into this analysis. Not once was there a need to discuss the morality of doing sit-ups, going to work, cooking dinner, sailing to Nova Scotia, or singing duets in order to discuss the existence of such practices. Individual agents are able to give themselves plans and create personal practices simply by engaging in the mental act of adoption and then by following through. Likewise, members of groups are able to give themselves plans and create JIAs by adopting certain plans, committing themselves to act in accordance with and because of meshing subplans, and following through on such commitments. The positivist theory is, therefore, parsimonious in that it does not rely on extra entities to explain

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

the existence of legal systems but rather treats legal activity like any other garden-variety practice. Facts about individuals and collectives suffice.

The position advocated here can be seen as continuous with and an extension of Hart's account in *The Concept of Law*. There Hart argued that the fundamental rules of a legal system come into existence in virtue of the fact that legal officials follow such standards from "the internal point of view."[61] When legal systems are recognized to be joint intentional activities, not only is Hart's account vindicated, but its claims can be reinterpreted and brought into sharper focus. In the framework we have been developing, the Hartian critical attitude is seen to be an intention to do one's part in the joint legal activity. Legal participants take the internal point of view towards their practice when they possess the intention to apply the same rules that pass the same criteria of validity as well as the intention to resolve disputes that might arise with respect to the application of such rules.

This account also suggests that we must distinguish, in the way that Hart did not, between the critical attitudes that most legal participants must *possess* from those that they must merely *profess*. The first internal point of view would consist of the intentions mentioned in the previous paragraph, that is, intentions to contribute to the joint legal activity and to mesh one's subplans with those of one's fellow participants. Most participants must satisfy these commitments if they are to create and maintain the fundamental rules of the system. In addition, because a legal system must hold itself out as an SCAA, participants must act as though they are uncoerced and are committed to mutual compensation. A legal system may exist even if most participants do not view their practice from this second, even more internal, point of view, but they must act as though they do.

Let us now return to the challenge raised at the beginning of the essay regarding the positivists' attempt to ground authority structures on social facts. As we saw, in order for the positivist strategy to be successful, the positivists must show that the social facts in question are capable of generating rights and duties. Austin's approach of reducing sovereignty to habits of obedience failed because habits are not the right sort of things that can impose obligations.

Fortunately, plans are the right sort of things. Plans create and define roles within practices and generate the rights and duties that attach to such roles. Plans have this normative power in virtue of their capacity to organize the behavior of participants. As I have argued, they are capable of guiding, reassuring, and coordinating participants so that each may contribute to the joint outcome in light of the cooperative relationship. I am not claiming, however, that the existence of plans suffices to create or maintain *morally legitimate* authority structures. The powers and obligations generated by plans are *internal* to legal practice. Whether participants are actually morally obligated to act according to their legal obligations depends on the goals of

61. Hart, *supra* note 47, at 55–57.

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

the legal system and the contributory relation between the obligations and the goals. If the goals are morally important and the obligations are reasonably tailored to effectuate such goals, then the authority structure will be morally legitimate and its claim to impose moral obligations will be vindicated.

While general conformity to an authority structure motivated by the appropriate intentions is not sufficient to create legitimacy, it is nevertheless necessary. First, as we saw earlier, it is a conceptual truth that one cannot have a duty, moral or otherwise, to play one's part in a joint activity when others do not do their part. Second, and more interesting, regimes are morally legitimate only when doing one's part is reasonably related to the production of morally significant goods. Insofar as the promotion of any morally significant good would likely require a high degree of organization and cooperation, it is rather unlikely that less-than-general conformity could produce them.

B. The Old Debate Revisited

The negative part of the argument in favor of legal positivism points out the pernicious consequences of natural-law theory. If the principles of morality are required for the existence of legal practice, it would follow that wicked legal systems are not really legal systems. We are thus thrust back into the old debate between legal positivists and natural lawyers, a debate that the legal positivists won long ago. Moreover, I don't see how the natural lawyer can restrict his claim to legal practice and not apply it more generally to other plans and practices that have a significant moral dimension. Did the Unabomber not have a plan to send bombs through the mail? Did Timothy McVeigh and Terry Nichols not have a plan to blow up the Alfred P. Murrah building in Oklahoma City? Did the Nazis not have a plan to exterminate the Jews? These plans were morally horrific, but these individuals and groups had them. To deny their existence is absurd.

The natural-law position also implies that legal participants are radically not in control of their own practices. For it is doubtful that legal participants would want each other to determine the existence or content of plans by looking directly to moral theory. First, it is unlikely that participants will agree on the ends that morality requires them to pursue. This is certainly true in societies such as our own that are strongly pluralistic. Even if the participants might concur on shallower goals—for example, that democratic self-rule is an important moral good that should be promoted—having officials look to comprehensive moral theories would lead to the promotion of very deep political objectives that would fail to mesh. The organizational difficulties that would attend such a strategy would likely be a nightmare.

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

Second, it is important to note that the natural-law approach presupposes a certain extreme conception of the participants' cooperative relationship. It assumes that the participants are highly cooperative, for it gives each official a tremendous amount of discretion in determining the existence or content of the fundamental plans of the system. Competitive individuals would not rationally give each other such license. More interesting, it presupposes that the participants trust each other implicitly, for it reposes in each the competence and good will to arrive at the morally correct solution.

The natural-law position, therefore, thrusts upon the participants a strategy that would likely be unproductive and very much unwanted. Such a strategy is appropriate only in those rare situations where the participants share an extremely cooperative and trustful participant relationship. To claim that legal participants are forced to employ such a method implies that they lack the basic control over the management of their cooperative relationship.

Ironically, natural-law accounts of legal authority fail for much of the same reason that the earlier legal-positivist approach did. They both overlook the fact that legal practice is a cooperative activity. Natural lawyers neglect the possibility that legal participants might work together to achieve outcomes that are morally horrendous. And they ignore the fact that the cooperative relationships in different legal systems might vary greatly and that participants would not want morality to play a central role in the determination of the practice.[62]

62. Natural lawyers often argue that the controversiality of fundamental legal rights and duties indicates that they are not determined by convention, for controversy, they claim, is inconsistent with conventionality. Though space considerations prevent a close examination of this objection, I would like to show how it can be answered by very briefly providing an alternative framework for understanding the nature of legal disagreement that is compatible with positivist strictures.

The existence of disagreement about the content of plans indicates that the content of intentions, like the content of all propositional attitudes, is not necessarily transparent to the adopter. The content of shared intentions in a JIA is ascertained by the adopter deliberating about the goals of the practice in light of the cooperative relationship that exists between the participants. A plan imposes certain duties and powers just in case the ascription of this content best furthers the goals of the system in light of the level of competition and trust that exists between the participants.

Because the goals of a system and its cooperative relationship are encoded within the structure of a system, participants can recover the goals and cooperative relationship through a rational reconstruction of the behavior of the designers of the system. Disagreements about the content of a plan can therefore take at least three forms: Participants can disagree about which goals the system is designed to promote, the nature of the cooperative relationship the designers presupposed, or which plan would best further such goals in light of the cooperative relationship.

Insofar as the goals of a legal system may be morally unacceptable, the method described is not extensionally equivalent with the natural-law approach. Nor need the result of the rational reconstruction attribute to the designers a cooperative and trusting relationship. Someone may have authority because it furthers the wicked goals of the system in light of the competitive and distrustful relationship that exists between the participants.

My intention has not been to prove that this interpretation of legal disagreement is the correct one. I offer it only as a means of deflecting the natural lawyer's objection by showing that it is possible to explain disagreement without resorting to the principles of morality. I explore these issues in greater detail in Interpretation and the Economy of Trust (unpublished manuscript on file with the author).

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press

XII. CONCLUSION

As Leslie Green noted in a recent article, conventionalism is not an essential feature of legal positivism.[63] While every legal positivist believes that the law has social foundations, only conventionalists believe that those foundations are best understood as conventions. Green has urged positivists to reject the conventionalist interpretation of positivism and seek alternative understandings of the law's social foundations.

Green is quite explicit that his repudiation of conventionalism is directed against those theories that regard the fundamental legal conventions as coordination conventions. He therefore implicitly leaves open the possibility that one can be a conventionalist positivist but with an alternate theory of convention. My aim in this paper has been to develop such an alternate account.

Yet, even though I have argued that the general plans that map out the authority structure of a legal system are best understood as conventions, nothing of much importance rides on the truth of this claim. I agree with Green that the essence of legal positivism is the claim that the law has social foundations. What we call these foundations matters little in the end.

There is one sense, however, in which I fear that labels do matter. Although the account I offered is technically a conventionalist one, I have shied away from referring to it as such in the body of this paper. To describe one's theory as conventionalist gives the impression that the fact of general conformity plays a central explanatory role in the theory, as is certainly the case for those accounts that treat the fundamental legal conventions as Lewisian conventions. If the choice of an authority structure were always a coordination problem, then any explanation of why legal participants choose one structure over another must foreground the fact of general conformity to the chosen structure. Nothing else could explain why participants act as they do.

However, on the account I presented, the "conventionality" of general plans does little explanatory work.[64] As I argued in Section VIII, general conformity to a set of plans is merely a conceptual precondition for adoption. If participants have intentions to participate in a practice, that practice must first exist, and practices exist only when most members of a group follow the plans that structure such practices. General conformity to a set of plans, therefore, plays only a small, supporting role in the explanation of official action.

Because the conventionality of plans is trivially entailed by any account that treats legal participants as having participatory intentions, it should not be given pride of place in any such theory. The term "conventionalist," though, has the tendency to play up conventionality and thus to obscure more important concepts such as cooperation, participation, and organization. For this reason, the label should be used with caution.

63. Leslie Green, *Positivism and Conventionalism,* 12 CAN. J. JURISPRUDENCE 35 (1999).

64. The discussion in the text has been influenced by Bratman's comment to my paper, where he makes a similar point.

https://doi.org/10.1017/S135232520208401Xh Published online by Cambridge University Press